

**CHARACTERIZATION OF FINAL STATE INTERACTION
STRENGTH IN PLASTIC SCINTILLATOR BY
MUON-NEUTRINO CHARGED CURRENT CHARGED
PION PRODUCTION**

by

Brandon M. Eberly

B.S. in Physics, Lebanon Valley College, 2006

M.S. in Physics, University of Pittsburgh, 2007

Submitted to the Graduate Faculty of
the Kenneth P. Dietrich School of Arts and
Sciences in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2014

UNIVERSITY OF PITTSBURGH
KENNETH P. DIETRICH SCHOOL OF ARTS AND
SCIENCES

This dissertation was presented

by

Brandon M. Eberly

It was defended on

March 17, 2014

and approved by

Steven Dytman, Ph.D., Professor

Vittorio Paolone, Ph.D., Associate Professor

James Russ, Ph.D., Professor of Physics at Carnegie Mellon University

Eric Swanson, Ph.D., Associate Professor

Andrew Zentner, Ph.D., Associate Professor

Dissertation Director: Steven Dytman, Ph.D., Professor

Copyright © by Brandon M. Eberly
2014

ABSTRACT

CHARACTERIZATION OF FINAL STATE INTERACTION STRENGTH IN PLASTIC SCINTILLATOR BY MUON-NEUTRINO CHARGED CURRENT CHARGED PION PRODUCTION

Brandon M. Eberly, PhD

University of Pittsburgh, 2014

Precise knowledge of neutrino-nucleus interactions is increasingly important as neutrino oscillation measurements transition into the systematics-limited era. In addition to modifying the initial interaction, the nuclear medium can scatter and absorb the interaction by-products through final state interactions, changing the types and kinematic distributions of particles seen by the detector. Recent neutrino pion production data from MiniBooNE is inconsistent with the final state interaction strength predicted by models and theoretical calculations, and some models fit best to the MiniBooNE data only after removing final state interactions entirely.

This thesis presents a measurement of $\frac{d\sigma}{dT_\pi}$ and $\frac{d\sigma}{d\theta_\pi}$ for muon-neutrino charged current charged pion production in the MINER ν A scintillator tracker. MINER ν A is a neutrino-nucleus scattering experiment installed in the few-GeV NuMI beam line at Fermilab. The analysis is limited to neutrino energies between 1.5-10 GeV. Dependence on invariant hadronic mass W is studied through two versions of the analysis that impose the limits $W < 1.4$ GeV and $W < 1.8$ GeV. The lower limit on W increases compatibility with the MiniBooNE pion data.

The shapes of the differential cross sections, which depend strongly on the nature of final state interactions, are compared to Monte Carlo and theoretical predictions. It is shown that the measurements presented in this thesis favor models that contain final state interactions. Additionally, a variety of neutrino-nucleus interaction models are shown to successfully reproduce the thesis measurements, while simultaneously failing to describe the shape of the MiniBooNE data.

TABLE OF CONTENTS

PREFACE	xiii
1.0 INTRODUCTION	1
1.1 The Standard Model	2
1.1.1 Fundamental Particles	2
1.1.2 Fundamental Interactions	4
1.2 Neutrino Oscillations	5
1.2.1 Neutrino Oscillation Theory	5
1.2.2 Long Baseline Neutrino Experiments	6
1.2.3 Example of a Nuclear Effect: Pion Absorption in Final State Interactions	8
1.3 Modeling Charged Current Neutrino-Nucleus Pion Production	9
1.3.1 Charged Current Neutrino-Nucleon Interactions	10
1.3.2 Nuclear Medium Modifications	13
1.3.3 Final State Interactions	14
1.4 Neutrino Charged Current Charged Pion Production Data	15
2.0 THE NUMI BEAMLINe	19
2.1 The Proton Beam	19
2.2 The Neutrino Beam	20
2.3 NuMI Flux Measurement	24
3.0 THE MINERvA DETECTOR	28
3.1 Detector Design	28
3.1.1 Inner Detector	29
3.1.1.1 Detector Modules	30

3.1.1.2	Scintillator Planes	31
3.1.2	Outer Detector	34
3.1.3	Helium Target and Veto Wall	35
3.1.4	Light Collection	35
3.1.5	MINOS Detector	38
3.1.6	Test Beam Detector	38
3.2	Detector Readout	39
3.2.1	Front End Boards	39
3.2.2	Data Acquisition	42
3.3	Detector Simulation	43
3.3.1	Particle Propagation	43
3.3.2	Data Overlay	44
3.3.3	Optical and Electronics Models	45
3.3.4	MINOS Simulation	46
4.0	CALIBRATION OF THE MINERνA DETECTOR	47
4.1	Energy Calibration	47
4.1.1	Electronic Pedestals	48
4.1.2	FEB Calibration	50
4.1.3	PMT Gains	51
4.1.4	Attenuation	52
4.1.5	Channel-to-Channel Variations	52
4.1.6	Detector Energy Scale	56
4.2	Timing Calibration	56
4.3	Plane Alignment	57
4.4	PMT Cross Talk	59
5.0	EVENT RECONSTRUCTION	62
5.1	Hit Grouping	62
5.1.1	Time Slicing	63
5.1.2	Clustering	63
5.2	Tracking and Vertexing	64

5.2.1	Long Track Pattern Recognition	65
5.2.2	Short Track Pattern Recognition	68
5.2.3	Event Building	70
5.3	Muon Reconstruction	74
5.3.1	MINOS Track-Matching	74
5.3.2	Muon Charge and Momentum	74
5.4	Hadron Particle Identification	75
5.4.1	dE/dx Particle Identification	76
5.4.2	Michel Electron Particle Identification	79
5.5	Neutrino Energy Reconstruction	79
6.0	NEUTRINO SIMULATION	83
6.1	Neutrino Cross Sections	83
6.1.1	Quasi-elastic Scattering	84
6.1.2	Resonance Production	84
6.1.3	Deep Inelastic Scattering	85
6.1.4	Coherent Pion Production	86
6.2	Nuclear Effects	87
7.0	MUON-NEUTRINO CHARGED CURRENT CHARGED PION PRODUCTION IN PLASTIC SCINTILLATOR	88
7.1	Event Selection	89
7.1.1	Vertex Selection	89
7.1.2	Muon Selection	90
7.1.3	Kinematic Constraints	92
7.1.4	Charged Pion Selection	93
7.2	Cross Section Calculation	98
7.3	Background Subtraction	99
7.4	Detector Resolution Correction	103
7.5	Efficiency Correction	108
7.6	Cross Section Normalization	110
7.7	Systematic Errors	110

7.7.1	Beam Flux	113
7.7.2	GENIE	115
7.7.3	Geant4	115
7.7.4	Michel Selection	116
7.7.5	Vertex Position and Particle Angles	117
7.7.6	Muon Tracking Efficiency	117
7.7.7	Hadron Tracking and Identification Efficiency	117
7.7.8	MINOS Muon Momentum	118
7.7.9	Detector Mass	118
7.7.10	Energy Loss Model	120
7.7.11	Energy Response	120
7.7.11.1	Charged Pions, Kaons, and Protons	120
7.7.11.2	Neutrons	120
7.7.11.3	Electromagnetic	122
7.7.11.4	Cross Talk	122
7.7.11.5	Other Particles	123
7.7.12	Birks Constant	123
7.7.13	Calorimetry Versions	124
7.7.14	Unfolding	124
8.0	RESULTS AND CONCLUSIONS	125
8.1	Results	125
8.2	Model Comparisons	133
8.3	Conclusions	133
APPENDIX A.	PMT GAIN CALIBRATION	136
A.1	Gain Measurement Procedure	136
A.2	Gain Measurement and Electronics Read-Out	138
A.3	Statistical Uncertainty	140
APPENDIX B.	N-PION ANALYSIS FIGURES	142
APPENDIX C.	MINERνA COLLABORATION	152
BIBLIOGRAPHY		154

LIST OF TABLES

1.1	Properties of bosons.	3
1.2	Properties of fermions.	3
1.3	T2K ν_e appearance systematic uncertainties	10
2.1	Pion and kaon decays.	23
2.2	Data used to constrain the NuMI flux measurement	26
3.1	Summary of MINER ν A material composition	34
3.2	MINER ν A scintillator plane composition	34
3.3	Geant4 physics modules	44
5.1	Pion tracking thresholds	69
5.2	Passive material corrections	81
6.1	Mesons produced by KNO scaling	86
7.1	Background Summary	101
7.2	Muon tracking efficiency corrections and errors.	111
7.3	Uncertainties on the NuMI beamline model	114
7.4	MINER ν A detector mass uncertainties	119
7.5	Detector energy response uncertainties	121
8.1	χ^2 /D.O.F. between the differential cross section measurements and GENIE	127
8.2	Systematic error groupings	127
8.3	χ^2 /D.O.F. between the differential cross section shape measurements and GENIE	131

LIST OF FIGURES

1.1	A $CC\nu$ interaction	8
1.2	Example of pion absorption	9
1.3	Comparison of ANL and BNL pion production data	16
1.4	MiniBooNE data compared to models	17
1.5	MiniBooNE data compared to GENIE and GiBUU with and without final state interactions	18
2.1	A diagram of the NuMI beamline.	21
2.2	A diagram of the NuMI target.	21
2.3	A diagram of the NuMI focusing horns.	22
2.4	A diagram of the NuMI absorber region.	23
2.5	A comparison of NA49 and FTFP_BERT cross sections	25
2.6	Ratio of weighted to unweighted flux	27
2.7	NuMI flux measurement	27
3.1	Diagram of the MINER ν A detector	29
3.2	MINER ν A scintillator plane orientations	31
3.3	MINER ν A nuclear targets region	32
3.4	Edgewise view of a scintillator plane	33
3.5	MINER ν A photomultiplier tube fiber mapping	37
3.6	Momentum vs. time-of-flight of test beam particles	40
4.1	Schematic diagram of a MINER ν A optical readout channel	49
4.2	Pedestal data for a single channel	50
4.3	MINER ν A gain measurement on April 1, 2010	53

4.4	Fitted rock muon peak energy per length for each plane	55
4.5	Calibration of detector energy scale	56
4.6	Time slewing vs. number of photoelectrons	58
4.7	MINER ν A time resolution	58
4.8	Alignment of two scintillator planes	60
4.9	Data-simulation comparison of the nearest neighbor cross talk fraction	61
5.1	Long track pattern recognition procedure	66
5.2	Tracking position resolution	69
5.3	Tracking example	71
5.4	Muon tracking efficiency	72
5.5	Pion tracking efficiency	72
5.6	Muon and pion angle resolution	73
5.7	Muon momentum resolution	76
5.8	Simulated dE/dx profiles compared to Bethe-Bloch	77
5.9	Pion kinetic energy resolution	78
5.10	Reconstructed Michel electron energy	80
5.11	Calorimetric energy resolution	82
7.1	Reconstructed q/p (one-pion)	91
7.2	Reconstructed E_ν (one-pion)	92
7.3	Reconstructed W_{exp} (one-pion)	94
7.4	Number of reconstructed hadron tracks (one-pion)	95
7.5	Pion score (one-pion)	96
7.6	Number of Michel electron views (one-pion)	97
7.7	Reconstructed pion distributions (one-pion)	98
7.8	Simulated background predictions (one-pion)	99
7.9	W_{exp} distributions before background fit.	102
7.10	W_{exp} distributions after the background fit	102
7.11	Background fit results	104
7.12	Adjusted background predictions (one-pion)	105
7.13	Background-subtracted distributions (one-pion)	105

7.14	Migration matrices (one-pion)	107
7.15	Unfolded T_π and $\theta_{\pi\nu}$ distributions (one-pion)	107
7.16	Reconstruction efficiency (one-pion)	109
7.17	Flux systematic uncertainties	114
7.18	MINOS Momentum Reconstruction Uncertainty	119
7.19	Pion energy response uncertainty	121
7.20	Electromagnetic response uncertainty	122
7.21	E_{recoil} systematic offset	124
8.1	Measurement of ν_μ charged current π^\pm production differential cross sections with respect to T_π and $\theta_{\pi\nu}$	126
8.2	Error summaries for the measurements of $\frac{d\sigma}{dT_\pi}$ and $\frac{d\sigma}{d\theta_{\pi\nu}}$	128
8.3	Shape measurement of ν_μ charged current π^\pm production differential cross sections with respect to T_π and $\theta_{\pi\nu}$	130
8.4	Shape error summaries for the measurements of $\frac{d\sigma}{dT_\pi}$ and $\frac{d\sigma}{d\theta_{\pi\nu}}$	132
8.5	Comparisons of the one-pion measurements to models	134
A.1	Statistical significance of PMT gain fluctuations	141
B.1	Reconstructed q/p (N-pion)	142
B.2	Reconstructed E_ν (N-pion)	143
B.3	Reconstructed W_{exp} (N-pion)	144
B.4	Number of reconstructed hadron tracks (N-pion)	145
B.5	Pion score (N-pion)	146
B.6	Number of Michel electron views (N-pion)	147
B.7	Reconstructed pion distributions (N-pion)	148
B.8	Simulated background predictions (N-pion)	148
B.9	Adjusted background predictions (N-pion)	149
B.10	Background-subtracted distributions (N-pion)	149
B.11	Migration matrices (N-pion)	150
B.12	Unfolded T_π and $\theta_{\pi\nu}$ distributions (N-pion)	150
B.13	Reconstruction efficiency (N-pion)	151

PREFACE

One does not successfully navigate a doctoral program without the help of many people along the way, and this is doubly true in the field of particle physics. First, I would like to thank my advisor, Dr. Steven Dytman, for his guidance through almost seven years of research. He challenged me to understand my research's role within its subfield and the broader landscape of particle physics, and yet was also able to provide helpful advice on the most specific and technical of problems.

I would be remiss if I did not thank every single person who was involved with the design, construction, calibration, and analysis of the MINER ν A detector. This dissertation would not have been possible without your hard work! I'd like specifically thank Debbie Harris, Kevin McFarland, and Jorge Morfin for their leadership as cospokespersons. Thank you as well to Mike Kordosky, Dave Schmitz, Gabe Perdue, and Laura Fields for their help and guidance as I learned about calibration, reconstruction, analysis, and especially good C++ coding practices.

I occasionally had to step back from this work in order to preserve my sanity, and I am lucky to have a number of friends that provided me with these opportunities. Thank you to everyone at the Pittsburgh Philharmonic for the privilege of serving as your principal trombonist. Also, thank you to the Edgewood Symphony for the trial-by-fire of learning to play bass trombone. When I wasn't blowing out my frustrations with trombone, I was kicking them out in karate. Thank you Dan for being a great karate teacher and for being the occasional punching bag. Thank you as well to Brian for all of our epic road trips and vacations; they helped me to discover a love for hiking and nature (not to mention excellent BBQ) that will stay with me for the rest of my life. To Brad and Sarah: you were selflessly supportive in the best and worst of times. I cannot imagine having better friends than you.

To end, I'd to thank all of my family. Mom, dad, Ryan, and Kirsten - thank you for all of your love and support. To Joanne and Joe, thank you for welcoming me to your family and caring for

me as though I were your own son. Finally, thank you to my wife Laura, whose patience, love, and support made it possible for me to survive every stress-filled and sleepless late night of writing and coding. I truly could not have finished this work without you.

1.0 INTRODUCTION

The Standard Model of particle physics seeks to describe the universe by cataloguing the fundamental particles and the forces that act upon them. It has withstood an unrelenting barrage of experimental scrutiny for almost 40 years with impressive success, but currently lacks the experimental information needed to fully describe neutrino mass and oscillations. The experimental neutrino physics community is answering this challenge with an exhaustive array of experiments to fully understand the nature of the neutrino. The success of this program requires precision knowledge of neutrino-nucleus interactions.

This chapter briefly describes the Standard Model and its limitations, then motivates the need for precision neutrino-nucleus scattering models to support neutrino oscillation experiments. Approaches to building neutrino-nucleus interaction models are reviewed, and current models are compared to recent neutrino scattering data. The comparisons reveal a confusing situation in which none of the models agree well with data, but the same data appears to be consistent with the absence of an important nuclear effect. The measurement presented in this thesis provides data that will help to resolve this puzzle.

The remainder of this thesis proceeds as follows:

- **Chapters 2-3** describe the neutrino beam and detectors.
- **Chapter 4** explains how the MINER ν A detector is calibrated.
- **Chapter 5** details the reconstruction techniques used by this analysis.
- **Chapter 6** discusses the neutrino interaction simulation.
- **Chapter 7** describes the analysis of ν_μ charged current charged pion production.
- **Chapter 8** concludes with the presentation and discussion of the analysis results.

1.1 THE STANDARD MODEL

The Standard Model is a quantum field theory that describes the 17 known fundamental particles and their interactions via three of the four known forces: strong, weak, and electromagnetic. Each particle has a corresponding field that permeates spacetime; a particle is actually an excitation in its respective field. Particle interactions are mediated by four particles known as gauge bosons, with each boson associated with a specific force.

The Standard Model has faced unprecedented experimental scrutiny since its formulation in the 1970s, but there are currently only a handful of experimental results, all of which are inconclusive, that directly conflict with its predictions¹. Many of its predictions have been experimentally verified to remarkable precision; the consistency of measurements of the fine structure constant α , for example, is better than one part in one hundred million [2]. However, the Standard Model does exclude some phenomena, gravity and neutrino mass being among the most notable.

1.1.1 Fundamental Particles

The fundamental particles consist of fermions and gauge bosons, and the fermions can be further divided into quarks and leptons. Additionally, each particle has a matching antiparticle with the same mass but opposite quantum numbers. The fermions have half-integer spin while the bosons have a spin of 0 or 1. There are six species of quarks – up (u), down (d), strange (s), charm (c), bottom (b), and top (t) – arranged into three pairs called generations: ud , cs , and tb . The leptons are comprised of three neutrinos and three charged leptons and are also arranged into three generations: $e\nu_e$, $\mu\nu_\mu$, and $\tau\nu_\tau$.

There are four spin-1 vector gauge bosons. The W and Z bosons mediate the weak force, the photon γ mediates the electromagnetic force, and the gluon g mediates the strong force. The fifth boson is a spin-0 scalar called the Higgs boson (H) and is responsible for providing itself and other particles with mass. Tables 1.1 and 1.2 summarize the properties of the fundamental particles.

¹One example is the BABAR measurement of the ratio of \overline{B} semileptonic decays, which disagrees with the Standard Model at 3.4σ [1].

Table 1.1: Properties of bosons.

Particle	Force or Effect	Mass (GeV/c ²)	Charge (e)
Gluon	Strong	0	0
Photon	Electromagnetic	0	0
W^{\pm}	Weak	81.8	± 1
Z	Weak	92.6	0
H	Provides mass to particles	126	0

Table 1.2: Properties of fermions.

Particle	Mass (MeV/c ²)	Lifetime (s)	Charge (e)	Weak Force?	Strong Force?
ν_e	undefined	∞	0	yes	no
e	0.511	∞	-1	yes	no
ν_{μ}	undefined	∞	0	yes	no
μ	105.7	2.2×10^{-6}	-1	yes	no
ν_{τ}	undefined	∞	0	yes	no
τ	1784	3.3×10^{-13}	-1	yes	no
u	1–5	∞	2/3	yes	yes
d	3–9	∞	-1/3	yes	yes
c	$1.15\text{--}1.35 \times 10^3$	–	2/3	yes	yes
s	75–170	–	-1/3	yes	yes
t	174.3×10^3	–	2/3	yes	yes
b	$4\text{--}4.4 \times 10^3$	–	-1/3	yes	yes

1.1.2 Fundamental Interactions

Each type of interaction can only affect those particles that carry a particular quantum number that is conserved by the interaction. The most well-known example of this is electric charge, which is the quantum number required for particles to interact via the electromagnetic force. The charged leptons, quarks, and W bosons all carry electric charge. The strong force analog to electric charge is called color charge and is only carried by quarks and gluons. Finally, the weak force requires participating particles to carry weak isospin, which is only carried by the W bosons and left-handed fermions.

The weak force is mediated by the electrically charged W and neutral Z bosons. Interactions mediated by a Z are referred to as neutral current, while W interactions are called charged current. The charged current interaction changes the charge of the participating particles – a charged lepton is converted to its partner neutrino, or *vice versa*, and a quark is converted to another quark of +1 or -1 electric charge. The quark can be converted to the other quark in its generation or to a quark in another generation. This arises because the quark states that couple to the weak interaction Hamiltonian are not identical to the quark mass states, which are eigenstates of the free particle Hamiltonian. This “mixing” of states is conventionally attributed to the down-type quarks d , s , and b . The relationship between the two bases is described by the Cabibbo-Kobayashi-Maskawa (CKM) matrix [3]

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & U_{cs} & V_{cb} \\ V_{td} & U_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix} \quad (1.1)$$

where d' , s' , and b' are the weak eigenstates and couple with u , c , and t respectively. The matrix elements describe the couplings between the quark mass states. The matrix V is a unitary matrix with four independent parameters, one of which is a CP-violating phase.

1.2 NEUTRINO OSCILLATIONS

Neutrino oscillation, a phenomenon in which neutrinos of one flavor are later measured in another flavor, is not included in the Standard Model. The effect arises from the quantum interference of neutrino flavor states and implies that neutrinos have mass. Oscillations were conclusively discovered by the year 2000 [4, 5], and since then there has been a substantial experimental program to measure the neutrino masses and oscillation parameters. This section will review the theory of neutrino oscillations, then motivate and discuss current efforts to measure neutrino oscillation parameters.

1.2.1 Neutrino Oscillation Theory

Analogous to quark mixing, the neutrino states that participate in the weak interaction, ν_e , ν_μ , and ν_τ , differ from the neutrino states of definite mass ν_1 , ν_2 , and ν_3 . The two sets of states are related by the leptonic equivalent of the CKM matrix, called the PMNS matrix [6, 7]:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} e^{i\frac{\alpha_1}{2}} & 0 & 0 \\ 0 & e^{i\frac{\alpha_2}{2}} & 0 \\ 0 & 0 & 1 \end{pmatrix} \times \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{cp}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{cp}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{cp}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{cp}} & -c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{cp}} & c_{23}c_{13} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \quad (1.2)$$

where $s_{ij} = \sin(\theta_{ij})$ and $c_{ij} = \cos(\theta_{ij})$. The PMNS matrix is expressed in terms of the mixing angles θ_{12} , θ_{13} , and θ_{23} , along with a CP-violating phase δ_{cp} . The phases α_1 and α_2 are only physical if neutrinos are Majorana particles².

The time-evolution of a neutrino mass state in the absence of an external potential is given by Schrödinger's equation to be

$$|\nu_i(t)\rangle = e^{-i(E_i t - p_i L)} |\nu_i(0)\rangle, \quad (1.3)$$

²A Majorana particle is its own anti-particle. There are a few experiments that are searching for neutrino-less double β -decay, which is a process that can only occur if neutrinos are Majorana.

where E_i is the neutrino energy, p_i is the neutrino momentum, and L is the distance traveled by the neutrino from $t = 0$. In the limit $p_i \gg m_i$, Equation 1.3 simplifies to

$$|\nu_i(t)\rangle = e^{\frac{-im_i^2 L}{2E_i}} |\nu_i(0)\rangle. \quad (1.4)$$

Neutrino flavor survival and disappearance probabilities are calculated using Equations 1.2 and 1.4. For example, the probability $P(\nu_\mu \rightarrow \nu_\mu)$ that a ν_μ remains a ν_μ , choosing distance and energy scales such that only the ν_μ and ν_τ states participate, is calculated to be

$$P_{\mu \rightarrow \mu} = 1 - \sin^2(2\theta_{23}) \sin^2 \left(1.27 \Delta m_{23}^2 \frac{L}{E} \right), \quad (1.5)$$

in which $\Delta m_{23}^2 = |m_3^2 - m_2^2|$, L is in units of km, and E is in units of GeV.

1.2.2 Long Baseline Neutrino Experiments

Neutrino oscillations were conclusively observed in the year 2000 [4, 5] after years of hints to their existence. Impressively, the Δm_{ij}^2 and θ_{ij} have all been measured since then. The value of δ_{cp} is the single remaining unknown neutrino mixing parameter. Its measurement is important for a few reasons. If non-zero, it will be the first measurement of CP-violation amongst leptons. Additionally, CP-violation is one of the three Sakharov conditions [8] needed to explain the observed dominance of matter over antimatter in the universe. However, the CP-violation observed in the quark sector is too small to account for the matter-antimatter asymmetry. Leptonic CP-violation may be large enough to explain the difference.

Long baseline neutrino oscillation experiments have made important contributions [9, 10] to the measurements of the Δm_{ij}^2 and θ_{ij} , and the proposed long baseline experiment LBNE [11] is one of the most promising avenues to a measurement of δ_{cp} [12]. In these experiments, a massive detector, called the far detector, is installed hundreds of miles from a few-GeV ν_μ beam³. The beam is aimed through the Earth⁴ at the detector. The ratio of the beam energy E to the baseline distance L is determined such that the detector sees maximal oscillations. The detector must be massive in order to maximize the observed number of neutrinos, given the small neutrino interaction cross

³Few-GeV generally indicates ~ 700 MeV up to ~ 5 GeV.

⁴For a particle whose mean free path through lead is approximately 1 light-year, traversing a few hundred miles of rock and dirt is not a challenge.

section and the spreading of the beam over the long distance. Typically this is achieved by using nuclear targets like water, iron, or liquid argon.

These experiments measure oscillation parameters by first measuring the energy distribution of a particular neutrino flavor at the far detector. The observed energy distribution is divided by the expected energy distribution in the absence of neutrino oscillations. The resulting ratio is the neutrino oscillation probability, like the example in Equation 1.5. The energy spectrum ratio is fit to the oscillation probability function and the oscillation parameters are extracted from the fit.

The neutrino beams used in long baseline experiments are fairly broadband, requiring that the neutrino energy be reconstructed because it is not known *a priori*. There are two strategies used to reconstruct the neutrino energy. Both strategies begin by selecting charged current interactions through the detection of the outgoing charged lepton, such as the one diagrammed in Figure 1.1, thereby tagging the incoming neutrino flavor. It is generally straightforward to measure the outgoing lepton energy, *e.g.* using the curvature of the lepton track in a magnetic field. Calorimetric detectors like MINOS [13] calculate the neutrino energy by summing the lepton energy with the reconstructed hadronic recoil:

$$E_{recoil} \equiv E_\nu - E_l, \quad (1.6)$$

where E_ν is the reconstructed neutrino energy, E_l is the measured outgoing lepton energy, and E_{recoil} is measured as the calorimetric sum of the remaining non-leptonic energy in the detector. Other experiments, such as T2K [14], elect to reconstruct the neutrino energy by limiting the selected events to interactions in which the hadronic recoil is a single nucleon. Such interactions are called quasi-elastic⁵. Given this two body final state and assuming that the struck nucleon is at rest, the incoming neutrino energy can be calculated using just the reconstructed lepton angle θ_l and E_l :

$$E_\nu = \frac{m_p^2 - (m_n - E_b)^2 - m_l^2 + 2(m_n - E_b)E_l}{2(m_n - E_b - E_l + p_l \cos \theta_l)}. \quad (1.7)$$

E_b is the nuclear binding energy. Both of these E_ν reconstruction strategies rely on knowledge of

⁵T2K and other experiments that use this approach do not actually detect the recoil nucleon, instead opting to select this topology by vetoing mesons. Therefore, the hadronic recoil may contain multiple nucleons; such events are called *quasi-elastic-like*. The validity of Equation 1.7 depends on the physical processes that cause the multi-nucleon ejection.




Figure 1.1: A $CC\nu$ interaction. A W boson is exchanged with the nucleus. There is a charged lepton, recoil nucleus, and hadron recoil X in the final state.

the hadronic recoil and are therefore sensitive to nuclear effects that modify the interaction and its final state particle content.

1.2.3 Example of a Nuclear Effect: Pion Absorption in Final State Interactions

Final state interactions (FSI), which are the interactions between neutrino interaction products and other nucleons inside the target nucleus, can significantly modify the measured E_ν by altering the particles observed in the detector. Pion absorption is one kind of FSI that is important for neutrino oscillation experiments. Figure 1.2 shows a cartoon example of pion absorption in a nucleus. A pion produced in the initial interaction, typically from the decay of a baryon resonance, is absorbed on a pair of nucleons before it can exit the nucleus. The pion is not observed in the detector and the reconstructed E_ν is incorrect, particularly if it is assumed that the neutrino interaction is quasi-elastic and Equation 1.7 is used. For example, a 1 GeV neutrino that produces a Δ resonance and a muon at 10° is calculated to be only ~ 650 MeV by Equation 1.7, so the effect of pion absorption on E_ν reconstruction can be quite large. In carbon at $E_\nu \sim 1$ GeV, pions produced in baryon resonance decays are absorbed approximately 20% of the time.

The current uncertainties on nuclear effects, including FSI, are quite large. As an example,




Figure 1.2: A cartoon example of pion absorption. The neutrino excites a proton to a Δ^{++} resonance. The Δ^{++} decays to $p\pi^+$. The π^+ is absorbed on a correlated np pair and is not seen in the final state.

Table 1.3 summarizes the systematic errors from a recent T2K measurement of the oscillation parameter θ_{13} [9]. Nuclear effects are responsible for the largest systematic uncertainties. This situation is even more dire for LBNE, which will be a systematics-limited experiment. Neutrino oscillation measurements are transitioning from the discovery era to the precision era, but a better understanding of neutrino nucleus interactions is needed to make the transition a success.

1.3 MODELING CHARGED CURRENT NEUTRINO-NUCLEUS PION PRODUCTION

The previous section used pion absorption as an example of a final state interaction that can significantly affect neutrino oscillation measurements. To account for this effect, experimenters need to be able to predict the rate at which pions are produced by charged current neutrino-nucleon and neutrino-quark interactions, as well as the rate of pion absorption in the nucleus. The latter requirement necessitates knowledge of the pion energy spectrum in the nucleus, as pion absorption is a highly energy-dependent process. In sum, experimenters need a complete model of pion

Table 1.3: T2K ν_e appearance systematic uncertainties from [9]. The first three entries are almost entirely ν cross section and nuclear effect uncertainties because the beam flux component of the ND280 (T2K near detector) constraint is small.

Error Source	% Error ($\sin^2 2\theta_{13} = 0.1$)
Beam flux & ν interactions (ND280)	5.0
ν interactions (external constraints)	7.4
Final state interactions	2.3
Far detector	3.0
Total	9.9

production from charged current neutrino-nucleus interactions.

The naively ideal way to construct such a pion production model would be to calculate the nuclear wave function and use scattering theory, but this is a prohibitively difficult many body problem. Instead, the typical approach is to model individual components of the interaction separately, then piece these models together. The three components are:

- Model of pion production in charged current neutrino-nucleon interactions.
- Modifications to the primary neutrino interaction due to the nuclear medium.
- Final state interaction model.

There are no precision measurements of any of these model components using neutrinos at energies relevant to neutrino oscillation experiments. However, they can be partially constrained by electron, photon, and pion scattering data.

1.3.1 Charged Current Neutrino-Nucleon Interactions

Pions are produced by neutrino-nucleon interactions primarily in two ways:

- **Resonance Production** – The neutrino excites the nucleon to a baryon resonance state, which then can decay to a nucleon and pion.
- **Deep Inelastic Scattering** – The neutrino interacts with one quark in the nucleon. The quark is ejected and hadronizes, often forming one or more pions.

Most neutrino oscillation experiments use the Rein-Sehgal model [15] to predict neutrino baryon resonance production. In this model, the cross section to produce a single resonance of finite width is given by

$$\frac{d\sigma}{dq^2 d\nu} = \frac{G_F^2}{4\pi^2} \left(\frac{-q^2}{P^2} \right) \kappa (u^2 \sigma_L + v^2 \sigma_R + 2uv \sigma_s), \quad (1.8)$$

where G_F is the Fermi constant, $\nu = E_\nu - E_l$ is the energy transfer, q is the 4-momentum transfer, P^2 is the 3-momentum transfer squared, and

$$\begin{aligned} u &= (E_\nu + E_l + P) / 2E_\nu, \\ v &= (E_\nu + E_l - P) / 2E_\nu, \\ \kappa &= (W^2 - M_n^2) / 2M_n, \end{aligned} \quad (1.9)$$

in which W is the resonance mass and M_n is the nucleon mass. σ_L , σ_R , and σ_s represent the partial cross sections to produce a resonance of mass W from an incident virtual vector boson of left-handed, right-handed, or zero helicity respectively. The σ terms contain the dynamics of the interaction and are calculated using the Feynman-Kislinger-Ravndal model [16], which treats the nucleon as three valence quarks coupled as a relativistic harmonic oscillator. Each σ term also contains a Breit-Wigner factor

$$\frac{1}{2\pi} \frac{\Gamma}{(M - W)^2 + \Gamma^2/4}, \quad (1.10)$$

that provides the resonance with a width Γ . The total cross section is obtained by integrating Equation 1.8 over appropriate M and q .

Rein and Sehgal calculate the production cross sections for eighteen resonances with $W < 2\text{GeV}$, and resonances with identical spin and orbital angular momentum are made to interfere when calculating the cross section to produce a particular final state post-decay. For example, $\nu_\mu p \rightarrow \mu^- n \pi^+$ results from the interference of the $\Delta(1232)$, $\Delta(1920)$, and $N(1720)$ resonances. In

addition to the resonances, a non-resonant isospin $\frac{1}{2}$ background is added incoherently to improve agreement with experiment. It is at the discretion of the experiment whether to use Rein-Sehgal's non-resonant prediction or to extend a model of deep inelastic scattering to lower $-q^2$ and invariant mass W .

The neutrino-nucleon deep inelastic scattering cross section, written in terms of structure functions F_1 , F_2 , and F_3 , is

$$\frac{d^2\sigma^{\nu,\bar{\nu}}}{dx dy} = \frac{G_F^2 M_n E_\nu}{\pi(1 + Q^2/M_{W,Z}^2)^2} \times \left[\frac{y^2}{2} 2xF_1(x, Q^2) + \left(1 - y - \frac{xyM_n}{2E_\nu}\right) F_2(x, Q^2) \pm y\left(1 - \frac{y}{2}\right) xF_3(x, Q^2) \right], \quad (1.11)$$

under the approximation that the outgoing lepton mass is small compared to the nucleon mass. The kinematic variable definitions are the same as those used in Equations 1.8 and 1.9, with $Q^2 = -q^2$, $x = \frac{Q^2}{2M_n\nu}$, and $y = \frac{\nu}{E_\nu}$. x can be understood as the fraction of the total nucleon momentum carried by the struck quark in the infinite momentum frame, and y is the fraction of the neutrino energy carried by the outgoing hadronic system⁶. The structure functions are functions of parton distribution functions (PDFs), which describe the probability that a quark of species q carries fraction x of the nucleon momentum at Q^2 . PDFs are typically obtained through fits to experimental data rather than theory.

