

Fermilab

Three-flavour neutrino oscillations with MINOS and CHIPS

FERMILAB-THESIS-2017-41

This manuscript has been authored by Fermi Forward Discovery Group, LLC under Contract No. 89243024CSC000002 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

Three-flavour neutrino oscillations with MINOS and CHIPS

Andrew John Perch
University College London

Submitted to University College London in fulfilment
of the requirements for the award of the
degree of **Doctor of Philosophy**

January 17, 2017

Declaration

I, Andrew John Perch confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Andrew Perch

Abstract

MINOS was a long-baseline neutrino oscillation experiment comprising two functionally identical detectors that observed Fermilab's NuMI neutrino beam in its low energy tune, at distances of 1 km and 735 km. When NuMI switched to a higher-power medium energy tune in 2012, the MINOS detectors continued to operate as MINOS+. Since its commissioning in 2003, the MINOS Far Detector has also been able to detect atmospheric neutrinos. Atmospheric neutrino oscillations are sensitive to the mass splitting Δm_{32}^2 and mixing angle θ_{23} , and are also subject to the Matter Effect as the neutrinos pass through the Earth, which affects neutrinos and antineutrinos differently in a way that depends upon the mass hierarchy. This thesis presents the first atmospheric neutrino analysis using data from the MINOS+ era, and the first dedicated MINOS atmospheric neutrino analysis to use a full three-flavour mixing model. It includes 10.79 kiloton years of new data and encompasses almost an entire period of the 11-year solar cycle, from 2003 to 2014.

The CHIPS experiment aims to reduce construction costs of large water Cherenkov detectors to \$200-300k per kiloton by submerging detectors with a lightweight structure in bodies of water on the surface of the Earth. Such detectors could reach masses of 1 Mton and would assist with the search for CP violation in the neutrino sector by measuring the rate of ν_e appearance in a ν_μ beam. A detailed reconstruction framework for CHIPS has been developed, incorporating a novel method based on the timing of PMT hits. This framework has been used to study the performance of different designs for a 10 kiloton CHIPS R&D module, and to demonstrate

that ν_e events can be identified in a sparsely-instrumented detector with a 6% coverage of 3" PMTs.

Acknowledgements

There are very many people without whom this thesis would never have come into being, first and foremost my parents Ken and Anne who encouraged and supported me through every stage of my education.

I'd also like to thank Jenny for being a dedicated and passionate supervisor, Tricia and Ryan for coordinating the Standard Oscillations group ahead of Neutrino 2014, and Andy Blake for sharing his wealth of experience with the atmospheric neutrino framework.

Throughout the course of my PhD I've met a great many people and made several long-lasting friendships. Included in that list are all of the 2012 HEP cohort at UCL: Becca, Ben, Laura and Xin Ran, my postdoc Leigh for providing advice and good humour in equal measure, and my C17 office-mates past and present. From my time at Fermilab I'd like to thank Ashley my flatmate and fellow MINOS LTA, Michelle and Joe who shared the Standard Oscillations workload, and all the other members of Young MINOS for making the Collaboration such a friendly, welcoming and successful one.

Lastly, I'd like to thank everyone with whom I shared two fantastic summers in Northern Minnesota, with a special mention for Dave, Denise, Jack, Jerry and everyone else working at Soudan, both for tirelessly keeping the Far Detector running and for all their help with CHIPS-M.

It's been an incredible four years.

Contents

List of figures	11
List of tables	17
1 Neutrino oscillations: history and theory	19
1.1 Discovery of the neutrino	20
1.1.1 Other neutrino flavours	21
1.2 Neutrino oscillations	22
1.2.1 The Solar Neutrino problem	22
1.2.2 Oscillations in the atmospheric sector	24
1.3 Formalism	25
1.3.1 Neutrino mixing formalism	26
1.3.2 Oscillations in vacuum	28
1.3.3 Neutrino oscillations in matter	30
1.3.4 Matter effect in two flavour mixing	32
2 Measuring neutrino oscillations	34
2.1 Neutrino interactions with matter	34
2.2 Two neutrino approximation	36
2.3 Solar oscillation parameters: Δm_{21}^2 and $\sin^2(\theta_{12})$	36
2.4 Atmospheric and accelerator neutrinos: Δm_{32}^2 and $\sin^2(\theta_{23})$	37
2.5 Short-baseline reactor antineutrinos and θ_{13}	39
2.6 Open questions	40
2.6.1 Octant of θ_{23}	40
2.6.2 Mass hierarchy	41
2.7 CP-violation	44
2.8 Summary of oscillation parameters	47

3	The MINOS Far Detector	49
3.1	Overview	49
3.2	The MINOS Far Detector	50
3.2.1	Scintillator planes	51
3.3	Instrumentation	53
3.4	Readout electronics	54
3.4.1	Overview	54
3.4.2	The VA chip and VA Front-end Board	56
3.4.3	The VMM and VARC	56
3.4.4	Data acquisition and triggering	57
3.5	Calibration	58
3.5.1	FD Calibration Stages	59
3.5.2	Measuring linearity: LI and CI systems	60
3.5.3	Strip-to-strip calibration with cosmic muons	61
3.5.4	Energy scale calibration: Muon Energy Units	62
4	Monte Carlo simulation	67
4.1	Cosmic rays	67
4.2	Atmospheric neutrino MC	70
4.3	Cosmic muon MC	72
4.4	Detector simulation	75
5	Data Quality	77
5.1	Run quality	77
5.2	High Voltage status	78
5.3	Magnet coil current	78
5.4	GPS timing uncertainty	79
5.5	Data quality results	79
5.6	Summary of data used for the analysis	79
6	Veto shield	82
6.0.1	Veto criteria	83
7	Timing calibration	86
7.1	Timing fits to through-going cosmics	87
7.2	Time-walk	88

7.3	Hardware swaps	88
7.4	Strip to strip calibration	89
7.5	Validation	90
8	Event selection	93
8.1	Event classification	93
8.2	Preselection	94
8.2.1	Definition of good events:	94
8.2.2	Track-like and shower-like events	94
8.2.3	Refining track-like and shower-like events	95
8.3	Containment	97
8.4	Selecting Contained Vertex Muons	98
8.4.1	Stage 1: Fiducial volume	98
8.4.2	Stage 2: Trace cut	98
8.4.3	Stage 3: Topology cuts	99
8.4.4	Stage 4: Futher topology requirements	100
8.5	Selecting CV showers	103
8.5.1	Stage 1: Fiducial Volume	103
8.5.2	Stage 2: Trace Cut	103
8.5.3	Stage 3: Topology Cuts	104
8.6	Neutrino Induced Muons	106
8.6.1	Classifying event directions	107
8.6.2	Selecting neutrino induced muons: event sample	108
8.6.3	Stage 1: Track Length	108
8.6.4	Stage 2: Direction from Timing	108
8.7	Avoiding double-counting	111
8.8	Summary of selected events	112
9	Oscillation Fit	114
9.0.1	Fit method	114
9.0.2	Oscillation templates	115
9.1	Systematic errors for the atmospheric sample	117
9.1.1	Normalisation uncertainties	118
9.1.2	Spectral shape uncertainties	120
9.1.3	Reconstruction uncertainties	121
9.1.4	Summary of systematic errors	122

10 Fit results	124
10.0.5 Atmospheric neutrino spectra	128
11 CHIPS	131
11.0.6 Water Cherenkov detectors	131
11.0.7 Location	133
11.1 CHIPS-M	133
11.2 CHIPS-10	135
11.2.1 Novel analysis strategies for water Cherenkov detectors	136
12 Reconstruction for CHIPS	138
12.1 Reconstruction considerations	138
12.1.1 Hough transformation for ring reconstruction	139
12.2 Reconstruction method	144
12.3 Calculating the likelihood	145
12.3.1 Track parameters	145
12.4 The charge component of the likelihood	146
12.4.1 Predicting the number of photons at the PMT: isotropic point-like source	146
12.4.2 Predicting photons for an extended, non-isotropic source	147
12.5 Calculating the components of the predicted charge	152
12.5.1 Number of photons $\Phi(E)$	152
12.5.2 Transmission function, $T(R)$	152
12.5.3 PMT Solid Angle, $\Omega(R)$	155
12.5.4 PMT efficiency	156
12.5.5 Evaluating the sums	157
12.5.6 Converting the charge prediction into a likelihood	159
12.6 The time component of the likelihood	161
12.6.1 Predicting the PMT hit time	161
12.6.2 Method for predicting PMT hit times	163
12.6.3 Emission profiles for timing	163
12.6.4 Predicting the arrival time	164
12.6.5 Smoothly parametrising the arrival times	168
12.6.6 Accounting for hits with multiple photoelectrons	170
12.6.7 Propagation time of particles	173
12.6.8 Calculating the time contribution to the likelihood	173

12.7	Minimising the log-likelihood	175
12.8	Extending to multiple tracks	176
12.8.1	Seeding π^0 events	177
13	CHIPS Reconstruction Results	180
13.1	Comparison of 3" and 10" tubes	180
13.2	Comparison of 10% and 6% coverage	185
13.3	Particle identification: Simple Cuts	188
13.4	ANN-based particle identification	189
13.4.1	ANN to separate CCQE ν_e and ν_μ events	193
13.4.2	ANN to separate CCQE ν_e events from NC events	193
13.4.3	Combining PID networks	194
13.5	Fitting pions	197
14	Summary and conclusion	199
	Bibliography	201

List of figures

2.1	Feynman diagrams for different neutrino interaction types	35
	(a) NC elastic scattering	35
	(b) CC quasielastic scattering	35
	(c) CC resonant scattering	35
	(d) CC deep inelastic scattering	35
2.2	Illustration of the two mass hierarchies	42
2.3	Oscillogram comparing neutrino and antineutrino oscillations in the Normal and Inverted Hierarchies	43
2.4	Oscillation probabilities for a beam of muon neutrinos	46
2.5	Toy MC showing oscillations applied to a ν_μ beam with a Gaussian initial spectrum	47
3.1	Photograph of MINOS Far Detector	51
3.2	Diagram of MINOS scintillator strip	53
3.3	Schematic of FD readout electronics	55
3.4	Bethe-Bloch equation for energy loss of charged particles in different materials	64
3.5	Definition of track window used to calculate MEU for energy scale calibration	65
4.1	Cosmic ray flux for different primary particles	68