Equation 1.11 describes deep inelastic scattering in terms of kinematic variables that summarize the hadronic system in a general way – namely, its total energy and momentum. As such, it does not describe the particle content to be found in the hadronic system. These predictions are made by a hadronization model that takes the outgoing quark and its invariant mass W as input. PYTHIA [17] is a well-known example of a hadronization model. It uses the phenomenological Lund string fragmentation prescription [18] and works well for $W \gtrsim 2$ GeV, but requires modifications to provide an alternative prediction for the non-resonant inelastic background used in the Rein-Sehgal baryon resonance model.

⁶ y is often referred to as the inelasticity.

1.3.2 Nuclear Medium Modifications

The procedures described above to calculate resonance and deep inelastic scattering production of pions by neutrinos must be modified for nucleons bound in a nucleus. The nuclear binding energy lowers the effective nucleon mass, and the nucleons have momentum relative to the nucleus rest frame. The nucleon momentum distribution is often modeled as a relativistic fermi gas (RFG) [19], though this introduces some inaccuracy. In the RFG model, the protons and neutrons are separately considered to be non-interacting fermions confined to a potential well with walls at the nuclear radius. The nucleons have a flat momentum distribution that is abruptly cut-off at the Fermi momentum p_F , which is typically ~ 250 MeV/c. This sets the scale of the mean field binding energy to approximately 30 MeV, which is confirmed by electron scattering data. Additionally, the RFG model requires that the nucleons obey the Pauli exclusion principle. In neutrino scattering, this manifests as a suppression at low Q^2 because interactions that impart the nucleon with momentum less than p_F are forbidden.

The RFG is inaccurate because it ignores interactions between nucleons. These interactions can give the nucleons momentum greater than p_F . Bodek and Ritchie provide a prescription that modifies the RFG momentum distribution with a tail at high-momentum, derived from comparisons of structure functions from deep inelastic scattering on nuclear targets to free proton and neutron structure functions [20]. Benhar *et al.* present an alternative approach whereby the RFG ground state is modified by two-body correlation operators using perturbation theory [21].

The nuclear environment further complicates the neutrino nucleon interactions at low Q^2 by allowing the neutrino to interact with a correlated pair of nucleons, most often an np pair. Such interactions are called meson exchange currents (MEC) because the nucleon-nucleon interaction can be approximated by the exchange of a meson. The effect of MECs on the total cross section is well-modeled by enhancing the transverse component of the cross section in a manner consistent with electron-nucleus scattering data [22]. Also, MECs can affect the final state content of the interaction by ejecting an extra nucleon from the nucleus⁷.

⁷Final state interactions can affect the outgoing nucleons, so the presence of a single extra nucleon is not guaranteed.

1.3.3 Final State Interactions

After the neutrino interacts with a bound nucleon, the products of the interaction must traverse the nucleus, possibly interacting with other nucleons along the way. Such interactions are called final state interactions (FSI). Pions, in particular, are affected by a number of FSI processes:

- **Elastic Scattering** ($\pi N \rightarrow \pi N$) – The pion and nucleon kinetic energy and momentum are unmodified. The pion may change direction, though this is unlikely because elastic scattering is forward-peaked.
- **Inelastic Scattering** ($\pi N \rightarrow \pi NX$) – The pion loses kinetic energy and there can be an outgoing shower X of additional hadrons from break-up of the target nucleon.
- **Charge Exchange** (e.g. $\pi^+ n \rightarrow \pi^0 p$) – The pion’s charge is changed by the interaction.
- **Absorption** (e.g. $\pi^+ np \rightarrow pp$) – There is no pion in the final state. This happens most often in interactions on a correlated pair of nucleons, analagous to pion absorption on deuterium. This process is discussed within the context of neutrino oscillation measurements in Section 1.2.3.

Pions are a very useful particle for studying FSI because they are the lightest of the mesons. Their small mass, relative to the nucleon mass, allows their momentum magnitude and direction to be easily changed when scattering off the much heavier nucleon.

There are a number of approaches used to model FSI. The simplest approach is to approximate the interaction using nuclear scattering data. For example, the GENIE neutrino event generator [23] uses π^+ and p scattering data on ^{56}Fe . Predictions on other nuclei are made by scaling the cross section according to $A^{2/3}$, where A is the atomic mass number. Isospin symmetry is used to make predictions for n , π^0 , and π^- . This approach has some disadvantages. First, the neutrino interaction deposits particles uniformly throughout the nucleus; this is not equivalent to the π^+p scattering data, where the interaction occurs on the nuclear surface. Second, this method cannot be used for particles that do not have nuclear scattering data. Finally, this method also relies on external data to set the 4-momenta of the outgoing particles. In the absence of data, the 4-momenta can be set in a very simple way by uniformly sampling from the available phase space.

Another approach to modeling FSI is to directly simulate each particle’s path through the nucleus using a Monte Carlo method. This is usually referred to as a cascade simulation. Each

particle is stepped through the nucleus, where the size of each step size is some fraction of the nuclear radius. After each step, a random number is used to determine if the particle interacts. In the case of interaction, random numbers also determine the interaction process and the 4-momenta of the outgoing particles. The interaction cross sections can be determined by free-nucleon scattering data or calculated. This method is ultimately limited by the availability of data and calculations for the relevant nucleon cross sections, as well as the validity of those cross sections when applied to bound nucleons.

The GiBUU [24] group calculates FSI using the Boltzmann-Uehling-Uhlenback (BUU) equation [25], which describes the time evolution of the phase space density of each particle species participating in FSI. All particles are subject to a mean field potential that approximates the nuclear medium, representing an improvement relative to the cascade simulation approach. The particle species are coupled through the mean field potential and “collision” terms that represent particle interactions and decays. The result is a system of coupled integro-differential equations, one for each particle species. Solving these equations can be computationally expensive, so the GiBUU method is not ideal for use by oscillation experiments that require fast simulations.

The previous few paragraphs presented a few techniques to model FSI in approximate order of increasing complexity, but increased complexity does not necessarily imply better accuracy. Ultimately, neutrino oscillation experiments require an accurate reproduction of FSI and the “correctness” of the physics used to build the model is of secondary concern. Each of these methods have been tuned, to some extent, to agree with charged lepton, pion, and photon nuclear scattering data. Nevertheless, this does not guarantee consistency with neutrino scattering data. In order to reduce the FSI systematic uncertainties in neutrino oscillation experiments, FSI models must be compared against and made to agree with neutrino scattering data that is sensitive to FSI.

1.4 NEUTRINO CHARGED CURRENT CHARGED PION PRODUCTION DATA

Neutrino interaction models are built on the foundation of a neutrino-nucleon interaction model. There are numerous baryon resonance pion production models available, including the Rein-Sehgal model described in Section 1.3.1, and the experimentalist paradigm is that the best model choice




Figure 1.3: The ANL [26] and BNL [27] measurements of the cross section for $\nu_\mu p \rightarrow \mu^- p \pi^+$ on a deuterium target.

marries convenience with good data agreement. Experimental uncertainties are set by the disagreement of the model with data and the size of the uncertainties in the data. Unfortunately, data for neutrino-nucleon baryon resonance production is both limited and inconsistent. The best measurements of the total cross section were made by ANL [26] and BNL [27] deuterium bubble chamber experiments in the 1980s. Figure 1.3 shows the ANL and BNL results for $\nu_\mu p \rightarrow \mu^- p \pi^+$. Both of these measurements have large systematic uncertainties – approximately 10-20%. Additionally, the measurements are inconsistent because the normalization disagrees by up to 40%. This situation results in large systematic uncertainties on the cross sections predicted by the baryon resonance production model regardless of the choice of model.

Models of nuclear effects must also withstand experimental scrutiny. One of the best ways to test FSI models is through pion production data because the pion is sensitive to a variety of FSI processes. The MiniBooNE experiment [28] published the first complete set of single and double differential cross sections for ν_μ charged current single pion production on a nuclear target (mineral oil – CH_2) [29]. Figure 1.4 compares the MiniBooNE result for the differential cross section with respect to pion kinetic energy, $d\sigma/dT_\pi$, to three calculations and three simulations. None of the models agree particularly well with the data, though different models agree with




Figure 1.4: The MiniBooNE $d\sigma/dT_\pi$ measurement [29] compared to three theoretical calculations [24, 30, 31] and three simulations [23, 32, 33]. Reprinted with permission from [34].

certain specific features. The prediction for the total cross section, *i.e.* the integral of $d\sigma/dT_\pi$, is primarily determined by

1. the neutrino-nucleon interaction model
2. nuclear effects that modify the width of the baryon resonance
3. the strength of pion absorption.

Despite not including resonance-modifying nuclear effects, GENIE is the only model that agrees with MiniBooNE's prediction for the total cross section.

The $d\sigma/dT_\pi$ shape disagreements are more interesting because the shape is uniquely affected by FSI. Figure 1.5 shows the disagreements within the context of the FSI models for the GiBUU calculation and GENIE simulation. Both GiBUU and GENIE predict that FSI deplete the cross section between ~ 100 -300 MeV, though this effect is much more pronounced for GiBUU. The depletion is due to pion absorption. The MiniBooNE data sees modest depletion in the absorption region and in this regard compares more favorably with GENIE. Additionally, GiBUU and GENIE predict that FSI cause the cross section to peak at lower energy, with GENIE also predicting enhancement to the cross section at low energy. This enhancement can be understood as inelastic pion scattering, which reduces the energies of the pions that exit the nucleus. GiBUU and GENIE

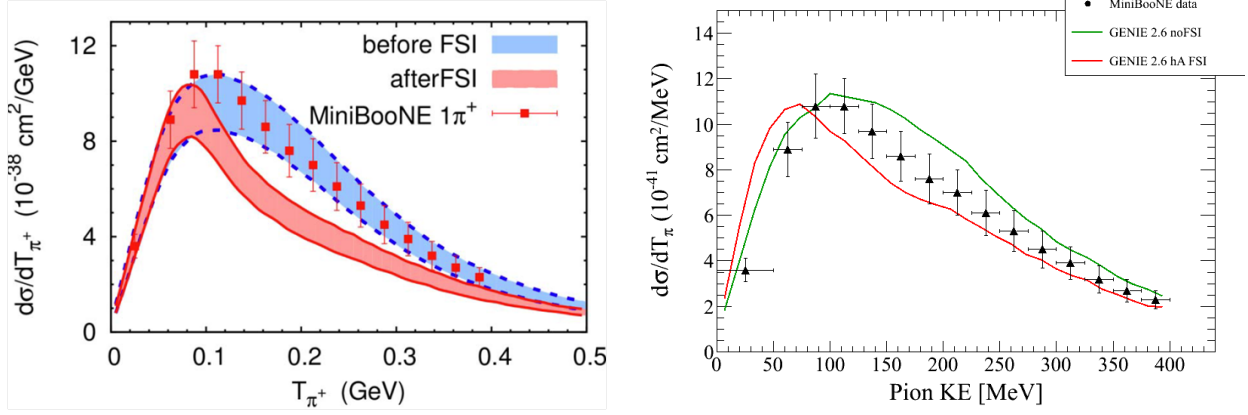


Figure 1.5: The MiniBooNE $d\sigma/dT_\pi$ measurement [29] compared to GiBUU (left) and GENIE (right) with and without final state interactions. The GiBUU figure is reprinted from [35], Copyright (2013) with permission from the American Physical Society. The GENIE figure is courtesy of Steve Dytman.

best reproduce the peak of the MiniBooNE data distribution by removing the FSI model. Strangely enough, GiBUU is remarkably consistent with the MiniBooNE distribution in both shape and magnitude when it does not include the effect of FSI. GENIE performs better than GiBUU when the FSI models are included. This is somewhat unexpected because GiBUU contains nuclear effects, such as the resonance width modification, that are not included in GENIE. Also, as described previously, GiBUU uses a more sophisticated FSI model.

The reason for the disagreement between the MiniBooNE data and each model is currently not understood. It is also not understood why the models that include the most nuclear effects, like GiBUU and Valencia [31], have the most difficulty reproducing the data. Since MiniBooNE is the only measurement of its kind, additional data is needed to resolve these puzzles. This motivates the central result in this thesis: the measurement of the shape of $d\sigma/dT_\pi$ in the MINERνA detector.

2.0 THE NUMI BEAMLINE

The NuMI (Neutrinos and the Main Injector) beamline [36] is a ν_μ and $\bar{\nu}_\mu$ beam located at the Fermi National Accelerator Laboratory (Fermilab) in Batavia, IL. It is used by the MINER ν A, MINOS, and NO ν A [37] experiments as a neutrino source for oscillation and cross section measurements. This chapter describes the design of the NuMI beam and the measurement of its ν_μ and $\bar{\nu}_\mu$ flux.

2.1 THE PROTON BEAM

The creation of the NuMI beam actually begins with a container of hydrogen gas. The hydrogen is ionized to form H^- ions, which are then accelerated to 750 KeV by a Cockcroft-Walton¹ generator. Next, the H^- ions are accelerated to 400 MeV by a series of linear accelerators. The beam then passes through carbon foil, converting the H^- to protons, and the protons are injected into the Booster ring.

The Booster is a synchrotron accelerator with a radius of 75 m. It accelerates the protons to 8 GeV in about 67 ms, requiring $\sim 40,000$ laps around the ring. The protons are then injected into another synchrotron, called the Main Injector, that accelerates the protons to 120 GeV. While the Tevatron was operational, the Main Injector also served as a proton source for antiproton production and a storage ring for the antiprotons. It separates the proton beam into six groups called batches. The Tevatron only required one of the six batches; the remaining five were delivered to the NuMI beamline. After the Tevatron was decommissioned, all six batches were used by NuMI.

¹In 2012, after the data used in this thesis was collected, Fermilab replaced the Cockcroft-Walton with a new radio frequency quadrupole accelerator.

The protons are extracted to the NuMI target once every 2.2 seconds using a single turn extraction. Each extraction consists of one five or six batch beam spill. A five batch spill generally contains $\sim 3.5 \times 10^{13}$ protons and lasts $8.4 \mu\text{s}$; a six batch spill is $10 \mu\text{s}$ and contains $\sim 4.2 \times 10^{13}$ protons. Before colliding with the NuMI target, the proton beam position is focused to a 1 mm profile by a baffle. The proton beam is monitored and its intensity is measured to 2%. The intensity measurement places a direct limit on how well the neutrino flux can be measured.

2.2 THE NEUTRINO BEAM

The NuMI neutrino beam is created from the decay of charged mesons produced by the collision of the 120 GeV proton beam with the NuMI target. The beamline is diagrammed in Figure 2.1. The NuMI target, shown in Figure 2.2, consists of 47 graphite (carbon) fins mounted to two stainless steel water cooling pipes along the beam direction. Each fin is 20 mm in the beam direction, with a 0.3 mm gap between each fin, giving the target a total length of 95.38 cm (approximately 2 interaction lengths). The fins present a 15 mm by 6.4 mm cross section to the beam; this narrow profile is chosen to reduce the probability that hadrons produced from the target reinteract in the fins. There is a 48th fin, used as a Budal monitor [38], that is electrically isolated from the rest of the target. It measures the shape of the beam profile and the beam intensity through detection of delta-ray charges created by beam interactions in the fin. The entire target is enclosed in an air-tight aluminum casing which in turn is mounted within a helium-filled steel canister. Approximately 89% of the neutrinos seen in the MINER ν A detector are created from interactions on the graphite target; the remaining neutrinos are attributed to interactions in the casings, cooling system, and helium gas.

The proton interactions in the NuMI target create mesons, mostly pions and kaons. These particles are focused into a beam by two 3 m magnetic parabolic focusing horns, shown in Figure 2.3. The position of the target can be adjusted relative to the horns, changing the beam energy by varying the meson momenta that are focused. The electrical current in the horns is nominally 185 kA; changing its sign allows the choice of focusing positive or negative mesons, which in turn create a neutrino or antineutrino beam. The magnitude of the current can also be changed to




Figure 2.1: A diagram of the NuMI beamline. Reprinted from [39], Copyright (2005) with permission from R. Zwaska.




Figure 2.2: A diagram of the NuMI target. Reprinted from [40], Copyright (2002) with permission from Fermi National Accelerator Laboratory.




Figure 2.3: A diagram of the NuMI focusing horns. Reprinted from [41], Copyright (2006) with permission from Elsevier.

adjust the focal point of the horns. In all cases, the resulting meson beam has two components: the mesons focused by the horns, and high momentum mesons that are produced collinear with the target and do not require focusing.

After the focusing horns, the meson beam enters a 2 m diameter, 675 m long decay pipe. The pipe is filled with helium, rather than air, to minimize interactions. Most of the pions and kaons decay to muons and muon neutrinos in the decay pipe. Table 2.1 summarizes the common decay modes for each meson that yield neutrinos. The neutrinos created by the decays are boosted forward in the laboratory frame, creating a neutrino beam. About 97% of the neutrinos that enter the MINER ν A are produced from pion decays, and the remainder are produced from charged kaons.

Any protons and mesons that reach the end of the decay pipe are met by a hadron monitor and hadron absorbers. The hadron monitor is used to monitor the beam at the downstream end of the decay pipe. The hadron absorbers, composed of concrete, remove all hadronic content from the beam, leaving just muons and neutrinos. A series of three muon monitors, each separated by dolomite rock, follow the hadron absorbers. The muon monitors measure the muon energy spectrum, which can be used to measure the neutrino flux [42] *in situ*. There is 240 m of dolomite separating the hadron absorber from the detector hall, which is sufficient to stop all muons in the beam, leaving only neutrinos. Figure 2.4 shows the layout of the NuMI absorber region.

Table 2.1: Decay modes of π^+ and K^+ that contribute significant neutrino flux to the NuMI beam. π^- and K^- decays are charge conjugates.

Mode	Fraction
$\pi^+ \rightarrow \mu^+ \nu_\mu$	99.99%
$K^+ \rightarrow \mu^+ \nu_\mu$	63.44%
$K^+ \rightarrow \pi^0 e^+ \nu_e$	4.98%
$K^+ \rightarrow \pi^0 \mu^+ \nu_\mu$	3.32%

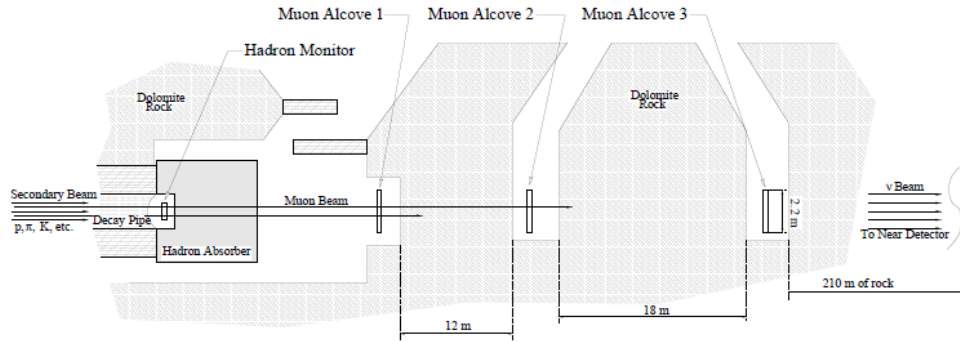


Figure 2.4: A diagram of the NuMI absorber region. Reprinted from [41], Copyright (2006) with permission from Elsevier.

2.3 NUMI FLUX MEASUREMENT

Neutrino beam flux measurements are a notoriously difficult but necessary component of any accelerator-based neutrino cross section measurement. The NuMI flux is measured using G4numi, a NuMI-specific implementation of Geant4 9.2.p03 [43], with external data constraints. G4numi contains a description of the NuMI beamline geometry and a map of the magnetic field generated by the focusing horns. Hadron production from 120 GeV p interactions on the graphite target is simulated using the FTFP_BERT model included in Geant4. FTFP_BERT uses the FTFP [44] model to simulate inelastic hadron-nucleus scattering above 4 GeV. In FTFP, QCD strings are formed between partons in the incoming hadron and target nucleus using the FRITIOF [45] model. The strings are hadronized according to the Lund fragmentation model [18] and the excited nucleus is de-excited by G4Precompound. The Bertini intranuclear cascade model [46] simulates interactions below 4 GeV.

The pC interaction products are propagated through the NuMI beamline geometry, subject to the magnetic field produced by the focusing horns, and allowed to reinteract in the beamline material. G4numi also simulates particle decay.

The FTFP_BERT prediction for 120 GeV pC interactions does not compare well with data; see, for example, the comparison of FTFP_BERT to NA49 data [47] in Figure 2.5. The relevant variables in these comparisons are:

- x_F – The Feynman scaling variable, $x_F = 2p_L/\sqrt{s}$, where p_L is the outgoing particle’s longitudinal momentum and \sqrt{s} is the total center of mass energy.
- p_T – The transverse momentum of the outgoing particle.
- $f(E, x_F, p_T)$ – The invariant production cross section, $f = E d^3\sigma/dp^3$.

For the data analyzed in this thesis, pions focused by the magnetic horns generally fall within the range $p_T = (0.1 - 0.6)$ GeV/ c and $x_F = (0.05 - 0.15)$, while the higher momentum unfocused mesons have the kinematic range $p_T = (0 - 0.2)$ GeV/ c and $x_F = (0.2 - 0.5)$.

Data agreement is improved by reweighting the FTFP_BERT predictions to match external data from three experiments, summarized in Table 2.2. The weight applied to the simulation for




Figure 2.5: A comparison of NA49 data and FTFP_BERT simulation of the cross sections for the process $pC \rightarrow \pi^+ X$.

Table 2.2: A summary of data used to constrain the NuMI flux measurement.

Data	Process	Proton Energy (GeV)	Kinematics
NA49 [49]	$pC \rightarrow \pi^\pm X$	158	$x_F < 0.5$
NA49 [50]	$pC \rightarrow K^\pm X$	158	$x_F < 0.2$
NA49 [51]	$pC \rightarrow pX$	158	$x_F < 0.95$
MIPP [52]	$pC \rightarrow K^\pm X$	120	$0.2 < x_F < 0.5$
Barton [53]	$pC \rightarrow \pi^\pm X$	100	$x_F > 0.5$

data collected at energy E_{data} is

$$weight = C(E_{data}) \cdot \frac{f(E_{data}, x_F, p_T)_{data}}{f(E_{g4numi}, x_F, p_T)_{g4numi}}, \quad (2.1)$$

where $C(E_{data})$ is derived from an alternate hadron production simulation called FLUKA [48] and scales the data to the simulated energy $E_{g4numi} \sim 120 GeV$:

$$C(E_{data}) = \frac{f(E_{g4numi}, x_F, p_t)_{fluka}}{f(E_{data}, x_F, p_t)_{fluka}}. \quad (2.2)$$

Simulated events that are not covered by the data in Table 2.2 are not reweighted. The proton production data is used to constrain simulation of the proton beam attenuation in the NuMI target. Figure 2.6 shows the ratio of weighted to unweighted simulation; the weights can be as large as 50% at higher E_ν . Figure 2.7 shows the reweighted flux prediction used by the analysis in this thesis.




Figure 2.6: The ratio of weighted to unweighted flux as a function of neutrino energy. “Right sign” means that the neutrino is produced from a charged meson that is focused by the focusing horns.




Figure 2.7: The weighted ν_μ flux measurement used in this thesis analysis. It is compared to the unweighted prediction.

3.0 THE MINER ν A DETECTOR

The MINER ν A experiment is a collaboration of approximately 65 nuclear and particle physicists from over 20 institutions around the world (see Appendix C). The collaboration's goals include the measurement of inclusive and exclusive neutrino-nucleus cross sections and the study of nuclear and nucleon structure using neutrinos as a probe. The MINER ν A detector, installed 100 meters underground in the NuMI beamline at Fermilab between 2009-2010, is designed with these goals in mind. This chapter describes the design and simulation of the MINER ν A detector and its electronics. More detail can be found in References [54, 55].

3.1 DETECTOR DESIGN

As shown in Figure 3.1, the main body of the detector is a regular hexagonal cylinder with a length of 5 m and an apothem of approximately 1.7 m. It consists of an inner detector of apothem 1.07 m that is radially enclosed by an outer detector. The central portion of the inner detector is a fine-grained, fully-active scintillator tracker that is capable of resolving charged particle tracks in high multiplicity neutrino interactions. The upstream (with respect to the direction of the neutrino beam) edge of the tracker is adjacent to the nuclear target region, which contains a series of carbon, iron, lead, and water targets that provide a means to study the nucleon number dependence of nuclear effects. There is a liquid helium target upstream of the nuclear target region, as well as a veto wall to reject charged particle backgrounds that enter the front of the detector. The remainder of the tracker and the sides of the nuclear target region are surrounded by electromagnetic and hadronic calorimeters (ECAL and HCAL respectively) for containment of electromagnetic and hadronic showers.

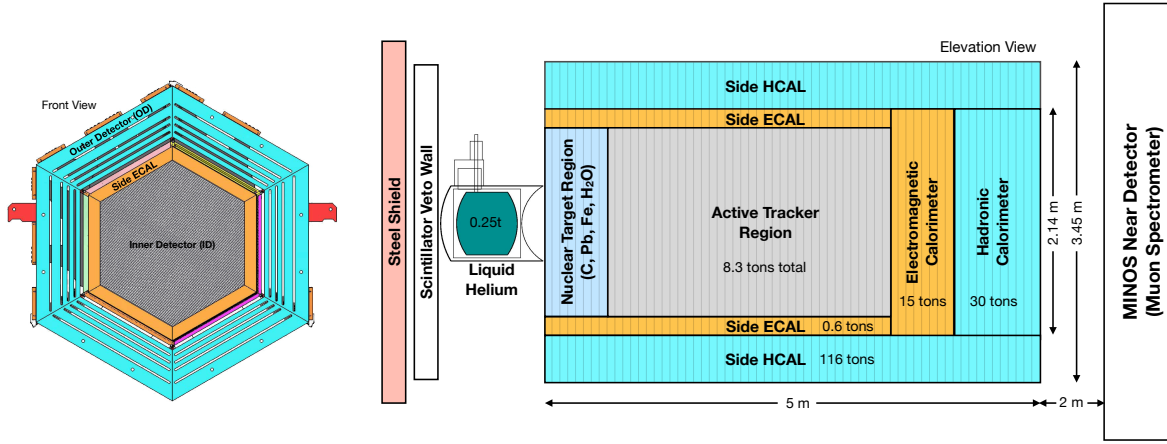


Figure 3.1: Left: front view of a single tracking module. Right: elevation view of the entire detector. Reprinted from [54], Copyright (2014) with permission from Elsevier.

In total, MINER ν A contains approximately 32000 active scintillator strips (“channels”). The detector is located 2.1 meters upstream of the MINOS near detector, hereafter referred to as the MINOS detector or simply MINOS. The MINER ν A detector is unmagnetized and incapable of containing most muons produced inside of it, so MINOS is used as a muon catcher and spectrometer.

The detector coordinate system is right-handed and defined such that the positive y-axis is directed away from the center of the Earth. The positive z-axis lies along the detector’s central symmetry axis and points towards MINOS, with the coordinate system origin defined such that $z = 1200$ cm is the front face of MINOS. The central axis of the NuMI beam lies in the y-z plane and is directed downward at an angle of 3.34° with respect to the z-axis¹.

3.1.1 Inner Detector

The inner detector consists of four regions, listed here in order from upstream to downstream: the nuclear target region, tracker, downstream ECAL, and downstream HCAL. Additionally, it

¹The beam is directed downward because it is aimed at the MINOS far detector in the Soudan Mine in northern Minnesota.

contains the side ECAL, which is located between the edges of neighboring scintillator planes in the nuclear target and tracker regions. The inner detector is further subdivided along the z-axis into 120 units called modules. There are four types of modules: tracking modules, ECAL modules, HCAL modules, and passive nuclear targets. The tracker, downstream ECAL, and downstream HCAL are constructed exclusively from their respective module types and contain 62, 10, and 20 modules respectively. The nuclear target region contains 5 solid passive nuclear targets, a water target, and 22 tracking modules. The modules are installed orthogonal to the z-axis and there is a ~ 2.5 mm air gap between each module.

3.1.1.1 Detector Modules A tracker module contains two hexagonal scintillator planes. Each scintillator plane is 1.7 cm thick and contains 127 triangular scintillator strips. Additionally, a hexagonal lead ring (otherwise known as the side ECAL) of thickness 2 mm, inner apothem 90 cm, and outer apothem 105 cm is located upstream of each scintillator plane. Each plane is rotated in one of three orientations, called views, with respect to the coordinate system. The orientation of the strips in each view are shown in Figure 3.2. In the X view, the scintillator strips are parallel to the y-axis, thus providing position information along the x-axis. The U and V views are rotated in the x-y plane by 60° clockwise and counterclockwise, respectively, relative to the X view. The downstream plane in each tracker module is in the X view, while the upstream plane can be either U or V. The modules are ordered such that the planes alternate throughout the tracker region according to the pattern U X V X. The use of three views enables MINER ν A to reconstruct certain track configurations that would be ambiguous in two views: for example, two tracks that have the same projection along the z-axis or overlap in one view.

ECAL modules contain two scintillator planes in the same configuration as the tracker modules, but missing a side ECAL collar. Instead each a 2 mm thick hexagonal lead plane with apothem identical to the scintillator planes is installed upstream of each scintillator plane in the ECAL. The full lead planes help contain electromagnetic showers that would otherwise escape the downstream end of the detector. The U X V X scintillator plane pattern is maintained in the ECAL.

HCAL modules contain an upstream hexagonal iron absorber of thickness 2.54 cm and a downstream scintillator plane. The modules are arranged so that the scintillator planes are ordered in an X U X V pattern. In the low energy NuMI beam configuration, most hadrons that originate from




Figure 3.2: The X (left), U (center), and V (right) scintillator plane orientations as seen by an upstream observer. The inner lines show the direction of the scintillator strips. This figure is provided by Laura McCarthy.

neutrino interactions in the tracker region are contained by the HCAL.

The nuclear target region is unique in that it contains more than one type of detector module. Lead, iron, water, and carbon targets are separated by tracker modules as shown in Figure 3.3. The analysis presented in this thesis does not use the nuclear target region and is negligibly affected by it, so it will not be described further. A more detailed description can be found in Ref. [54].

Each detector module includes six trapezoidal outer detector frames, to be described in more detail in Sec. 3.1.2. The frames are attached to the module in the pattern shown in figure 3.1, maintaining the hexagonal shape. With the exception of the thickness, the frames are identical across all regions of the detector. Downstream HCAL frames are slightly thicker than the others due to the thickness of the iron absorber.

3.1.1.2 Scintillator Planes Each scintillator plane is a 1.7 cm thick regular hexagon with a 107 cm apothem. The planes are composed of 127 scintillator strips glued together with 3M-DP190 translucent epoxy. Sheets of Lexan are attached to each plane with 3M-DP190 gray epoxy to increase rigidity and reduce exposure to ambient light. Black PVC electrical tape is used to seal joints in the Lexan and cover any remaining light leaks, allowing MINER ν A to be sensitive to single photons generated in the scintillator. Figure 3.4 shows an edgewise view of a scintillator




Figure 3.3: The MINERνA nuclear targets region. The beam direction is from left to right.
 Reprinted from [54], Copyright (2014) with permission from Elsevier.




Figure 3.4: Edgewise view of the strips in a scintillator plane. Each plane contains 127 triangular scintillator strips fitted side by side as shown. This figure is provided by Laura McCarthy.

plane and demonstrates how the strips are fitted together. The triangular shape ensures that every particle that crosses a scintillator plane will intersect at least two strips. This configuration improves position resolution relative to equally-sized rectangular strips by comparing the amount of scintillation light produced in each strip intersected by the particle.

The scintillator strips vary in length according to their position in the plane, and have a triangular cross section with a base of 3.3 cm and height of 1.7 cm. The scintillator is Dow Styron 663 W polystyrene ($(C_8H_8)_n$) doped with 1% 2,5-diphenyloxazole (PPO) and 0.03% 1,4-bis(5-phenyloxazole-2-yl) benzene (POPOP) by weight. The strips are co-extruded with a 0.25 mm reflective coating of polystyrene and 15% (by weight) TiO_2 . For readout, each strip contains a 2.6 mm diameter hole that runs the length of the strip. The hole is centered on the base and is located at a height of 0.85 cm above the base. A 1.2 mm diameter green wavelength shifting (WLS) fiber is put in the hole and fixed in place by optical epoxy (Epon Resin 815C and Epicure 3234).

The WLS fiber is a 175 ppm Y-11 doped, S-35, multicladd fiber made by the Kuraray Corporation. It shifts the blue scintillation light to green in order to best match the photosensitivity of the MINERνA photomultiplier tubes (PMTs). Only one end of the fiber is read out. The other end is mirrored via an “ice-polishing” technique [56], followed by the deposition of aluminum and Red Spot UV Epoxy layers. The mirroring directs additional light to the readout end of the fiber.

The material composition of the planes must be well understood in order to interpret cross section results on scintillator. The fractional mass density of elements in each material is determined through direct measurement (pure scintillator and coated strips), assayed compositions (coated strips and epoxy), and manufacturer data sheets (Lexan and electrical tape). Table 3.1 summarizes

Table 3.1: Density and composition by mass percentage of scintillator plane materials [54].

Material	Density (g/cm ³)	H	C	N	O	Al	Si	Cl	Ti
Scintillator	1.043 \pm 0.002	7.6	92.2	0.06	0.07	-	-	-	-
Coating	1.52	6.5	78.5	-	6.0	-	-	-	9.0
Lexan	1.2	6.7	66.7	-	26.7	-	-	-	-
PVC tape	1.2	4.8	38.7	-	-	-	-	56.5	-
Transl. Epoxy	1.32	10.0	69.0	2.6	17.0	-	-	0.5	-
Gray Epoxy	1.70	5.0	47.0	1.7	27.0	6.0	6.0	0.05	-

each material’s fractional mass density and Table 3.2 summarizes the material assay for a single MINER ν A scintillator plane.

3.1.2 Outer Detector

The outer detector is a hadronic calorimeter consisting of a steel frame instrumented with scintillator strips. Additionally, it serves as the detector support structure and constrains the plane alignment. Each outer detector frame is trapezoidal and contains eight scintillator bars of varying length that run parallel to the trapezoid base. In contrast to the inner detector, the outer detector

Table 3.2: Elemental composition of scintillator strips and planes by mass percentage [54].

Component	H	C	O	Al	Si	Cl	Ti
Strip	7.59%	91.9%	0.51%	-	-	-	0.77%
Plane	7.42%	87.6%	3.18%	0.26%	0.27%	0.55%	0.69%

scintillator bars have a rectangular cross section. The bars in the nuclear target, tracker, and downstream ECAL regions have a base of 19.0 mm and height of 16.6 mm. The downstream HCAL bars are larger due to the thickness of the iron absorber. The outer detector bars also have a larger fiber hole than their inner detector counterparts, but otherwise contain similar WLS fiber technology.

3.1.3 Helium Target and Veto Wall

From upstream to downstream, the veto wall consists of a 5 cm thick steel plate, a 1.9 cm thick scintillator plane, a 2.5 cm thick steel plate, and another 1.9 cm thick scintillator plane. It is designed to shield the detector from low energy hadrons and tag muons created by neutrino interactions in the surrounding cavern rock (such particles are referred to as “rock muons” in this thesis). The helium target is a 2300 liter aluminum cryostat located between the veto wall and detector. The analysis presented in this thesis is not affected by the presence of these components so they will not be described further. Ref. [54] contains a more detailed description of the helium target.

3.1.4 Light Collection

The WLS fibers are fed into Fujikura-DDK connectors in groups of eight. Each connector is matched to a bundle of eight clear optical fibers that are sealed in light tight material and also terminated by DDK connectors. The clear optical fibers are used to transport scintillation light to the PMTs. The fibers are 1.2 mm diameter S-35 multicladd clear optical fiber manufactured by Kuraray and each bundle is cut to one of four lengths: 1.08 m, 1.38 m, 3.13 m, or 6 m. The length of the bundle is determined by the location of the WLS fiber connector on the detector. The clear optical fiber attenuation length and mean DDK connector transparency (the percentage of incident light transmitted by the connector) are measured to be 7.83 m and 87.0% respectively.

MINERvA uses 507 Hamamatsu Photonics H8804MOD-2 multi-anode PMTs to amplify the scintillation light. These are essentially the same PMTs used by MINOS [57]. Each multi-anode PMT is a collection of 64 individual PMTs distributed in an 8x8 grid measuring 4 cm². Henceforth, the multi-anode PMT will be referred to as PMT and the component channels will be called pixels. The pixels consist of a bialkali photocathode with a borosilicate glass window and a twelve stage dynode amplification chain. The photocathode quantum efficiency is required to be at least 12%

at 520 nm and the maximum to minimum pixel gain ratio can be no more than three. The gain of the dynode chain, defined as the number of electrons collected at the anode divided by the number of photoelectrons arriving at the first dynode, is $\sim 5 \times 10^5$. The scintillation light from a minimum ionizing particle typically produces a few photoelectrons at the photocathode, resulting in a few-hundred fC electrical signal at the anode.

Each PMT is mounted onto a base circuit board that contains the Cockcroft-Walton high voltage (HV) power source. The operating high voltage is generally between 750-850 V and is distributed amongst the dynodes and anode according to the ratios 3:2:2:1:1:1:1:1:1:1:2:5, where the first 12 numbers correspond to the 12 dynodes and the last number refers to the anode. The PMT and base circuit board are installed inside a 2.36 mm thick steel cylindrical box that provides protection from ambient light, dust, and residual magnetic fields. The PMT boxes are mounted onto racks directly above the detector.

Eight clear optical fiber bundles, numbering 64 fibers total, are connected to the faceplate of each PMT box. In the interior of the box, the light is delivered from the faceplate connector to each pixel by clear optical fibers. An 8x8 “cookie”, mounted onto the face of the PMT, ensures the alignment of each fiber with its corresponding pixel. The fibers are mapped such that the light from adjacent scintillator strips is not fed to adjacent pixels in the PMT. This minimizes the effect of PMT cross talk, the process by which signal in one pixel can induce a signal in neighboring pixels, on event reconstruction. Figure 3.5 diagrams the fiber mapping.

The PMT boxes also have two optical fiber ports with diffusers on the front face plate. These ports are used to inject a calibration light source into the PMT; see Sec. 4.1.3. In this case, the light is not transported directly to the pixels by fibers; rather, the light is spread by polypropylene diffusers and arrives at the pixels through spaces in the cookie.

The MINOS detector magnetic coil creates magnetic fields in the vicinity of MINER ν A that can be as large as 30 gauss. The performance of the PMTs is adversely affected by magnetic fields higher than 5 gauss, so shielding is necessary. The PMT box itself provides some magnetic shielding. Additionally, the PMTs are oriented perpendicular to the residual field and the 40 PMT boxes closest to the MINOS detector are fitted with a high permeability metal shielding. Finally, a steel “mirror” plane, installed between the MINOS and MINER ν A detectors, provides additional shielding.




Figure 3.5: The MINERνA PMT fiber map. Clear optical fibers connect to the PMT face plate and interior fibers are routed to the cookie. The cookie is mounted on the face of the PMT pixel array. Reprinted from [54], Copyright (2014) with permission from Elsevier.

3.1.5 MINOS Detector

MINOS is a long baseline neutrino oscillation experiment with a near detector at Fermilab and a far detector installed in the Soudan Mine in northern Minnesota. The MINOS near detector is a one kiloton magnetized course-grained scintillator detector and calorimeter located 2.1 meter downstream relative to the MINER ν A detector. It is used by MINER ν A to measure the charge and momentum of forward-going muons. This section contains a brief description of the MINOS detector design. A more complete description can be found in Ref. [13].

The MINOS detector is composed of 282 one inch thick steel planes. Throughout the entire detector, every fifth steel plane is followed by a fully-instrumented scintillator plane. In addition, each of the first 120 steel planes that is not paired with a fully-instrumented scintillator plane is instead followed by a partially-instrumented scintillator plane. The scintillator planes are 1 cm thick and are composed of 4.1 cm wide rectangular scintillator strips. The scintillator composition is identical to that used by MINER ν A. The strips are oriented $\pm 45^\circ$ with respect to the positive y-axis, with the orientation alternating in successive planes. The two orientations, referred to as the MINOS U and V views, differ from the U and V views used by MINER ν A.

The MINOS detector is magnetized by a magnetic coil that runs the entire length of the detector (along the beam direction) inside a hole through the detector interior. The magnetic field inside the detector is toroidal with an average strength of 1.3 T. Muon momentum and charge can be measured from the curvature of the muon track through MINOS. Additionally, the momentum can be measured from the track range if the track is contained in the detector.

3.1.6 Test Beam Detector

In 2010, a miniature version of the MINER ν A detector was installed in a hadron test beam at the Fermilab MTest facility [58] in order to study hadron response in the scintillator and calorimeters. The test beam detector scintillator planes are 1 m² and contain 63 scintillator strips. The scintillator composition is identical to the full MINER ν A detector. Test beam tracker, ECAL, and HCAL modules are constructed from scintillator planes, lead absorbers, and iron absorbers. The scintillator is read out by MINER ν A multi-anode PMTs connected directly to the WLS fibers; one PMT services one plane. Data was collected in two configurations: 20 tracker planes followed by

20 ECAL planes, and 20 ECAL planes followed by 20 HCAL planes.

The hadron test beam is a tertiary beam. 120 GeV protons from the Main Injector strike an Al target, producing mostly pions that are focused into a secondary beam. The secondary beam interacts with a copper target, creating tertiary charged pions, protons, and electrons. Tertiary particle momentum is measured with two dipole magnets and four wire chambers upstream of the detector. The particle species is identified by comparing the reconstructed momentum to the time-of-flight through the wire chambers.

Figure 3.6 plots the momentum vs. time-of-flight for particles entering the test beam detector. The detector does not see pions and protons below ~ 400 MeV/c. 400 MeV/c pions are contained in the ECAL+HCAL configuration, but not the Tracker+ECAL. For this reason, the test beam pion data is not used to study the pion particle identification algorithms described in this thesis. Instead, the pion and proton data are used to constrain the uncertainty on the MINER ν A detector calorimetric response to these particles. In addition, a sample of non-interacting protons that stop in the tracker planes are used to measure the non-linearity in the conversion of energy to scintillation light (see Sec. 3.3.3).

3.2 DETECTOR READOUT

The raw data retrieved from the detector consists of the timing and magnitude of the charges collected at the PMT anode during each beam spill for each of the ~ 32000 channels. Each PMT is serviced by a front end board (FEB), mounted directly onto the PMT, that is tasked with organizing and collecting data from the PMT. The FEBs are read out by a computer-based data acquisition system (DAQ), facilitated by rack-mounted electronics that handle communications between the DAQ computers and the FEBs. Ref. [55] contains more details.

3.2.1 Front End Boards

The FEBs serve two roles: supply high voltage to the PMTs via a Cockroft-Walton generator and read out the PMT anode charge. Before read out, capacitive dividers split the charge from




Figure 3.6: Momentum vs. time-of-flight of particles in the tertiary test beam. The “accidentals” are particles whose time-of-flight reconstruction is confused by unrelated activity induced by the previous radio frequency bucket in the Main Injector batch structure.

each pixel into low, medium, and high gain channels. This splitting increases dynamic range while maintaining fine resolution on signals that are only a few photoelectrons (PEs). Read out is handled by six 32-channel Application-Specific Integrated Circuit (ASIC) chips, called TriP-t chips [59], that integrate, time-stamp, and digitize the anode charge. Digitization is handled by 12 bit Analog-to-Digital Converters (ADCs). FEB operation is controlled by a Spartan 3E Field-Programmable Gate Array (FPGA) chip, which also decodes incoming timing signals.

The first four TriP-t chips, numbered 0-3, read out the high and medium gain channels, while TriP-t chips 4-5 handle the low gain channels. The high and medium gain channels of a single pixel are always read out by the same TriP-t chip; under this requirement, each of the chips 0-3 read out 16 high gain channels and 16 medium gain channels. The result of this configuration is that two TriP-t chips are used to read out the three gain ranges of a single pixel. In order to avoid ambiguity during the read out process, the TriP-t chips are organized into two triplets: chips (0,1,4) and chips (2,3,5). Each triplet contains full low, medium, and high gain information for 32 pixels.

The fundamental unit of time on a TriP-t chip is called a clock tick, approximately 9.4 ns. The clock tick is set by a 53.1 MHz crystal contained within a master timing module (described in Sec. 3.2.2) that distributes a timing signal to all other read out electronics. Additionally, during charge integration, a quadrature circuit is used to achieve quarter-tick timing resolution, approximately 2.4 ns. Charge is time-stamped by the channel's discriminator, which triggers when the integrated charge exceeds a certain threshold. In practice, only the high gain discriminators are used because any charge that triggers a pixel's high gain discriminator will also trigger its medium and low gain discriminators. Time on each FEB is measured relative to a local counter. The counters are initialized taking into account differences in signal propagation time, synchronizing the timing across all FEBs to within one clock tick.

Read out is organized into gates. A gate is a 1702 clock tick ($\sim 16 \mu\text{s}$) window that starts 500 ns before the arrival of the beam spill and ends $5.5 \mu\text{s}$ after the spill ends. The TriP-t chip triplets are able to independently integrate and time-stamp their channels' charges up to seven times per gate. A triplet begins integration when one of its high gain discriminators is triggered. The integration window lasts for 16 clock ticks ($\sim 150 \text{ ns}$), then the integrated charge is pushed into the chip's ADC pipeline; each instance of integrated charge in the ADC pipeline is called a hit. All channels on the triplet are time-stamped with T_0 , the clock tick of the earliest triggered high gain discriminator.

Subsequent discriminator triggers during the integration window are time-stamped with delay and quarter ticks relative to T_0 . The integrators are then reset in preparation for the next integration window. The reset requires 20 clock ticks (~ 188 ns). After the end of the gate, all hits in the ADC pipelines are digitized.

There are two sources of electronics dead time that arise from FEB operation. The first occurs during the reset period after an integration window. Any charge that arrives at the TriP-t chip triplet during reset is not integrated. This primarily affects read out of decay activity and neutrino interactions that are close in time and spatially overlap. The other form of dead time occurs when a triplet reaches its limit of seven time-stamped integration windows within a gate. All subsequent charge is integrated and pushed into the ADC pipeline at the end of the gate, but does not receive a time stamp. Untimed hits are not used in this thesis analysis.

3.2.2 Data Acquisition

The FEBs are daisy-chained into groups of nine or ten with standard ethernet cables. Both ends of the chain are connected to a custom VME module called the Chain Readout Controller (CROC), which can service up to four chains. The CROC reads out its FEBs after each gate and is also responsible for sending instructions and timing information. CROCs receive timing information from another custom VME module called the CROC Interface Module (CRIM). The CRIMs collect timing information from MINOS, which is used for event-matching between the two detectors, and the MINER ν A Timing Module (MvTM), which receives timing information from the Main Injector and contains the 53.1 MHz timing synchronization crystal. The signal from the MvTM controls the start of each gate relative to the beam spill. Additionally, one CRIM serves as the master CRIM, and is tasked with distributing timing triggers to the other CRIMs when collecting light injection data (see Sec. 4.1.3). In full detector running, MINER ν A uses two VME crates, each of which contains a CAEN V2718 Crate Controller, two CRIMs, and seven or eight CROCs. The crates are controlled and read out by separate computers installed with a CAEN A2818 PCI Card that communicates with the V2718 Crate Controller.

3.3 DETECTOR SIMULATION

Simulation is an important component of modern particle physics experiments. It helps experimentalists to understand backgrounds, detector resolutions and efficiencies, and sensitivity to the physics that will be studied. The MINER ν A detector simulation has two components. The first is the simulation of particles traveling through the detector, which provides a prediction of the energy deposited in the active scintillator. The second component is a simulation of the detector optics and electronics read out. It converts energy deposited in the scintillator to light, propagates the light to the PMTs, and simulates the digitization of the electrical signal produced at the PMT anodes. A data overlay technique is used to simulate the effects of pile-up², unsimulated backgrounds, detector dead time, dead channels, and time-depended detector effects.

3.3.1 Particle Propagation

Geant4 9.4.p02 [43] is used to simulate the propagation of particles through the MINER ν A detector. However, Geant4 is not used to simulate the initial neutrino interaction. Instead, the GENIE neutrino event generator [23] simulates neutrino-nucleus interactions, including nuclear effects and final state interactions. Chapter 6 discusses the MINER ν A implementation of GENIE, focusing on the physics models that it contains. GENIE outputs the particle species and 4-vectors seen in the final state, after final state interactions. Geant4 uses the GENIE output as its input and propagates the particles through the detector.

Geant4 propagates particles through MINER ν A in discrete steps. After each step, it calculates the energy deposited in the detector due to ionization and radiation. Energy deposited in the active scintillator is converted to a digitized electric charge by the optical and electronics simulations. Additionally, after each step, Geant4 uses interaction models to determine whether any of the particles interacted in the detector. If an interaction occurs, Geant4 selects the outgoing particle species and 4-vectors according to the interaction model.

The QGSP_BERT model is used to simulate the strong inelastic interactions between hadrons and detector nuclei. It is an intranuclear cascade model that is most accurate for interactions at

²Pile-up refers to multiple neutrino interactions in the same read out gate.

Table 3.3: Geant4 physics modules used by the MINER ν A detector simulation.

Physics Process	Geant4 Module	Notes
Inelastic hadronic	QGSP_BERT	Intranuclear cascade, best below 10 GeV
Elastic hadronic	Hadron_Elastic	Glauber model [60] including multiple scattering
Electromagnetic	EmStandard	Bremsstrahlung, e^+e^- production, multiple scattering
Decays	DecayPhysics	-

energies below 10 GeV. The analysis in this thesis is restricted to neutrino energies below 10 GeV and is therefore insensitive to inaccuracies in the QGSP_BERT model above 10 GeV. Standard Geant4 physics modules are used to simulate other physics interactions; some are listed in Table 3.3.

3.3.2 Data Overlay

The data overlay pairs each simulated event with a data gate from the run period that is simulated. Dead channel and electronics dead time information is applied to the simulated event, masking simulated activity that occurs in a dead channel. Certain time-dependent calibrations are also applied to the simulation, according to the time of the overlayed data. Additionally, some of the data activity, *i.e.* the times and energies of hits in the data gate, is added to the simulated event. Only the activity that occurs between the time of the first and last simulated hits is added. All overlayed data activity is propagated through the rest of the simulation in order to simulate its effect on reconstruction and analysis.

3.3.3 Optical and Electronics Models

The MINER ν A optical model converts energy deposited in the active scintillator to photons. Non-linearity is accounted for with Birks' law [61]:

$$N_\gamma = C \frac{dE}{1 + k_B dE/dx}. \quad (3.1)$$

N_γ is the number of photons produced in the scintillator, dE is the deposited energy in path length dx , C scales energy to number of photons, and k_B , called Birks' constant, determines the size of the non-linearity. k_B is set to 0.133 mm/MeV in the simulation, but is measured to be 0.0905 ± 0.012 mm/MeV by the MINER ν A test beam detector [62]. This difference is accounted for as a systematic uncertainty. Minimumly ionizing particles (MIP) in MINER ν A scintillator lose energy at the rate of ~ 2.5 MeV/cm, resulting in a $\sim 2\%$ non-linearity. The effect is more noticeable for low energy charged hadrons, which can lose energy at a much higher rate than a MIP. A proton that stops in MINER ν A loses 15 MeV/cm or more at the end of its path, resulting in a non-linearity greater than 10%.

The light N_γ is propagated along the WLS and clear optical fibers to the PMT, accounting for attenuation along the fibers and reflection from the mirrored end of the scintillator strip. The effect of scintillator aging is modeled by applying a time-dependent scale factor that forces the simulated distribution of photons to match data. The timing of simulated hits is initially determined by Geant4 and the light propagation time along the optical fibers. It is much narrower in the simulation than data, so a data-driven model is used to smear the simulated hit times.

The quantum efficiency of the PMT cathode is simulated by using a Poisson distribution to calculate the number of photoelectrons (PEs) produced by the cathode. The first two stages of the PMT dynode chain are modeled separately as amplifications with a Poisson width. The remaining stages are treated as a single amplification with Gaussian width. The PMT gains are drawn from the gain calibrations used in the overlaid data gate. The PMT model also includes optical cross talk. Hit digitization is handled by a simple model of the front end electronics.

3.3.4 MINOS Simulation

Simulated particles that exit MINER ν A and point towards the MINOS detector are the input to a Geant3 [63] simulation of the MINOS near detector provided by the MINOS collaboration. Geant3 simulates particle energy loss, interactions, and decays in MINOS. It also curves the trajectories of charged particles according to the magnetic field. Energy deposited in the scintillator is converted to a digitized signal by models of MINOS optics and electronics. The MINOS simulation does not include a data overlay procedure.

4.0 CALIBRATION OF THE MINER ν A DETECTOR

The MINER ν A detector measures the position, energy, and time of charged particles traveling through the scintillator strips. The position is known from the location of scintillator strip that registers a hit. The energy is stored as a digitized PMT anode charge, and the time is an FEB clock tick relative to the start of the read out gate. The actual scintillation energy and its time must be calculated from the digitized data stored by the detector. Additionally, misalignments in detector modules must be accounted for when determining the position of the scintillator strips. Detector calibrations are used to make these corrections. This chapter summarizes the calibration of various MINER ν A detector components. Ref. [54] and Appendix A contain additional detail.

4.1 ENERGY CALIBRATION

The goal of the energy calibration is to convert the digitized PMT anode charge in channel i , denoted ADC_i ¹, to the energy E_i deposited in the scintillator strip. The series of effects that create the measured ADC_i are as follows:

1. E_i is converted to scintillation light.
2. The scintillation light is attenuated as it travels along the WLS and clear optical fibers to the PMT. Discrete attenuation also occurs in the connectors that transfer light from one fiber to another.
3. The PMT photocathode converts the scintillation light to photoelectrons (PEs), which are amplified to a measurable electric signal by the PMT dynode chain.

¹ADC refers to analog-to-digital conversion.

4. The charge at the PMT anode is digitized by the FEB to an ADC charge ADC_i

Figure 4.1 provides a visualization of this list. The energy calibration procedure corrects for these effects in reverse order. The full calibration chain is summarized by the following equation:

$$E_i = Q_i(ADC_i - ADC_i^0) \cdot \frac{1}{g_i(t)} \cdot e^{l_i/\lambda_{clear}} \cdot \eta_i^{att} \cdot S_i(t) \cdot C(t). \quad (4.1)$$

The components of Equation 4.1 are:

- $ADC_i^0(t)$ – The quantity registered by the FEB ADC for zero electric charge. This is called the electronics pedestal.
- $Q_i(ADC)$ – A function that converts a pedestal-subtracted ADC charge to electric charge measured at the PMT anode.
- $g_i(t)$ – The gain (amplification) of the PMT dynode chain. It converts the PMT anode charge to PEs. It is a time-dependent quantity.
- l_i – The clear optical fiber length from the PMT to the scintillator strip.
- λ_{clear} – The clear optical fiber attenuation length. It is measured to be 7.83 m by a dedicated test stand measurement described in Ref. [54].
- η_i^{att} – Attenuation within the scintillator strip. It is not initially known where the particle intersected the scintillator strip, so this correction is initially made to the center of the strip. Three dimensional tracking, described in 5.2.1, provides the position along the strip. The attenuation correction is then adjusted accordingly.
- $S_i(t)$ – A relative time-dependent correction that normalizes the response of channel i to the other channels. This accounts for variations in channel response caused by factors not included in the other calibration corrections. It is only applied to inner detector channels.
- $C(t)$ – A time-dependent correction that converts the channel-to-channel normalized light to energy deposited in the scintillator strip.

4.1.1 Electronic Pedestals

The FEB ADCs register a non-zero ADC value for an input of zero electric charge. This offset, $ADC_i^0(t)$, is the electronics pedestal and must be subtracted from the ADC measurement. The




Figure 4.1: Schematic diagram of a MINERνA optical readout channel. The energy calibration procedure converts the ADC charge read out by the DAQ to the energy deposited in the scintillator strip. Reprinted from [54], Copyright (2014) with permission from Elsevier.

electronics pedestal for every channel on MINERvA is measured *in situ* every ~ 10.5 hours using 750 empty gates collected between beam spills. Figure 4.2 is a histogram of pedestal data collected over ~ 30 minutes in a single channel. The outliers in the pedestal distributions are created by discrete background processes such as cosmic rays and radioactivity. Outliers are removed using Peirce's Criterion [64]. The remaining distribution has a small width due to time variation in the channel's electronic noise. The mean and standard deviation of the outlier-subtracted pedestal distribution is measured for each channel. The measured mean is used as the pedestal correction $ADC_i^0(t)$ for data collected in the subsequent ~ 10.5 hours until the next pedestal measurement. The standard deviation is used by the PMT gain measurement described in 4.1.3.




Figure 4.2: Pedestal gates collected in ~ 30 minutes in a single channel. The outlier at ~ 555 is removed by Peirce's Criterion before measuring the pedestal mean and RMS. Reprinted from [54], Copyright (2014) with permission from Elsevier.

4.1.2 FEB Calibration

The function $Q_i(ADC)$ is determined *ex situ* for the high, medium, and low gain ADCs of each FEB using data from a dedicated test stand at the University of Pittsburgh. A series of known

electric charges were injected into each ADC and the corresponding ADC response was recorded. The response is not linear; this is accounted for by fitting the ADC vs. charge curves to a trilinear function². The fit results reproduce the test stand data to within 1% over the dynamic range of the ADCs.

The trilinear fit parameters for each ADC are stored in an offline database and are used to construct the function $Q_i(ADC)$. The calculation first selects the highest gain ADC that is not saturated, maximizing charge resolution. Next, the pedestal for the chosen ADC is subtracted from the measured ADC value. Finally, the trilinear fit parameters are used to calculate the PMT anode charge.

4.1.3 PMT Gains

PMT response changes with time due to temperature fluctuations and aging. Additionally, the time dependence varies for each PMT. For these reasons, MINER ν A monitors PMT response using *in situ* calibration data collected once per day. This is sufficient to monitor response variations; noticeable changes to the response occur over a span of weeks rather than days. PMT calibration data is collected between beam spills so that neutrino data collection is not interrupted.

LED light from a light injection (LI) system is the calibration source. The LI system pulses 23 blue (472 nm) LEDs when triggered by the DAQ system. Each LED flashes 50 optical fibers, and two fibers are routed to each PMT box. Sec. 3.1.4 describes the interface between the LI optical fibers and PMT. This system is sufficient to deliver at least a few PEs to most pixels, and the average light delivered to a pixel does not change significantly with time. However, the exact amount of light delivered is not known *a priori* and the pixel-to-pixel variation in the light is large, about 30%. This situation requires a calibration technique that is insensitive to the amount of calibration light delivered to the pixel.

The PMT response that MINER ν A measures is the PMT gain g , defined as

$$g = \frac{\overline{Q}}{\lambda_e}, \quad (4.2)$$

²The trilinear function is three connected line segments. The free parameters are the slope and starting point of each line.

where \overline{Q} is the mean of the anode charge distribution, λ is the mean number of PEs arriving at the first dynode from the LI source, and e is the electron charge. The quantum efficiency of the photocathode and the collection efficiency of the first dynode are not included in this definition of the gain. These cannot be measured with LI data and are accounted for in the relative channel-to-channel ($S_i(t)$) and energy scale ($C(t)$) calibrations. By assuming that λ is constant in time and using the Rademacker model [65] for PMT dynode chain amplification, the gain is measured independent of λ with the following equation:

$$g = \frac{\sigma^2 - \sigma_p^2}{\overline{Q}(1 + w(g)^2)e}. \quad (4.3)$$

Here, σ and σ_p^2 are the standard deviations of the anode charge distribution and electronics pedestal, and $w(g)^2$ is a slowly-varying function of the gain.

The measured gains are used according to Equation 4.1. Figure 4.3 shows the measured distribution of gains in MINER ν A for April 1, 2010; most gains are between 4×10^5 and 8×10^5 electrons per PE. Appendix A contains more detail on the gain calibration procedure, including the derivation and implementation of Equation 4.3.

4.1.4 Attenuation

The attenuation of light along the clear optical fibers is calculated using measured values of the fiber length l_i and the attenuation length λ_{clear} . The optical attenuation along the WLS fiber in each scintillator strip is measured *ex situ*. Before detector installation, each detector module was scanned with a Cs-137 source by the MINER ν A “module mapper”. The source was systematically moved along each scintillator strip in the module while a PMT recorded the response. This measured the attenuation as a function of longitudinal position along the strip and is used to calculate the η_i^{att} in Equation 4.1. An alternative *in situ* attenuation measurement using rock muons (see Sec. 3.1.3) agrees with the Cs-137 scan results, but is not as precise.

4.1.5 Channel-to-Channel Variations

After applying the calibrations described above, there are still channel-to-channel variations in the measured light level. These variations are caused by effects not considered in the other calibrations;




Figure 4.3: The gain measurement for every pixel using LI calibration data collected on April 1, 2010. Reprinted from [54], Copyright (2014) with permission from Elsevier.

for example, the optical coupling between clear fiber connectors is not uniform across channels, and the scintillator composition can vary slightly from batch to batch. The remaining variations can also be time-dependent, *e.g.* changing the PMT installed on the detector changes the photocathode quantum efficiencies and the coupling between the clear optical fibers at the PMT face plate.

MINER ν A uses a minimum-ionizing rock muon sample to normalize the channel-to-channel response in the inner detector. First, the plane alignment corrections described in Sec. 4.3 are applied. Next, the procedure calculates a tail-truncated mean energy, corrected for path length, deposited by the rock muons in each scintillator strip i . This quantity is represented as x_i . A calibration constant A_i defined as

$$A_i = \frac{\frac{1}{x_i}}{\frac{1}{u_i}} \quad (4.4)$$

is applied to each strip to remove variations in x_i . u_i is the mean value of x_i for all strips in the detector. If a strip registers zero light for more than 30% of the rock muons that pass through it, it is marked as “dead” and is not used to calculate u_i . Dead channels identified in this way are also not used by the neutrino event reconstruction in data and simulation. The definition of A_i is such that the average value of A_i is one, leaving the overall detector energy scale unchanged.

The desired result is that the path length-corrected peak energy, rather than x_i , is the same in each channel. This cannot be enforced directly because the rock muon statistics are too low to fit for the peak energy response in individual strips. Instead, the peak energy response is measured in each plane after applying the A_i factors to the individual strips. Another factor P_j is calculated for each plane j that normalizes the fitted peak energy response. P_j , like A_i , is constructed such that the overall detector energy scale does not change. The channel-to-channel calibration constant $S_i(t)$ from Equation 4.1 is the product of A_i and P_j . It is recalculated every time the detector hardware changes³ to account for time-dependent variations in the relative response. Figure 4.4 shows the peak energy response to rock muons across all planes in the detector after applying the channel-to-channel calibrations.

³PMTs and FEBs occasionally malfunction and must be replaced.




Figure 4.4: The fitted rock muon peak energy per length for each plane. The distribution is flat due to the channel-to-channel calibration. The non-statistical fluctuations, such as those in the ECAL, can be reduced by additional iterations on the fits that calculate P_j , but these fluctuations are already approximately the size of the statistical uncertainty after one iteration. Reprinted from [54], Copyright (2014) with permission from Elsevier.




Figure 4.5: Comparison of muon energy in data and simulation (left) and the fit of a fifth-order polynomial to the peak in data (right). Reprinted from [54], Copyright (2014) with permission from Elsevier.

4.1.6 Detector Energy Scale

The detector energy scale correction $C(t)$ is calculated from a rock muon sample by comparing the energy distribution, after all other energy calibrations are applied, in data and simulation. The reconstructed muon positions and momenta (Sec. 5.3.2) in the data sample are used to create the simulated sample in order to avoid kinematic differences. The data and simulated muon energy distributions are fitted to a fifth-order polynomial as shown in Figure 4.5. The physics that describes muon energy loss in scintillator is well-understood and well-simulated, so $C(t)$ is defined as the ratio of the simulated fitted peak to the data fitted peak. $C(t)$ is calculated in two-day intervals to account for scintillator aging.

4.2 TIMING CALIBRATION

Each energy deposit in a scintillator strip is time-stamped with an FEB clock tick that indicates the time at which the corresponding PMT anode charge triggered the channel's high gain discriminator.

The timing calibration converts the clock tick to the time at which the energy was deposited in the strip. This correction must account for propagation time along the optical fibers, time slewing, and channel-to-channel timing offsets. The time slewing is primarily caused by variances in the time required to produce scintillation light and is a PE-dependent effect. The channel-to-channel offsets are mostly due to difference in propagation time along cables for FEBs on one chain and timing offsets between different chains. The timing calibration described below is redone after each hardware change to account for changes in the channel-to-channel offsets.

The optical fiber propagation time is calculated using the fiber lengths and the speed of light in the fiber. The time slewing and channel-to-channel offset effects are measured using a rock muon sample. The raw time of hits along each muon track are corrected for optical fiber propagation and muon time of flight. These corrected times are used to calculate a truncated mean time for each muon. The difference between each hit time and the truncated mean time of its muon is calculated and averaged over all muons as a function of number of PEs, shown in Figure 4.6. This distribution is fit to a polynomial and the fit result is used to apply a time slewing correction for each hit. The channel-to-channel offset is calculated for each group of channels read out by the same high gain TriP-t chip. The time slewing and channel-to-channel corrections are applied iteratively, calculating new truncated mean muon times with the corrections from the previous iteration, until convergence.

Figure 4.7 compares the calibrated hit time to the truncated mean time after the final iteration. The central peak is fitted to a Gaussian with a width of 3 ns. The non-Gaussian asymmetric tails are due to the variance in time slewing, which is a stochastic effect.

4.3 PLANE ALIGNMENT

Detector modules can be perturbed from their ideal position during installation, resulting in misalignment. The misalignments are small, generally a few mm translationally and a few mrad rotationally. MINER ν A uses rock muon tracks to measure the translational misalignment along each plane's measurement axis and the rotational misalignment about the z-axis, treating the scintillator strips in each plane as a rigid unit. The average energy deposited in each strip in a plane,




Figure 4.6: Time slewing vs. number of photoelectrons. This is fit to a polynomial to derive the time slewing correction. Reprinted from [54], Copyright (2014) with permission from Elsevier.




Figure 4.7: Time resolution after applying the timing calibration. Reprinted from [54], Copyright (2014) with permission from Elsevier.

corrected for normal incidence, is calculated as a function of the muon intersection point at the base of the strip. The resulting distribution is fitted to the triangular shape of the scintillator strip to find the peak energy. The peak offset from zero is the size of the translational shift for the plane. The peak offsets are then recalculated as a function of the longitudinal position along the strip. The slope indicates the rotational shift. Figure 4.8 shows the distributions used to extract the alignment parameters for two planes.

4.4 PMT CROSS TALK

Cross talk describes the process by which signal in one channel induces a response in another channel. There are three types of cross talk in MINER ν A:

- **PMT optical cross talk** – PEs produced by a pixel’s photocathode leak to another pixel’s first dynode. This most commonly occurs for adjacent pixels.
- **PMT dynode cross talk** – Electrons produced from a dynode in one pixel leak to the dynode in a different pixel. Like optical cross talk, this effect is strongest in neighboring pixels.
- **Electrical cross talk** – Cross talk between two channels on an FEB.

The two variations of PMT cross talk are by far the most significant seen by MINER ν A. Ideally, PMT cross talk is measured on an *ex situ* test stand that systematically injects a known quantity of light into each pixel while measuring the response in the other pixels. This was not done before PMT installation, so instead a sample of rock muons is used to calculate an average cross talk fraction per PMT.

Each rock muon is associated with a time slice, defined in Sec. 5.1.1, that groups the hits on the rock muon track with other hits that are coincident in time. Hits that are not on the rock muon track are considered to be cross talk candidates if they register in pixels that are on the same PMT as a rock muon track hit. This safely avoids confusion of cross talk with other muon-induced activity, like delta electrons, because the PMT fiber weave described in Sec. 3.1.4 ensures that neighboring scintillator strips are not serviced by neighboring PMT pixels. The PMT average cross talk fraction is defined as the ratio of the energy of cross talk candidate hits to the energy of on track hits. A




Figure 4.8: The distributions used to calculate the alignment shifts for module 50, plane 2 (top) and module 61, plane 1 (bottom). In the left plots, the offset of the peak in the triangular fit from zero indicates the translational shift in the plane. In the right plots, the rotational shift is calculated from the slope. Reprinted from [54], Copyright (2014) with permission from Elsevier.

version of this quantity for nearest-neighbor (adjacent) pixels, called $f_{xt,NN}$ is also calculated. $f_{xt,NN}$ is typically less than 4%; it is plotted for all PMTs and compared to the MINER ν A optical cross talk simulation in Figure 4.9.




Figure 4.9: Data-simulation comparison of the nearest neighbor cross talk fraction. Reprinted from [54], Copyright (2014) with permission from Elsevier.

5.0 EVENT RECONSTRUCTION

Event reconstruction is the collection of techniques and algorithms that use the location and timing of energy deposits in the detector to determine the trajectories, energies, and species of the particles that caused the detector activity. Reconstruction is particularly critical for understanding the nature of neutrino scattering because the neutrino energy is not known *a priori*. The presence of the neutrino, as well as its energy and the kinematics of the interaction, can only be inferred from the interaction by-products. This chapter describes the reconstruction used in the analysis of charged pion production in MINER ν A.

5.1 HIT GROUPING

The NuMI beam is a very intense neutrino source and typically produces multiple neutrino-induced interactions¹ per read out gate. Before attempting to reconstruct particle trajectories, hits are grouped by time in order to separate unrelated neutrino interactions. The groups of hits formed by this procedure are called “time slices.” A time slice in MINER ν A usually has a one-to-one correspondence with a single neutrino-induced interaction. The most common exception occurs for muon decay. The muon lifetime is $2.2 \mu\text{s}$, so the decay electron, called a Michel electron, is often placed in a separate time slice. Sec. 5.4.2 describes the procedure used to identify Michel electrons and reassociate them with their parent time slice.

Hits within the same time slice are also organized spatially to form “clusters”, leveraging the knowledge that a charged particle crossing a scintillator plane must intersect and deposit energy in at least two strips. Additionally, the hit energy is used to identify likely PMT cross talk candidates.