4.2	Combined ‘all-particle’ cosmic ray flux	69
4.3	Atmospheric neutrino flux at the Soudan Underground Laboratory	73
5.1	Summary of data quality checks for new atmospheric data	80
6.1	Schematic and photograph of veto shield	83
6.2	Time to nearest shield hit	84
6.3	Veto shield efficiency over time	85
7.1	Variables used for timing calibration fit	88
7.2	Example of timing jump	89
7.3	Comparison of uncalibrated and calibrated East-West time differences	91
7.4	Mean East-West time difference by calibration period	92
	(a) 1st Feb. - 31st Jul. 2012	92
	(b) 1st Aug. 2012 - 31st Jan. 2013	92
	(c) 1st Feb. - 31st Jul. 2013	92
	(d) 1st Aug. - 31st Jan. 2014	92
8.1	Z-trace variable for selecting CV ν_μ events	99
8.2	Distribution of $\langle \Delta_{UV} \rangle$ and $\langle \Delta_{UV}^2 \rangle$ variables for selecting CV ν_μ events	101
8.3	Distribution of ΔR_{max} and Q_{vtx} variables for selecting CV ν_μ events	101
8.4	Distribution of Q_{vtx} against $\cos \theta_Z$ used for selecting CV ν_μ events .	102
8.5	Shower z-trace variable used for selecting CV showers	104
8.6	Moment of inertia variables used for selecting CV showers	106
8.7	Distribution of $\frac{1}{\beta}$ demonstrating the separation of downward cosmic muons from upward NIM events	110

8.8	Distribution of reconstructed zenith angle for events with well-defined directions	111
8.9	Rate at which candidate atmospheric neutrinos are selected over time	113
9.1	Four-layer model of the Earth used for calculating oscillation probabilities	116
9.2	Fluxes and uncertainties for muon neutrinos and antineutrinos . . .	118
9.3	Uncertainty in the up/down and up/horizontal flux ratios	122
10.1	Two dimensional contours for Δm_{32}^2 and $\sin^2(\theta_{23})$	125
10.2	One dimensional profiled likelihood surface for Δm_{32}^2	126
10.3	One dimensional profiled likelihood surface for $\sin^2(\theta_{23})$	127
10.4	Comparison of detected and predicted CV ν_μ and $\bar{\nu}_\mu$ spectra at the best fit point	129
10.5	Comparison of detected and predicted NIM spectra at the best fit point	130
11.1	Photographs of the CHIPS site at the Wentworth 2W Pit	133
	(a) Wentworth 2W Pit (aerial photo)	133
	(b) Wentworth 2W Pit	133
11.2	Photograph of CHIPS-M prototype shortly before deployment . . .	134
11.3	Early design drawing of CHIPS-10 structure	135
11.4	Energy spectra for leading muons and electrons at the CHIPS site .	137
12.1	Alternative parametrisation of a straight line used for Hough transformations	140
12.2	Illustration of a Hough transformation for ring finding	142
	(a)	142

(b)	142
(c)	142
12.3	Output from a Hough transform for photons with different separation angles	143
(a)	Ring separation 90°	143
(b)	Ring separation 30°	143
(c)	Ring separation 10°	143
12.4	Calculating the linear emission profile ρ	148
12.5	Predicted number of photons without the angular emission profile g	149
12.6	Interpretation of the angular emission profile, for three different bins	150
12.7	Emission profiles for 1 GeV electrons and muons	151
12.8	Number of Cherenkov photons produced by MC electrons, as a function of their energy	153
12.9	Transmission length used in the fitter	154
12.10	Comparison of absorption lengths and Cherenkov wavelength spectrum	155
12.11	Diagram of PMT protruding through the inner liner	156
12.12	Model of PMT simulation	160
12.13	Event display showing charge and time of PMT hits	162
12.14	Timing emission profiles	165
12.15	Diagram showing the calculation of the predicted hit time for a single step along the track	166
12.16	Predicted arrival times at a PMT from steps along the track, and the weights	167
12.17	Example of Catmull-Rom spline used to interpolate between four equidistant points	169

12.18	PDF describing the arrival time of 3 photoelectron, compared to a toy MC	172
12.19	Overlap between predicted arrival time PDF and a Gaussian describing the PMT resolution	174
12.20	Situation in which the two-track minimiser may effectively fit a single track	178
13.1	Reco - true vertex coordinates for ν_e CCQE events	181
	(a) Vertex X-coordinate	181
	(b) Vertex Y-coordinate	181
	(c) Vertex Z-coordinate	181
13.2	Reco - true vertex coordinate for ν_μ CCQE events	182
	(a) Vertex X-coordinate	182
	(b) Vertex Y-coordinate	182
	(c) Vertex Z-coordinate	182
13.3	Angle, time and energy reco - true distributions for ν_e CCQE events	183
	(a) Angle between reconstructed and true directions	183
	(b) Track energy	183
	(c) Vertex time	183
13.4	Angle, time and energy reco - true distributions for ν_μ CCQE events	184
	(a) Angle between reconstructed and true directions	184
	(b) Track energy	184
	(c) Vertex time	184
13.5	Reco - true distributions for ν_e CCQE events, comparing 10% and 6% coverage	186
	(a) Distance from true vertex	186

	(b) Angle to true track	186
	(c) Track energy	186
	(d) Vertex time	186
13.6	Reco - true distributions for ν_μ CCQE events, comparing 10% and 6% coverage	187
	(a) Distance from true vertex	187
	(b) Angle to true track	187
	(c) Track energy	187
	(d) Vertex time	187
13.7	Comparison of charge and time likelihood components between electron and muon fits to ν_e and ν_μ events	188
13.8	Efficiencies for selecting electrons and rejecting muons using simple cuts on the likelihood components	189
	(a) Electron selection efficiency	189
	(b) Muon rejection efficiency	189
13.9	Purity of selected ν_e CCQE sample as a function of selection efficiency	190
13.10	Output of ANN to separate ν_e CCQE from ν_μ CCQE events	194
13.11	Output of ANN to separate ν_e CCQE events from NC events	195
13.12	Efficiency \times purity figure of merit as a function of ANN cuts	196
13.13	Distribution of ANN selection variables for different event types	196
	(a) CCQE ν_μ vs ν_e classifier	196
	(b) CCQE vs NC classifier	196
13.14	Selection efficiency and purity for different event types	197
13.15	Invariant mass of two photons when a sample of clean π^0 events is reconstructed	198

List of tables

2.1	Best fit oscillation parameters their 1σ uncertainties from the global analysis [1] performed by the NuFit group [2]. Note that Δm_{3i}^2 refers to the larger of atmospheric mass splittings (Δm_{31}^2 in the NH and Δm_{32}^2 in the IH).	48
5.1	Summary of good data used for atmospheric neutrino analysis	81
8.1	Number of selected events of each category compared to no-oscillation prediction	112
8.2	Number of selected events events of each category compared to the prediction with nominal oscillation parameters	112
9.1	Grid points used to generate oscillated templates for GhostFitter	116
9.2	Summary of the sources of systematic error and their uncertainties	123
10.1	Best fit oscillation parameters and $-2\Delta \ln \mathcal{L}$ in each hierarchy	127
10.2	Systematic nuisance parameters at the best fit	128
13.1	Comparison of resolutions for a 10 kton detector with 10% coverage of 10 inch or 3 inch PMTs	184
13.2	Comparison of resolutions for 10% and 6% coverage	186
13.3	Relative rate of different interaction types assumed when training the particle ID	190
13.4	Variables used in PID neural network	192

Chapter 1

Neutrino oscillations: history and theory


The field of neutrino physics has a rich history, from the proposal of a mysterious, undetected particle to explain the behaviour of beta decays in the 1930s through to the resolution of a 30-year problem with the confirmation of oscillations in the early 2000s, and on into the precision era of the coming decades. Neutrino oscillations were first discussed in 1957, when Bruno Pontecorvo proposed a model [3] in which neutrinos could oscillate into antineutrinos and back, similar to the mixing of the neutral kaon system [4]. In the decades that followed, it was instead shown that neutrinos oscillate from one kind (or flavour) to another, and that antineutrinos do the same.

Experiments that measure neutrino oscillations probe quantum phenomena with subatomic particles propagating over distance scales ranging from the human to the astronomical, providing a beautiful demonstration of quantum interference. They encompass a range of techniques and technologies, and use neutrinos from solar, nuclear, atmospheric and accelerator sources. This chapter presents a brief history of the field, and a description of the relevant theoretical framework underpinning neutrino oscillations.

1.1 Discovery of the neutrino

The neutrino was first postulated in 1931 by Wolfgang Pauli to account for the apparent non-conservation of energy and angular momentum in the beta decay of atomic nuclei. This was known to produce a nucleus with a different atomic number, and an electron. Assuming a simple two-body decay with both nuclei at rest, the electron would be expected to carry away a fixed kinetic energy corresponding the difference in the binding energies of the two nuclei. However, a continuous spectrum of electron energies was observed and prompted Pauli's 'desperate remedy', namely the invention of a second particle, also emitted in the decay, which escaped undetected. This light, neutral, spin- $\frac{1}{2}$ fermion was proposed in a letter to the 'Radioactive ladies and gentlemen' attending a meeting in Tübingen in 1930 [5], and was christened the 'neutron'. Following Chadwick's discovery in 1932 of the heavier nucleon now called the neutron, Pauli's particle was renamed the 'neutrino', from the Italian for 'little neutral one'.

The existence of the (anti)neutrino was experimentally confirmed by Reines and Cowan [6] in 1956, using inverse beta decay in a 200l tank of water to detect antineutrinos produced by a nuclear reactor at Los Alamos. Here, the $\bar{\nu}_e$ from the reactor interacted with protons in the water to produce a positron, which quickly annihilated with an electron in the water yielding two back-to-back 511 keV gamma rays. Shortly afterwards, Goldhaber *et al.* [7] demonstrated that electron neutrinos have negative helicity, using the interaction:



The Eu and Sm nuclei are spin-0 whilst the Sm* is spin-1. In order to conserve angular momentum, if the neutrino spin component along its direction of travel is $\mp\frac{1}{2}$, the photon polarisation must be ± 1 . For resonance scattering where the photon and neutrino are preferentially emitted back-to-back, the photon direction measures the neutrino direction, and its polarisation measures the neutrino spin component along that direction. The neutrino helicity must then be same as that of the photon, and was found to be -1 in agreement with the V-A theory of the weak interaction.

1.1.1 Other neutrino flavours

The Reines-Cowan experiment involved placing a detector a short distance away from an intense source of reactor antineutrinos and detecting them through interactions that produced antielectrons. In 1962, Lederman, Schwartz and Steinberger demonstrated that neutrinos originating primarily from pion decays at the Brookhaven AGS produced not electrons but muons [8], detected as single long tracks in a spark chamber. This demonstrated the existence of a second type of neutrino, the muon neutrino, and was recognised by the 1988 Nobel Prize.