¹A neutrino-induced interaction is most often a rock muon or a neutrino scattering event in the detector.

The energy scale calibration implicitly corrects for the energy lost as cross talk, so it is not correct to add this energy back to the source channel. Instead, cross talk hits are simply ignored by the reconstruction.

5.1.1 Time Slicing

The time slicing algorithm selects hits that triggered their high gain discriminator and sorts them by calibrated time, correcting the time to the center of the scintillator strip. An 80 ns window walks along the collection of hits until the sum of photoelectrons² for the hits contained within the window exceeds 10. The hits within the window form a time slice. The window continues to walk forward, adding additional hits to the time slice, until the sum of photoelectrons in the window is less than 10. This procedure is repeated, forming additional time slices, until the window has walked along the entire collection of hits. Hits without a triggered discriminator are added to a time slice if they share a TriP-t chip with another hit already in the slice, and if their calibrated time is consistent with the hits in the slice.

5.1.2 Clustering

A cluster is a collection of hits in a time slice that are in consecutive scintillator strips in the same plane. Every hit is added to one and only one cluster; a hit that does not neighbor other hits is the sole member of its cluster. Clusters are classified topologically according to the number and energies of hits they contain. The classifications are designed to be indicative of the physical process or type of particle(s) that is responsible for the hits in the cluster. The definitions are as follows:

- **Cross talk** – This cluster classification supersedes the others. Each hit is compared to hits in the neighboring PMT pixels using the cross talk fraction measurement described in Sec. 4.4. If the number of photoelectrons in the hit is consistent with cross talk, the hit is placed in its own cluster. Cross talk clusters can only contain one hit and are ignored by subsequent reconstruction.

²Recall that a photoelectron is an electron produced by the photocathode of a PMT.

- **Trackable** – The cluster topology is consistent with a single minimum ionizing particle crossing the plane at an angle $\lesssim 70^\circ$. It requires that the total cluster energy is between 1 and 12 MeV, the cluster contain no more than four hits, and there is at least one hit with energy greater than 0.5 MeV. Additionally, hits with more than 0.5 MeV of energy must all be adjacent in the plane.
- **Heavy-Ionizing** – This is a higher-energy version of the trackable cluster that is consistent with a number of topologies, *e.g.* high-angle particles³, spatially-overlapping particles, or the large energy deposits of non-minimumly ionizing particles. The requirements to be classified as heavy-ionizing are similar to trackable, but without the upper limit on cluster energy. The trackable classification supersedes heavy-ionizing.
- **Supercluster** – Wide clusters that are consistent with hadron or electromagnetic particle showers rather than individual particle trajectories. Any cluster with at least five hits is automatically a supercluster.
- **Low Activity** – This is a type that was originally intended to select cross talk activity. Any cluster with total energy less than 1 MeV is classified as low activity. There are very few of these because this type is overridden by the cross talk type.

Each cluster is assigned a position and time. The position is the energy-weighted position of the hits contained in the cluster, and the time is the time of the highest energy hit.

5.2 TRACKING AND VERTEXING

A *track* is a reconstructed object that approximates a portion of a charged particle’s true trajectory through the detector. Typically, only one track is needed to reconstruct a particle trajectory. Multiple tracks may be needed if the particle undergoes a large-angle scatter or decays into other charged particles. This analysis uses tracks to tag and identifying muons and charged hadrons.

A *vertex* is the reconstructed start or end position of one or more tracks. Each vertex has at least one track associated with it, and every track starts and ends at a vertex. This analysis uses

³As the particle angle with respect to the z-axis increases, so does the path length through the plane, resulting in a higher energy cluster.

vertices in two ways. The vertex that is identified as the neutrino interaction point, called the *primary vertex*, is used to select interactions in the tracker. Additionally, the end vertex of each hadron track candidate is used as a search point for Michel electrons.

This analysis uses multiple pattern recognition schemes to find tracks. A master event building algorithm systematically applies the various track pattern recognitions to find muon and hadron tracks. It associates groups of tracks with common vertices, using a Kalman filter fit [66] to reconstruct the best position for the vertex. Tracks are also subjected to “cleaning” algorithms that remove extra energy from the track to make it available to future iterations of the track pattern recognition. This is most useful for resolving events where two or more tracks overlap in one or more views, thus sharing clusters at the point of overlap. It is important to note that the algorithms described in this section are restricted to use clusters from the same time slice.

5.2.1 Long Track Pattern Recognition

Figure 5.1 contains a cartoon that illustrates the following procedure. The *long track* pattern recognition scheme begins by using a set of trackable and heavy-ionizing clusters, provided by the master event building algorithm, to build all possible track *seeds*. A seed is a collection of three clusters, all in the same view (X, U, or V), that meet the following requirements:

1. No two clusters belong to the same scintillator plane.
2. The clusters must be in consecutive⁴ scintillator planes.
3. The clusters must satisfy a fit to a two-dimensional straight line.

Note that a cluster may belong to more than one seed, and the requirement that the seed is composed of clusters from the same view implies that a seed is a two-dimensional object. Seeds are the basic pattern recognition unit used to find long tracks. By requiring that seeds contain only one cluster per plane, this pattern recognition is limited to finding long tracks with a polar angle no larger than 70° .

Next, seeds within the same view are merged to form larger two-dimensional track-like objects called *candidates*. The merger of two seeds requires that the slope of the seeds’ linear fits are consistent, the seeds share at least one cluster, and the seeds do not contain different clusters in

⁴Planes in the same view are ordered according to position along the z-axis.

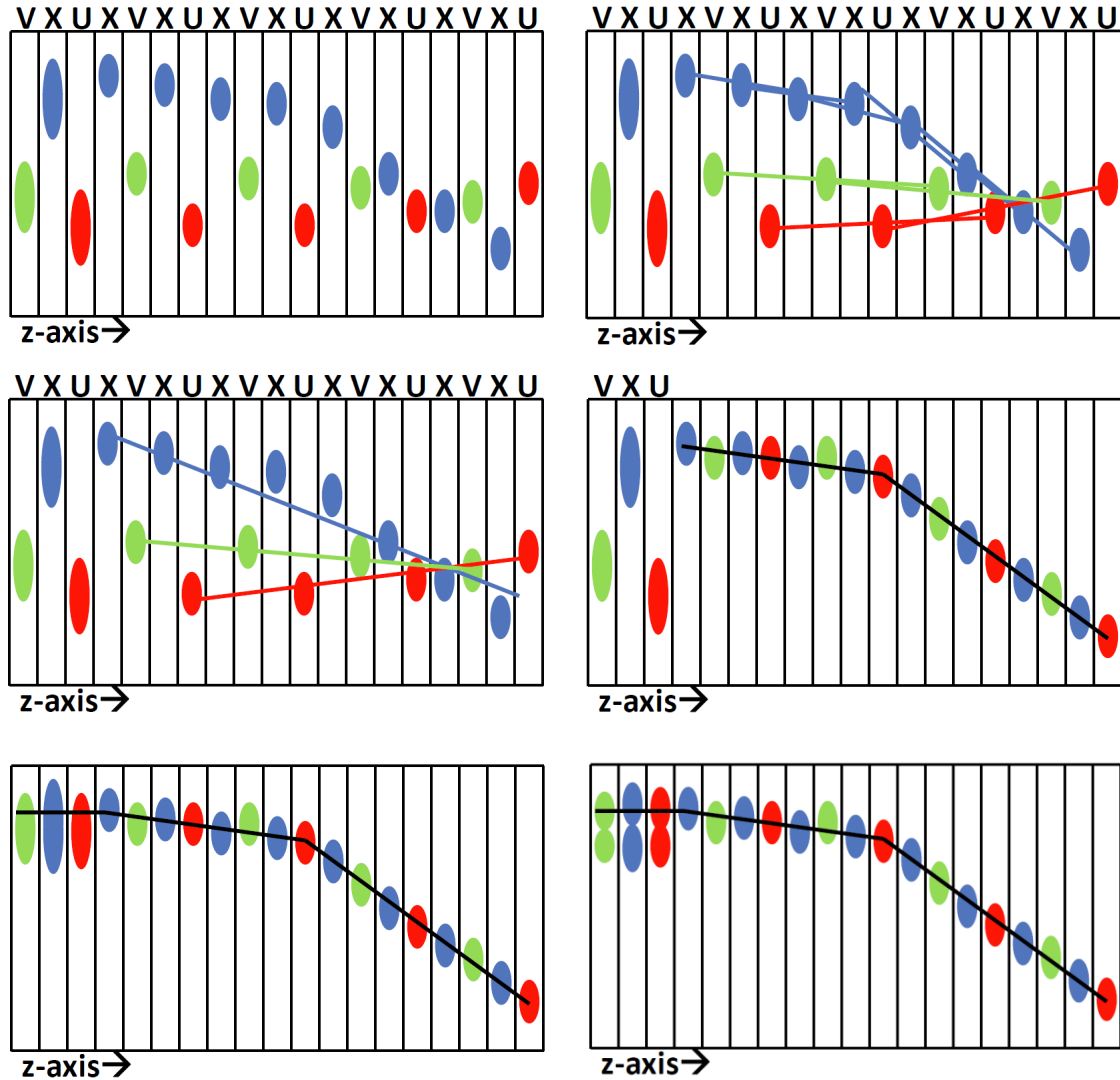


Figure 5.1: A cartoon example of the three-view long tracking. (Top left) Clusters are formed from hits in each plane. (Top right) Track seeds are formed in each view. Superclusters are not yet used. (Middle left) Seeds are merged into candidates. (Middle right) A three-dimensional track is formed from a candidate in each view. (Bottom left) The track is projected backwards, claiming the superclusters intersected by the projection. (Bottom right) Superclusters are broken by the master event building algorithm's cleaning procedure.

the same scintillator plane. If the algorithm merges two seeds into a candidate, it will attempt to merge additional seeds to the candidate using the same criteria. A seed may only be added to one candidate, so this stage of the algorithm is sensitive to the order of merging attempts.

After all candidates are built, they may be merged into larger candidates using criteria similar to those for seeds: the slopes and intercepts must be consistent and there cannot be two different clusters in the same scintillator plane. Merged candidates are not required to share clusters, allowing candidates to cross a scintillator plane without containing a cluster in that plane. This allows the candidates to follow particle trajectories that intersect dead regions in the detector.

The next step of the algorithm attempts to combine candidates from different views into three-dimensional track objects. This is attempted by two sub-algorithms: three-view tracking and two-view tracking. Both algorithms reject candidates that contain only one seed with a polar angle greater than 60° ; such candidates typically correspond to random energy deposits that happen to fit a straight line. Three-view tracking is unable to form tracks that intersect fewer than 11 planes, while the looser two-view tracking is able to form tracks that intersect as few as 9 planes.

Three-view tracking is attempted first. It examines all possible combinations of three candidates such that there is one candidate in each view. Such a combination of candidates are formed into a track if they fully overlap along the z-direction and are mutually consistent with the same three-dimensional line. The algorithm also searches for a particular topology in which the particle trajectory bends, but the bend is only visible in two views. In this instance, the combination of candidates will fail the overlap requirement because the candidate in the view without the bend will be longer. If this topology is detected, the longer candidate is broken into two shorter candidates.

Two-view tracking is run with all remaining candidates that were not used to form three-view tracks. It forms all possible pairs of two candidates in which the candidates are in different views. If the pair fully overlaps along z, they are used to construct a three-dimensional line. Then, the remaining view is searched for unused (not claimed by another track) clusters that have a position consistent with the pair of candidates. If a sufficient number of clusters are found in the remaining view, the two candidates are formed into a track. This algorithm is particularly useful for tracking particle trajectories that are obscured by detector activity in one of the three views.

All tracks are fit by a Kalman filter fit routine [67] that includes multiple scattering [68]. Each cluster on the track receives a three-dimensional position and direction from the fit. Figure 5.2

shows the tracking position resolution after the fit. The fit is used to add additional clusters to the track by searching nearby planes in which the track does not contain a cluster. Superclusters may be used here, allowing the track to project into a region of shower-like activity (as in the case of a muon track that is partially obscured by a hadron shower in all three views). Also, multiple tracks are allowed to claim the same cluster, each taking a fraction of the cluster energy. This is done to accomodate overlapping particles.

5.2.2 Short Track Pattern Recognition

Each cluster has a two-dimensional position: transverse (X, U, or V) and longitudinal (Z) coordinates. A three-dimensional line is completely specified by four parameters (X, Y, dX/dZ , dY/dZ). Therefore, a track in MINER ν A must contain at least four clusters: at least two X, and two of U or V⁵. In other words, a MINER ν A track must intersect at least four planes. This places a limit on the polar angle of the track⁶ and a species-dependent limit on the minimum trackable particle momentum. The long track pattern recognition cannot construct a track that intersects fewer than 9 planes, so a *short track* pattern recognition is employed to track lower-momentum particles. Table 5.1 summarizes the pion tracking thresholds for each pattern recognition.

The long tracking threshold is ultimately the result of using a two-dimensional track seed. The short track pattern recognition instead uses a three-dimensional seed constructed from four clusters in consecutive planes: two X clusters, one U, and one V. The clusters used by the short track pattern recognition are supplied by the master event building algorithm. All possible combinations of XUXV⁷ clusters are formed, requiring that the polar angle of the combination is less than 55°.

Next, the short tracking algorithm attempts to merge seeds into tracks. The merging requirements are that the two seeds share at least one cluster, have similar polar angles, fit well to a straight line, and pass a Kalman filter fit with a χ^2 per degree of freedom less than 10. The merged track may be merged with additional seeds or other merged tracks. Merging continues until all possible mergers are made. The remaining unmerged four-clusters seeds are not used by this analysis: every analyzed track must intersect at least five planes.

⁵Any combination of two X clusters and two U or V clusters are always consistent with a three-dimensional line, provided each cluster is in a different plane

⁶A 90° particle is contained in one plane.

⁷The combinations UXVX, XVXU, and VXUX are also considered.




Figure 5.2: Residual of fitted positions along a track relative to the measured cluster positions for a rock muon sample. The resolution is 3 mm. This figure is from Ref. [54]. Reprinted from [54], Copyright (2014) with permission from Elsevier.

Table 5.1: Pion tracking thresholds calculated for normal-incidence trajectories. These are estimated from a tabulation of muon range in polystyrene [69] and noting that particles with the same $\beta\gamma$ (p/m) have the same range.

Pattern Recognition	Planes	Range (mm)	Kinetic Energy (MeV)
Long three-view	11	170	85
Long two-view	9	136	80
Short	5	68	50

5.2.3 Event Building

The track pattern recognition schemes are used by a master event building algorithm designed to reconstruct high-multiplicity final states while also imposing restrictions that greatly reduce the number of low quality tracks that are found. In each step of the master algorithm, a unique combination of clusters is supplied to the pattern recognition. The full algorithm is as follows:

1. Run the long track pattern recognition with all trackable and heavy ionizing clusters in the current timeslice. Select the track that has the longest projection along the z-axis. This track is called the *anchor* track and its clusters are marked as used. All other tracks are deleted. If there are no tracks or the anchor contains fewer than 25 clusters, do not proceed further.
2. Rerun the long track pattern recognition using all unused trackable and heavy ionizing clusters. Use the startpoint of the anchor track to estimate the primary vertex and select tracks that are consistent with emerging from the primary vertex. Use these tracks to fit the position of the primary vertex via a Kalman filter technique [66]. Delete any tracks that are inconsistent with the primary vertex. Mark the clusters on the retained tracks as used.
3. Repeat step 2, but do not check for consistency with the primary vertex. Instead, iteratively check the end vertex of each track connected to the primary vertex. This finds tracks for particle trajectories that abruptly change direction due to scattering in the detector.
4. Repeat steps 2 and 3, but use the short track pattern recognition. Also, search for additional short tracks at the end vertex of every short track that is found.

Figure 5.3 illustrates this procedure with a series of event displays. Figures 5.4 and 5.5 show long tracking efficiencies for muons and pions. Pions often scatter, absorb, or charge exchange through interactions in the detector. These interactions decrease pion tracking efficiency relative to muons.

Each track found with the long track pattern recognition is submitted to a cleaning procedure before searching for additional tracks. The purpose of the cleaning procedure is to remove energy that is likely to be unrelated to the tracked particle. This is accomplished by breaking the offending cluster into two clusters. One is kept by the track and the other is marked as unused so that it may be used to find additional tracks. The cleaning procedure varies according to whether the track is the anchor track.




Figure 5.3: (Top left) V-view event display of a neutrino interaction candidate in data. The vertical axis is strip number and the horizontal axis is module number. Each triangle is a strip and the color indicates the energy (gray and yellow are lower energy, red and blue are higher energy). (Top right) In step one, the master event building algorithm finds the anchor track. (Bottom left) In step two, additional tracks are found at the primary vertex. Step three searches for long tracks at the end of the primary long tracks but doesn't find any. (Bottom right) The short track pattern recognition is run in step four and finds a track at the end of one of the long tracks.




Figure 5.4: Muon tracking efficiency as a function of momentum (left) and angle with respect to Z-axis (right), predicted by simulation. The shape of the efficiencies is primarily a function of the number of planes intersected by the muon trajectory. Backwards muons are tracked less often because their momentum is usually less than 1 GeV/c. Reprinted from [54], Copyright (2014) with permission from Elsevier.

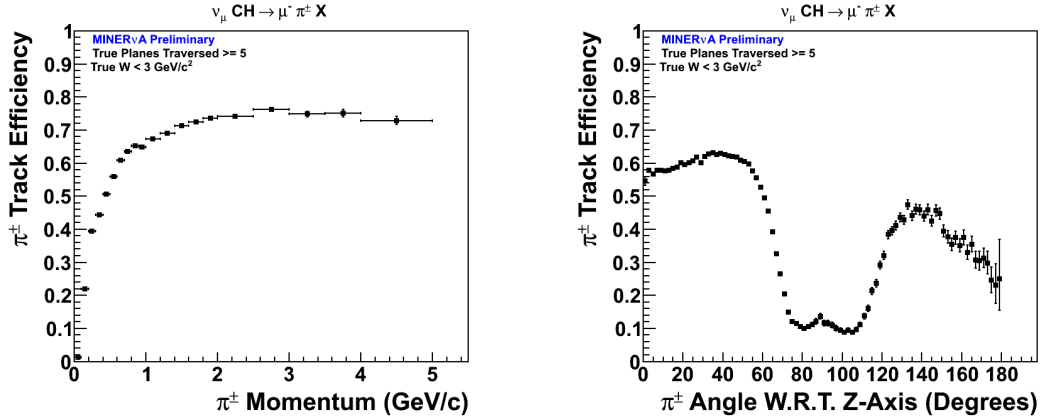


Figure 5.5: Pion tracking efficiency as a function of momentum (left) and angle with respect to Z-axis (right), predicted by simulation. Pion tracking is not as efficient as muon tracking because pions often scatter in the detector, complicating reconstruction.




Figure 5.6: Simulated muon (left) and pion (right) angular resolution, defined as reconstructed angle minus true angle. The distributions are fit to the sum of two Gaussians. The pion resolution is a bit worse than the muon resolution due to scattering in the detector.

- **Anchor Track** – The anchor track is almost always a minimumly ionizing muon if the neutrino interaction was charged current. Any energy near the primary vertex that is inconsistent with a minimumly ionizing particle (MIP) is removed from the track. Superclusters are always broken, with the anchor track retaining a portion of the supercluster consistent with a MIP.
- **Non-anchored Track** – This is usually a proton or pion. All superclusters are broken, with the track only claiming energy from the hits in the supercluster that intersect with the fitted track trajectory.. If the supercluster is near the track's endpoint, the track will take as much energy as possible in order to not disturb the expected rise in the energy loss profile. Otherwise, the track will claim an energy equal to its mean cluster energy.

The track fits are used to reconstruct the angle of the tracked particle with respect to the neutrino beam. The muon angle resolution is 0.6° and the pion angle resolution is 3.3° . The pion resolution is worse due to pion-nucleus scattering in the detector. Figure 5.6 shows the resolution distribution for samples of simulated muons and pions. The long tails in the pion distribution are caused by particularly severe pion scattering.

5.3 MUON RECONSTRUCTION

This analysis studies charged current ν_μ interactions, which are tagged by the presence of a muon. In MINER ν A, for neutrino interactions in the tracker, the muon can be sufficiently identified by a track that exits the detector; it is extremely rare for a charged hadron to travel through the ECAL and HCAL without stopping or interacting. However, MINER ν A is not magnetized, so it cannot measure the momentum or charge of exiting muons. The charge is needed to reject $\bar{\nu}_\mu$ interactions, and the momentum is needed to reconstruct E_ν . These quantities are measured by requiring that the muon is reconstructed in the MINOS near detector. For neutrino interactions in the tracker, this limits acceptance to muons with momentum (E_μ) greater than 1.5 GeV/c and angle with respect to the NuMI beam (θ_μ) less than 25° .

5.3.1 MINOS Track-Matching

First, an exiting muon track in MINER ν A must be matched to a reconstructed muon track in MINOS. The MINER ν A track is required to contain a cluster in one of the five most downstream MINER ν A planes, and the MINOS track must contain a cluster in one of the four most upstream MINOS planes. Additionally, the tracks must occur within 200 ns of each other.

The matching is done in two ways. First, the MINOS track is projected to the last active MINER ν A plane, and the MINER ν A track is projected to the first active MINOS plane. The difference between the projection and the other track's activity in the plane must be less than 40 cm in both cases in order to make a match. If this fails to yield a match, then the point of closest approach along the projection is used. This latter criterion is helpful in cases where the muon scattered in the passive material between the two detectors.

5.3.2 Muon Charge and Momentum

The muon charge is determined by the deflection of the matched MINOS track in the MINOS magnetic field. The direction of the deflection indicates the sign of the charge. The MINOS muon momentum can be measured in two ways. The first method uses a track fitting algorithm that

relates the momentum to the track curvature using the equation

$$p_\mu = 0.3BR, \quad (5.1)$$

where p_μ is the muon momentum where it enters MINOS (MeV/c), B is the MINOS magnetic field (kGauss), and R is the radius of curvature (cm). The second method calculates the momentum directly from the range of the track, but this can only be done for lower energy muons contained in MINOS. The range method is used on contained muons, and the curvature method on uncontained muons.

To calculate the muon momentum at the neutrino interaction point, the reconstructed muon energy in MINOS is added to the energy lost by the muon in the passive material between the detectors and in the material intersected by the MINERνA muon track. This energy is calculated using the Bethe-Bloch equation [70]

$$-\frac{dE}{dx} = KZ^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 T_{\max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right]. \quad (5.2)$$

The parameter definitions are the same as Ref. [70]. Figure 5.7 shows the muon momentum resolution for a sample of simulated muons. The resolution is $\sim 4\%$.

5.4 HADRON PARTICLE IDENTIFICATION

Charged hadrons are identified as the non-muon tracks at the primary vertex. The fitted track provides the angle of the hadron with respect to the neutrino beam, but additional reconstruction is needed to determine the hadron species and its momentum. This analysis uses two particle identification algorithms to identify charged pions and reconstruct their energy. The first algorithm separates pions from protons using the energy loss (dE/dx) profile of the track. The second algorithm searches for the Michel electron that is the result of the consecutive decays $\pi^+ \rightarrow \mu^+ \nu_\mu$ and $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$.




Figure 5.7: Simulated muon momentum fractional resolution, defined as $(p_{reco} - p_{true})/p_{true}$. The distribution is fit to the sum of two Gaussians.

5.4.1 dE/dx Particle Identification

The rate that charged particles lose energy in matter is well-determined by the Bethe-Bloch equation (Eq. 5.2). Particles of equal $\beta\gamma$ (p/m) lose energy at equal rates⁸. This is used to identify particle species by comparing the rate of energy loss along the track to the track's range. Figure 5.8 demonstrates this by comparing the energy loss profile of a simulated proton and pion to the profiles predicted by the Bethe-Bloch equation.

In practice, the separation is done by finding the best fit dE/dx profile for each track under pion and proton hypotheses. The fit is constrained to be consistent with the range of the track, although it is allowed to ignore the last cluster on the track or extend up to two planes beyond the end of the track without penalty to the fit. This is done to account for possible mis-reconstruction of the track end point location. Each fit returns the best fit momentum under its particle hypothesis and calculates the χ^2 between the measured and fitted dE/dx profiles. The χ^2 are used to construct

⁸There is also a dependence on the particle mass m , but this is negligible for typical hadron energies in MINERνA




Figure 5.8: Simulated pion (left) and proton (right) dE/dx profile compared to Bethe-Bloch (solid curves). Reprinted from [54], Copyright (2014) with permission from Elsevier.

a pion score s_π

$$s_\pi = 1 - \frac{\chi_{\pi,DOF}^2}{\sqrt{\chi_{\pi,DOF}^4 + \chi_{p,DOF}^4}} \quad (5.3)$$

that quantifies whether the track is better fit as a pion or proton. $\chi_{\pi,DOF}^2$ is the pion fit χ^2 per degree of freedom and $\chi_{p,DOF}^2$ is the proton fit χ^2 per degree of freedom. Sec. 7.1.4 demonstrates how s_π is used to select pions in this analysis. The best fit pion momentum is used as the reconstructed pion momentum. Figure 5.9 shows the pion kinetic energy resolution. The resolution is $\sim 10\%$ in the peak region. The low side tail is composed of pions that scatter, absorb, or charge-exchange in the detector, which generally reduces the pion's visible range and causes the kinetic energy to be under-reconstructed. These interactions also affect the particle identification of the pion track by altering the measured dE/dx profile from the expected Bethe-Bloch shape.

The track cleaning procedure described in Sec. 5.2.3 does not attempt to remove extra vertex energy from hadron tracks. This energy can bias the dE/dx fit, so a simple algorithm is used to remove that energy from consideration. The algorithm searches for sudden large decreases in the rate of energy loss near the beginning of the track, indicating the end of overlapping vertex activity. All clusters before the rate decrease are ignored by the dE/dx fit.




Figure 5.9: Simulated pion kinetic energy fractional resolution, defined as $(T_{reco} - T_{true})/T_{true}$. The low side tail are pions where the reconstructed kinetic energy is too small due to interactions in the detector.

5.4.2 Michel Electron Particle Identification

Michel electrons are found by searching later time slices for clusters that are spatially consistent with the end of the hadron track. The search is done separately in each view within a 35×25 cm² (transverse \times longitudinal) box centered on the track's end position. If clusters are found in multiple views, they are combined to form a three dimensional Michel candidate.

The search box is fairly wide and will often return clusters from unrelated neutrino-induced activity later in the gate, so a few quality control selections are applied. The total energy of the Michel clusters must be less than 55 MeV and the total number of hits contained in the clusters must be less than 35. These selections are motivated by the well-understood kinematic limit imposed on the Michel electron. Additionally, Michels in only one view must originate from a time slice that contains less than 100 MeV total energy. Figure 5.10 compares the reconstructed Michel energy of data and simulated samples selected by this analysis. The simulated spectrum is slightly lower energy than the data spectrum, but this is consistent within the uncertainty on the detector response to electromagnetic particles.

5.5 NEUTRINO ENERGY RECONSTRUCTION

Ideally, the neutrino energy E_ν can be reconstructed by individually measuring and summing the energy of each particle in the final state. This is not practical for most neutrino interactions in MINER ν A, however, because neutral particles do not directly produce scintillation light. Instead, for charged current events, E_ν is calculated as the sum of the reconstructed muon plus the calorimetric sum of all non-muon (hadronic) energy seen in the detector. The sum ignores cross talk clusters. This method introduces some neutrino interaction and final state interaction model-dependence to the measurement because it requires a prediction of the mean amount of energy carried by neutral particles in the recoil system.

First, each hadronic energy deposit is scaled to account for the undetected energy deposited in passive material. This is simply the expected energy deposit in the active and passive material




Figure 5.10: Reconstructed Michel electron energy in data and simulation. The particle species refers to the true primary particle associated with the Michel. A “stopping” particle is one that does not interact hadronically with the detector. The simulation is area-normalized to the data to compare shape.

Table 5.2: Passive material corrections.

	Tracker	ECAL	HCAL
C^{sd}	1.222	2.013	10.314

divided by the visible energy, or

$$C^{sd} = \frac{E_{abs}^{sd} + \frac{1}{f}E_{vis}}{E_{vis}}, \quad (5.4)$$

where C^{sd} is the correction factor for a particular sub-detector, E_{abs}^{sd} is the expected energy loss in one absorber plane in the sub-detector, $f = 0.8185$ is the fraction of active material in a scintillator plane, and E_{vis} is the measured energy. E_{abs}^{sd} is calculated by assuming that the energy deposited per g/cm² in scintillator is also deposited in the neighboring absorber plane. Table 5.2 summarizes the values of C^{sd} that are most important for this analysis.

Next, the hadronic recoil energy, defined as $E_{recoil} \equiv E_\nu - E_\mu$, is calculated to be

$$E_{recoil} = \beta \left(\alpha \sum_i C_i^{sd} E_i \right), \quad (5.5)$$

where the sum is over sub-detectors, E_i is the total visible energy measured in the sub-detector, and α and β are model-dependent scales that account for undetected neutral particles and charged particles that are not contained in the detector. α , determined to be 1.568, is calculated by minimizing the difference between the sum term and the true value of E_{recoil} . β is an energy-dependent scale that increases agreement with true simulated E_{recoil} . The fractional resolution on reconstructed E_{recoil} is shown in Figure 5.11. Resolution is worse at lower energy because the missing energy from a single neutral or uncontained particle is a larger fraction of the total hadronic energy. More detail is included in Ref. [54].




Figure 5.11: Calorimetric energy resolution $\sigma/E_{\text{recoil}}^{\text{true}}$ for ν_{μ} interactions. Reprinted from [54], Copyright (2014) with permission from Elsevier.

6.0 NEUTRINO SIMULATION

MINER ν A uses the GENIE neutrino event generator [23], version 2.6.4, to simulate neutrino interactions. GENIE generates neutrino interactions in a geometrical description of the MINER ν A detector (Sec. 3.1) with the simulated NuMI flux (Sec. 2.3). The output of the GENIE simulation is the species and four-momentum of each particle seen in the final state, as well as a record of the interaction. The output is fed into the detector simulation (Sec. 3.3) to simulate detector response. The simulation is used to predict backgrounds, estimate detector efficiencies, and evaluate systematic uncertainties. This chapter briefly describes how GENIE simulates neutrino-nucleus interactions.

6.1 NEUTRINO CROSS SECTIONS

GENIE first simulates the underlying neutrino-nucleon or neutrino-quark interaction, then modifies the interaction using models of various nuclear effects. The relevant neutrino interactions for this thesis are the ν_μ charged current versions of the following¹:

- **Quasi-elastic Scattering** – The processes $\nu_\mu n \rightarrow \mu^- p$ and $\bar{\nu}_\mu p \rightarrow \mu^+ n$. The former is a source of background in this analysis when the final state proton is mis-identified as a pion.
- **Resonance Production** – The neutrino interacts with a nucleon and excites it to a Δ or N^* resonance. This is signal when the baryon resonance decays to a pion and a charged pion escapes the nucleus. Higher mass resonances can produce multiple pions.

¹There are essentially no neutral current interactions selected by this analysis.

- **Deep Inelastic Scattering** – The neutrino interacts with a single quark. The quark is ejected from the nucleon and hadronizes. This is signal if the hadronization yields a charged pion.
- **Coherent Pion Production** – The neutrino interacts with the nucleus such that a pion is produced but the nucleus is unchanged by the interaction. This is another signal channel.

The simulation determines whether an interaction occurs by using the total interaction cross section $\sigma(E_\nu)$, which is tuned to data. If an interaction occurs, the type of interaction is determined randomly according to branching fractions. The kinematics of the interaction are decided by the models.

6.1.1 Quasi-elastic Scattering

The Llewellyn-Smith formalism [71] is used to model quasi-elastic scattering. The hadronic current is not calculated from a dynamical quark model; instead, it is parameterized in terms of the most general Lorentz-invariant form factors. GENIE models the form factors using experimental results. Two vector form factors are related to electromagnetic form factors, measured in electron scattering experiments, by using the conserved vector current hypothesis. GENIE uses the BBBA2005 parametrization [72] of the electromagnetic form factors. The form of the axial form factor $F_A(Q^2)$ ² is not well-known, but its value at $Q^2 = 0$ is measured to be -1.267 by neutron decay measurements; GENIE models $F_A(Q^2)$ as a dipole:

$$F_A(Q^2) = \frac{F_A(0)}{(1 + Q^2/M_A^2)^2}. \quad (6.1)$$

The parameter M_A , called the axial mass, is measured by numerous bubble chamber experiments and is set to $0.99 \text{ GeV}/c^2$. The pseudo-scalar form factor is related to F_A by the partially conserved axial current hypothesis [73]; this relationship is confirmed by experiment [74].

6.1.2 Resonance Production

GENIE uses the Rein-Sehgal model as described in Sec. 1.3.1 but makes a few changes. Two resonances calculated by Rein-Sehgal are omitted by GENIE because the PDG [70] considers

²Recall that $Q^2 = -q^2$, where q is the four momentum of the virtual boson.

them to be ambiguous. GENIE also does not allow neighboring resonances to interfere, and does not include lepton mass terms in the differential cross section calculation. The σ terms from Eq. 1.8 contain axial and vector form factors, analogous to quasi-elastic scattering. The axial form factor is taken to be a dipole with axial mass $M_A^{RES}=1.12 \text{ GeV}/c^2$.

GENIE simulates the decay of Δ s produced by resonance interactions. The pion angular distribution $W_\pi(\cos\theta)$ in Δ decay has the form

$$W_\pi(\cos\theta) = 1 - p(\frac{3}{2})P_2(\cos\theta) + p(\frac{1}{2})P_2(\cos\theta), \quad (6.2)$$

where θ is the pion angle with respect to the Δ angular momentum quantization axis in the Δ rest frame, P_2 is the second order Legendre polynomial and $p(\frac{3}{2})$ and $p(\frac{1}{2})$ are coefficients for the $\frac{3}{2}$ and $\frac{1}{2}$ angular momentum states. GENIE incorrectly simulates an isotropic decay, where $p(\frac{3}{2}) = p(\frac{1}{2}) = 0.5$, but provides event weights that reweight to the anisotropic decay predicted by Rein and Sehgal ($p(\frac{3}{2}) = 0.75$ and $p(\frac{1}{2}) = 0.25$). This analysis uses an anisotropic decay distribution that is “half way” between isotropic and the Rein-Sehgal anisotropic model by modifying the GENIE weights according to

$$w_{ana} = 1 + \frac{w_{RS} - 1}{2}, \quad (6.3)$$

in which w_{ana} is the weight used to analyze data and w_{RS} is the weight provided by GENIE to reproduce the Rein-Sehgal prediction for anisotropy.