With the discovery of the τ lepton in 1975, a natural question was whether there existed a corresponding third type of neutrino. This was finally discovered in 2000 by the DONUT Collaboration [9] using an emulsion cloud chamber and 800 GeV protons from the Tevatron, with the neutrinos produced in the decay of hadrons created in a tungsten beam dump. Four candidate interactions were identified based on an isolated vertex leading to a kinked track (i.e. the appearance and subsequent decay of a τ) with an expected background of 0.34 events.

These observations demonstrated the existence of three different types of neutrino. When detected through their charged current interactions close to where they were produced, one type produces only electrons, one produces only muons, and one produces only taus. They can therefore be categorised as three ‘flavours’, the ν_e , ν_μ , and ν_τ .

The strongest constraint on the number of active neutrino flavours comes from collider physics. Combined measurements of the line shape of the Z^0 resonance performed by the ALEPH, DELPHI, L3 and OPAL experiments at LEP, in particular the ratio of partial decay widths Γ_{inv}/Γ_{ll} for invisible and charged lepton (e and μ) final states, indicate that the number of active neutrino flavours is 2.984 ± 0.008 [10]. This suggests that any as-yet-undiscovered neutrino states must be sterile (i.e. not coupled to the weak interaction) or be heavier than $0.5m_Z$. Cosmology can also be used to constrain the number of neutrinos, with a fit performed using data from the Planck satellite suggesting 3.30 ± 0.27 species of relativistic neutrino-like particles [11].

1.2 Neutrino oscillations

Throughout the development of the Standard Model of Particle Physics, neutrinos were presumed to be massless and to occur only in the three flavour eigenstates. However, several anomalies developed with this description. Neutrino oscillations, the quantum interference phenomenon by which a neutrino of one flavour can convert to a neutrino of a different flavour, emerged as the explanation for these anomalies, and non-zero masses followed as a direct consequence.

1.2.1 The Solar Neutrino problem

The nuclear fusion processes occurring in the Sun provide a rich source of electron neutrinos. The amount and energy of the neutrinos produced depends upon the specific reaction that produces them. Various fusion reactions combine to form the pp-chain, the overall result of the which is to fuse four protons and produce a ${}^4\text{He}$ nucleus:



where E_ν is the energy of the two neutrinos, and $(26.73\text{MeV} - E_\nu)$ is the energy emitted as photons. Characteristic neutrino fluxes are also produced at individual steps along this chain.

In the late 1960s, Davis and Bahcall set out to use the solar neutrino flux to test the Standard Solar Model (SSM) describing nuclear fusion in the Sun. A 100,000 gallon tank of perchloroethylene dry-cleaning fluid was installed in the Homestake Mine, detecting neutrinos from boron-8 decay via the charged-current neutrino capture reaction:



The tank was periodically flushed with helium to remove and measure the amount of argon produced, and hence the number of detected neutrinos. The initial results were consistent with background, setting an upper limit of 3.0SNU^1 [12] compared

¹1 Solar Neutrino Unit = 1 interaction per ${}^{37}\text{Cl}$ per second

to predictions ranging from 4.4 to 21 SNU, depending on model parameters [13]. This deficit persisted into the 1990s, with the final Homestake results showing $2.56 \pm 0.16(\text{stat.}) \pm 0.16(\text{syst.})$ SNU compared to updated predictions of $9.3_{-1.4}^{+1.2}$ SNU [14], 6.4 ± 1.4 SNU [15], and 7.7 SNU [16], again depending on solar model parameters.

A similar deficit was observed using the Kamiokande-II water Cherenkov detector. This was capable of inferring the neutrinos' directions and demonstrating that they did indeed originate from the Sun, but measured a flux only $0.46 \pm 0.13(\text{stat.}) \pm 0.08(\text{syst.})$ of the SSM prediction [17].

Because of its energy threshold, ν_e capture on chlorine is mostly sensitive to neutrinos from ${}^8\text{B}$ interactions, but capture on ${}^{71}\text{Ga}$ instead of ${}^{37}\text{Cl}$ has a lower energy threshold of 233 keV and is sensitive to the more-abundant pp neutrinos as well. Deficits in this channel were observed by SAGE, which initially reported a rate of 20_{-20}^{+15} SNU [18] and GALLEX which measured 83 ± 21 SNU [19], compared to a SSM prediction of 132 SNU [20]. Again this deficit persisted, with later results of $65.4_{-3.0}^{+3.1}(\text{stat.})_{-2.8}^{+2.6}(\text{syst.})$ and $77.5 \pm 6.2(\text{stat.})_{-4.7}^{+4.3}(\text{syst.})$ from SAGE [21] and GALLEX [22] respectively.

Neutrino oscillations were one suggested mechanism of explaining this deficit: if the ν_e were somehow converting to ν_μ or ν_τ the electrons in the detector media of the radiochemical experiments would be unable to capture them (and solar neutrinos have insufficient energy to produce a μ or τ which would allow their detection in a Cherenkov detector). This was finally conclusively demonstrated by the SNO experiment in 2001 and 2002.

SNO featured a tank containing 1 kton of heavy water (water containing ${}^2\text{H}$ isotopes), and was able to detect solar neutrinos via three different, complementary channels:

$$\nu_e + d \rightarrow p + p + e^- \quad (1.4)$$

$$\nu_x + d \rightarrow p + n + \nu_x \quad (1.5)$$

$$\nu_x + e^- \rightarrow \nu_x + e^- \quad (1.6)$$

where d is a deuterium nucleus. The first channel is a charged current (CC) interaction and is sensitive only to electron neutrinos (solar neutrinos have energies $\sim\text{MeV}$ so can produce electrons but not heavier muons or taus). The second channel is a

neutral current (NC) interaction between a neutrino of any flavour and a deuterium nucleus, and is detected using the gamma rays produced when the dislodged neutron thermalises and is captured by another nucleus. The final channel is elastic scattering of the neutrino from an electron, and is again sensitive to all three neutrino flavours. Comparing the fluxes measured with the three channels and dividing them into ϕ_e , the component attributable to ν_e , and $\phi_{\mu\tau}$, the component attributable to ν_μ and ν_τ , SNO demonstrated a 5.3σ signal of oscillations from ν_e to the other flavours.

1.2.2 Oscillations in the atmospheric sector

Atmospheric neutrinos were first observed experimentally in 1965 by the Kolar Gold Fields [23] and CWI/SAND [24] experiments which operated cosmic ray observatories in extremely deep mines in India (7500 metres water equivalent) and South Africa (8890 m.w.e.) respectively. At such large overburdens, the flux of cosmic rays at high zenith angles is negligibly small. Even if the timing resolution is insufficient to distinguish upward trajectories from downward ones, a horizontal track is unambiguously horizontal. Hence, when horizontal muon events were detected, both experiments correctly realised that these were muons induced by atmospheric neutrinos interacting in the Earth and propagating into the detector.

Oscillations of these atmospheric neutrinos were uncovered by experiments searching for the unrelated (and still undiscovered) phenomenon of proton decay. To search for decays with long lifetimes, detectors with large masses and low backgrounds are required. This led to the construction of large underground water Cherenkov detectors such as IMB and Kamiokande. Atmospheric neutrinos were an important background to such experiments.

At low energies, atmospheric neutrinos are primarily produced from the interactions of cosmic ray protons with nuclei, N , in the atmosphere by:

$$p + N \rightarrow \pi^+ + X \quad (1.7)$$

$$\pi^+ \rightarrow \mu^+ \nu_\mu \quad (1.8)$$

$$\mu^+ \rightarrow e^+ \bar{\nu}_e \bar{\nu}_\mu \quad (1.9)$$

$$\Rightarrow p + N \rightarrow e^+ + \bar{\nu}_e + \bar{\nu}_\mu + \nu_\mu + X \quad (1.10)$$

The charge conjugate applies in the case of π^- production. In the absence of oscillations this leads to a ratio $R = \frac{\nu_\mu + \bar{\nu}_\mu}{\nu_e + \bar{\nu}_e}$ that approaches 2. The absolute flux of each flavour is subject to a number of uncertainties (see Section 9.1.1), many of which cancel when this flavour ratio is taken.

When evaluating their systematic errors, both IMB [25] and Kamiokande [26] reported a deficit of ν_μ and $\bar{\nu}_\mu$ events relative to ν_e and $\bar{\nu}_e$ when comparing data with Monte Carlo. Later dedicated atmospheric neutrino analyses by experiments including Super-Kamiokande [27], MACRO [28] and Soudan 2 [29] also demonstrated a deficit of atmospheric muon neutrinos as a function of their zenith angle, and interpreted it in terms of neutrino oscillation.

1.3 Formalism

The change in flavour of propagating neutrinos can be explained by supposing that the flavour eigenstates (through which neutrinos couple to the weak W and Z bosons) are different from the energy eigenstates for the material through which the neutrinos are travelling. The flavour eigenstates can therefore be expressed in terms of the energy states, and oscillations naturally arise through interference as the different energy states evolve differently with time.

For free particles travelling through a vacuum with momentum \mathbf{p} , energy eigenstates can be treated as plane waves with energies $E = \sqrt{p^2 + m^2}$. Hence in a vacuum the energy eigenstates are states of well-defined mass. Significantly, this implies the existence of (at least) three mass states with different quantum numbers, hence at least two of them must have non-zero masses.

The general case of neutrino oscillations in a medium can be understood in terms of:

- A flavour basis: orthonormal states consisting of ν_e , ν_μ , and ν_τ
- An orthonormal energy basis, consisting of three energy eigenstates for that medium
- Eigenvalues of these energy states
- A unitary matrix that can be used to convert between the two bases

Oscillation experiments typically probe the relation between flavour and mass eigenstates. When neutrinos propagate through a medium of known electron density, the additional transformation required to convert into the mass basis gives rise to departures from vacuum mixing through the so-called Matter Effect. These departures can also be used to constrain the vacuum mixing parameters.

1.3.1 Neutrino mixing formalism

The simplest example of neutrino oscillations are those that occur in a vacuum. Here, the only two bases involved are the flavour and mass bases. We can describe the mixing of mass states to produce flavour states with a matrix U_{ij} so that, adopting notation where mass eigenstates have Roman indices and flavour eigenstates have Greek indices:

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i} |\nu_i\rangle \quad (1.11)$$

If we require that $\langle \nu_\alpha | \nu_\beta \rangle = \delta_{\alpha\beta}$ then the mixing matrix satisfies:

$$\sum_i \sum_j \langle \nu_j | U_{j\alpha}^\dagger U_{\beta i} | \nu_i \rangle = \delta_{\alpha\beta} \quad (1.12)$$

$$\sum_i \sum_j U_{j\alpha}^\dagger U_{\beta i} \langle \nu_i | \nu_j \rangle = \delta_{\alpha\beta} \quad (1.13)$$

$$\sum_i \sum_j U_{j\alpha}^\dagger U_{\beta i} \delta_{ij} = \delta_{\alpha\beta} \quad (1.14)$$

$$\sum_i U_{i\alpha}^\dagger U_{\beta i} = \delta_{\alpha\beta} \quad (1.15)$$

$$\Rightarrow U^\dagger U = \mathbb{I} \quad (1.16)$$

In the case of exactly three neutrino flavours, the 3×3 unitary mixing matrix is known as the PMNS matrix (named for Pontecorvo, Maki, Nakata and Sakawa) [30] [31]. It

parametrises the mixing of mass states to produce flavour ones as follows:

$$\begin{pmatrix} |\nu_e\rangle \\ |\nu_\mu\rangle \\ |\nu_\tau\rangle \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \\ |\nu_3\rangle \end{pmatrix} \quad (1.17)$$

The PMNS matrix is analogous to the CKM (Cabibbo-Kobayashi-Maskawa) matrix that describes quark mixing, whose unitarity has been regularly and precisely tested. Combined fits to the CKM matrix currently are consistent with unitarity [32].

Under these assumptions, U_{PMNS} is a unitary, complex, 3×3 matrix. Like all $n \times n$ unitary complex matrices it can, in general, be described with n^2 independent parameters, of which $\frac{n(n-1)}{2}$ are angles and $\frac{n(n+1)}{2}$ are phases. Hence U_{PMNS} can be rewritten in terms of three mixing angles, θ_{12} , θ_{23} , and θ_{13} , and six complex phases. Because it describes the mixing between particle fields, some of these phases can be freely removed without affecting any physical processes by absorbing them into the neutrino fields.

If neutrinos are Dirac particles then five phases can be absorbed in this manner, leaving a single CP-violating phase δ_{CP} behind. For Majorana neutrinos only three of the phases can be removed, leaving δ_{CP} and two Majorana mass phases α_{21} and α_{31} which have no effect on neutrino oscillations [33].

This gives the standard formulation of the PMNS matrix:

$$U_{\text{PMNS}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{\text{CP}}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{\text{CP}}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (1.18)$$

$$= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-\delta_{\text{CP}}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{\text{CP}}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{\text{CP}}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{\text{CP}}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{\text{CP}}} & c_{23}c_{13} \end{pmatrix} \quad (1.19)$$

where $s_{ij} = \sin \theta_{ij}$ and $c_{ij} = \cos \theta_{ij}$

In the case of Majorana neutrinos the mixing matrix becomes:

$$U_{\text{PMNS}}^{\text{Dirac}} \times \text{diag}(1, e^{i\frac{\alpha_{21}}{2}}, e^{i\frac{\alpha_{31}}{2}}) \quad (1.20)$$

1.3.2 Oscillations in vacuum

To calculate the oscillation probabilities for neutrinos produced in a given flavour state, we must use the mixing matrix to convert into the basis of energy eigenstates, evolve these in time, and project the evolved states back into the flavour basis using the mixing matrix once more. This yields an amplitude which can then be squared to produce the oscillation probability.

1.3.2.1 Approximate energy of a neutrino propagating in vacuum

Consider a neutrino of some flavour α produced at time $t = 0$ with a well-defined momentum p_α . Following [34], we can write its initial wavefunction as a plane wave, choosing coordinates such that the wave propagates along the x-direction:

$$\psi(x, t = 0) = \sum_{i=1}^3 U_{\alpha i} e^{ip_\alpha x} \quad (1.21)$$

And after a time t' has elapsed:

$$\psi(x, t = t') = \sum_{i=1}^3 U_{\alpha i} e^{(ip_\alpha x)} e^{(-iE_i t)} \quad (1.22)$$

But $E_i = \sqrt{p_\alpha^2 + m_i^2}$ which can be expanded under the assumption $m_i \ll p_\alpha$ to give:

$$E_i \approx p_\alpha + \frac{m_i^2}{2p_\alpha} \quad (1.23)$$

The assumption of $m_i \ll p_\alpha$ also implies that the neutrino states are travelling at approximately c ($c = 1$ in these units) and hence that $x(t) = ct = t$. So we can replace t with x in the above, yielding:

$$\psi(x, t) \approx \sum_i U_{\alpha i} e^{(-i[m_i^2/2p_\alpha]x)} \quad (1.24)$$

A more complete treatment of neutrino oscillations would use a wavepacket formalism to account for the slightly different propagation speeds of components with different masses but the same momentum. Under this formalism the wavepackets slowly separate, and the two mass states cease being coherent after propagating a distance $L \geq L_{coh} = d/\Delta\beta$, where d is the initial length of the wavepacket and $\Delta\beta$ is the difference in their propagation speeds as a fraction of c [35]. Full wavepacket treatments considering the length of the neutrino source [36] and detector [37] have been performed, but for distances less than an extremely large number of oscillation lengths the results are the same as when using the approximation above.

Using this approximation, we can now calculate the oscillation probability in vacuum, using the PMNS matrix.

1.3.2.2 Vacuum oscillation probability

In a vacuum, if we produce a state $|\psi(t=0)\rangle = |\nu_\alpha\rangle = \sum_j U_{\alpha j} |\nu_j\rangle$ then:

$$|\psi(t)\rangle = \sum_j U_{\alpha j} e^{-i[m_j^2/2p_\alpha]x} |\nu_j\rangle \quad (1.25)$$

By multiplying by the inverse of U_{PMNS} we can also express mass eigenstates in terms of flavour ones:

$$|\nu_j\rangle = \sum_\gamma U_{\gamma j}^* |\nu_\gamma\rangle \quad (1.26)$$

Substituting $|\nu_j\rangle$ into Equation 1.25, the amplitude for an oscillation from flavour α at $x = 0$ to flavour β at distance x is:

$$\langle \nu_\beta | \psi(x) \rangle = \sum_j \sum_\gamma \langle \nu_\beta | U_{\alpha j} e^{-i[m_j^2/2p_\alpha]x} U_{\gamma j}^* |\nu_\gamma\rangle \quad (1.27)$$

$$= \sum_j \sum_\gamma U_{\gamma j}^* U_{\alpha j} e^{-i[m_j^2/2p_\alpha]x} \delta_{\beta\gamma} \quad (1.28)$$

$$= \sum_j U_{\alpha j} U_{\beta j}^* e^{-i[m_j^2/2p_\alpha]x} \quad (1.29)$$

So the general expression for the oscillation probability is:

$$P(\nu_\alpha \rightarrow \nu_\beta) = |\langle \nu_\beta | \psi(x) \rangle|^2 \quad (1.30)$$

$$= \sum_i \sum_j U_{\beta i}^* U_{\alpha j} U_{\beta i} U_{\alpha j} e^{i[(m_i^2 - m_j^2)/2p_\alpha]x} \quad (1.31)$$

Converting to more conventional units, we have:

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sum_i \sum_j U_{\beta j}^* U_{\alpha j} U_{\beta i} U_{\alpha i} e^{(i \times 2.534 \Delta m_{ij}^2 L/E)} \quad (1.32)$$

Where $\Delta m_{ij}^2 = m_i^2 - m_j^2$ in units of $\frac{eV^2}{c^4}$, $\frac{L}{E}$ is the ratio of the experiment's baseline to the energy of the neutrino source (which is approximately equal to p_α) in $\frac{\text{km}}{\text{GeV}}$ or $\frac{\text{m}}{\text{MeV}}$, and the numerical factor converts to these non-SI units.

Written in this format the oscillatory pattern becomes evident. The oscillations arise from the superposition of mass eigenstates, and their period depends on the difference between the squared masses. The size of the oscillations depends both on this interference and directly on the elements of the PMNS matrix.

Note also that neutrino oscillations require the mass eigenstates to have different masses (otherwise the oscillatory part goes as $e^0 = 1$), and hence an observation of oscillations described by this formalism requires that neutrinos are not (all) massless.

1.3.3 Neutrino oscillations in matter

Neutrino oscillations in matter can be described in a similar fashion, with additional complexity due to the presence of electrons in the medium through which the neutrinos propagate. While all neutrino flavours can scatter off the electrons through the exchange of a Z boson, the electron neutrino component can also interact through coherent forward scattering via a W boson.

This can be accounted for by introducing a potential to the Hamiltonian describing a neutrino's propagation. Let \hat{H}_0 be the Hamiltonian describing neutrino oscillations

in vacuum. In the mass basis, \hat{H}_0 is given by:

$$\hat{H}_0 = \frac{1}{2E} \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta m_{21}^2 & 0 \\ 0 & 0 & \Delta m_{31}^2 \end{pmatrix} + \frac{m_1^2}{2E} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (1.33)$$

This can be rewritten in the flavour basis via a transformation using U_{PMNS} :

$$\hat{H}_0 = U_{\text{PMNS}} \begin{pmatrix} 0 & 0 & 0 \\ 0 & \frac{\Delta m_{21}^2}{2E} & 0 \\ 0 & 0 & \frac{\Delta m_{31}^2}{2E} \end{pmatrix} U_{\text{PMNS}}^\dagger + \frac{m_1^2}{2E} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (1.34)$$

In the flavour basis it is then simple to include the matter effect by adding a potential in the e-e channel:

$$\hat{H}_{\text{matter}} = \hat{H}_0 + \begin{pmatrix} \sqrt{2}G_F n_e & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (1.35)$$

Where G_F is the Fermi constant, and n_e is the electron density of the matter through which the neutrino is propagating.

To solve for the neutrino oscillation probabilities, we must find the matrix M that allows \hat{H}_{matter} to be decomposed into the form $\hat{H}_{\text{matter}} = M\Lambda M^{-1}$, where Λ is a diagonal matrix of eigenvalues of \hat{H}_{matter} and the rows in M^{-1} are its eigenvectors. Then M is the analogue of U_{PMNS} for a material of constant density, and the algebra just performed has allowed us to describe it in terms of $\theta_{12}, \theta_{23}, \theta_{13}$, and δ_{CP} . We can then reuse the result in Equation 1.31, substituting $U_{\text{PMNS}} \leftrightarrow M$ and $e^{i[(m_k^2 - m_j^2)/2p_\alpha]x} \leftrightarrow e^{i[\Lambda_k - \Lambda_j]x}$.