6.1.3 Deep Inelastic Scattering

GENIE models all non-resonant inelastic interactions as deep inelastic scattering using a procedure by Bodek and Yang [75]. The Bodek-Yang model extends deep inelastic scattering to low Q^2 and E_ν such that it can be used to predict the inelastic background to resonance production. This is somewhat dangerous, as the Bodek-Yang model double counts some interactions predicted by Rein-Sehgal. GENIE avoids double-counting by first limiting the Rein-Sehgal model to hadronic invariant mass $W_{cut} < 1.7 \text{ GeV}/c^2$. The Bodek-Yang prediction is reduced for $W_{cut} < 1.7 \text{ GeV}/c^2$ to agree with measurements of one-pion production, two-pion production, and total cross section measurements.

Table 6.1: Meson production probabilities in KNO scaling.

Meson Species	Probability
π^\pm	62.66%
π^0	31.33%
$K^{\pm,0}$	6%

Struck quarks are hadronized according to the so-called AGKY model [76]. AGKY uses PYTHIA/JETSET [17] at $W > 3 \text{ GeV}/c^2$ and KNO scaling [77] at $W < 2.3 \text{ GeV}/c^2$. The transition region is handled by linearly increasing the probability of using PYTHIA/JETSET from 2.3 to $3 \text{ GeV}/c^2$, while simultaneously decreasing the probability of using KNO scaling. Hadronization at low W is particularly relevant for the analysis in this thesis because it determines the hadronization of quarks, producing pions, in the analysis signal region. At low W , the struck quark is usually hadronized to one baryon and any number of π and K that are kinematically allowed. The average charged hadron multiplicity is calculated from an empirical expression tuned to data. KNO scaling determines the actual multiplicity via its prediction for the multiplicity dispersion. The meson particle species are determined according to the probabilities in Table 6.1.

6.1.4 Coherent Pion Production

Coherent pion production is simulated using another model from Rein and Sehgal [78] modified with lepton mass terms [79]. The partially conserved axial current hypothesis relates this process to pion-nucleus scattering at $Q^2=0$, then assumes a dipole $F_A(Q^2)$ with an axial mass of $1 \text{ GeV}/c^2$ for non-zero Q^2 . The pion-nucleus scattering cross sections are determined with data.

6.2 NUCLEAR EFFECTS

GENIE simulates the nucleus as a relativistic Fermi gas with a Bodek-Ritchie high-momentum tail, as described in Sec. 1.3.2. Pauli-blocking is applied to quasi-elastic and elastic scattering by requiring that the outgoing nucleon have momentum greater than the Fermi momentum; resonance baryon production is not subject to Pauli-blocking. The deep inelastic scattering structure functions are modified to account for shadowing, anti-shadowing, and the EMC effect by the Bodek-Yang procedure. Final state interactions are simulated using π^+p scattering data on ^{56}Fe according to the procedure described in Sec. 1.3.3.

7.0 MUON-NEUTRINO CHARGED CURRENT CHARGED PION PRODUCTION IN PLASTIC SCINTILLATOR

This chapter describes an analysis of ν_μ charged current charged pion production in the MINER ν A scintillator tracker. The goal is to measure the shape of the differential cross sections with respect to pion kinetic energy T_π and angle $\theta_{\pi\nu}$ in order to study nuclear final state interaction (FSI) strength in the target. The analysis signal is defined as

$$\nu_\mu + A \rightarrow \mu^- + N\pi^\pm + X, \quad (7.1)$$

where A is a nucleus in the tracker (composition defined in Tables 3.1 and 3.2), the \rightarrow symbol represents the primary neutrino-nucleon interaction and subsequent FSI, N is the number of charged pions, and X is the remaining hadron and nuclear recoil. Charged current coherent pion interactions are signal; for such interactions, $X = A$, $N = 1$, and the charged pion is a π^+ . The neutrino energy E_ν is restricted to $1.5 \text{ GeV} < E_\nu < 10 \text{ GeV}$ because of poor muon acceptance below 1.5 GeV and highly uncertain neutrino flux knowledge above 10 GeV.

The analysis is subdivided into two analyses: one-pion and N-pion analyses. The one-pion analysis is Δ resonance-dominated and is more readily compared to theoretical calculations, while the N-pion analysis is less model dependent and includes higher order nucleon resonances. They are partly distinguished by W , the total invariant mass of the hadronic recoil. In this analysis, W is calculated before FSI and does not include correlated nucleons¹. The equation for W is

$$W^2 = m_n^2 + (p_\nu - p_\mu)^2 + 2p_n \cdot (p_\nu - p_\mu). \quad (7.2)$$

¹Nuclear effects make the definition of an invariant hadronic mass somewhat ambiguous. Recall from Sec. 1.3.2 that meson exchange currents include another nucleon in the initial neutrino interaction. Also, FSI can impart energy to nucleons without ejecting them from the nucleus, thus leaving the nucleus in an excited state. For simplicity, W is defined without considering these effects.

The four-vectors p_ν , p_μ , and p_n correspond to the incoming neutrino, outgoing muon, and incoming nucleon, respectively, while m_n is the effective mass of the incoming nucleon. The one-pion analysis requires that there is exactly one charged pion in the final state ($N = 1$) and $W < 1.4$ GeV/c². The N-pion analysis allows for one or more charged pions in the final state ($N > 0$) and requires $W < 1.8$ GeV/c². In the interest of brevity, this chapter will generally only show plots and figures from the 1-pion analysis. The corresponding N-pion figures are included in Appendix B.

7.1 EVENT SELECTION

Events are selected from 2.99×10^{20} protons on target of ν_μ -mode² data collected by MINER ν A between 2010-2012. The event selection for each analysis has four components. The first is the primary vertex selection that constrains the target nucleus to be within the scintillator tracker. The second component is the μ^- identification, which tags charged current ν_μ interactions. The third component is a kinematical selection that constrains E_ν and W to fall within the limits imposed by the signal definition. The final component is π^\pm identification.

7.1.1 Vertex Selection

The reconstructed primary vertex is required to be within detector modules 27-79, while the full tracker spans modules 23-84. This selection creates a sufficient buffer between the tracker fiducial volume, upstream nuclear targets, and downstream ECAL, eliminating events that originate in the passive nuclear targets or the ECAL lead. Additionally, the transverse position of the vertex is required to fall within a regular hexagonal area of apothem 85 cm that is centered on the z-axis. This is done to remove events that originate in the lead of the side ECAL, which has an inner apothem of 90 cm. The vertex position reconstruction is not perfect and otherwise fine signal events with a true interaction vertex just outside of the fiducial volume are regularly included by this selection. The fiducial volume of true signal events is defined to be slightly larger, within

²In ν_μ -mode, the NuMI focusing magnets focus positively-charged mesons, resulting in a ν_μ -dominated neutrino beam.

modules 25-81 in a regular hexagon with apothem 86.5 cm, in order to avoid labeling such events as background.

Occasionally, multiple neutrino interactions occur in the detector at approximately the same time, creating multiple primary vertices³. In these cases, the time slicing algorithm is unable to resolve the interactions into separate physics events. When this occurs, this analysis chooses the primary vertex with a MINOS-matched track as the interaction vertex, ignoring the other interactions. If multiple primary vertices contain MINOS-matched tracks, then the analysis chooses the vertex with the largest number of outgoing tracks. Any remaining ambiguity is resolved by selecting the most upstream primary vertex.

7.1.2 Muon Selection

A muon is identified as a single track that exits the downstream end of MINER ν A and is matched to a reconstructed track in the MINOS detector. The MINOS-matching requirements are described in Sec. 5.3.1. This analysis requires one MINOS-matched track attached to the interaction vertex. This is an extremely pure ($>99\%$) selection of charged current events because it is very unlikely for any other particle to traverse the entire downstream ECAL and HCAL without scattering or losing all of its kinetic energy. The single exception is a neutral current interaction that yields a high energy pion that promptly decays in flight to a muon, but this is extremely rare. Muons that exit MINER ν A and are not matched to MINOS are not used because their momentum cannot be reconstructed. Also, it is very difficult to differentiate muons and pions that stop in MINER ν A, so MINER ν A-contained muons are also excluded from the analysis.

In addition to the track selection, this analysis requires that the reconstructed muon charge significance, q/p , is less than zero. This selects negatively-charged muons, corresponding to a charged current ν_μ interaction. Figure 7.1 shows the reconstructed q/p distribution.

The MINOS-match requirement imposes strict limits on the muon acceptance, resulting in a charged current event selection efficiency of about 40%. MINOS is unable to track muons with momentum less than 150 MeV/c [80] and the muon must have enough energy to traverse the MINER ν A tracker and downstream calorimeters without stopping. This places a lower limit on

³ $\sim 0.5\%$ of time slices contain multiple neutrino interactions.




Figure 7.1: The reconstructed q/p distribution. Events below zero contain a muon with negative charge. $q/p > 0$ is primarily background from anti-neutrino interactions (in gray).

the accepted muon momentum of about 1-1.5 GeV/c. This analysis requires a signal event to have a minimum neutrino energy of 1.5 GeV, which implies a minimum muon momentum with adequate acceptance. The MINOS-match requirement also restricts the accepted muon angle because transverse muons tend to exit the sides of MINERνA and not enter MINOS. The maximum accepted angle is a function of the location of the interaction vertex in MINERνA and for this analysis is generally $\sim 20^\circ$. The muon angle is not restricted in the signal definition. This simplifies the signal definition for easy model comparisons, but also introduces some model dependence to the measured cross sections, which must be corrected for the missing high angle muons. This issue is discussed further in Chapter 8.

It is also necessary to impose a track quality cut to remove rock muon background. Rock muons can be selected by the analysis when the upstream portion of the rock muon track is masked by detector dead time. This causes the rock muon to appear to be a muon from a neutrino interaction in the detector fiducial volume. Rock muon background is removed by using the track fit to project the muon track upstream and count the number of dead channels that intersect the projection. The analysis requires that there are fewer than two dead channels in the upstream projection.




Figure 7.2: The reconstructed E_ν distribution. The contamination from $E_\nu > 10$ GeV is shown in the light blue fill and is greatly reduced by the selection.

7.1.3 Kinematic Constraints

The one-pion (N-pion) analysis signal is restricted to $1.5 \text{ GeV} < E_\nu < 10 \text{ GeV}$ and $W < 1.4 \text{ GeV}/c^2$ ($W < 1.8 \text{ GeV}/c^2$). The E_ν selection is made by requiring that reconstructed E_ν is between 1.5-10 GeV. E_ν is reconstructed according to

$$E_\nu = E_\mu + E_{recoil}, \quad (7.3)$$

where E_μ is the muon energy reconstructed via the procedure described in Sec. 5.3.2 and E_{recoil} is hadronic recoil energy defined and calculated according to Sec. 5.5. Low E_ν background is already suppressed by muon acceptance and resolution on reconstructed E_ν is good for larger neutrino energies, so this selection is highly efficient and pure. The reconstructed E_ν distribution is shown in Figure 7.2.

As described by Equation 7.2, W is a quantity that describes the hadronic recoil at the primary interaction vertex before FSI and excluding correlated nucleons. As such, it is sensitive to the initial momentum of the struck nucleon and to modifications of the interaction by the mean nuclear

potential. W is therefore not a quantity accessible to direct experimental measurement. Instead, this analysis identifies a quantity W_{exp} that is closely related to W . To define W_{exp} it is first necessary to define Q^2 :

$$Q^2 = 2E_\nu(E_\mu - |\vec{p}_\mu| \cos(\theta_{\mu\nu})) - m_\mu^2. \quad (7.4)$$

W_{exp} is defined to be

$$W_{exp}^2 = m_p^2 - Q^2 + 2m_p E_{recoil}. \quad (7.5)$$

Equation 7.5 is identical to Equation 7.2 under the additional assumption that the struck nucleon is a free proton at rest. FSI do not play a role in differentiating W_{exp} and W because the quantity E_{recoil} is constructed to account for the simulation's FSI prediction. The reconstructed W_{exp} distribution is shown in Figure 7.3.

Simulation is used to model the difference between W_{exp} and W in order to evaluate the signal selection efficiency and background contamination of these cuts. A significant background at larger W remains in both analysis samples after this cut; section 7.3 discusses the procedure used to constrain the size of this background and subtract it from the T_π and $\theta_{\pi\nu}$ distributions.

7.1.4 Charged Pion Selection

Charged pions are identified by looking for additional, non-muon tracks at the primary vertex and processing the tracks with a couple of particle identification (PID) algorithms. The one-pion analysis requires that the reconstruction find one or two hadron track candidates at the event vertex, while the N-pion analysis requires that there is at least one hadron track candidate. The number of hadron track candidates is plotted in Figure 7.4. Additionally, at least one hadron track candidate must be fully contained in the inner detector, excluding the side ECAL, so that the PID algorithms and momentum reconstruction can be done reliably.

Most hadron track candidates are protons and charged pions, with rare matches to charged kaons, knock off deltas from the muon, and photons from prompt π^0 decays. Therefore, it is most important that the PID algorithms are able to separate charged pions from protons. The first PID algorithm uses the rate of energy loss along the track, dE/dx , to determine whether a track is




Figure 7.3: The reconstructed W_{exp} distribution. The $W_{exp} > 1.4$ GeV/c² background leaks into the signal region due to fluctuations in the calorimetric reconstruction of E_{recoil} . The entries at zero are events in which most of the hadronic recoil is neutral and not seen in the detector; these events are also removed by the analysis.




Figure 7.4: The number of reconstructed hadron tracks at the event vertex. The signal events that are removed by this selection are mostly events where the outgoing pion and nucleon are low energy and below tracking threshold.

more likely a proton or a pion. The dE/dx algorithm is the full procedure described in Sec. 5.4.1. This analysis requires that the pion score of the hadron track candidate, shown in Figure 7.5, is greater than 0.6. The cut value is tuned to maximize the product of the charged pion selection efficiency and purity. Approximately 80% of the tracks that pass this cut correspond to charged pions, while the remaining 20% are predominantly protons that interact in a way that confuses the PID algorithm. The best fit momentum returned by the dE/dx PID algorithm is used as the reconstructed pion momentum.

This analysis also requires that a Michel electron candidate is matched to the end of at least one reconstructed track that also passes the dE/dx PID selection. The Michel electron matching and selection is done according to the procedure described in Sec. 5.4.2. This selection is done for two reasons. First, it removes a majority of the remaining non-pion backgrounds – a proton would have to interact in the scintillator tracker and produce a trackable, decaying π^+ in order to be associated with a Michel electron. Second, the Michel selection removes tracks that are matched to pions that absorb, charge exchange, or otherwise interact within the scintillator in such




Figure 7.5: The pion score variable used to separate pion and proton tracks. The stopping categories correspond to particles that are contained in the detector and lose all of their energy through ionization and atomic excitation before stopping or decaying. Interacting pions and protons either interact with a detector nucleus or decay before losing all of their energy. Stopping pions and protons are well separated by this score. Separation is worse for interacting particles because their energy loss profiles do not match the Bethe-Bloch equation prediction. However, most interacting pions pass this selection because the score (Eq. 5.3) is designed to indicate whether a hadron is more likely a stopping pion or proton, rather than the likelihood that the particle is a stopping pion.




Figure 7.6: The number of detector views spanned by selected Michel candidates. The shape of this distribution is well-modeled by simulation.

a way that invalidates the momentum reconstruction. This reduces sensitivity to uncertainties in the simulation’s hadron interaction models when correcting from reconstructed to true pion kinetic energy. Figure 7.6 shows the selection.

The final analysis selection is specific to the one-pion analysis and simply requires that only one of the two allowed hadron track candidates pass the previously-described PID criteria. This removes some of the background from interactions that produce multiple pions. The full set of analysis selections yield 3474 event candidates in the one-pion analysis and 5410 event candidates in the N-pion analysis. Figure 7.7 shows the reconstructed T_π and $\theta_{\pi\nu}$ distributions for events passing the analysis selection. The error bars on the simulation reflect the large uncertainties on the flux measurement and the neutrino-nucleon cross sections predicted by GENIE⁴; the difference in normalization between the data and simulated distributions is likely due to the poor simulation of one of these.

⁴Recall from Sec. 1.4 that the experimental data for neutrino baryon resonance production on free nucleons does not provide strong constraints for models.

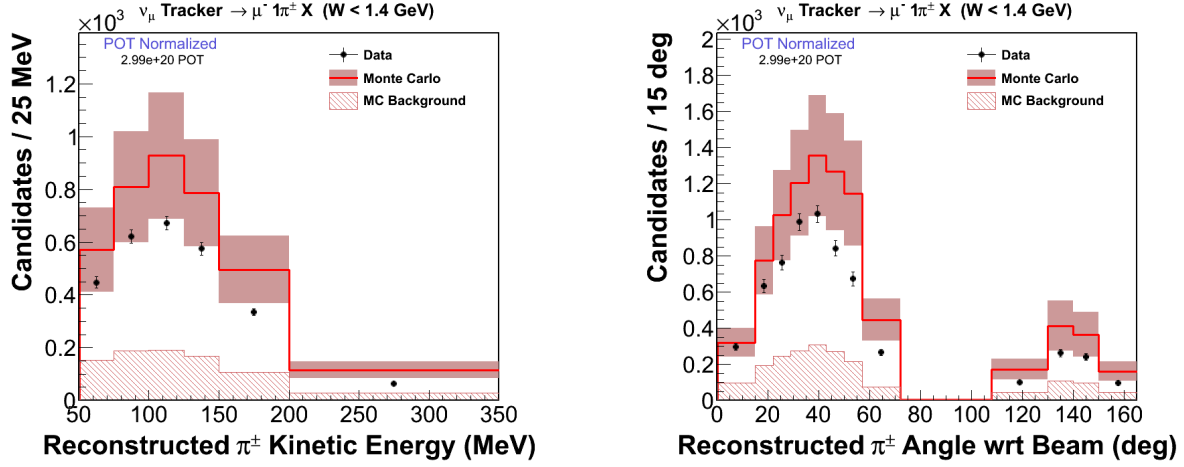


Figure 7.7: The one-pion reconstructed T_π and $\theta_{\pi\nu}$ distributions compared to simulation. The data error bars are statistical and the error bars on the simulation include the full set of systematic errors as described in section 7.7.

7.2 CROSS SECTION CALCULATION

The T_π flux-integrated differential cross section per nucleon, in bin i , is

$$\left(\frac{d\sigma}{dT_\pi}\right)_i = \frac{\sum_j U_{ij} (N_j - N_j^{bg})}{\epsilon_i T \Phi \Delta_i(T_\pi)}, \quad (7.6)$$

where j is the index of a reconstructed T_π bin, U_{ij} is an unfolding function that calculates the contribution to true bin i from reconstructed bin j , N_j is the number of selected events, N_j^{bg} is the predicted number of background events, ϵ_i is the signal event reconstruction efficiency, T is the number of tracker nucleons in the fiducial volume, Φ is the ν_μ flux prediction integrated between 1.5-10 GeV, and $\Delta_i(T_\pi)$ is the width of T_π bin i . The $\theta_{\pi\nu}$ flux-integrated differential cross section per nucleon is defined in a similar manner:

$$\left(\frac{d\sigma}{d\theta_{\pi\nu}}\right)_i = \frac{\sum_j U_{ij} (N_j - N_j^{bg})}{\epsilon_i T \Phi \Delta_i(\theta_{\pi\nu})}. \quad (7.7)$$




Figure 7.8: The one-pion simulated background distributions divided according to background type. The largest background category is events with true $W > 1.4 \text{ GeV}/c^2$.

The event selection yielding N_j was described in 7.1. The background estimation that determines N_j^{bg} is discussed in 7.3. The unfolding function U_{ij} is explained in 7.4 and the efficiency correction ϵ_i is described in 7.5. The normalization corrections T and Φ do not change the shape of the differential cross sections, but are briefly discussed in 7.6.

7.3 BACKGROUND SUBTRACTION

Figure 7.8 shows the components of the total simulated background predictions in Figure 7.7. Table 7.1 lists each background component as a fraction of the total selected sample. The rock muon background and backgrounds from neutrino activity originating in the outer detector and downstream calorimeters are not simulated. These are estimated from a visual scan of 1200 selected data events, finding one rock muon and one event originating in a calorimeter. This background is subtracted from data before subtracting the rest of the background.

The largest background in both versions of the analysis corresponds to events with large hadronic invariant mass W . Rather than use the simulation to directly predict the N_j^{bg} in Equations

7.6 and 7.7, data is used to adjust the simulated background prediction. This reduces sensitivity to uncertainties in GENIE’s cross section and FSI models, but increases sensitivity to uncertainties in the detector energy response and statistical fluctuations in data. Ultimately, the data-constrained procedure outline below reduces the total uncertainty on the background prediction from $\sim 32\%$ to $\sim 24\%$ in the one-pion analysis, with a similar reduction in uncertainty in the N-pion version.

The first step in the background subtraction procedure is to create data and simulated distributions of reconstructed W_{exp} after applying the entire analysis event selection except for the W_{exp} cut. The simulated distributions are then divided into components. The one-pion analysis distribution has three components: signal events, background events with $W < 1.7 \text{ GeV}/c^2$, and background events with $W > 1.7 \text{ GeV}/c^2$. The N-pion analysis distribution is divided into two components: signal and background. In both analyses, the background components include all possible sources of background except for those subtracted by the visual scan. The W_{exp} distributions, summed over all bins of T_π and $\theta_{\pi\nu}$, are shown in Figure 7.9.

The background and signal components are used as fit templates. The data distribution is fit to the sum of the signal and background templates, with the normalizations of the templates serving as the fit parameters. The fits are done in two ways: separately in each bin of T_π and $\theta_{\pi\nu}$, and summed over all bins. The fitting algorithm is a maximum likelihood technique using Poisson statistics described in Ref. [81]. Figure 7.10 shows the reconstructed W_{exp} distribution after fitting. The fit improves the shape agreement between data and simulation, though there is a small discrepancy near the peak of the distribution. This discrepancy is consistent with the systematic error on the hadron recoil energy reconstruction.

The fit returns the best relative normalizations for each signal and background template. The normalizations are used to construct weights w_{jk} :

$$w_{jk} = \frac{f_{jk}^{fit}}{f_{jk}^{sim}}, \quad (7.8)$$

where j is an index referring to a bin of T_π or $\theta_{\pi\nu}$, k is an index indicating a particular template, and f^{sim} and f^{fit} are, respectively, the simulated and fitted predictions for the fraction of events in bin j that correspond to template k . The index j is dropped for the version of the fit that sums over all bins of T_π and $\theta_{\pi\nu}$.

The predicted number of signal events $N_j - N_j^{bg}$ in bin j of the T_π or $\theta_{\pi\nu}$ distribution is

Table 7.1: This table summarizes the remaining background fractions after all event selections are applied. Background categories are list from top to bottom in order of decreasing precedence; events that qualify for more than one background category are only counted in the category that takes precedence. The Proton and Other Particle backgrounds refer to events where a proton or other particle are misidentified as a charged pion.

Background	Percentage of One-Pion Sample	Percentage of N-Pion Sample
Rock Muon	0.08	0.08
Outside Fiducial Volume	0.16	0.17
Not $\text{CC}\nu_\mu$	0.05	0.1
$W > 1.4 \text{ GeV}/c^2$ ($1.8 \text{ GeV}/c^2$)	16.7	6.05
$E_\nu > 10 \text{ GeV}$	0.45	0.84
Multiple Charged Pions	1.61	N/A
Proton	3.32	4.47
Other Particles	0.8	2.2
Total	23.2	13.9

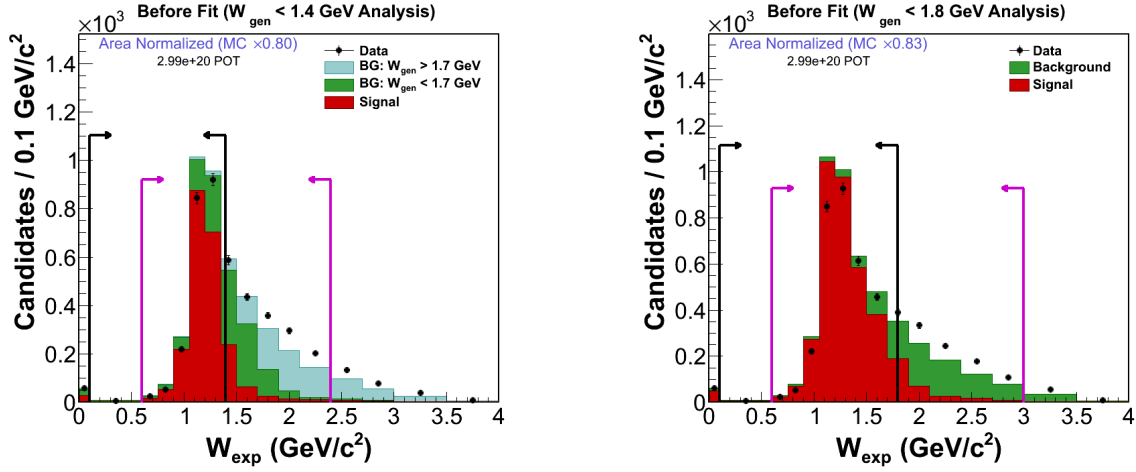


Figure 7.9: W_{exp} distributions summed over all bins of T_π and $\theta_{\pi\nu}$ and with all cuts except for the W_{exp} cut applied. The one-pion analysis is on the left and N-pion is on the right. The black arrows indicate the location of the analysis cut, and the magenta arrows indicate the fitted region. The sum of the signal and background components are fit to the data. The relative normalization of the components are the fit parameters.

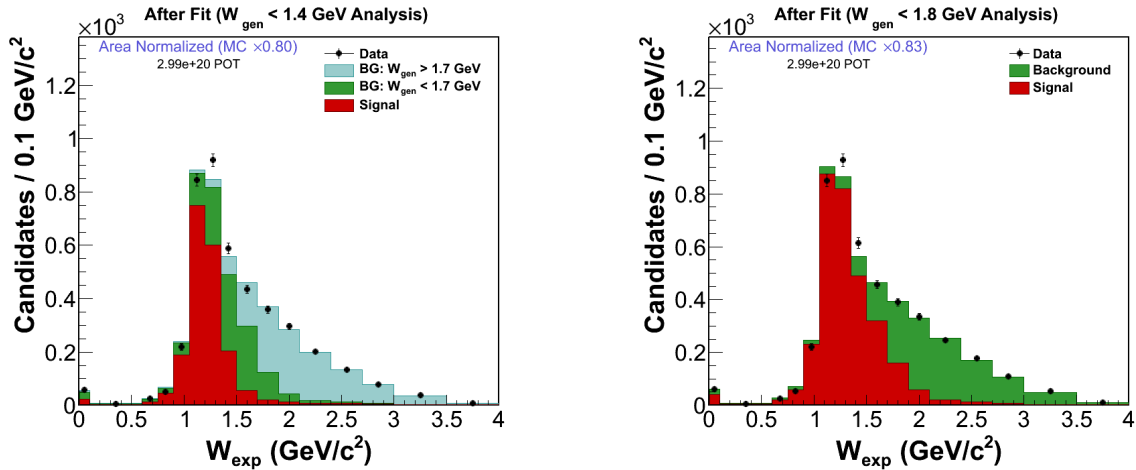


Figure 7.10: The W_{exp} distributions after fitting.

calculated to be

$$N_j - N_j^{bg} = \left(1 - \sum_{k \neq \text{signal}} \frac{N_{jk, \text{sim}}}{N_{j, \text{sim}}} r_{jk} \right) N_j. \quad (7.9)$$

r_{jk} is a function of w_{jk} or w_k , depending on the version of the fit; it modifies the nominal simulated fraction of template k events in bin j , denoted $\frac{N_{jk, \text{sim}}}{N_{j, \text{sim}}}$. The term in parentheses on the right hand side of Equation 7.9 can be understood as a sum over background templates in which each nominal background fraction is adjusted by r_{jk} . This sum is the *adjusted total background fraction* and is subtracted from one to calculate the *adjusted signal fraction*.

The r_{jk} do not equal the fit weights w_{jk} because the fit region includes events that are not selected by the analysis cut on W_{exp} (see Figure 7.9). The r_{jk} are determined by applying the w_{jk} to the events that pass all analysis cuts and calculating new fractions with the passed events, analogous to Equation 7.8. The result, after some algebra and simplification, is

$$r_{jk} = \frac{N_{j, \text{sim}} w_{jk}}{\sum_s w_{js} N_{js, \text{sim}}}, \quad (7.10)$$

where the sum in the denominator is over all templates using the dummy index s . If the fit is summed over all bins, w_{jk} is replaced with w_k in Equation 7.10. Two versions of r_{jk} , calculated from w_{jk} and w_k , are plotted in Figure 7.11. When calculated from the w_{jk} , the r_{jk} show dependence on T_π and $\theta_{\pi\nu}$ that is not consistent with statistical fluctuations. For this reason, only the bin-by-bin fits are used in this analysis.

Figure 7.12 shows the one-pion background predictions after applying the fit weights. The one-pion background-subtracted distributions $N_j - N_j^{bg}$ are plotted in figure 7.13. The corresponding N-pion distributions are in Appendix B.

7.4 DETECTOR RESOLUTION CORRECTION

Reconstructed quantities like T_π and $\theta_{\pi\nu}$ generally do not exactly equal the corresponding true quantity. This happens for a variety of reasons. The detector itself may be incapable of a perfectly precise measurement. For example, the amplification of photons by a PMT is a probabilistic pro-

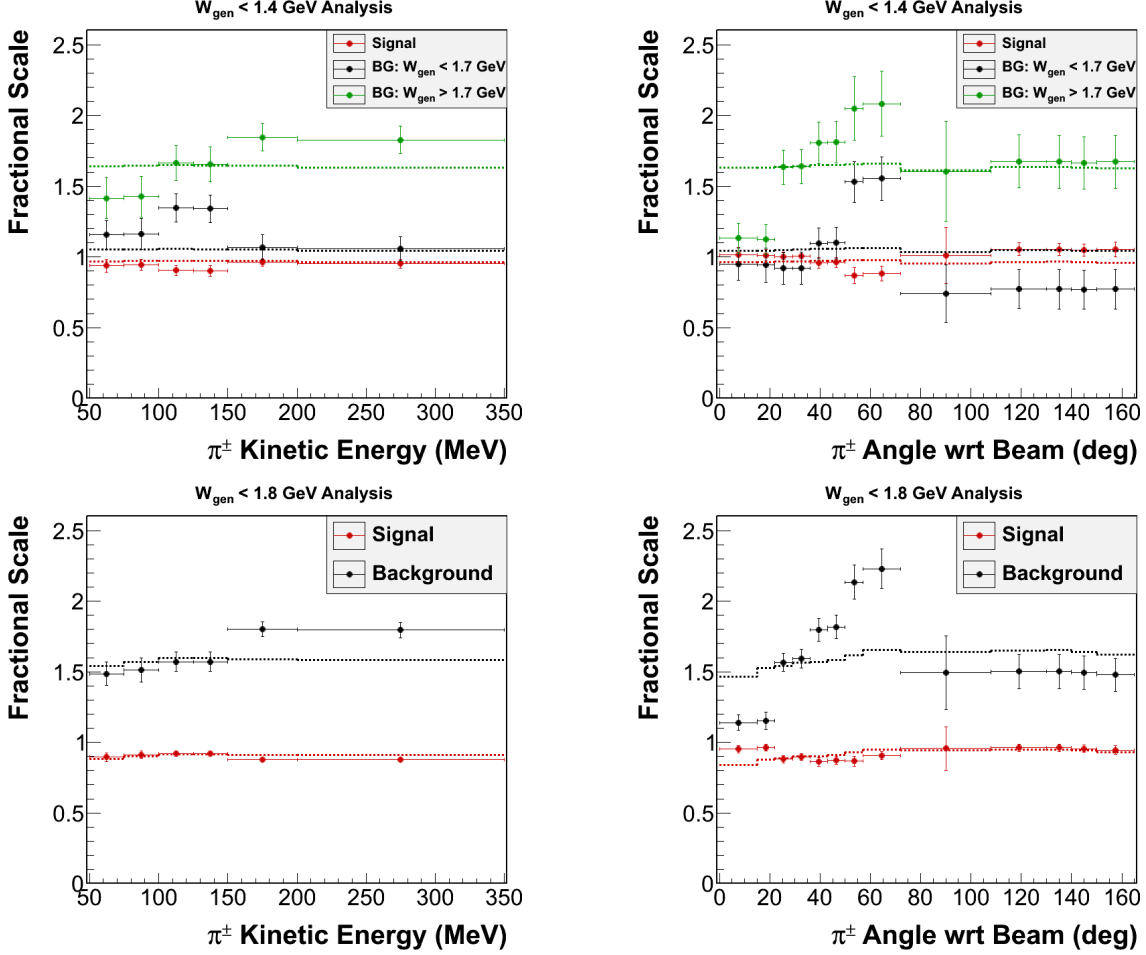


Figure 7.11: The background fit weights r_{jk} as functions of T_π and $\theta_{\pi\nu}$ for the one-pion (top) and N-pion (bottom) analyses. The points are calculated using bin-by-bin fits and include the fit statistical error. The dashed lines are calculated from a fit summed over all bins T_π and $\theta_{\pi\nu}$. In order to increase statistics, bins are paired together for the bin-by-bin fitting, resulting in the paired structure of the points.

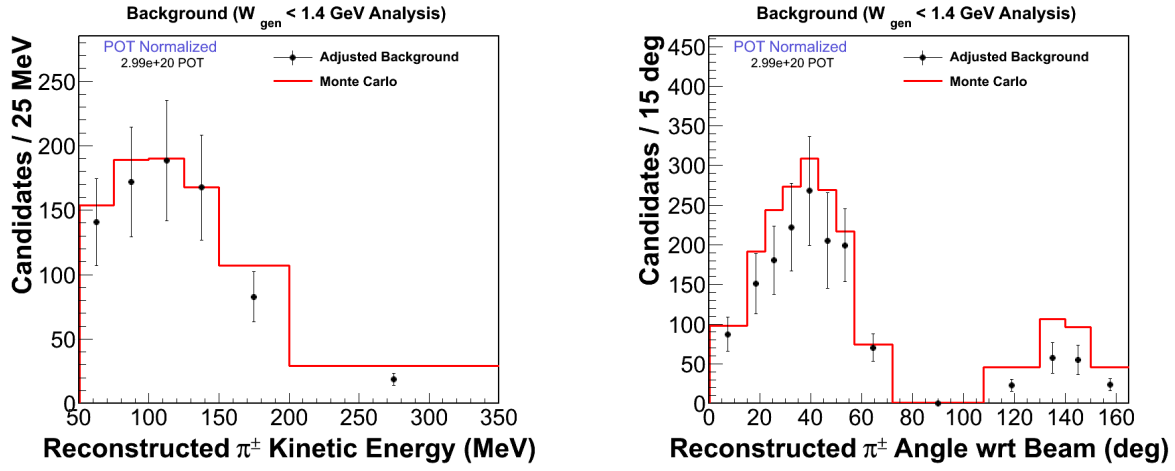


Figure 7.12: The one-pion adjusted background predictions compared to the nominal simulated predictions.