Diagonalising \hat{H}_{matter} is simple, albeit algebraically tedious, and is performed exactly by the MINOS software when calculating oscillation probabilities in matter. The following section describes the matter effect in the more simple case of only two neutrino flavours to give an indication of its behaviour.

1.3.4 Matter effect in two flavour mixing

It is instructive to consider the matter effect as it applies to two flavour neutrino mixing between ν_a and ν_b . In this case, there is a single mass splitting Δm^2 , and the vacuum mixing matrix becomes:

$$U = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \quad (1.36)$$

Multiplying out the two-dimensional equivalent of Equation 1.34, and neglecting the term proportional to the identity matrix (which has no effect on oscillations), we find that:

$$\hat{H} = \begin{pmatrix} V - \frac{\Delta m^2}{4E} \cos 2\theta & \frac{\Delta m^2}{4E} \sin 2\theta \\ \frac{\Delta m^2}{4E} \sin 2\theta & \frac{\Delta m^2}{4E} \cos 2\theta \end{pmatrix} + \frac{\Delta m^2}{4E} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad (1.37)$$

Again, the second component is a diagonal matrix and can be disregarded. The diagonalisation becomes more convenient if we also subtract a term equal to the identity matrix multiplied by $V/2$ to finally obtain:

$$\hat{H} = \frac{\Delta m^2}{4E} \begin{pmatrix} A - \cos 2\theta & \sin 2\theta \\ \sin 2\theta & \cos 2\theta - A \end{pmatrix} + \text{const.} \times \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad (1.38)$$

with $A = \frac{2\sqrt{2}G_F n_e E}{\Delta m^2}$.

Diagonalising \hat{H} we find:

$$\hat{H} = U_M \Lambda U_M^\dagger \quad (1.39)$$

with:

$$U_M = \begin{pmatrix} \cos \theta_M & \sin \theta_M \\ -\sin \theta_M & \cos \theta_M \end{pmatrix} \text{ and } \Lambda = \begin{pmatrix} 0 & 0 \\ 0 & \Delta m_m^2 \end{pmatrix} \quad (1.40)$$

where:

$$\sin^2 2\theta_M = \frac{1}{C} \sin^2 2\theta \quad (1.41)$$

$$\Delta m_M^2 = C \Delta m^2 \quad (1.42)$$

and:

$$C = \sqrt{(\cos 2\theta - A)^2 + \sin^2 2\theta} = \sqrt{\left(\cos 2\theta - \frac{2\sqrt{2}G_F n_e E}{\Delta m^2}\right)^2 + \sin^2 2\theta} \quad (1.43)$$

Inspecting Equation 1.42 shows that $\sin^2 2\theta_M = 1$, and hence the effective mixing becomes maximal, when $\cos 2\theta = A$, i.e. $\sqrt{2}G_F n_e = \frac{\Delta m^2}{2E} \cos 2\theta$. Neutrinos produced in the core of the Sun pass through regions of varying n_e as they propagate to the surface, and undergo this resonant enhancement of their oscillations despite the non-maximal value of θ_{12} which drives the solar oscillations. This effect is known as the MSW (Mikheyev-Smirnov-Wolfenstein) Effect [38] [39]. When n_e is zero, we find $C = 1$ and the formula for mixing in a vacuum is recovered.

For antineutrinos, the potential term V acquires a negative sign, and so the resonance occurs when $\sqrt{2}G_F n_e = -\frac{\Delta m^2}{2E} \cos 2\theta$. This allows the sign of the Δm^2 terms to be determined, and resonant enhancement in the Sun is the reason that the sign of the solar mass splitting is known [40] [41]. The Matter Effect also proves relevant for atmospheric neutrinos and long-baseline accelerator experiments (especially those measuring ν_e appearance) as in both cases the neutrinos travel long distances through the Earth.

Chapter 2

Measuring neutrino oscillations

The goal of neutrino oscillation experiments is to infer the values of the various components of the PMNS matrix through fits to the period and magnitude of oscillations, based on the number and energy distribution of detected neutrinos. To make such measurements it is important to understand which parameters are relevant to which oscillation channels.

This chapter gives an overview of the current status of each mass splitting and PMNS mixing angle, paying particular attention to those relevant to the MINOS atmospheric neutrino analysis, and to the appearance of ν_e in a ν_μ beam which is relevant for CHIPS.

2.1 Neutrino interactions with matter

In the Standard Model, neutrinos interact only through the weak force. These interactions can be divided into two classes depending on the weak vector boson that mediates them: charged current (CC) interactions are mediated by W^\pm bosons and neutral current (NC) interactions are mediated by the neutral Z^0 boson.

Figure 2.1 shows Feynman diagrams for various classes of neutrino interaction. The simplest is NC elastic scattering, where the neutrino and nucleon simply scatter off one another. Charged current quasielastic (CCQE) scattering is similar, but the interaction is mediated by a W boson and so the neutron is converted into a proton and the ν_e to an electron. This is the dominant channel below around 1 GeV and produces a clean signal with a single charged lepton track whose energy can be used

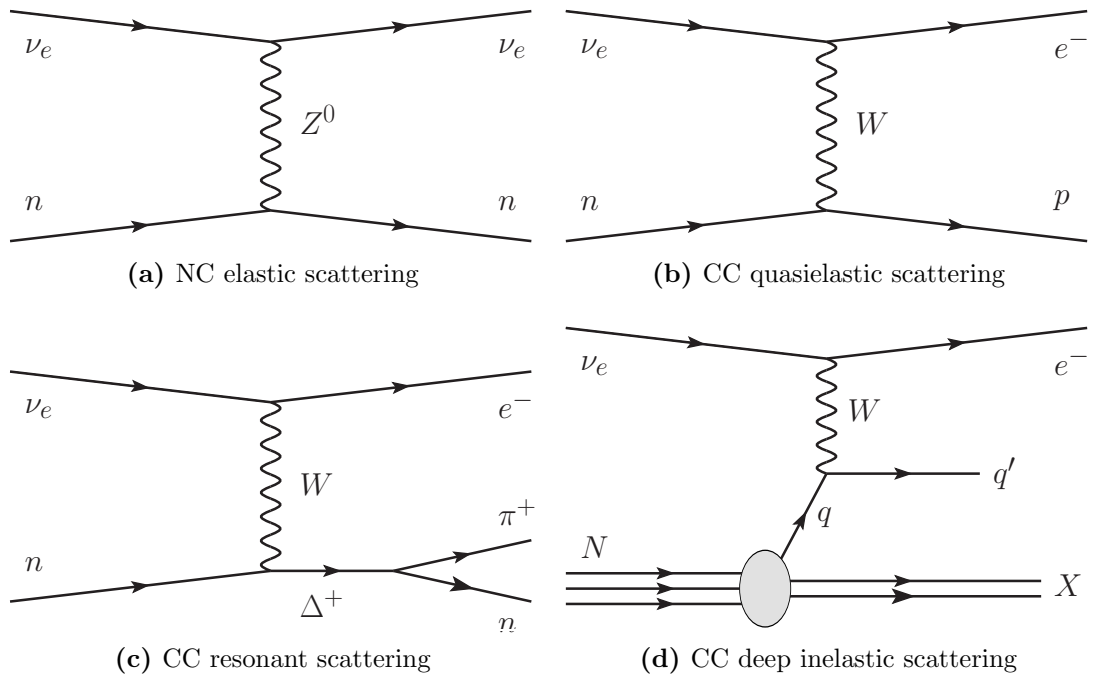


Figure 2.1: Feynman diagrams for different types of neutrino interactions.

to reconstruct that of the neutrino. A CC interaction may also involve the production of a resonance, typically a Δ state, whose subsequent decay produces some hadronic activity as well as the outgoing charged lepton. This channel is most significant from 1-3 GeV. At higher energies, deep inelastic scattering dominates. Here, the neutrino interacts with an individual quark inside the nucleon, producing substantial hadronic activity due to both the outgoing quark and the nucleon remnants.

Resonances can also accompany NC interactions, and coherent NC interactions also occur when the neutrino transfers a sufficiently small amount of energy that the target nucleus remains intact. Both of these event categories produce hadronic activity without there being a charged lepton in the final state, and can be a source of neutral pions. Interactions between neutrinos and electrons in the target medium occur with lower probability; neutrinos of all flavours can undergo NC elastic scattering from electrons, while electron neutrinos can also elastically scatter via a CC interaction.

2.2 Two neutrino approximation

While the oscillation probabilities can be calculated exactly using matrix methods, useful insights and intuitions can be gained by expanding them in terms of the mixing angles θ_{12} , θ_{13} and θ_{23} .

Historically neutrino oscillations were analysed using a two-flavour model because the mass splittings responsible for atmospheric oscillations, $\Delta m_{32}^2 \approx \Delta m_{31}^2$, are approximately 30 times larger than the Δm_{21}^2 splitting that drives solar oscillations. For this reason, the effect of considering a full three flavour treatment was small compared to the statistical and systematic errors involved. The two flavour paradigm describes oscillations measured by accelerator and atmospheric experiments operating on L/E scales of $\sim \text{km/GeV}$ reasonably well.

Under the approximation of a single mass splitting, Δm^2 and a single mixing angle θ , we find the survival probability for ν_μ is:

$$P(\nu_\mu \rightarrow \nu_\mu) \approx 1 - \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 L}{4E}\right) \quad (2.1)$$

where the approximate equality would become exact if there truly were only two neutrino flavours. The combined MINOS limits from accelerator and atmospheric neutrinos and antineutrinos yielded $\Delta m^2 = 2.41_{-0.10}^{+0.09} \times 10^3 \text{ eV}^2$ and $\sin^2 2\theta = 0.950_{-0.036}^{+0.035}$ [42].

With the discovery that θ_{13} was both non-zero and, in fact, reasonably large [43], and as larger data sets cause the uncertainty in Δm^2 to approach the level of Δm_{21}^2 , the validity of the two-flavour approximation begins to break down.