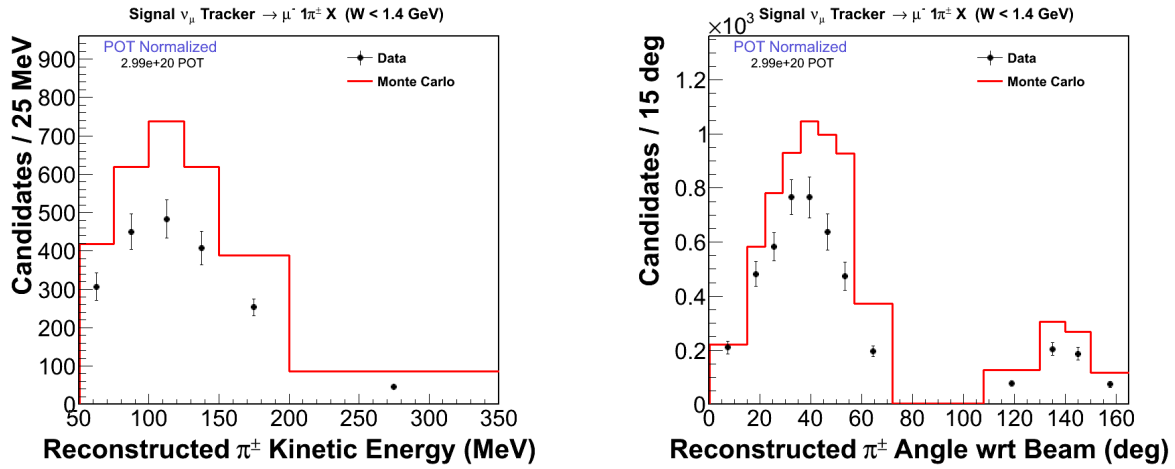


Figure 7.13: The one-pion background-subtracted distributions, $N_j - N_j^{bg}$, compared to simulation.

cess such that a single incoming photon results in a distribution of measured current at the PMT anode, thus broadening any energy measurement made with the PMT. Imperfections in a reconstruction algorithm can also smear a measurement, as would happen when the dE/dx particle identification procedure is used to reconstruct the momentum of a π^\pm that charge-exchanges to a π^0 . Ideally, the probability that a measurement is smeared by some amount is only a function of the detector and reconstruction design, remaining independent of the underlying models of neutrino interactions in nuclei. Unfortunately, this is not entirely true; for example, the energy deposited by low energy hadrons near the interaction vertex, which is predicted by GENIE, can affect a tracking algorithm's ability to reconstruct the direction of a high energy particle track that extends from that vertex.

In order to compare to theoretical predictions, the reconstructed distributions must be corrected to the underlying true distribution by removing all smearing effects. This process is often referred to as “unfolding.” This analysis unfolds the reconstructed T_π and $\theta_{\pi\nu}$ by using a Bayesian procedure [82] with four iterations. The input to the unfolding procedure is the background-subtracted data distribution and a migration matrix, shown in Figure 7.14. The migration matrix is derived from simulation and determines the probability that the true value of a quantity is smeared into a reconstructed value. It is fairly insensitive to FSI effects because the T_π and $\theta_{\pi\nu}$ distributions are corrected to the true kinematics at the point where the π^\pm exits the nucleus. Also, because it is constructed from probabilities, the migration matrix is not sensitive to the large differences in the resonance pion production normalization in data and simulation. This in turn minimizes the unfolding procedure's sensitivity to any normalization uncertainty.

The unfolded T_π and $\theta_{\pi\nu}$ distributions are shown in Figure 7.15. The unfolding generally migrates events from low to high T_π bins, accounting for the reconstruction's tendency to report a momentum that is too small for pions that interact inelastically in the scintillator. Unfolding has a small effect on the $\theta_{\pi\nu}$ distribution because the pion angular resolution is only a couple degrees. The events added to the zero-acceptance 90° bin correspond to initially transverse pions that scatter near the interaction vertex and are reconstructed in a different bin.

The off-diagonal elements in the upper left corner of the $\theta_{\pi\nu}$ migration matrix are true backwards pions that are mis-reconstructed as forwards. This can occur when the backwards pion is heavily-ionizing and nearly 180° relative to the muon. The tracking algorithm reconstructs the

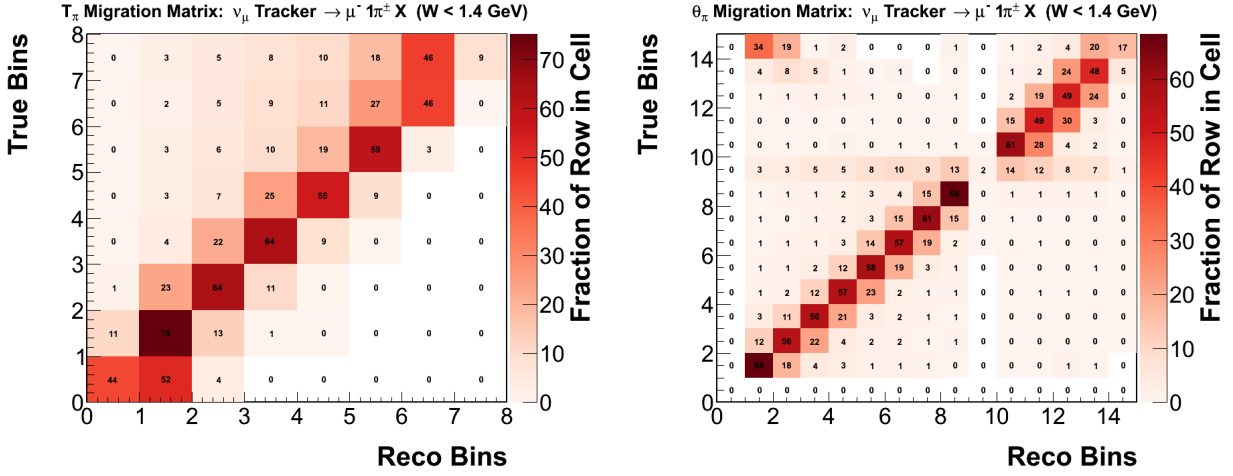


Figure 7.14: The one-pion migration matrices. The bin number indicates the bin number in the T_π or $\theta_{\pi\nu}$ distribution. Bin 0 is the underflow bin and the last bin is the overflow bin. Perfect detector resolution would correspond to a diagonal matrix; i.e., the reconstructed value equals the true value. The matrix is row-normalized - the contents of each row sum to one.

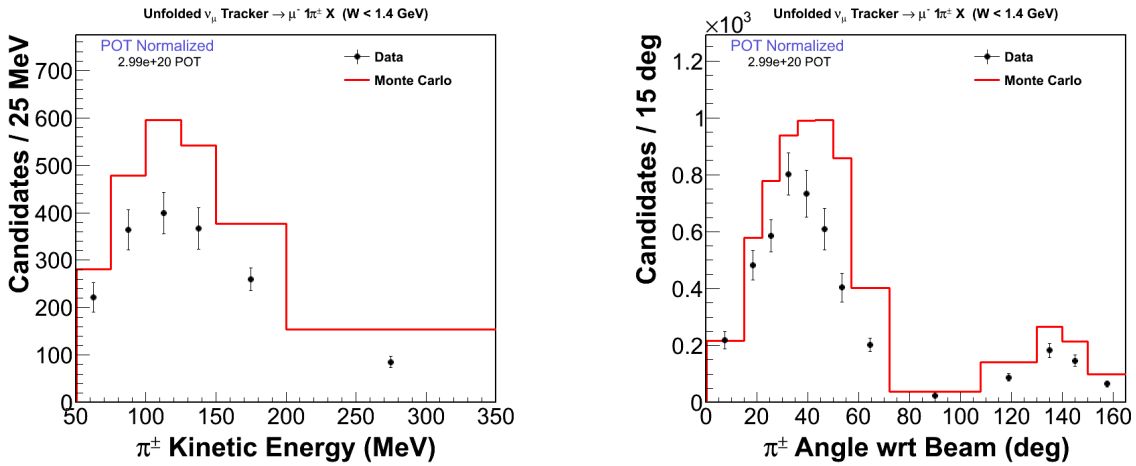


Figure 7.15: The one-pion unfolded T_π and $\theta_{\pi\nu}$ distributions.

muon and pion as one track, and the track cleaning procedure identifies the heavily-ionized portion of the track as a forwards pion overlapping the muon. Large off-diagonal elements can have a deleterious effect on the unfolding procedure, so an extra check is done to show that there is no problem. A new migration matrix is formed by moving the contents of bins (1,14) and (2,14) to bin (14,14), effectively making it so that the backwards pions are correctly reconstructed. The data is unfolded with the new migration matrix and compared to the nominal result. The difference is only a couple percent in bins 1 and 2, large in bin 14, and negligible elsewhere. In other words, the difference is only in the bins directly affected by the mis-reconstruction, as is expected for a well-behaved unfolding procedure.

7.5 EFFICIENCY CORRECTION

Only a fraction of signal channel events are detected and reconstructed. Furthermore, the analysis event selection rejects a portion of signal interactions. It is necessary to correct for this inefficiency in order to calculate a cross section. The efficiency ϵ_i to select an event in a true bin i is defined as

$$\epsilon_i = \frac{w_i N_i^S}{N_i^T}, \quad (7.11)$$

where N_i^S is the number of selected signal events predicted by simulation and N_i^T is the total number of signal events as predicted by simulation. w_i is a weight that can modify the efficiency, and is used to correct the simulated prediction of the efficiency using data-driven efficiency measurements. In general, w_i can vary from bin to bin, but this analysis only applies a flat correction. Figure 7.16 plots the efficiency as a function of T_π and $\theta_{\pi\nu}$.

This analysis uses data measurements to correct the simulated MINER ν A and MINOS muon tracking efficiencies. These efficiencies are sensitive to the beam intensity due to overlapping neutrino-induced detector activity that can obscure the muon track or confuse the tracking algorithm. Therefore, the efficiencies are separately measured for each data run period. w_i is a time-dependent product of the two corrections, and the correction factor is calculated by taking the ratio of the efficiency measured in data to the simulated efficiency.




Figure 7.16: The one-pion reconstruction efficiency as a function of T_{π} and $\theta_{\pi\nu}$. The error bar contains all systematic errors.

The MINER ν A tracking efficiency is derived from a data sample of muon tracks in MINOS that point back towards the MINER ν A fiducial volume. The study also required that the MINER ν A detector contained muon-like activity in order to increase the likelihood that the MINOS track was created from a muon that traveled inside MINER ν A. The MINER ν A muon tracking efficiency is then measured as the fraction of the events in this sample that have a reconstructed track in MINER ν A.

The MINOS track efficiency is measured in a similar manner using a sample of muon tracks in the MINER ν A fiducial volume that point towards MINOS. This sample is divided into two small samples, corresponding to greater than and less than 3 GeV/c, by selecting on the magnitude of the track's deflection in the downstream calorimeters due to multiple scattering. The MINOS tracking efficiency is measured in each subsample by calculating the fraction of events with a track in MINOS. The MINER ν A and MINOS efficiency corrections as a function of time are summarized in Table 7.2.

Other reconstruction efficiencies are measured in data, but they do not differ significantly from simulation. For example, Sec. 7.7.4 describes a study of rock muons used to measure the Michel electron detection efficiency. Also, the efficiency of the detector dead time selection described

in 7.1.2 is measured separately in data and simulation by dividing the number of events passing the selection by the number of events before the selection. The data and simulated dead time selection efficiency vary by only a few tenths of one percent. Remaining selection and reconstruction efficiencies are not explicitly measured in data, but their uncertainty is constrained by various systematic studies described in Sec. 7.7.

7.6 CROSS SECTION NORMALIZATION

A cross section is a rate-independent quantity and is derived from an unfolded and efficiency-corrected event distribution by dividing by the number of targets and the beam flux. Cross sections are often represented as a function of the incoming particle energy, but this analysis measures flux-integrated differential cross sections in order to remove dependence on the flux shape uncertainty. The total ν_μ flux prediction, shown in Fig. 2.7, is calculated by integrating from 1.5-10 GeV. The total number of target nucleons in the fiducial volume is calculated from a careful accounting of the materials used to construct each scintillator planes. Since both of these numbers are constant across bins of T_π and $\theta_{\pi\nu}$, they do not affect the central result of this thesis: measurement of the differential cross section shapes.

7.7 SYSTEMATIC ERRORS

Each step in the data analysis is sensitive to certain aspects of the simulation. For example, the simulation is used to predict the shapes of the background fit templates, the migration matrices, and the efficiency of the event selection. Unfortunately, the simulation is not perfect and many of its parameters are only known to within a finite uncertainty, denoted in this section as σ . This leads to a systematic error on the measured result.

Systematic errors are evaluated by shifting a parameter in the simulation within its uncertainty σ to produce a new simulated sample. The differential cross sections are then remeasured using each shifted universe. The nominal simulation is referred to as the central value universe and each

Table 7.2: Corrections to the muon reconstruction efficiency in MINER ν A and MINOS. This table is reprinted from [83], Copyright (2014) with permission from B. Tice.

Dataset (m/d/y)	Source of Efficiency Correction	Correction
<hr/>		
3/22/10–6/12/10 (high intensity)	MINER ν A tracking	0.980 ± 0.01
	MINOS tracking ($p_{\mu}^{MINOS} < 3 \text{ GeV}/c$)	0.936 ± 0.002
	MINOS tracking ($p_{\mu}^{MINOS} > 3 \text{ GeV}/c$)	0.982 ± 0.001
<hr/>		
3/22/10–6/12/10 (low intensity)	MINER ν A tracking	0.985 ± 0.01
	MINOS tracking ($p_{\mu}^{MINOS} < 3 \text{ GeV}/c$)	0.970 ± 0.007
	MINOS tracking ($p_{\mu}^{MINOS} > 3 \text{ GeV}/c$)	0.996 ± 0.002
<hr/>		
5/6/11–5/13/11	MINER ν A tracking	0.985 ± 0.01
	MINOS tracking ($p_{\mu}^{MINOS} < 3 \text{ GeV}/c$)	0.984 ± 0.009
	MINOS tracking ($p_{\mu}^{MINOS} > 3 \text{ GeV}/c$)	0.995 ± 0.002
<hr/>		
6/11/11–6/21/11	MINER ν A tracking	0.985 ± 0.01
	MINOS tracking ($p_{\mu}^{MINOS} < 3 \text{ GeV}/c$)	0.956 ± 0.006
	MINOS tracking ($p_{\mu}^{MINOS} > 3 \text{ GeV}/c$)	0.995 ± 0.002
<hr/>		
10/18/11–11/17/11	MINER ν A tracking	0.98 ± 0.01
	MINOS tracking ($p_{\mu}^{MINOS} < 3 \text{ GeV}/c$)	0.942 ± 0.004
	MINOS tracking ($p_{\mu}^{MINOS} > 3 \text{ GeV}/c$)	0.987 ± 0.001
<hr/>		
11/17/11–3/12/12	MINER ν A tracking	0.98 ± 0.01
	MINOS tracking ($p_{\mu}^{MINOS} < 3 \text{ GeV}/c$)	0.942 ± 0.002
	MINOS tracking ($p_{\mu}^{MINOS} > 3 \text{ GeV}/c$)	0.987 ± 0.001
<hr/>		
3/12/12 – 3/17/12	MINER ν A tracking	0.98 ± 0.01
	MINOS tracking ($p_{\mu}^{MINOS} < 3 \text{ GeV}/c$)	0.934 ± 0.006
	MINOS tracking ($p_{\mu}^{MINOS} > 3 \text{ GeV}/c$)	0.985 ± 0.002
<hr/>		
3/17/12 – 4/30/12	MINER ν A tracking	0.98 ± 0.01
	MINOS tracking ($p_{\mu}^{MINOS} < 3 \text{ GeV}/c$)	0.945 ± 0.002
	MINOS tracking ($p_{\mu}^{MINOS} > 3 \text{ GeV}/c$)	0.987 ± 0.001
<hr/>		

shifted simulation is a shifted universe. Each parameter may be shifted many times. If shifted once or twice, the $\pm 1\sigma$ uncertainties are used; otherwise, the shift is drawn from a Gaussian distribution with a width equal to σ . The covariance matrix that describes the systematic error from a single parameter is

$$C_{ij} = \frac{1}{N} \sum_n (x_{n,i} - u_i) (x_{n,j} - u_j), \quad (7.12)$$

where N is the total number of shifts, n is an index referring to a particular shifted universe, i and j indicate bins of the differential cross section, and $x_{n,i}$ is the measurement of the differential cross section in bin i using shifted universe n . The definition of u_i changes according to the number of shifted universes. If there is only one shift, then it is the value of the differential cross section in bin i predicted by the central value simulation. In the case of multiple shifts, u_i is the mean of the shifted universes in bin i . The total covariance matrix is the sum of the covariance matrices for each systematic uncertainty. The systematic errors shown in the results in Chapter 8 are equivalent to the square root of the diagonal entries in the total covariance matrix.

There are two general categories of systematic errors, vertical and lateral, which differ in how the shifted universes are produced. Vertical systematics are equivalent to the uncertainty on the probability that an event occurs. For example, a $\sigma = 1\%$ uncertainty on the Michel electron selection efficiency directly causes a 1% uncertainty on the number of Michel electrons in the final selected sample. Vertical shifted universes are created by weighting appropriate events according to the ratio

$$w = \frac{P_{shift}}{P_{CV}}. \quad (7.13)$$

P_{shift} is the probability of observing a certain event in the shifted universe and P_{CV} is the central value probability. In the example of the Michel electron selection efficiency, the +1% uncertainty corresponds to $w = 1.01$, and w is applied only to events with a selected Michel electron.

Lateral systematics change the value of an observed quantity and are typically associated with uncertainties on detector resolutions and calibration. Unlike vertical systematics, lateral systematics can move events from one measured bin to another and change whether an event passes an analysis cut. If the uncertainty is due to a reconstruction bias or scale uncertainty, *e.g.* uncertainty

on the detector energy scale, two shifted universes are produced by shifting the affected quantities symmetrically about σ . In the case that the uncertainty is on a detector resolution, then one shifted universe is produced in which the affected quantities are modified by adding a number drawn from a normal distribution with width equal to the uncertainty σ .

This analysis measures the absolutely-normalized differential cross sections and their shapes. The shape errors are calculated by normalizing the measurement in each shifted universe so that the integrated cross sections in the shifted and central value measurements are equal. The remaining differences between the central value and shifted measurements are the shape uncertainties.

The remainder of this section will describe how the σ for each systematic error is evaluated.

7.7.1 Beam Flux

The beam flux systematic is derived from uncertainties in hadron productions models, beamline modeling, and the data constraint technique described in Sec. 2.3. These three sources are mutually uncorrelated. The data constraint uncertainty is equal to the experimental errors of the NA49 measurement, with an additional negligible contribution [84] from the uncertainty on scaling the data to NuMI proton beam energy. The hadron production model uncertainty covers the uncertainty on the production of hadrons from the NuMI target for particles that are not constrained by data. Three models are used to predict this component of the hadron production spectrum: FTFP_BERT and QGSP_BERT as implemented in Geant4 [43], and Fluka[48]. The hadron production model uncertainty is the maximum spread between the three models.

The beamline uncertainty is measured by MINOS [85] and contains contributions from the alignment of the magnetic horn, horn current, target density, and protons-on-target counting. The size of these effects is summarized in Table 7.3. The errors are propagated to the flux measurement using the Fluka simulation. The beamline uncertainty is about 2% over most of the flux spectrum, but increases to close to 10% at the trailing edge of the focusing peak due to the horn current uncertainties. The total flux error and its components are summarized in Figure 7.17. This is one of the dominate systematics on the absolutely-normalized cross section measurements, but the uncertainty on the cross section shape is negligibly small.

There is one additional flux systematic: the uncertainty on the angle of the neutrino beam

Table 7.3: Uncertainties on the NuMI beamline model. The horn current distribution systematic is evaluated as the difference between simulating with a 6 mm and infinite skin depth (δ). This information is reprinted from [85], Copyright (2008) with permission from Ž. Pavlović.

Source	Uncertainty
Number of protons on target	2.0%
Horn transverse misalignment	1.0 mm
Horn tilt	0.2 mrad
Horn current miscalibration	1.0%
Horn current distribution	$\delta = 6 \text{ mm} / \delta = \infty$
Baffle scraping	0.25%
Misalignment of shielding blocks	1.0 cm
Target density	2.0%

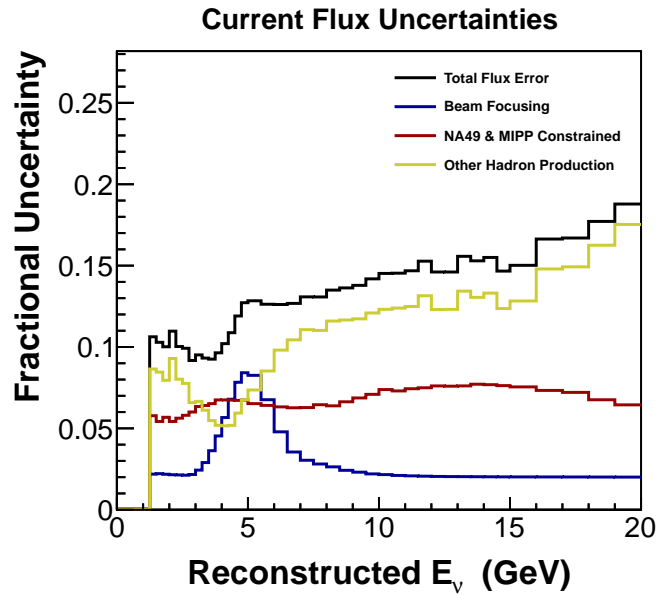


Figure 7.17: Flux systematic uncertainties as a function of neutrino energy.

relative to MINER ν A. This affects the measurement of the muon and pion angles with respect to the beam. The former is used to calculate Q^2 and W_{exp} and therefore can affect the efficiency of the kinematic event selections, while the latter is one of the quantities for which the differential cross section is measured. The uncertainty is found to be less than 6 mrad and is evaluated by measuring the angles of electrons in an electron-neutrino scattering sample and comparing to the well-understood theoretical calculation[86].

7.7.2 GENIE

This measurement uses the GENIE neutrino-nucleus event generator to simulate neutrino events in the MINER ν A detector. This introduces the unfortunate possibility that the measurement is sensitive to the same physics that it attempts to study, though the dependence is limited by restricting the measurement to the differential cross section shape. Model uncertainties are evaluated by using GENIE’s reweighting tools to shift the model parameters by $\pm 1\sigma$. The σ are determined by GENIE studies and are listed in [87], with the exception of the uncertainty on the anisotropy of the Δ resonance decay (see Sec. 6.1.2). For this uncertainty, the one-sigma errors are taken to be the half the difference between the isotropic and Rein-Sehgal anisotropic decay models. The total GENIE error is calculated by adding the component errors in quadrature, assuming that all parameters are uncorrelated. The GENIE authors consider this to be a safe assumption [88].

There are a few parameters in GENIE that cannot be reweighted. They control the effective size of the nucleus and quark hadronization time. These are studied by producing independent simulated samples, one for each change of a non-reweightable parameter, using a procedure similar to one used by MINOS[89]. In all cases, the deviations from the central value measurement are within statistical uncertainties, so no uncertainty is applied to the measurement for these parameters.

7.7.3 Geant4

Geant4 simulates the transport, energy deposition, and interaction of particles in the MINER ν A detector. In particular, uncertainties on the models used by Geant4 to simulate pion secondary interactions with detector nuclei can have a significant effect on the prediction of the tracking and particle identification efficiencies. Even small uncertainties can be important because Geant4

predicts that 70% of pions in MINER ν A experience at least one secondary interaction.

Charged pions can scatter elastically, resulting in a "kinked" trajectory that is more difficult to track. They can also interact inelastically, sometimes resulting in an untrackable shower of particles. Pions can also absorb or charge-exchange to the neutral π^0 ; if this occurs close to the primary vertex, then the particle trajectory before the interaction may be below tracking threshold. Each of these interactions also affects the performance of the particle ID algorithms, which assume that the pion has stopped and decayed in MINER ν A without experiencing a hadronic interaction.

The uncertainty on the Geant4 interaction model is estimated by modifying the pion and proton total inelastic cross sections by $\pm 10\%$, and modifying the ratio of the pion disappearance to pion scattering cross section by $\pm 10\%$. The 10% errors are determined by comparing Geant4 cross section predictions to data, and the cross sections are modified by a reweighting technique [90]. The systematic uncertainty is largest at high pion kinetic energy, about 7% (5% in shape). The uncertainty on the pion angle differential cross section is no more than 4% (2% in shape).

7.7.4 Michel Selection

The systematic uncertainty on the Michel selection efficiency is evaluated by comparing a data rock muon sample to simulation. The rock muons are required to stop in the tracker fiducial volume and be the only activity in their time slice. The difference between data and simulation in the efficiency to detect the decay electron is found to be 1%. The uncertainty on the false tag rate (background) is estimated by looking for Michels at random positions and times in both data and simulation. Any Michels found by the random search are background, and the uncertainty is the difference in the false Michel identification rate. This study finds that false Michels are found in data at 2.5 times the rate in simulation. This large difference is due to a flaw in the method used to overlay data on simulated events. The data overlay is only applied at the times in the gate spanned by detector activity, so unrelated detector activity at later times (when the algorithm searches for Michels) is often not simulated. Fortunately, the background prediction itself is very small, so this large uncertainty has only a couple percent effect on the measured differential cross sections.

7.7.5 Vertex Position and Particle Angles

The vertex position and particle angles are measured by the tracking and vertexing algorithm, specifically by the Kalman filter used to fit for the track trajectory. The uncertainty on the resolution of the track fit is determined by a study of rock muons that enter the front of MINER ν A and are matched to a track in MINOS. The tracks are broken at their midpoint along the z-axis in MINER ν A and each half is independently refit. The two fits provide measurements of the track position and angle at the point where it is broken. The differences in the two fits are compared in data and simulation to find the uncertainty on the resolution. The difference on the (x,y) position resolution is found to be (0.91,1.25) mm, and this is used as the vertex resolution uncertainty. The uncertainty on the angular resolution is 2 mrad, and this is used as the uncertainty on the muon angle resolution. The pion angle resolution uncertainty is conservatively set to 10 mrad because the pion angle resolution is five times worse than for muons⁵. All three of these uncertainties have a negligible few tenths of a percent effect on the measured differential cross sections.

7.7.6 Muon Tracking Efficiency

The MINER ν A and MINOS muon tracking and matching efficiencies are measured together in the study described in Sec. 7.5. The uncertainties are the statistical errors on the measurement and are tabulated in Table 7.2.

7.7.7 Hadron Tracking and Identification Efficiency

The tracking and dE/dx particle identification efficiencies of protons are measured using test beam detector data [91]. Protons are identified by the beamline reconstruction and are required to be contained in the detector by vetoing activity in the downstream half of the detector. Inelastic interactions are avoided by limiting the number of scintillator planes with energy deposits that are inconsistent with a minimum ionizing proton. The tracking efficiency of all protons passing this selection is measured to differ by less than 1% between data and simulation, though the disagreement increases up to 5% for proton trajectories that cross fewer than 8 planes.

⁵Compare the two distributions in Figure 5.4.

This proton sample is also used to measure the rate at which protons are misidentified as pions by the dE/dx particle identification algorithm. The data misidentification rate is measured to be 60% greater than the simulated rate. The source of this difference is not yet determined, so this analysis conservatively applies the 60% difference as an uncertainty on the proton background fraction. Fortunately, the proton background is extremely small, so this large systematic uncertainty does not significantly affect the measurement.

7.7.8 MINOS Muon Momentum

The MINOS momentum reconstruction by range has an uncertainty of 2%, measured by MINOS [13]. The uncertainty on the MINOS momentum reconstruction by track curvature is measured using a high-statistics sample of rock muons that enter the front face of MINOS and stop in the fully instrumented region of the detector. The momentum is reconstructed by range and curvature, and the distribution of the difference $(1/P_{curv} - 1/P_{range})$ is formed, see Figure 7.18. The momentum by curvature systematic is the difference in the mean values of this distribution in data and simulation. It is added in quadrature with the 2% range uncertainty to find a 2.1% (3.1%) momentum uncertainty for muons that enter MINOS at a momentum above (below) 1 GeV/c.

7.7.9 Detector Mass

The uncertainties on the detector mass, shown in Table 7.4, affects the dE/dx particle identification algorithm, the muon and pion energy reconstruction, and the calculation of the number of target nucleons. The effect on the muon energy reconstruction is estimated to be 11 MeV on average for muons originating in the tracker and exiting the back of MINERνA. This is determined by changing the detector mass within the uncertainties and recalculating the energy lost by the muon in MINERνA.

The effect on the dE/dx particle identification and pion energy reconstruction is estimated individually for each hadron track candidate by propagating the mass uncertainties into the calculation of the best fit energy loss profile. The target nucleon number uncertainty is equivalent to the 1.4% scintillator plane mass uncertainty. It is a flat, fully-correlated uncertainty across all bins of T_π and $\theta_{\pi\nu}$, so it does not contribute an error on the shapes of the differential cross sections.




Figure 7.18: (Left) Residual $P_{curv} - P_{range}$ distribution of all CC inclusive contained MINOS μ^- events. (Right) Residual $1/P_{curv} - 1/P_{range}$ distribution for the same events. Reprinted from Ref. [54], Copyright (2014) with permission from Elsevier.

Table 7.4: MINER ν A detector mass uncertainties.

Material	Uncertainty
Scintillator Plane	1.4%
Lead (ECAL)	1.3%
Iron (HCAL)	0.5%

7.7.10 Energy Loss Model

The energy loss model, *i.e.* the MINER ν A implementation of the Bethe-Bloch equation, is used by the dE/dx particle identification algorithm and the muon and pion energy reconstruction. The uncertainty is estimated by comparing Bethe-Bloch predictions to the Groom muon range tables [69]. The differences are found to be $\sim 1\%$ for pure materials and $\sim 3\%$ for mixtures. The contribution to muon energy uncertainty is estimated to be 30 MeV on average. The effect on dE/dx particle identification and pion energy reconstruction is evaluated separately for each hadron track candidate.

7.7.11 Energy Response

The hadronic recoil energy E_{recoil} , which is used by the kinematic event selections, is sensitive to the uncertainty on the measured energies of hits in the detector. The uncertainty can vary according to the particle that deposited the energy and can come from numerous sources: Birks' parameter, the detector energy scale calibration, and the Geant4 interaction model, to name a few. The total uncertainty on the hadronic recoil energy is evaluated for each event by dividing E_{recoil} into particle-specific components, evaluating the uncertainty on each component, then summing the uncertainties in quadrature. Table 7.5 summarizes the fractional uncertainty of each particle component.

7.7.11.1 Charged Pions, Kaons, and Protons The energy response uncertainty for π^\pm is evaluated with test beam detector data [92]. The fraction of observed energy to total pion energy is measured as a function of total pion energy, shown in Figure 7.19. After evaluating systematic errors, it is determined that an additional 5% uncertainty is needed to account for the difference between data and simulation. A similar study for protons [93] measures a 6% uncertainty. There are not enough charged kaons in the test beam data to repeat this study, so it is assumed that the charge kaon energy response uncertainty is the same as for charged pions.

7.7.11.2 Neutrons Neutrons only deposit energy in the detector if they scatter and produce a charged particle, so the only important source of uncertainty on the neutron energy response is the

Table 7.5: Detector energy response uncertainties.

Source	Uncertainty
Charged pions and kaons	5%
Protons	6%
Neutron	10%
γ, π^0, e^\pm	3%
Cross talk	20%
Other	10 %




Figure 7.19: Fraction of pion energy seen in the test beam detector vs. pion energy. This figure is from Ref. [92].




Figure 7.20: Michel electron energy distribution compared to simulation. The difference in the means is 3%. Reprinted from Ref. [54], Copyright (2014) with permission from Elsevier.

Geant4 neutron scattering model. This uncertainty is evaluated by comparing the Geant4 prediction for the cross section of the process $nA \rightarrow pX$ to external measurements [94]. The simulated cross section is found to differ by as much as 25% when compared to data. This is translated to an uncertainty on the neutron energy response by adjusting the neutron mean free path by 25% using a reweighting procedure [95], then finding the average change in the neutron component of the hadronic recoil energy for a sample of inclusive charged current neutrino interactions. This results in a 10% uncertainty on the neutron energy response.

7.7.11.3 Electromagnetic γ , π^0 , and e^\pm particles deposit energy in the detector through electromagnetic showers. The uncertainty on this process is evaluated by comparing the mean Michel electron energy in data and simulation, shown in Figure 7.20. The means are found to differ by 3%. This uncertainty is also applied to the Michel energy quality cut described in Sec. 5.4.2.

7.7.11.4 Cross Talk Cross talk is not removed from the hadronic recoil system with 100% efficiency. The uncertainty on the amount of cross talk energy that remains in the recoil is determined

by comparing two independent measurements of PMT cross talk. One of the measurements is the *in situ* procedure described in Sec. 4.4. The other procedure measured cross talk by injecting PMTs with a light source on a test stand. The two methods disagree by 20%, and this is taken to be the uncertainty on the cross talk component of the hadronic recoil.

7.7.11.5 Other Particles If a particle is not assigned an energy response uncertainty by any of the procedures discussed above, it is given a default 10% uncertainty. This the neutron response uncertainty, chosen because it is the largest of the non-cross talk energy response uncertainties.

7.7.12 Birks Constant

Birks' law is given by Equation 3.1. Birks' parameter k_B is measured in the test beam detector [62] to be 0.0905 ± 0.012 mm/MeV. This differs significantly from the value used in the detector simulation, 0.133 mm/MeV, and this difference must be accounted for as a systematic uncertainty. The procedures used to estimate the energy response uncertainties already account for data-simulation differences in k_B ; therefore, only the effect on the dE/dx particle identification and pion energy reconstruction must be evaluated.

To evaluate the remaining uncertainty, the energy of each hit on the hadron track candidate is adjusted to be consistent with the measured value of k_B [96]. The track's modified energy loss profile is then used by the dE/dx particle identification procedure to determine the change in pion score and reconstructed pion energy. Each hit energy is modified by adding the $O(1)$ correction to the energy deposit for a shift δ in k_B , calculated by the Taylor expansion of Birks' Law:

$$\Delta E_{hit} = -\frac{E_{hit}^2}{x_{hit}}\delta + O(2). \quad (7.14)$$

x_{hit} is the scintillator strip path length of the particle that responsible for the energy E_{hit} of the hit. An additional factor of $(2.25 \text{ MeV/cm})^2 x_{hit} \delta$ is added to avoid changing the detector energy scale, since the energy scale uncertainty is accounted for by the mass and energy loss model uncertainties. δ is chosen to be -0.0545 mm/MeV, setting k_B to one standard deviation below the measured value.

7.7.13 Calorimetry Versions

The calorimetric constants α and β in Equation 5.5 were determined by a different version of the simulation than the one used in this thesis. When applied to the thesis version of the simulation, the reconstructed value of E_{recoil} is shifted by $\sim 2\%$ ($\sim 12\%$) above (below) 300 MeV, as shown in Figure 7.21. The shift is taken as an additional uncertainty on the measurement of E_{recoil} .

7.7.14 Unfolding

The error on the unfolding procedure is the difference between iterating four and five times. The difference is found to be less than 2% in all bins.




Figure 7.21: E_{recoil} systematic offset as a function of true E_{recoil} .

8.0 RESULTS AND CONCLUSIONS

8.1 RESULTS

Figure 8.1 shows the ν_μ charged current π^\pm production differential cross sections with respect to T_π and $\theta_{\pi\nu}$. The distributions are compared to the nominal GENIE prediction, using an isotropic Δ resonance decay model¹, with and without final state interactions (FSI). The χ^2 per degree of freedom ($\chi^2/\text{D.O.F.}$) is calculated, comparing both versions of GENIE to each measurement, using the data covariance matrices to account for bin-to-bin correlations in the uncertainties. The $\chi^2/\text{D.O.F.}$ are listed in Table 8.1. In every comparison of the data to GENIE, data favors GENIE with FSI included. It is also evident that the simulation over-predicts the total cross section by $\sim 30\%$, likely due to the simulation's highly uncertain models for the flux and neutrino baryon resonance production.

The statistical and systematic errors are summarized in Figure 8.2. There are many more error sources than can be practically shown on one plot, so similar errors are grouped together. Table 8.2 lists the component errors in each group. The largest sources of uncertainty are the beam flux, the detector energy response, and the GENIE model. The flux uncertainty is $\sim 9\%$ with very little fluctuation from bin to bin. The GENIE uncertainty is $\sim 5\text{--}13\%$ and is generally largest at low pion kinetic energy and small pion angles. The detector energy response uncertainty is larger in the one-pion analysis because this analysis selects lower multiplicity final states, which increases sensitivity to the simulation's model of the hadronic system and its calorimetric energy response.

In the one-pion analysis, the largest GENIE uncertainties are from the vector (4–6%) and axial

¹Recall that the central value measurement is made using an anisotropic Δ decay model. The measurement is compared to the nominal GENIE prediction for simplicity and consistency with other model comparisons shown later in this chapter.

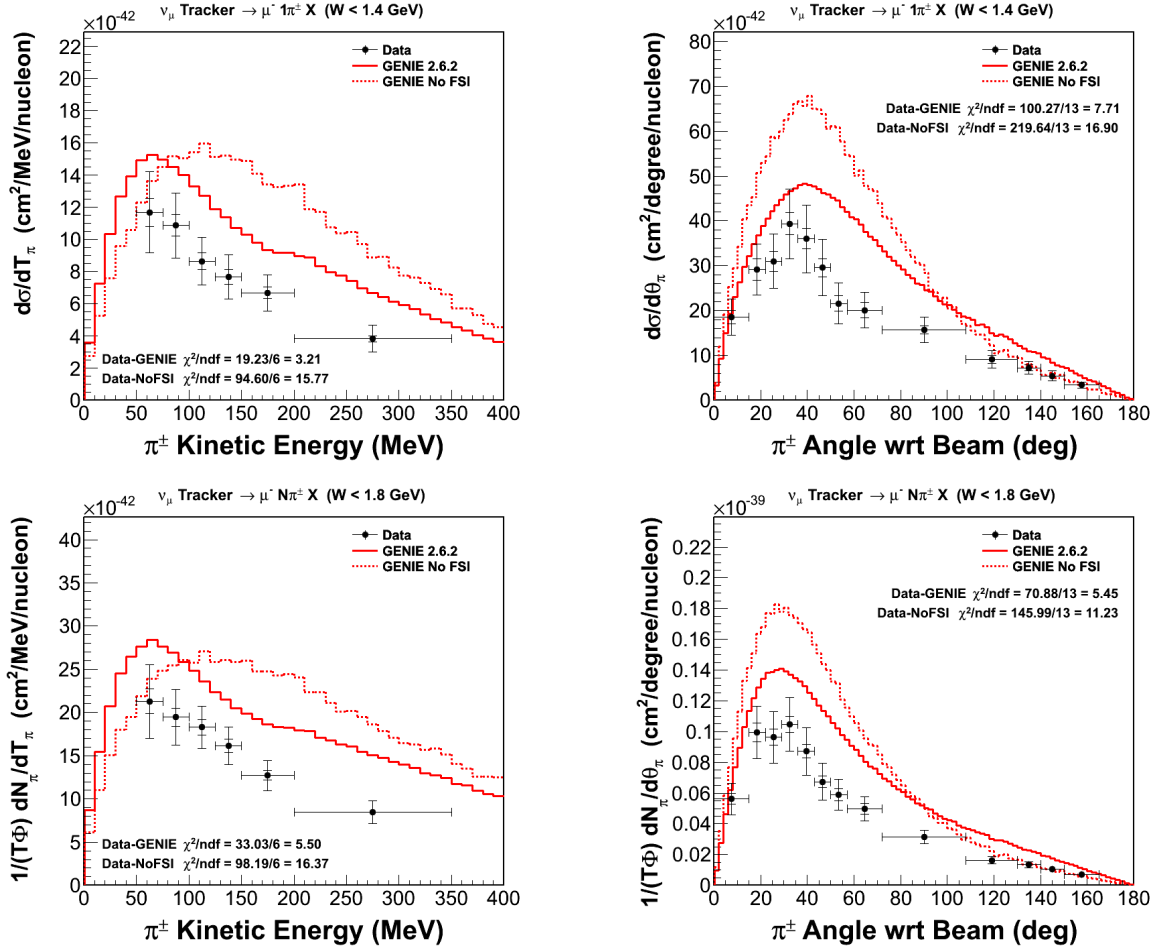


Figure 8.1: Measurement of ν_μ charged current π^\pm production differential cross sections with respect to T_π and $\theta_{\pi\nu}$. The one-pion distributions are on top, and the N-pion distributions are on the bottom.

Table 8.1: $\chi^2/\text{D.O.F.}$ between the differential cross section measurements and GENIE.

Measurement	$\chi^2/\text{D.O.F.}$ (with FSI)	$\chi^2/\text{D.O.F.}$ (no FSI)
One-pion $\frac{d\sigma}{dT_\pi}$	3.21	15.77
N-pion $\frac{d\sigma}{dT_\pi}$	5.50	16.37
One-pion $\frac{d\sigma}{d\theta_{\pi\nu}}$	7.71	16.90
N-pion $\frac{d\sigma}{d\theta_{\pi\nu}}$	5.45	11.23

Table 8.2: Systematic error groupings.

Group Name	Component Errors
Flux	All flux uncertainties except beam direction
Interaction Model	All GENIE model uncertainties
Detector Model	Geant4 cross section uncertainties
Michel	Michel selection efficiency and background rate
Angle & Vertex	Vertex and track angle resolutions, beam angle
Normalization	Tracking efficiency, mass (effect on rate)
Energy Response	Errors described in Sec. 7.7.8 through Sec. 7.7.13
Other	Unfolding, background fit uncertainties

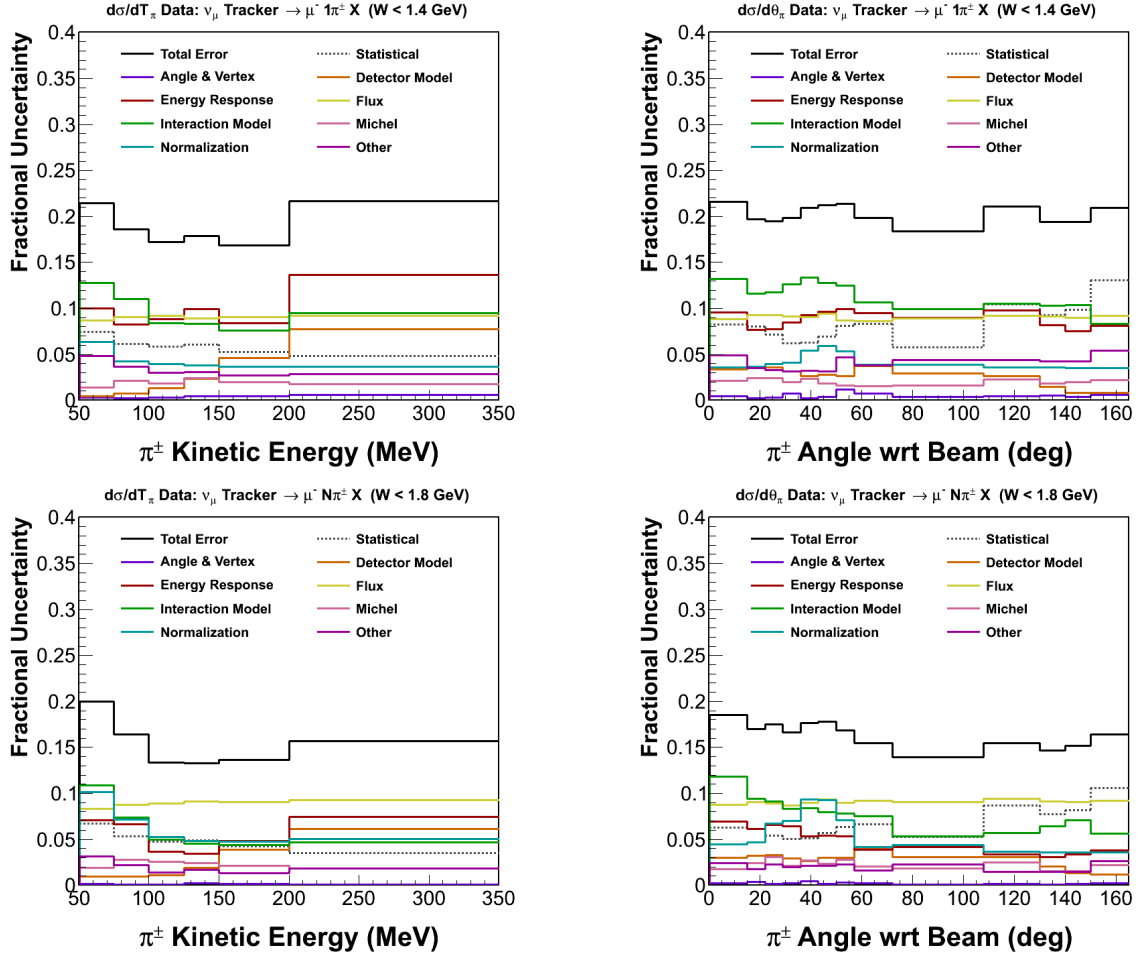


Figure 8.2: Error summaries for the measurements of $\frac{d\sigma}{dT_\pi}$ and $\frac{d\sigma}{d\theta_{\pi\nu}}$. The errors are shown as fractions of the measured value. The one-pion summaries are on top, and the N-pion summaries are on the bottom.

(6–10%) masses used in the Rein-Sehgal resonance model, which divide the Q^2 terms in their respective form factors (similar to Equation 6.1). The measurement is sensitive to these parameters because the cross sections are corrected for muon angles that are not accepted by MINOS. Different muon angles sample different Q^2 , and the form factor mass parameters affect Q^2 shape. If the cross sections are only corrected for muon angles less than 20° , the uncertainties due to the resonance vector and axial masses are reduced by half.

As discussed in Sec. 1.4, FSI sculpt the shape of the T_π differential cross section in a strong and unique way. Also, the shapes of the differential cross sections are not as sensitive to the large flux and GENIE cross section model uncertainties. Figure 8.3 shows the measurement of the differential cross section shapes compared to GENIE predictions. The $\chi^2/\text{D.O.F.}$ for the shape comparisons are in Table 8.3. The shape measurements all prefer GENIE with FSI, and this preference is even more significant than for the measurement with the absolute normalization included.

The shape uncertainties are summarized in Figure 8.4. Every shape measurement is dominated by the $\sim 4\text{--}12\%$ statistical uncertainty. The only significant systematic uncertainty in the $\frac{d\sigma}{d\theta_{\pi\nu}}$ measurement is the uncertainty on the GENIE cross section and FSI models. There is no GENIE uncertainty that dominates; collectively, there are a variety of cross section and FSI model parameters that contribute no more than 2–4% to the total GENIE uncertainty.

The $\frac{d\sigma}{dT_\pi}$ shape measurements only contain significant systematic errors below 75 MeV and above 200 MeV. The largest GENIE shape uncertainties are on parameters that determine the rate of pion scattering and absorption in FSI, but these do not rise above 2% except in the lowest energy bin in the N-pion analysis. The Geant4 pion interaction uncertainties contribute a $\sim 4\%$ shape uncertainty in the highest energy bin. The detector mass model is responsible for the largest energy response shape uncertainty and is greatest, $\sim 4\%$, in the highest energy bin.

These measurements can only be compared to the MiniBooNE pion production data [29] through a model. There are a few reasons for this. First, MiniBooNE uses a pure mineral oil (CH_2) target, while this measurement uses a $\sim 95\%$ polystyrene ($\sim \text{CH}$) target. Additionally, MiniBooNE has a lower energy neutrino beam, approximately 700 MeV compared to 3 GeV. Both MiniBooNE and the one-pion version of this measurement are dominated by pions from Δ decay, but the higher beam energy gives the Δ a larger average boost in this measurement. Finally, this measurement reaches higher Q^2 than MiniBooNE, again because of the higher beam energy. This

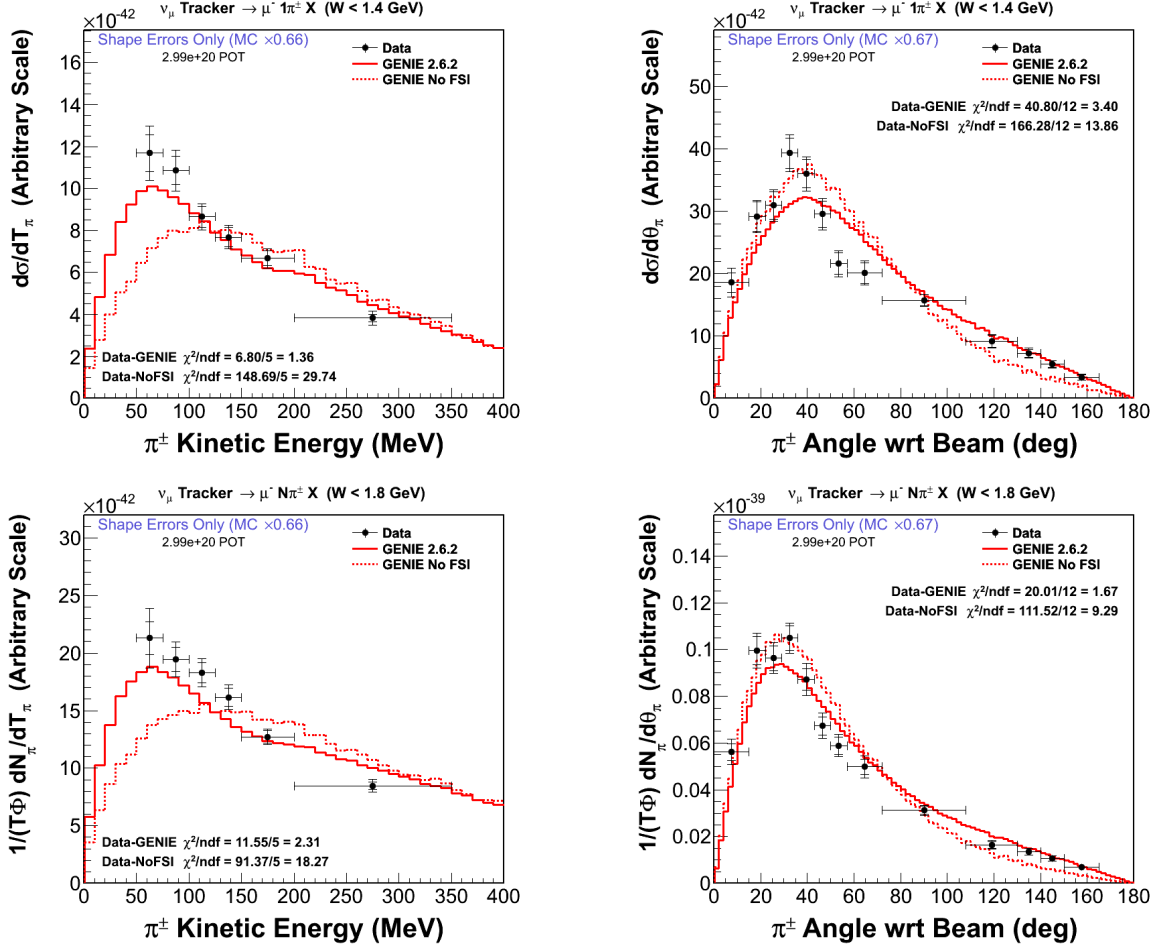


Figure 8.3: Shape measurement of ν_μ charged current π^\pm production differential cross sections with respect to T_π and $\theta_{\pi\nu}$. The one-pion distributions are on top, and the N-pion distributions are on the bottom. The simulations are scaled to match the normalization of the data in region of measurement.

Table 8.3: $\chi^2/\text{D.O.F.}$ between the differential cross section shape measurements and GENIE.

Shape Measurement	$\chi^2/\text{D.O.F.}$ (with FSI)	$\chi^2/\text{D.O.F.}$ (no FSI)
One-pion $\frac{d\sigma}{dT_\pi}$	1.36	29.74
N-pion $\frac{d\sigma}{dT_\pi}$	2.31	18.27
One-pion $\frac{d\sigma}{d\theta_{\pi\nu}}$	3.40	13.86
N-pion $\frac{d\sigma}{d\theta_{\pi\nu}}$	1.67	9.29

means that the two measurements have different sensitivity to a variety of Q^2 -dependent nuclear effects².

GENIE comparisons to this measurement and the MiniBooNE data are shown in Figures 1.5, 8.1, and 8.3. Both measurements see a modest depletion due to FSI pion absorption between 100–300 MeV. GENIE models the shape of the pion absorption region quite well for both measurements, but also is shifted to a lower kinetic energy relative to the MiniBooNE data. Given the size of the shift, it is unlikely that it can be explained by an offset in the MiniBooNE detector energy scale. This discrepancy can perhaps be explained by a nuclear effect not included in the GENIE model, such as multi-nucleon correlations or broadening of the Δ decay width. However, models that include some of these nuclear effects, such as GiBUU, are already unable to reproduce the MiniBooNE data. The reconciliation of the peak pion kinetic energy in these two measurements will be an important new goal for model builders.

Interestingly, while GENIE reproduces the thesis measurement shape better than the MiniBooNE shape, it greatly prefers the MiniBooNE normalization. This could be partially explained as a flux effect; the largest uncertainty on the MiniBooNE measurement is the neutrino flux, and the flux is among the largest uncertainties on the normalization of the measurement in this thesis.

²For example, the MiniBooNE measurement is more sensitive to multi-nucleon correlations due to the longer-wavelength virtual W probe.

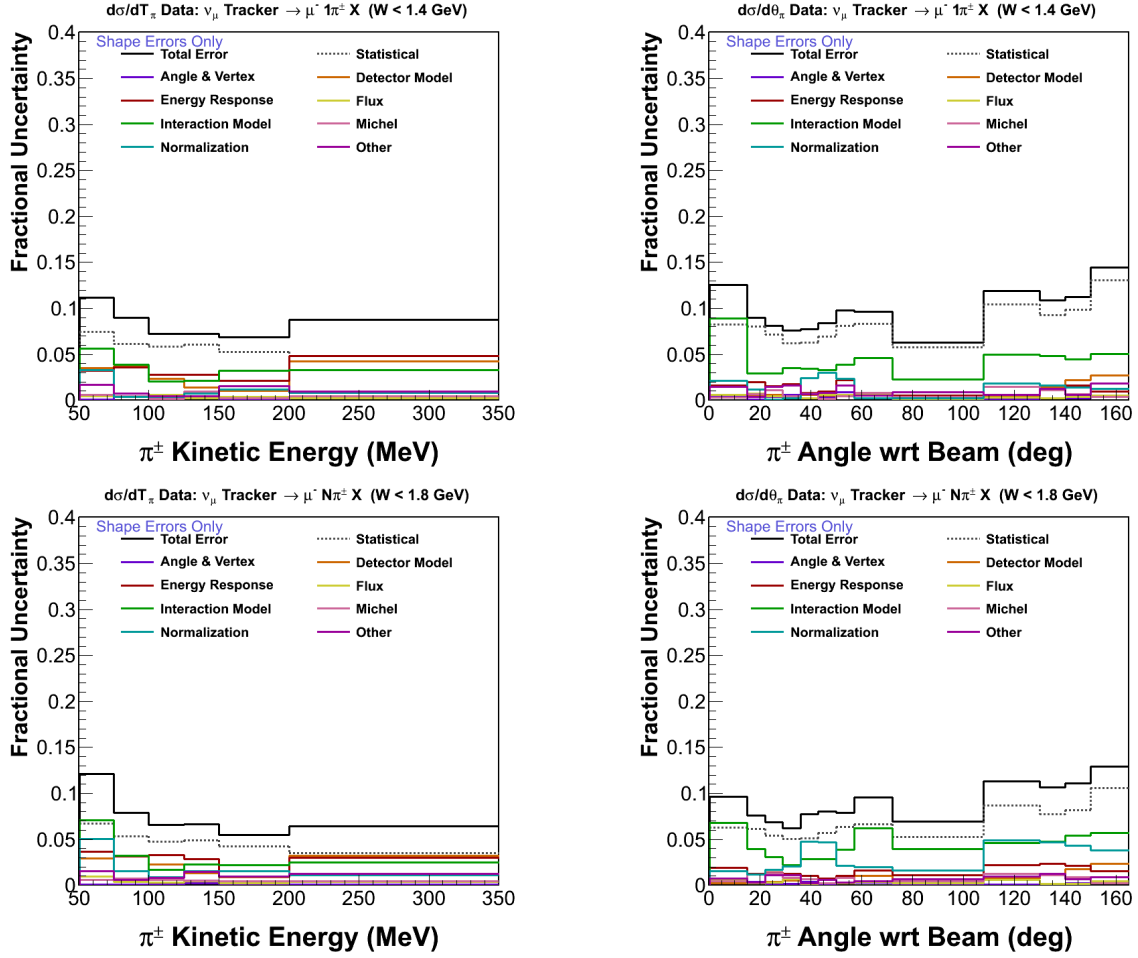


Figure 8.4: Shape error summaries for the measurements of $\frac{d\sigma}{dT_\pi}$ and $\frac{d\sigma}{d\theta_{\pi\nu}}$. The errors are shown as fractions of the measured value. The one-pion summaries are on top, and the N-pion summaries are on the bottom.

8.2 MODEL COMPARISONS

Figure 8.5 compares the one-pion measurements to the GENIE, NuWro [32], and Neut [33] event generators, as well as a theoretical calculation from Athar [30]. Neut and NuWro agree better with the measured total cross section than GENIE, and all three event generators reproduce the shape measurement very well. The GENIE prediction for the rate of baryon resonance production on deuterium is $\sim 20\%$ higher than Neut's prediction at the mean NuMI beam energy of 3 GeV, so it is very likely that this is responsible for most of the difference in the total cross section prediction³. The Athar calculation, which contains an incomplete FSI model that flatly suppresses the differential cross section, does not agree well with the shape of the data.

8.3 CONCLUSIONS

This thesis presented measurements of $\frac{d\sigma}{dT_\pi}$ and $\frac{d\sigma}{d\theta_\pi}$ for muon-neutrino charged current charged pion production in the MINER ν A scintillator tracker. The analysis was limited to neutrino energies between 1.5-10 GeV. Dependence on invariant hadronic mass W was studied through two versions of the analysis that impose the limits $W < 1.4$ GeV and $W < 1.8$ GeV. This is the first neutrino pion production measurement made with the MINER ν A detector. A number of reconstruction techniques, particularly in hadron tracking and particle identification, were developed to do this measurement.

The shapes of the differential cross sections, which depend strongly on the nature of nuclear final state interactions, were compared to simulation and theoretical predictions. The measurements were shown to favor models that contain final state interactions. Additionally, the GENIE model was used to compare these measurements to the MiniBooNE pion production data, which is the only other high-statistics neutrino pion production measurement available. GENIE was found to reproduce the shape of the thesis measurements while failing to reproduce the full shape of the MiniBooNE result. These results will be useful for constraining neutrino interaction models, par-

³Both Neut and GENIE use the Rein-Sehgal model for baryon resonance production. The difference comes from how the model is tuned to data. The data have large uncertainties (Fig. 1.3), so the tuning can vary quite a bit.

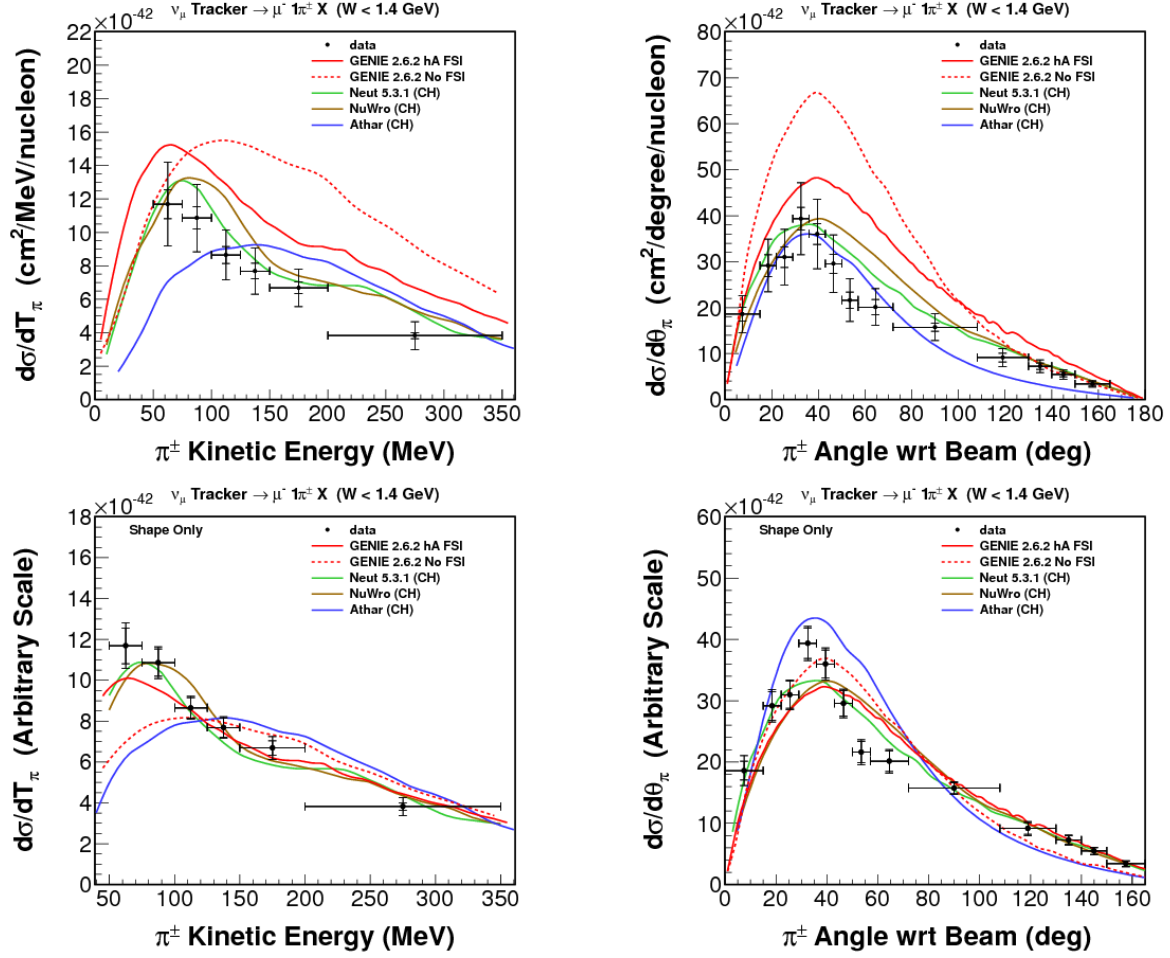


Figure 8.5: Comparisons of the one-pion measurements to various models [23, 30, 32, 33]. The cross sections are on the top, and the shape measurements are on the bottom.

ticularly implementations of final state interactions, which in turn will help to reduce systematic uncertainties for current and future neutrino oscillation experiments.

APPENDIX A

PMT GAIN CALIBRATION

A.1 GAIN MEASUREMENT PROCEDURE

The gain g of a PMT pixel is defined as

$$g = \frac{\overline{Q}}{\lambda e}, \quad (\text{A.1})$$

where \overline{Q} is the mean of the pixel's pedestal-subtracted anode charge distribution, λ is the mean number of photoelectrons (PEs) arriving at the first dynode, and e is the electron charge. Note that the quantum efficiency of the photocathode and the collection efficiency of the first dynode are not included in this definition of the gain. These effects cannot be measured with *in situ* light injection (LI) data and are accounted for in the relative channel-to-channel calibration.

The parameter λ must be measured to calculate the gain. MINER ν A adopts and modifies a technique from MINOS [57] that calculates λ as a function of the gain and the mean and width of the PMT anode charge distribution. For a PMT with N dynodes, the probability distribution $P_N(q)$ of the charge q measured at the PMT anode is equal to the probability distribution of the number of PEs observed at the anode convoluted with a Gaussian to account for electronics noise (the electronics pedestal). If each dynode amplifies according to a Poisson distribution, and the amplification is linear with the number of incoming PEs, $P_N(q)$ is

$$\begin{aligned}
P_N(q) = & \sum_{k_{N+1}=0}^{\infty} \sum_{k_N=0}^{\infty} \cdots \sum_{k_1=0}^{\infty} \\
& \frac{e^{-g_N k_N} (g_N k_N)^{k_{N+1}}}{k_{N+1}!} \frac{e^{-g_{N-1} k_{N-1}} (g_{N-1} k_{N-1})^{k_N}}{k_N!} \cdots \\
& \frac{e^{-g_1 k_1} (g_1 k_1)^{k_2}}{k_2!} \frac{e^{-\lambda} \lambda^{k_1}}{k_1!} \frac{1}{\sqrt{2\pi}\sigma_p} e^{\frac{(q - e k_{N+1})^2}{2\sigma_p^2}}, \tag{A.2}
\end{aligned}$$

where q is the charge measured at the anode, σ_p is the electronics pedestal width, and g_n is the mean gain of the n th dynode [65]. The variance σ_N^2 of $P_N(q)$ is

$$\sigma_N^2 = \sigma_p^2 + \lambda g^2 e^2 + \lambda g^2 e^2 w^2, \tag{A.3}$$

where the product of the dynode gains, $\prod_{n=1}^N g_n$, is identified as the gain g of the pixel, and w^2 is defined as

$$w^2 \equiv \sum_{j=1}^N \left(\prod_{i=1}^j \frac{1}{g_i} \right). \tag{A.4}$$

w^2 typically has values between 0.2 and 0.3 and is related to the fractional width of the one-PE peak. Equation A.3 has a simple physical interpretation: the variance of the PMT anode charge distribution is the sum of the variance of the electronics pedestal, the variance of the incoming PE distribution, and the variance due to the statistical broadening of the dynode chain.

Upon solving equations A.1 and A.3 simultaneously, the gain can be written as a function of w and the mean and standard deviation of the anode charge distribution:

$$g = \frac{\sigma_N^2 - \sigma_p^2}{Q(1 + w^2)e}. \tag{A.5}$$

Equation A.5 is not sufficient to calculate the pixel gains because the inclusion of w on the right hand side introduces N unmeasured parameters g_n . The number of unmeasured parameters can be reduced to one by using a model in which each dynode gain is proportional to a power of V_n , the potential difference that accelerates each PE arriving at the n th dynode:

$$g_n \propto V_n^\alpha. \tag{A.6}$$

The parameter α typically has values between 0.7 and 0.8 [97] and can be measured by collecting light injection data while varying the PMT high voltage. Instead, MINER ν A assumes $\alpha = 0.75$ for each pixel, which results in less than 1% systematic error on the gain measurement.

The high voltage applied to each MINER ν A PMT is divided amongst the dynodes according to the optimal ratios recommended by Hamamatsu: (3:2:2:1:1:1:1:1:1:2:5), where the final amplification stage corresponds to the potential difference between the last dynode and the anode. Using these ratios and equation A.6, each g_n can be written as a function of g_1 only. This leads to a relationship between g_1 and g :

$$g_1 = g^{\frac{1}{12}} \frac{177147^{\frac{\alpha}{12}}}{8}. \quad (\text{A.7})$$

w can now be expressed as a sum of non-integer powers of g and inserted into equation A.5. The resulting equation does not have an analytic solution, so MINER ν A uses a recursive procedure to simultaneously calculate w and g :

1. Calculate the gain with equation A.5 using a value for w supplied by the previous iteration of this procedure. $1 + w^2 = 1.24$ is used for the first iteration
2. Use equation A.7 and the gain from the previous step to calculate g_1
3. Use equations A.4, A.6, and the calculated value g_1 to calculate $1 + w^2$
4. If the w calculated in step 3 and the w used in step 1 differ by less than 0.001%, use the current value of w to calculate the gain. Otherwise, return to step 1 with the value of w calculated in step 3.

This procedure typically converges after two iterations.

A.2 GAIN MEASUREMENT AND ELECTRONICS READ-OUT

The gain measurement expects that there is one hit per channel per calibration gate, so it is necessary to know whether the LI data was collected with discriminators on or off¹. The discriminator

¹The discriminators are described in 3.2.1. If discriminators are on, there could be more than one hit in a channel due to noise or PMT after-pulsing triggering the discriminator. If discriminators are off, each channel has one hit with PMT anode charge summed over the entire gate.

status is not stored by the DAQ, so it must be extracted from the data. This is done by determining whether any hits fired their discriminator in the first LI gate of the data to be processed. If the discriminators are on, the algorithm must also determine whether all ADC pipelines were read out, or if only the deepest pipeline was read². The pipeline readout mode is determined by finding channels with more than one hit in the gate. If every channel has only one hit, then the data is deepest pipeline data.

The algorithm to calculate the gains is changed according to the discriminator status and the pipeline readout mode. If discriminators are off, each channel only has one hit per gate and these hits do not have timing information. In this case, it is assumed that each hit corresponds to an LI pulse and the calibrated charge of the hit is used to calculate the mean and variance of the anode charge distribution.

If the discriminators are on and all pipelines are read out, each channel may have more than one hit per gate. These hits may be formed from a few sources: LI pulses, PMT after-pulsing, and dark noise that fluctuates above the discriminator threshold. The charges of these hits cannot be added together to duplicate the discriminator-off result because of discriminator dead time. Instead, the discriminator timing information is used to find the time window during which the LI pulses arrives at the PMTs. This time window is set to be the clock tick that corresponds to the timing information of the greatest number of hits, plus and minus ten clock ticks. All hits outside of this time window is rejected, and the hit inside the window is used to calculate the mean and variance of the anode charge distribution. If there is no hit inside the LI time window, the channel is assigned a new hit with a pedestal-subtracted charge of zero fluctuated about a width equal to the pedestal width. This ensures that zeroes³ are counted correctly.