2.3 Solar oscillation parameters: Δm_{21}^2 and $\sin^2(\theta_{12})$

The dominant results in the Δm_{21}^2 and $\sin^2(\theta_{12})$ sector are due to the SNO and KamLAND experiments, measuring solar and reactor neutrinos respectively, and combining their results with a range of other solar neutrino experiments. Combined fits by SNO [44] to solar neutrino data collected from its three run configurations

using a three-flavour oscillation model resulted in values of $\tan^2(\theta_{12}) = 0.436_{-0.036}^{+0.048}$ and $\Delta m_{21}^2 = 5.13_{-0.58}^{+1.49} \times 10^{-5} \text{ eV}^2$.

KamLAND was a reactor neutrino experiment constructed to test neutrino oscillations at the L/E scale suggested by the solar neutrino experiments, but using reactor antineutrinos which have a much higher average energy so require a shorter oscillation distance. It consisted of a 1 kton liquid scintillator detector positioned between several nuclear reactors at a flux-averaged baseline of 180 km, and reported $\tan^2(\theta_{12}) = 0.56_{-0.07}^{+0.10}$ and $\Delta m_{21}^2 = 7.58_{-0.13}^{+0.14} \times 10^{-5} \text{ eV}^2$.

The Particle Data Group's 2015 Review of Particle Physics [45] quotes the results of a combined fit to solar neutrinos, Kamland, short-baseline reactor experiments, and accelerator experiments [46], which finds $\Delta m_{21}^2 = 7.53 \pm 0.18 \times 10^{-5} \text{ eV}^2$ and $\sin^2(\theta_{21}) = 0.846_{-0.021}^{+0.021}$.

2.4 Atmospheric and accelerator neutrinos: Δm_{32}^2 and $\sin^2(\theta_{23})$

The Δm_{32}^2 and $\sin^2(\theta_{23})$ sector is probed by complementary experiments using atmospheric neutrinos and accelerator neutrinos. A beam of primarily muon neutrinos can be produced at a particle accelerator by directing a proton beam towards a target, using magnets to focus the charged pions and kaons this produces, and allowing these to decay downstream.

The main channel relevant to these parameters is the survival of muon neutrinos. Writing out the sum in Equation 1.31, and using the identities $|U_{\mu 1}|^2 + |U_{\mu 2}|^2 + |U_{\mu 3}|^2 \equiv 1$ for a unitary matrix, and $(e^{i\alpha} + e^{-i\alpha}) \equiv 2 \cos \alpha \equiv 2(1 - 2 \sin^2(\alpha/2))$ gives:

$$P(\nu_\mu \rightarrow \nu_\mu) = 1 - 4|U_{\mu 1}|^2|U_{\mu 2}|^2 \sin^2\left(\frac{\Delta m_{21}^2 L}{2E}\right) \quad (2.2)$$

$$- 4|U_{\mu 1}|^2|U_{\mu 3}|^2 \sin^2\left(\frac{\Delta m_{31}^2 L}{2E}\right) \quad (2.3)$$

$$- 4|U_{\mu 2}|^2|U_{\mu 3}|^2 \sin^2\left(\frac{\Delta m_{32}^2 L}{2E}\right)$$

Assuming Δm_{21}^2 is negligible, its term can be ignored and Δm_{31}^2 can be set equal to Δm_{32}^2 . We find:

$$P(\nu_\mu \rightarrow \nu_\mu) \approx 1 - \sin^2(\theta_{23}) \cos^2 \theta_{13} (1 - \sin^2(\theta_{23}) \cos^2(\theta_{13})) \sin^2 \left(\frac{1.267 \Delta m_{32}^2 L}{E} \right) \quad (2.4)$$

Given that θ_{13} is also known to be small (see Section 2.5) we can set $\cos^2 \theta_{13} \approx 1$ to finally obtain the two-flavour approximation quoted in Equation 2.1:

$$P(\nu_\mu \rightarrow \nu_\mu) \approx 1 - \sin^2(2\theta_{23}) \sin^2 \left(\frac{1.267 \Delta m_{32}^2 L}{E} \right) \quad (2.5)$$

This demonstrates that the two-flavour approximation can be applied to the atmospheric neutrino regime, with the mixing angle and mass splitting roughly corresponding to θ_{23} and Δm_{32}^2 .

Two conclusions can be drawn from this process. Firstly, measuring the disappearance of ν_μ provides a sensitive probe of θ_{23} and Δm_{32}^2 . Secondly, in deriving this expression a number of reasonable assumptions were made that relied upon Δm_{21}^2 and θ_{13} being small. These allowed terms containing $U_{\mu 1}$ and $U_{\mu 2}$ (which carry factors of δ_{CP}) to be neglected, the difference between Δm_{31}^2 and Δm_{32}^2 to be ignored, and terms in $\sin^2(\theta_{23})$ to be converted into $\sin^2(2\theta_{23})$. Conversely, if our aim is to measure ‘three flavour effects’ such as the mass hierarchy (whether $|\Delta m_{32}^2|$ is larger or smaller than $|\Delta m_{31}^2|$), the octant of θ_{23} , the size of θ_{13} , and the presence or extent of CP-violation, then doing so via the ν_μ survival channel requires enough sensitivity to measure terms that were previously negligible.

Recent conventional measurements of atmospheric neutrinos have been reported by Super-K [47]. Neutrino telescopes, primarily searching for ultra-high energy neutrinos have also demonstrated sensitivity to atmospheric neutrino oscillations, with consistent results available from the IceCube [48] and ANTARES [49] experiments.

Measuring the same parameters using accelerator neutrinos was pioneered by K2K [50], with later results from MINOS [51] and T2K [52] substantially improving the constraints. In recent years, reactor electron antineutrino experiments aiming to measure θ_{13} have also contributed to the atmospheric mixing sector by measuring Δm_{31}^2 or a closely related effective mass splitting [53].

Because the ordering of the mass state m_3 relative to the other two states ($m_2 > m_1$) is not yet known, results for the mass splittings Δm_{32}^2 and Δm_{31}^2 are usually quoted for the normal hierarchy ($m_3 > m_2 > m_1$) and the inverted hierarchy ($m_2 > m_1 > m_3$) separately. The current limits from the Particle Data Group, based on a combined fit to the published results from MINOS, T2K and Daya Bay are $\Delta m_{32}^2 = 2.49 \pm 0.06 \times 10^{-3} \text{ eV}^2$ assuming the normal mass hierarchy, and $|\Delta m_{32}^2| = 2.42 \pm 0.06 \times 10^{-3} \text{ eV}^2$ assuming the inverted hierarchy. The PDG values for the mixing angle are $\sin^2(\theta_{23}) = 0.514_{-0.056}^{+0.055}$ (NH) and 0.511 ± 0.055 (IH) based on [52], which is consistent with the MINOS result of $0.41_{-0.06}^{+0.26}$ (NH) and $0.41_{-0.07}^{+0.28}$ (IH) [51].

2.5 Short-baseline reactor antineutrinos and θ_{13}

The θ_{13} sector is the most recent to be measured, but is now also the most tightly constrained. It was known that θ_{13} must be small because measurements of the two-flavour atmospheric mixing angle found near-maximal disappearance (so $\cos^2 \theta_{13}$ must be close to 1), but this means it is difficult to measure using solar or atmospheric neutrinos. Instead, $\sin^2(\theta_{13})$ is measured using electron antineutrinos produced in nuclear reactors.

The expression for the $\bar{\nu}_e$ survival probability is the same as in Equation 2.2 but uses the U_{ei} row instead of the $U_{\mu i}$ row of the PMNS matrix. However the different oscillation parameters mean that different assumptions about which parameters are small must be made. Neglecting the term in Δm_{21}^2 and assuming $\Delta m_{31}^2 \approx \Delta m_{32}^2$ (again for illustrative purposes only):

$$P(\bar{\nu}_e \rightarrow \bar{\nu}_e) \approx 1 - 4 \cos^2(\theta_{12}) \cos^2(\theta_{13}) \sin^2(\theta_{13}) \sin^2\left(\frac{\Delta m_{31}^2 L}{4E}\right) \quad (2.6)$$

$$\begin{aligned} & - 4 \sin^2(\theta_{12}) \cos^2(\theta_{13}) \sin^2(\theta_{13}) \sin^2\left(\frac{\Delta m_{31}^2 L}{4E}\right) \\ & \approx 1 - \sin^2(2\theta_{13}) \sin^2\left(\frac{\Delta m_{31}^2 L}{4E}\right) \end{aligned} \quad (2.7)$$

Three reactor-based experiments provide measurements of $\sin^2(2\theta_{13})$: Daya Bay [53], RENO [54] and Double Chooz [55]. There are also measurements with

higher uncertainties that use ν_e appearance in ν_μ accelerator beams, from T2K [56], MINOS [57], and such appearance has also been seen by NO ν A [58].

For θ_{13} the PDG reference value is $\sin^2(2\theta_{13}) = 0.085 \pm 0.005$, constructed from a weighted average of reactor experiments' results.

2.6 Open questions

Despite the excellent progress that has been made in characterising neutrino oscillations, a number of significant open questions remain. These generally relate to the transition from the two-flavour to the three-flavour experimental landscape, where it becomes desirable to measure parameters that were previously considered small and neglected. The three main unknowns concern the octant of θ_{23} , the ordering of the three neutrino mass states, and whether or not neutrino oscillations violate charge-parity (CP) symmetry.

2.6.1 Octant of θ_{23}

Atmospheric and accelerator experiments have determined that $\sin^2(2\theta_{23}) \approx 1$, and hence that $\theta_{23} \approx 45^\circ$. In the case of two-flavour mixing, $\sin^2(2\theta) = 1$ corresponds to maximal disappearance (i.e. at some point all ν_μ of a given energy will have oscillated away) although this ceases to be the case with non-zero θ_{13} .

However if $\theta_{23} \neq 45^\circ$, its octant (that is, whether $\theta_{23} < 45^\circ$ or $> 45^\circ$) is uncertain. While the degeneracy between θ and $\frac{\pi}{2} - \theta$ is exact for two-flavour oscillations, there are also approximate three-flavour degeneracies in terms depending on $\sin^2(\theta_{23})$ where fits involving other poorly-constrained parameters can produce similar oscillation probabilities with the 'wrong' best-fit.

For ν_μ survival, the vacuum oscillation probability can be expanded to first order in the form of two-flavour oscillations [59]:

$$P(\nu_\mu \rightarrow \nu_\mu) = 1 - \sin^2(2\theta_{\text{eff}}) \sin^2\left(\frac{1.267\Delta m_{\text{eff}}^2 L}{E}\right) \quad (2.8)$$

Where:

$$\Delta m_{\text{eff}}^2 = \Delta m_{32}^2 + \Delta m_{21}^2 (\sin^2 \theta_{12} + \cos \delta_{CP} \sin \theta_{13} \tan \theta_{23} \sin 2\theta_{12}) \quad (2.9)$$

$$\sin^2 2\theta_{\text{eff}} = 4 \sin^2 \theta_{23} \cos^2 \theta_{13} (1 - \sin^2 \theta_{23} \cos^2 \theta_{13}) \quad (2.10)$$

The MINOS beam disappearance analysis therefore has some access to the octant information, but its effect depends upon the level at which $\cos^2 \theta_{13}$ differs from 1.