If the discriminators are on but only the deepest pipeline is read out, the procedure is similar to the procedure where all pipelines are read out, with one important difference: pedestal hits are not added to channels that do not have a hit inside the LI time window. This is omitted due to an effect from integrating TriP-t chips in triplets that only occurs during deepest pipeline runs. The effect causes one or more of the earliest hits in the channels on a triplet of TriP-t chips to be lost

²If only the deepest pipeline is read, then the DAQ only stores the earliest hit in a channel. This is the default setting for LI data that is read out between beam gates. This is done to make calibration data collection as fast as possible.

³A “zero” is a gate in which zero PEs reach the first dynode.

during readout. With the hit lost, it is unknown whether it could have occurred during the LI time window or not. Therefore, a missing hit in the LI time window does not indicate that there was a pedestal hit. This difference between deepest pipeline and all pipeline data is very small; the gains measured are similar in both cases.

The light injection system is also used to calibrate PMTs on the test beam detector. The discriminator-off test beam LI data contains occasional hits with charges that are much larger than can be created by the LI system. These hits can have a large effect on the measurement of the gain by pulling the variance of the anode charge distribution. An outlier rejection scheme is implemented to remove these hits and is used on all data, regardless of discriminator status, because it introduces less than 1% systematic uncertainty.

The outlier rejection scheme works by sorting the hits in a channel in order of ascending charge. It then searches for the first occurrence of a consecutive pair of charges with a charge greater than the mean of the charge distribution and with a charge difference more than 300 fC. For a typical PMT pixel, this charge difference is equal to 3.75 PEs, and such a pair of charges is likely to be located in the very end of the high-side tail of the anode charge distribution. If such a pair is found, all hits with charges equal to or greater than the charge of the higher-charged hit in the pair are removed from the anode charge distribution.

A.3 STATISTICAL UNCERTAINTY

The statistical uncertainty is calculated by taking the derivative of [A.5](#) and then assuming that each source of statistical uncertainty is uncorrelated and Gaussian. The resulting equation is

$$\sigma_g = g \sqrt{\frac{2}{(\sigma^2 - \sigma_p^2)^2} \left(\frac{\sigma^4}{N-1} + \frac{\sigma_p^4}{N_p-1} \right) + \frac{1}{Q^2} \left(\frac{\sigma^2}{N} + \frac{\sigma_p^2}{N_p} \right)} \quad (\text{A.8})$$

where σ_N has been redefined as σ , N is the number of LI gates used to construct that anode charge distribution, and N_p is the number of gates used to measure the pedestal mean and variance. In calculating [A.8](#), it is assumed that w^2 is constant. This assumption is fine because w^2 only varies by a couple of percent over an order of magnitude of gains.

There are a couple more approximations that must be made when calculating the statistical error. In general, an LI charge distribution can include data from all three FEB gain ranges, and thus different data in the distribution are calibrated with different pedestal distributions. This creates some ambiguity in choosing a σ_p in equation A.8. Fortunately, for the majority of channels in the MINER ν A detector, most entries in the LI charge distribution fall within the high gain range. σ_p is taken to simply be the width of the high gain pedestal. A second approximation is made in converting the pedestal width from ADC units to units of electron charge. The conversion is made by using the FEB trilinear fit calibration constants. These constants are valid only for charges above pedestal, but here are used to calibrate a quantity that is partly determined by charges below pedestal. The systematic effect of this approximation is negligible. Figure A.1 shows the statistical significance $\Delta g / \sigma_{\Delta g}$ of the gain drift between consecutive gain measurements⁴. The width of the distribution is 0.97, indicating that the statistical error σ_g is slightly overestimated.




Figure A.1: Statistical significance of gain fluctuations between consecutive measurements. The mean and width of the distribution are 0.002 and 0.97 respectively. Reprinted from Ref. [54], Copyright (2014) with permission from Elsevier.

⁴Gains are typically measured once per day.

APPENDIX B

N-PION ANALYSIS FIGURES

This appendix contains the N-pion analysis plots that were omitted from Chapter 7.




Figure B.1: The N-pion reconstructed q/p distribution. Events below zero contain a muon with negative charge. $q/p > 0$ is primarily background from anti-neutrino interactions (in gray).




Figure B.2: The N-pion reconstructed E_ν distribution. The contamination from $E_\nu > 10$ GeV is shown in the light blue fill and is greatly reduced by the selection.




Figure B.3: The N-pion reconstructed W_{exp} distribution. The $W_{exp} > 1.4$ GeV/c² background leaks into the signal region due to fluctuations in the calorimetric reconstruction of E_{recoil} . The entries at zero are events in which most of the hadronic recoil is neutral and not seen in the detector; these events are also removed by the analysis.




Figure B.4: The number of reconstructed hadron tracks at the event vertex in the N-pion analysis. The signal events that are removed by this selection are mostly events where the outgoing pion and nucleon are low energy and below tracking threshold.




Figure B.5: The pion score variable used to separate pion and proton tracks in the N-pion analysis. The stopping categories correspond to particles that are contained in the detector and lose all of their energy through ionization and atomic excitation before stopping or decaying. Interacting pions and protons either interact with a detector nucleus or decay before losing all of their energy. Stopping pions and protons are well separated by this score. Separation is worse for interacting particles because their energy loss profiles do not match the Bethe-Bloch equation prediction.




Figure B.6: The number of detector views spanned by selected Michel candidates in the N-pion analysis. The shape of this distribution is well-modeled by simulation.

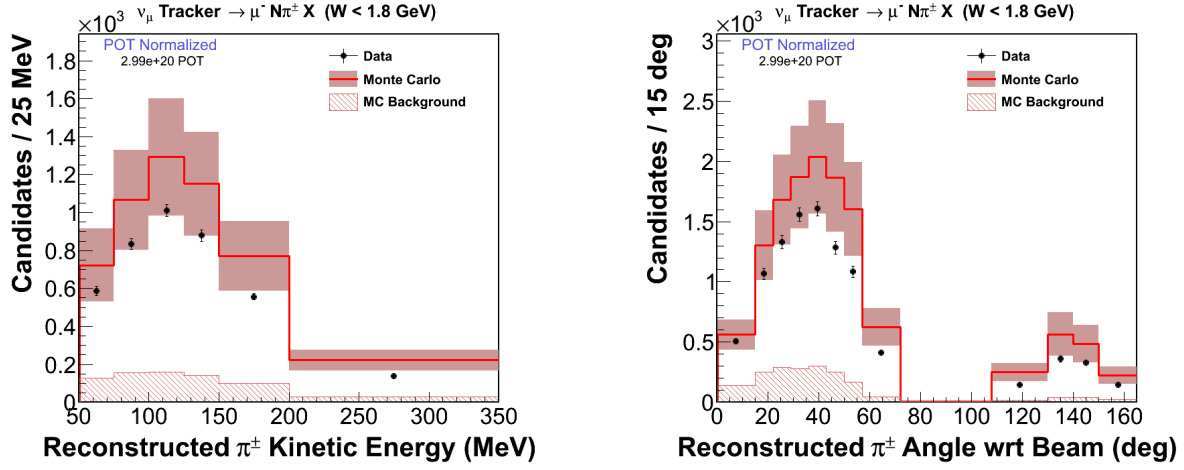


Figure B.7: The N-pion reconstructed T_π and $\theta_{\pi\nu}$ distributions compared to simulation. The data error bars are statistical and the error bars on the simulation include the full set of systematic errors as described in section 7.7.

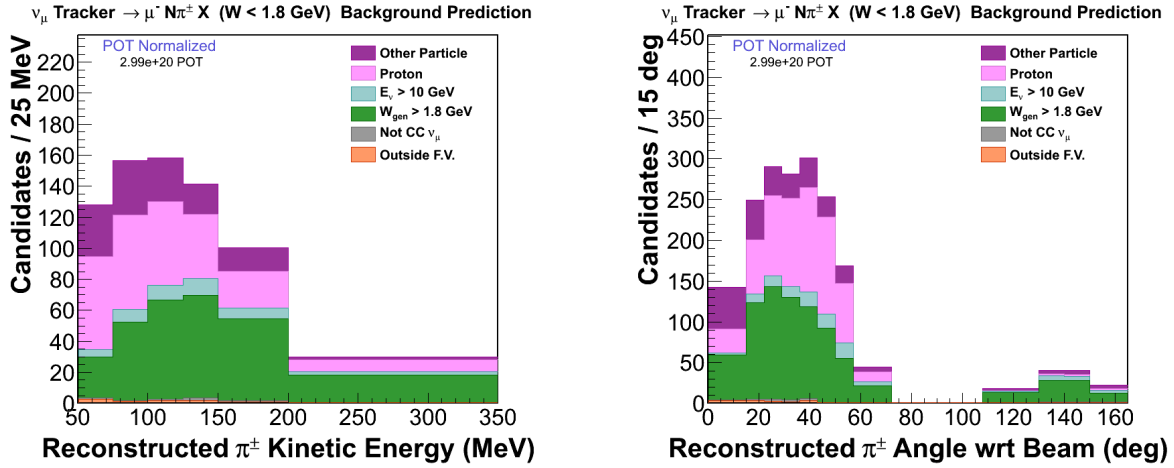


Figure B.8: The N-pion simulated background distributions divided according to background type. The largest background category is events with true $W > 1.4\text{GeV}/c^2$.

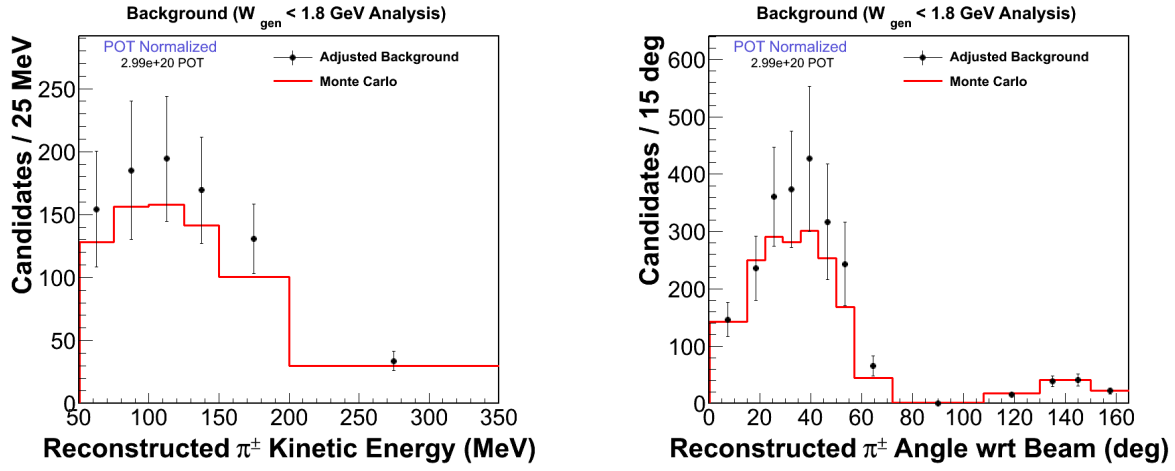


Figure B.9: The N-pion adjusted background predictions compared to the nominal simulated predictions.




Figure B.10: The N-pion background-subtracted distributions, $N_j - N_j^{bg}$, compared to simulation.

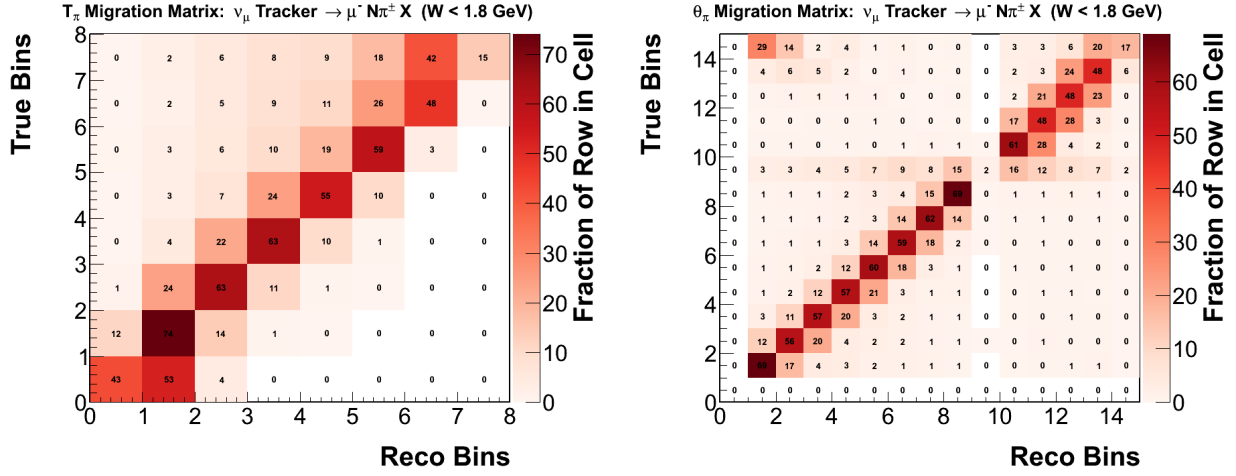


Figure B.11: The N-pion migration matrices. The bin number indicates the bin number in the T_π or $\theta_{\pi\nu}$ distribution. Bin 0 is the underflow bin and the last bin is the overflow bin. Perfect detector resolution would correspond to a diagonal matrix; i.e., the reconstructed value equals the true value. The matrix is row-normalized - the contents of each row sum to one.

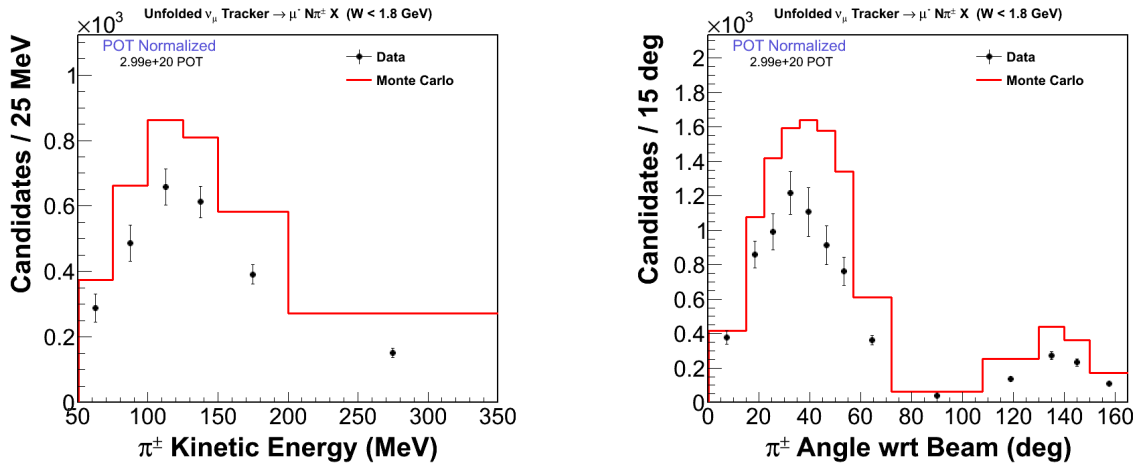


Figure B.12: The N-pion unfolded T_π and $\theta_{\pi\nu}$ distributions.




Figure B.13: The N-pion reconstruction efficiency as a function of T_π and $\theta_{\pi\nu}$. The error bar contains all systematic errors.

APPENDIX C

MINER ν A COLLABORATION

This is a list of past and present members of the MINER ν A collaboration.

University of Athens: G. Tzanakos

University of Texas at Austin: J. Cravens, M. Jerkins, S. Kopp, J. Ratchford

Centro Brasileiro de Pesquisas Físicas: D.A. Martinez Caicedo, M.F. Carneiro, C.M. Castromonte, H. da Motta, G.A. Fiorentini, K. Hurtado, J.L. Palomino

University of California at Irvine: C. Simon, B.P. Ziemer

University of Chicago: D.W. Schmitz

University of Florida: J. Mousseau, B. Osmanov, H. Ray

Fermilab: D. Boehnlein, R. DeMaat, N. Grossman, D.A. Harris, J.G. Morfín, J. Osta, E. A. Paschos, G.N. Perdue, L. Rakotondravohitra, P. Rubinov, F.D. Snider, J.S. Sobczyk, R. Stefanski

University of Geneva: A. Bravar, C. Martin Mari

Universidad de Guanajuato: Victor M. Castillo Vallejo, J. Felix, A. Higuera, Z. Urrutia, E. Valencia, G. Zavala

Hampton University: M.E. Christy, M. Datta, C.E. Keppel, W. Tan, T. Walton, L. Y. Zhu

Inst. Nucl. Reas. Moscow: A. Butkevich, S.A. Kulagin

James Madison University: G. Niculescu, I. Niculescu

Mass. Col. Lib. Arts: E. Maher

University of Minnesota at Duluth: R. Gran, M. Lanari

Universidad Nacional de Ingeniería: M. Alania, C.J. Solano Salinas

Northwestern University: L. Fields, B. Gobbi, J.A. Hobbs, C.E. Patrick, L. Patrick, H. Schellman

Otterbein University: N. Tagg

University of Pittsburgh: S. Boyd, S.A. Dytman, B. Eberly, Z. Isvan, C.L. McGivern, D. Naples, V. Paolone, L. Ren

Pontificia Universidad Catolica del Peru: M. J. Bustamante , G.A. Díaz , A.M. Gago, J.P. Velásquez

Purdue University Calumet: R. Napora

University of Rochester: S. Avvakumov, A. Bodek, R. Bradford, H. Budd, J. Chvojka, M. Day, H. Lee, L. Loiacono, S. Manly, C.M. Marshall, K.S. McFarland, A.M. McGowan, A. Mislivec, J. Park, P.A. Rodrigues, J. Wolcott

Rutgers University: G. J. Kumbartzki, T. Le, R.D. Ransome, E.C. Schulte, B.G. Tice

Universidad Técnica Federico Santa María: W.K. Brooks, G. Maggi, J. Miller, C. Peña

Tufts University: O. Altinok, H. Gallagher, T. Kafka, W.A. Mann, W. P. Oliver

College of William and Mary: L. Aliaga, J. Devan, M. Kordosky, J.K. Nelson, A. Norrick, J. Walding, D. Zhang

BIBLIOGRAPHY

- [1] J.P. Lees *et al.* Evidence for an Excess of $\overline{B} \rightarrow D^{(*)} \tau^- \overline{\nu}_\tau$ Decays. *Phys. Rev. Lett.*, 109: 101802, (2012).
- [2] M. Peskin and D. Schroeder. *An Introduction to Quantum Field Theory*, page 198. Westview Press, (1995).
- [3] M. Kobayashi and T. Maskawa. CP-Violation in the Renormalizable Theory of Weak Interaction. *Progress of Theoretical Physics*, 49(2):652–657, (1973).
- [4] Y. Fukuda *et al.* Evidence for Oscillation of Atmospheric Neutrinos. *Phys. Rev. Lett.*, 81: 1562–1567, (1998).
- [5] Q.R. Ahmad *et al.* Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory. *Physical Review Letters*, 89(1):011301, (2002).
- [6] Z. Maki, M. Nakagawa and S. Sakata. Remarks on the Unified Model of Elementary Particles. *Progress of Theoretical Physics*, 28(5):870–880, 1962.
- [7] D.J. Griffiths. *Introduction to Elementary Particles*, chapter 11. WILEY-VCH Verlag GmbH & Co., (2008).
- [8] A.D. Sakharov. Violation of CP Invariance, C Asymmetry, and Baryon Asymmetry of the Universe. *Soviet Physics Uspekhi*, 34(5):392, (1991).
- [9] K. Abe *et al.* Evidence of Electron Neutrino Appearance in a Muon Neutrino Beam. *Phys. Rev.*, D88:032002, (2013).
- [10] P. Adamson *et al.* Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS. *Phys. Rev. Lett.*, 110:251801, 2013.
- [11] LBNE collaboration. LBNE Homepage, (2014). URL <http://lbne.fnal.gov>.
- [12] C. Adams *et al.* Scientific Opportunities with the Long-Baseline Neutrino Experiment. *arXiv:1307.7335 [hep-ex]*, (2013).

- [13] D.G. Michael *et al.* (MINOS collaboration). The Magnetized Steel and Scintillator Calorimeters of the MINOS Experiment. *Nucl. Inst. and Meth., Phys. Res. Sect. A*, 596:190–228, (2008).
- [14] T2K collaboration. T2K Homepage, (2014). URL <http://t2k-experiment.org/>.
- [15] D. Rein and L.M. Sehgal. Neutrino-Excitation of Baryon Resonances and Single Pion Production. *Annals of Physics*, 133(1):79–153, (1981).
- [16] R.P. Feynman, M. Kislinger, and F. Ravndal. Current Matrix Elements from a Relativistic Quark Model. *Phys. Rev.*, D3(11):2706–2732, (1971).
- [17] T. Sjöstrand *et al.* . High-Energy-Physics Event Generation with PYTHIA 6.1. *Computer Physics Communications*, 135(2):238–259, (2001).
- [18] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand. Parton Fragmentation and String Dynamics. *Phys. Rept.*, 97:31–145, (1983).
- [19] R.A. Smith and E.J. Moniz. Neutrino Reactions on Nuclear Targets. *Nuclear Physics B*, 43: 605–622, (1972).
- [20] A. Bodek and J.L. Ritchie. Fermi-Motion Effects in Deep-Inelastic Lepton Scattering from Nuclear Targets. *Phys. Rev.*, D23(5):1070, (1981).
- [21] O. Benhar, A. Fabrocini, and S. Fantoni. The Nucleon Spectral Function in Nuclear Matter. *Nucl. Phys. A*, 505(2):267–299, (1989).
- [22] A. Bodek, H.S. Budd, and M.E. Christy. Neutrino Quasielastic Scattering on Nuclear Targets. *The European Physical Journal C*, 71(9):1–18, (2011).
- [23] C. Andreopoulos *et al.* The GENIE Neutrino Monte Carlo Generator. *Nucl. Instrum. Meth.*, A614:87–104, (2010).
- [24] O. Buss *et al.* Transport-Theoretical Description of Nuclear Reactions. *Phys.Rept.*, 512: 1–124, (2012).
- [25] E.A. Uehling and G.E. Uhlenbeck. Transport Phenomena in Einstein-Bose and Fermi-Dirac Gases. I. *Phys. Rev.*, 43:552–561, (1933).
- [26] G.M. Radecky *et al.* Study of single pion production by weak charged currents in low-energy neutrino d interactions. *Phys. Rev.*, D25:1161, (1982).
- [27] T. Kitagaki *et al.* Charged current exclusive pion production in neutrino deuterium interactions. *Phys. Rev.*, D34:2554, (1986).
- [28] MiniBooNE collaboration. MiniBooNE Homepage, (2014). URL <http://www-boone.fnal.gov/>.

- [29] A.A. Aguilar-Arevalo *et al.* Measurement of Neutrino-Induced Charged-Current Charged Pion Production Cross Sections on Mineral Oil at $E_\nu \sim 1$ GeV. *Phys. Rev.*, D83:052007, (2011).
- [30] M.S. Athar, S. Chauhan, and S.K. Singh. Theoretical Study of Lepton Events in the Atmospheric Neutrino Experiments at SuperK. *Eur.Phys.J.*, A43:209–227, (2010).
- [31] E. Hernandez, J. Nieves, and M.J. Vicente Vacas. Single π production in neutrino nucleus scattering. *Phys.Rev.*, D87:113009, (2013).
- [32] T. Golan, C. Juszczak, and J.T. Sobczyk. Final State Interactions Effects in Neutrino-Nucleus Interactions. *Phys.Rev.*, C86:015505, (2012).
- [33] Y. Hayato. A Neutrino Interaction Simulation Program Library NEUT. *Acta Phys. Pol.*, B40: 2477, (2009).
- [34] P.A. Rodrigues. Comparing Pion Production Models to MiniBooNE Data. *arXiv:1402.4709 [hep-ex]*, (2014).
- [35] O. Lalakulich and U. Mosel. Pion Production in the MiniBooNE Experiment. *Phys. Rev.*, C87:014602, (2013). URL <http://journals.aps.org/prc/abstract/10.1103/PhysRevC.87.014602>.
- [36] D.A. Crane *et al.* Status report: Technical design of neutrino beams for the Main Injector (NuMI). *FERMILAB-TM-1946, NUMI-B-92, FERMILAB-DESIGN-1995-03*, (1995).
- [37] R.B. Patterson. The NOvA Experiment: Status and Outlook. *Nucl. Phys. B, Proc. Suppl.*, 235:151–157, (2013).
- [38] K. Budal. Measurement of Charge Emission from Targets as a Means of Burst Intensity and Beam Intensity Monitoring. *IEEE Trans. Nucl. Sci.*, 14:1132, (1967).
- [39] R. Zwaska. *Accelerator Systems and Instrumentation for the NuMI Neutrino Beam*. PhD thesis, FERMILAB-THESIS-2005-73, University of Texas at Austin, (2005).
- [40] Fermi National Accelerator Laboratory. The NuMI Technical Design Handbook. (2005). URL http://www-numi.fnal.gov/numwork/tdh/tdh_index.html.
- [41] S. Kopp *et al.* Secondary Beam Monitors for the NuMI Facility at FNAL. *Nucl. Inst. and Meth., Phys. Res. Sect. A*, 568:503–519, (2006).
- [42] L. Loiacono. *Measurement of the Muon Neutrino Inclusive Charged Current Cross Section on Iron Using the MINOS Detector*. PhD thesis, FERMILAB-THESIS-2011-06, University of Texas at Austin, (2011).
- [43] S. Agostinelli *et al.* (Geant4 collaboration). G4—a Simulation Toolkit. *Nucl. Inst. and Meth., Phys. Res. Sect. A*, 506:250–303, (2003).

- [44] Description of FTFP QCD String Model, (2014). URL http://geant4.cern.ch/support/proc_mod_catalog/physics_lists/hadronic/FTFP.html.
- [45] B. Andersson, G. Gustafson, and B. Nilsson-Almqvist. A Model for Low-pT Hadronic Reactions with Generalizations to Hadron-Nucleus and Nucleus-Nucleus Collisions. *Nuclear Physics B*, 281(1):289 – 309, (1987).
- [46] Description of Bertini Cascade Model, (2014). URL http://geant4.cern.ch/support/proc_mod_catalog/physics_lists/hadronic/BertiniCascade.html.
- [47] NA49 collaboration. NA49 Homepage, (1997). URL <http://na49info.web.cern.ch/na49info/>.
- [48] A. Ferrari, P.R. Sala, A. Fasso, and J. Ranft. Fluka. *CERN-library in: http://fluka.web.cern.ch/fluka*, (2005).
- [49] C. Alt *et al.* Inclusive Production of Charged Pions in p+C Collisions at 158-GeV/c Beam Momentum. *Eur.Phys.J.*, C49:897–917, (2007).
- [50] T. Anticic *et al.* Inclusive Production of Charged Kaons in p+p Collisions at 158 GeV/c Beam Momentum and a New Evaluation of the Energy Dependence of Kaon Production up to Collider Energies. *Eur.Phys.J.*, C68:1–73, (2010).
- [51] B. Baatar *et al.* Inclusive Production of Protons, Anti-Protons, Neutrons, Deuterons and Tritons in p+C Collisions at 158 GeV/c Beam Momentum. *Eur.Phys.J.*, C73:2364, (2013).
- [52] A. Lebedev. *Ratio of pion kaon production in proton carbon interactions*. PhD thesis, FERMILAB-THESIS-2007-76, Harvard, (2007).
- [53] D.S. Barton *et al.* Experimental Study of the A Dependence of Inclusive Hadron Fragmentation. *Phys. Rev.*, D27:2580, (1983).
- [54] L. Aliaga *et al.* (MINERvA collaboration). Design, Calibration, and Performance of the MINERvA Detector. *Nucl. Inst. and Met., Phys. Res. Sect. A*, 743C:130–159, (2014).
- [55] G.N. Perdue *et al.* (MINERvA collaboration). The MINERvA Data Acquisition System and Infrastructure. *Nucl. Inst. and Met., Phys. Res. Sect. A*, 694(0):179 – 192, (2012).
- [56] E. Gallas and J. Li. Polishing Optical Fibers for the D0 ICD in Run II. FNAL-TM-2062, (1998).
- [57] N. Tagg *et al.* Performance of Hamamatsu 64-Anode Photomultipliers for use with Wavelength-Shifting Optical Fibres. *Nucl. Inst. and Meth., Phys. Res. Sect. A*, 539:668–678, (2005).
- [58] Fermilab Test Beam Facility. Fermilab test beam facility, (2014). URL <http://www-ppd.fnal.gov/ftbfb/>.

- [59] J. Estrada, C. Garcia, B. Hoeneisen, and P. Rubinov. MCM II and the TriP Chip. FERMILAB-TM-2226, (2002).
- [60] R.J. Glauber. *High Energy Physics and Nuclear Structure*. Plenum Press, (1970).
- [61] J.B. Birks. Scintillations from organic crystals: Specific fluorescence and relative response to different radiations. *Proceedings of the Physical Society. Section A*, 64(10):874, (1951).
- [62] R. Gran. Birks law parameter from minerva testbeam data. [MINERvA Internal Note TN37](#), (2013).
- [63] R. Brun, A.C. McPherson, P. Zancarini, M. Maire, and F. Bruyant. GEANT 3 : user's guide Geant 3.10, Geant 3.11. *CERN-DD-EE-84-01*, (1987).
- [64] S.M. Ross. Peirce's Criterion for the Elimination of Suspect Experimental Data. *Jour. Eng. Tech.*, (Fall 2003).
- [65] J. Rademacker. An Exact Formula to Describe the Amplification Process in a Photomultiplier Tube. *Nucl. Instrum. Meth.*, A484:432–443, (2002).
- [66] R. Luchsinger and C. Grab. Vertex reconstruction by means of the method of kalman filtering. *Computer Physics Communications*, 76(3):263 – 280, (1993).
- [67] R. Frühwirth. Application of kalman filtering to track and vertex fitting. *Nucl. Inst. and Meth., Phys. Res. Sect. A*, 262(2-3):444 – 450, (1987).
- [68] E.J. Wolin and L.L. Ho. Covariance Matrices for Track Fitting with the Kalman Filter. *Nucl. Inst. and Meth., Phys. Res. Sect. A*, 329:493–500, (1993).
- [69] D.E. Groom, N.V. Mokhov, and S.I. Striganov. Muon Stopping Power and Range Tables 10 MeV–100 TeV. *Atomic Data and Nuclear Data Tables*, 78(2):183–356, (2001).
- [70] J. Beringer *et al.* (Particle Data Group). Review of particle physics. *Phys. Rev. D*, 86:010001, (2012).
- [71] C.H. Llewellyn Smith. Neutrino Reactions at Accelerator Energies. *Phys. Rept.*, 3:261–379, (1972).
- [72] R. Bradford, A. Bodek, H. Budd, and J. Arrington. A new parameterization of the nucleon elastic form factors. *Nuclear Physics B-Proceedings Supplements*, 159:127–132, (2006).
- [73] S.L. Adler. Tests of the Conserved Vector Current and Partially Conserved Axial-Vector Current Hypotheses in High-Energy Neutrino Reactions. *Phys. Rev.*, 135(4B):B963, (1964).
- [74] G.T. Jones *et al.* Experimental Test of the PCAC-Hypothesis in Charged Current Neutrino and Antineutrino Interactions on Protons. *Zeitschrift für Physik C Particles and Fields*, 37 (1):25–38, (1987). URL <http://dx.doi.org/10.1007/BF01442065>.

- [75] A. Bodek and U.K. Yang. Higher Twist, ξ_w Scaling, and effective LO PDFs for Lepton Scattering in the Few GeV Region. *J. Phys.*, G29:1899–1906, (2003).
- [76] T. Yang, C. Andreopoulos, H. Gallagher, K. Hoffmann, and P. Kehayias. A Hadronization Model for Few-GeV Neutrino Interactions. *The European Physical Journal C*, 63(1):1–10, (2009).
- [77] Z. Koba, H.B. Nielsen, and P. Olesen. Scaling of Multiplicity Distributions in High-Energy Hadron Collisions. *Nucl. Phys.*, B40:317–334, (1972).
- [78] D. Rein and L.M. Sehgal. Coherent π^0 Production in Neutrino Reactions. *Nucl. Phys.*, B223:29, (1983).
- [79] D. Rein and L.M. Sehgal. PCAC and the Deficit of Forward Muons in π^+ Production by Neutrinos. *Phys. Lett.*, B657:207–209, (2007).
- [80] J.S. Marshall. *A study of muon neutrino disappearance with the MINOS detectors and the NuMI neutrino beam*. PhD thesis, FERMILAB-THESIS-2008-20, University of Cambridge, (2008).
- [81] R. Barlow and C. Beeston. Fitting Using Finite Monte Carlo Samples. *Comp. Phys. Comm.*, 77:219–228, (1993).
- [82] T. Adye. Unfolding Algorithms and Tests Using RooUnfold. *arXiv:1105.1160*, (2011).
- [83] B.G. Tice. *Measurement of Nuclear Dependence in Inclusive Charged Current Neutrino Scattering*. PhD thesis, Rutgers, The State University of New Jersey, (2014).
- [84] J. Wolcott. Uncertainties due to Feynman scaling of NA49 data to NuMI energy. [MINER \$\nu\$ A Internal Note TN15](#), (2013).
- [85] Ž. Pavlović. *Observation of Disappearance of Muon Neutrinos in the NuMI beam*. PhD thesis, FERMILAB-THESIS-2008-59, University of Texas at Austin, (2008).
- [86] J. Park. Beam angle bias estimation. [MINER \$\nu\$ A Internal Document docdb-8659v1](#), (2013).
- [87] H. Gallagher. Summary of minerva estimates of genie systematics. [MINER \$\nu\$ A Internal Document docdb-7578v3](#), (2012).
- [88] H. Gallagher. Personal Communication, (2013).
- [89] S.A. Dytman, H. Gallagher, and M. Kordosky. Hadronic Shower Energy Scale Uncertainty in the MINOS Experiment. *arXiv:0806.2119[hep-ex]*, (2008).
- [90] B. Eberly. Charged pion update. [MINER \$\nu\$ A Internal Document docdb-9523v1](#), (2014).
- [91] C. McGivern. Test beam tracking efficiencies for pion production systematic uncertainties. [MINER \$\nu\$ A Internal Document docdb-9725v3](#), (2014).

- [92] R. Gran. Testbeam pion calorimetry figures for summer 2012. [MINER \$\nu\$ A Internal Document docdb-7529v6](#), (2012).
- [93] J. Devan. Test beam mc upgrades and proton calorimetry. [MINER \$\nu\$ A Internal Document docdb-8701v1](#), (2013).
- [94] J.P. Velásquez, S.A. Dytman, and A.M. Gago. Geant4 hadronic uncertainties. [MINER \$\nu\$ A Internal Document docdb-8249v1](#), (2012).
- [95] L. Fields and K. McFarland. Neutron interaction reweighting procedure. [MINER \$\nu\$ A Internal Document docdb-7518v3](#), (2012).
- [96] B. Eberly. Birks uncertainty strategy. [MINER \$\nu\$ A Internal Document docdb-9526v3](#), (2014).
- [97] Hamamatsu Photonics K.K. *Photomultiplier Tubes: Basics and Applications*, (2007).