Electron neutrino appearance analyses have a better sensitivity to the octant because the leading order expansion of the oscillation probability goes as:

$$P(\nu_\mu \rightarrow \nu_e) = \sin^2(\theta_{23}) \sin^2(2\theta_{13}) \sin^2\left(\frac{1.267 \Delta m_{31}^2 L}{E}\right) \quad (2.11)$$

Although the overall rate of ν_e appearance is suppressed by $\sin^2(2\theta_{13})$, this can be well-constrained by reactor experiments, and so the size of the appearance signal gives a direct probe of $\sin^2(\theta_{23})$.

The situation regarding ν_μ disappearance is improved when atmospheric neutrinos are considered because of the additional matter effect. A similar two-flavour-like expansion for neutrinos passing through matter can be obtained by replacing θ_{13} in the definition of θ_{eff} from Equation 2.10 with [60]:

$$\sin^2(2\theta_{13}) \rightarrow \sin^2(2\theta_M) = \frac{\sin^2(2\theta_{13})}{(A - \cos(2\theta_{13}))^2 + \sin^2(2\theta_{13})} \quad (2.12)$$

With $A = \pm 2\sqrt{2}G_F n_e E / \Delta m_{31}^2$ and the plus (minus) sign applying to neutrinos (antineutrinos). As with two flavour mixing in matter, there exists a resonance condition. Here the resonance occurs for $A = \cos(2\theta_{13})$ which is satisfied for neutrinos travelling thousands of kilometres through the Earth with energies ~ 5 GeV, and counteracts the usual difficulty associated with small θ_{13} .

2.6.2 Mass hierarchy

The neutrino oscillation probabilities in vacuum depend only on terms proportional to $\sin^2(1.267 \Delta m_{ij}^2 L / E)$ and therefore have no sensitivity to the sign of Δm_{ij}^2 . The sign of Δm_{21}^2 is known to be positive, i.e. $m_2 > m_1$, because of MSW resonance

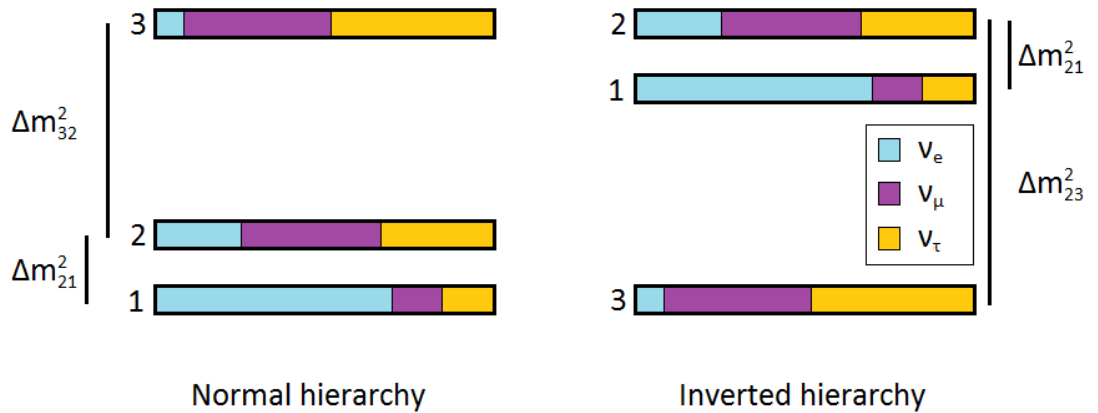


Figure 2.2: Illustration of the two possible mass hierarchies. In the normal hierarchy, m_3 is the mass of the heaviest state and Δm_{32}^2 is positive. In the inverted hierarchy, m_3 is the lightest mass, and so Δm_{32}^2 is negative. The ordering $m_2 > m_1$ is known from the solar oscillations and the MSW effect. The blue, purple and yellow shading indicates the approximate ν_e , ν_μ and ν_τ composition of each mass state.

effects in the Sun. The magnitude of the atmospheric mass splitting is ~ 30 times that of the solar mass splitting, so it can be deduced that m_1 and m_2 are close together, with a larger separation between themselves and m_3 .

This leaves a twofold ambiguity in the order of the mass states, which is illustrated in Figure 2.2. One possibility is that $m_3 \gg m_2 > m_1$. In this case, $|\nu_1\rangle$, the state with the largest ν_e component, is the lightest mass state. This is reminiscent of the charged lepton sector, and so is called the Normal Hierarchy (NH). The alternative is called the Inverted Hierarchy (IH), where $m_2 > m_1 \gg m_3$, causing Δm_{31}^2 and Δm_{32}^2 to be negative.¹

Atmospheric neutrinos are able to access the mass hierarchy through the matter effect. This can be seen by examining the expression for $\cos \theta_M$ in Equation 2.12, or by analogy with the exact form for two-flavour mixing given in Equation 1.42. Here, Δm_{31}^2 appears in a position that is sensitive to its sign. The energy required for resonance is:

$$E = \pm \Delta m_{31}^2 \left(\frac{\cos(2\theta_{13})}{2\sqrt{2}G_F n_e} \right) \quad (2.13)$$

¹Some discussions instead refer to Normal Ordering (NO) and Inverted Ordering (IO).




Figure 2.3: Oscillogram showing the difference in oscillation probabilities between full three-flavour mixing in matter and two-flavour vacuum mixing, given nominal oscillation parameters. The patterns are similar for neutrinos in the normal hierarchy and antineutrinos in the inverted hierarchy, with the most prominent differences occurring at low $\cos \theta_{\text{zenith}}$ where the neutrinos have travelled a long distance through the Earth and the matter effect is strong. Plot courtesy of J. Coelho.

As before, the plus sign applies to neutrinos and the minus sign applies to antineutrinos. Therefore, if the normal hierarchy is correct, then $\Delta m_{31}^2 > 0$ and the matter resonance can occur for neutrinos. Conversely, if the hierarchy is inverted then $\Delta m_{31}^2 < 0$ and we require the additional minus sign associated with antineutrinos to allow the resonance to take place. This is illustrated in Figure 2.3.

In summary, the mass hierarchy can be determined by comparing the oscillations of neutrinos and antineutrinos under the influence of the matter effect. Atmospheric neutrinos provide the matter effect, and magnetised experiments like MINOS are well-suited to comparing neutrinos and antineutrinos. A high statistics sample is, as always, also desirable.

2.7 CP-violation

A process is said to conserve C (charge) symmetry if it is invariant under charge conjugation, i.e. the replacement of all particles with their antiparticles and the reversal of all electromagnetic fields. Similarly, P (parity) symmetry corresponds to invariance under a parity transformation, inverting all coordinates through the origin so that $(x, y, z) \rightarrow (-x, -y, -z)$.

The weak interaction maximally violates both C and P symmetry: its V-A nature couples the W boson only to neutrino states with left-handed chirality, and to right-handed antineutrinos. Charge conjugation converts a ν_L state to a $\bar{\nu}_L$ whilst parity inversion converts a ν_L to a ν_R . Neither of these states couple to the W, and so the weak interaction is not symmetric under these transformations.

Applying a combined CP transformation requires using the Hermitian conjugate of U_{PMNS} , and so neutrino oscillations conserve CP if $U_{\text{PMNS}} = U_{\text{PMNS}}^\dagger$. For the elements that contain a term in $e^{\pm i\delta_{CP}}$, this can only be satisfied if $\delta_{CP} = 0$ or π . Hence a measurement of δ_{CP} that excludes these two values constitutes a demonstration that neutrino oscillations violate CP symmetry.

CP-violation is known to occur in the quark sector [61], but the existence or extent of CP-violation in the neutrino sector is unknown and is a major area of interest for the coming generation of oscillation experiments. Specifically, CP-violation is one of the Sakharov conditions required to generate the baryon charge asymmetry evident in the universe today [62]. The level of CP-violation in the quark sector is insufficient to explain this asymmetry, but at the high energies present in the early universe a lepton charge asymmetry could be converted into a baryon charge asymmetry by sphaleron processes [63], Standard Model processes which conserve (B-L) but violate (B+L). Thus a lepton charge asymmetry due to CP-violation in the neutrino sector could provide an explanation of baryogenesis.

In particular, if a model employing the Seesaw Mechanism introduces a heavy, right-handed (sterile), Majorana neutrino state N , with leptonic decays $N \rightarrow l^- + X^+$ and $N \rightarrow l^+ + X^-$, then CP-violation would allow these two decay modes to proceed with different rates. Attractively, such models could simultaneously account for leptogenesis and the extremely low mass of the conventional neutrinos relative to other Standard Model particles [64]. Detecting CP-violation would not confirm such

a model (it makes no statement about the existence of Majorana neutrinos, amongst other reasons), but would provide important constraints and insights.

Examining the PMNS matrix shows that measuring CP-violation is likely to be difficult because the $e^{\pm i\delta_{\text{CP}}}$ terms all contain a prefactor of $\sin\theta_{13}$. Any CP-violating contribution to oscillations will therefore be suppressed by a factor $\lesssim 0.147$.

In the absence of matter effects, there is no CP-violation in the $\nu_\mu \rightarrow \nu_\mu$ channel - this is because the CP-conjugate process $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ is also the CPT conjugate. Even in matter, Equation 2.10 shows that ν_μ disappearance is a particularly challenging channel as this suppression is compounded by Δm_{21}^2 and other trigonometric functions of the mixing angles.

CP-violation is most-easily probed via ν_e appearance in a ν_μ beam. An approximate form of this oscillation probability in matter, expanded to second order, is [65]:

$$\begin{aligned}
 P(\nu_\mu \rightarrow \nu_e) = & \sin^2(\theta_{23})\sin^2(2\theta_{13})\frac{\sin^2(\Delta_{31} - aL)}{(\Delta_{31} - aL)^2}\Delta_{31}^2 & (2.14) \\
 & + \left(\sin(2\theta_{23})\sin(2\theta_{13})\sin(2\theta_{12})\frac{\sin(\Delta_{31} - aL)}{(\Delta_{31} - aL)} \right. \\
 & \quad \left. \times \Delta_{31}\frac{\sin(aL)}{(aL)}\Delta_{21}\cos(\Delta_{31} + \delta_{\text{CP}}) \right) \\
 & + \cos^2(\theta_{23})\sin^2(2\theta_{12})\frac{\sin^2(aL)}{(aL)^2}\Delta_{21}^2
 \end{aligned}$$

Where $a = G_F n_e / \sqrt{2}$, $\Delta_{ij} = \Delta m_{ij}^2 L / 4E$, and L remains the experiment baseline. The CP conjugate probability, $P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)$, is obtained by replacing $a \leftrightarrow -a$ and $\delta_{\text{CP}} \rightarrow -\delta_{\text{CP}}$.

The ν_e appearance and ν_μ disappearance probabilities are illustrated in Figure 2.4, for a matter density of 1.36 mol cm^{-3} which is the same as is used in MINOS beam analyses [51]. As described, the ν_μ survival channel is not sensitive to δ_{CP} , but the ν_e appearance channel is. The effect of δ_{CP} becomes even more noticeable if the baseline is long enough to allow the second oscillation maximum to be resolved and compared with the first. Figure 2.5 illustrates the same effect in terms of event yields from a toy Monte Carlo, and a clear difference in event counts is apparent as

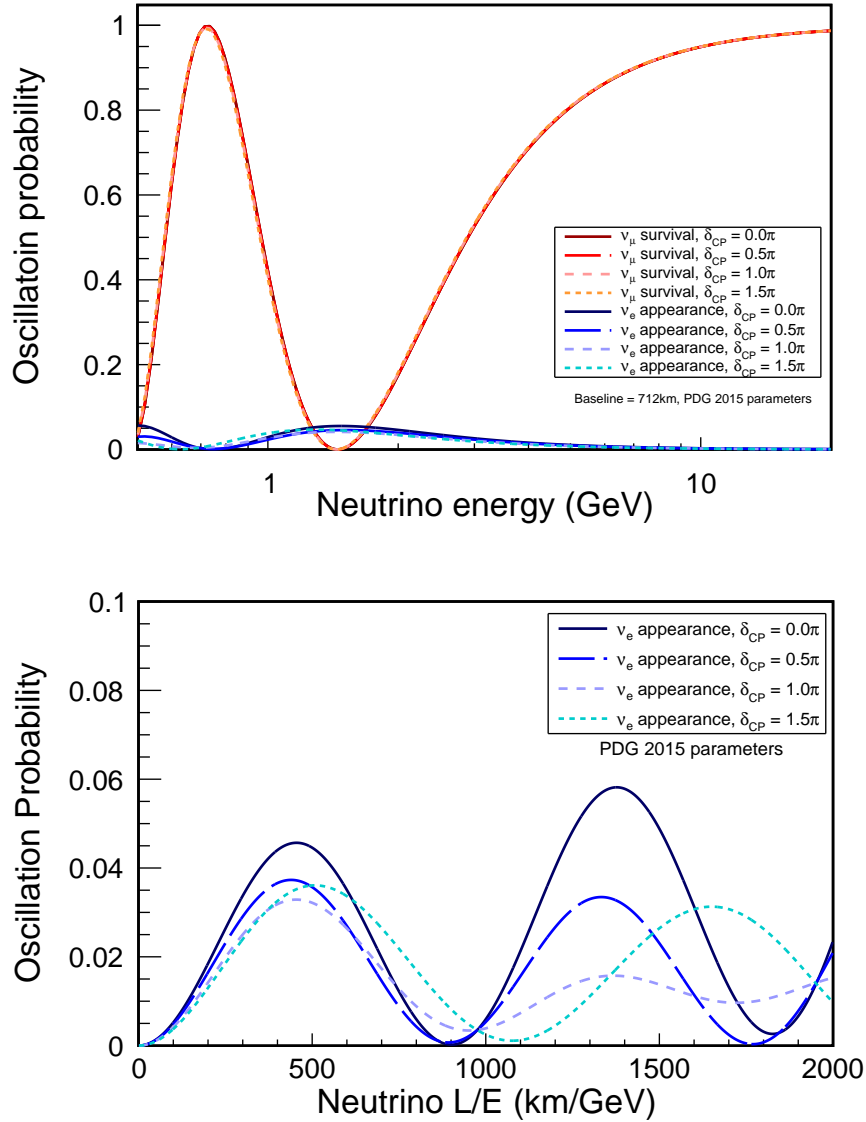


Figure 2.4: Oscillation probabilities for ν_e and ν_μ , assuming a beam initially composed of ν_μ . Probabilities are shown as a function of neutrino energy (top) assuming a 712 km baseline as δ_{CP} is varied from 0 (darkest solid line) to $\frac{3}{2}\pi$ (lightest colour, smallest dashes) in steps of $\frac{\pi}{2}$. The bottom figure shows the ν_e appearance probabilities as a function of L/E. Both plots assume the normal hierarchy and take all the oscillation parameters except δ_{CP} from the 2015 update to [45].




Figure 2.5: Toy Monte Carlo oscillated spectra showing the number of ν_e events produced when oscillations are applied to a sample of 10000 pure ν_μ events distributed according to a Gaussian with width 1.5 GeV and mean 3 GeV. A clear decrease in the number of detected events is observed as δ_{CP} is varied from 0 (dark blue solid line) to $\frac{3}{2}\pi$ (pale blue dotted line). Oscillations are calculated assuming the normal hierarchy and using world-average parameters from the 2015 update to [45] for all parameters except δ_{CP}

δ_{CP} is varied. While there is a spectral component, the most obvious effect is in the overall number of events detected.

Current constraints on δ_{CP} are weak, although recent publications from T2K [56] and NO ν A [58] provide tantalising hints of a non-zero value.

2.8 Summary of oscillation parameters

Global fits, using data from a range of experiments and attempting to simultaneously constrain all the PMNS angles and mass differences provide a concise summary of the present state of the field. The results from one such global analysis [1] are presented in Table 2.1.

Parameter	NH best fit	IH best fit
$\Delta m_{3l}^2/10^{-3} \text{ eV}^2$	$2.524^{+0.039}_{-0.040}$	$2.514^{+0.038}_{-0.041}$
$\Delta m_{21}^2/10^{-5} \text{ eV}^2$	$7.50^{+0.19}_{-0.17}$	$7.50^{+0.19}_{-0.17}$
$\sin^2(\theta_{12})$	$0.306^{+0.012}_{-0.012}$	$0.306^{+0.012}_{-0.012}$
$\sin^2(\theta_{23})$	$0.441^{+0.027}_{-0.021}$	$0.587^{+0.020}_{-0.024}$
$\sin^2(\theta_{13})$	$0.02166^{+0.00075}_{-0.00075}$	$0.02179^{+0.00076}_{-0.00076}$
$\delta_{\text{CP}}/^\circ$	261^{+51}_{-59}	277^{+40}_{-46}

Table 2.1: Best fit oscillation parameters their 1σ uncertainties from the global analysis [1] performed by the NuFit group [2]. Note that Δm_{3l}^2 refers to the larger of atmospheric mass splittings (Δm_{31}^2 in the NH and Δm_{32}^2 in the IH).

Chapter 3

The MINOS Far Detector

3.1 Overview

The MINOS (Main Injector Neutrino Oscillation Search) experiment is a long-baseline experiment designed to study neutrino oscillations. It consists of two functionally identical sampling tracking calorimeters which are exposed to the NuMI (Neutrinos from the Main Injector) beam, a wideband neutrino beam produced at the Fermi National Accelerator Laboratory in Batavia, IL.

The Near Detector has a mass of 980 ton and is situated at a baseline of 1.04 km downstream of the beam target, 100 m underground. At this distance, and assuming standard three-flavour mixing, oscillations are negligible so it is used to measure the energy spectrum of the beam. The larger Far Detector has a mass of 5.4 kton and is located 705 m underground at the Soudan Mine, a former iron mine in Northern Minnesota, at a baseline of 735 km. It is used to re-measure the neutrino energy spectrum so that the oscillation probabilities can be determined by extrapolating the Near Detector spectrum to the Far Detector and taking the ratio between the no-oscillation prediction and the measured spectrum. This two-detector method has the major advantage that many systematics, such as those relating to the beam composition and interaction cross-sections, approximately cancel in the ratio.

The atmospheric neutrino analysis only selects events from outside the beam spill windows and does not make use of the Near Detector due to its small size and prohibitively shallow overburden. The design and construction of both the MINOS Near and Far detectors are described in great detail in [66] and [67]. This chapter

provides a brief summary of the design, readout and calibration of the MINOS Far Detector.

3.2 The MINOS Far Detector

The MINOS Far Detector is a magnetised sampling tracking calorimeter made of steel and plastic scintillator. It began taking commissioning data with cosmic rays and atmospheric neutrinos in September 2002, and became fully operational the following year. The detector took cosmic ray and atmospheric neutrino data from July 2003 onwards, supplementing neutrino data from the NuMI beam, which came online in March 2005. Until 2016 data was still being taken under the name of MINOS+, reflecting the increased beam energy and intensity of NuMI, which has been reconfigured to provide neutrinos to the off-axis NO ν A experiment. MINOS+ running came to an end on 29th June 2016 and decommissioning of the Far Detector has since begun, while the Near Detector continues to operate a muon spectrometer for the MINER ν A experiment.

The detector itself has a total mass of 5.4kton and is located in Northern Minnesota, at latitude 47.82°N and longitude 92.24°W. It resides 705 m below the surface of the Earth (2020 m water-equivalent) in the Soudan Mine, a former iron mine now operated by the Minnesota Department of Natural Resources, at a net altitude 210m below sea level. Figure 3.1 shows a photograph of the detector from the downstream end, with the dimensions and major components highlighted.

The FD is an octagonal prism whose long axis (the z-axis) lies parallel to the NuMI beam direction. The octagonal faces have a diameter of 8 m between opposite sides, and the entire detector is 31 m long. It is divided up into two separate regions, called supermodules (abbreviated ‘SM’), each of which has its own separate magnetic coil. Supermodule 1 is 14.78 m long and is separated by 1.15 m from SM2, which is 14.10 m long.

The two supermodules contain a total of 486 steel planes, 249 in SM1 and 237 in SM2, each of which are 2.54 cm (one inch) thick. Interleaved between these planes are 484 layers of 1 cm thick plastic scintillator, made from extruded polystyrene. An




Figure 3.1: Photograph of the MINOS Far Detector, indicating the steel plate, the two orthogonal scintillator planes behind it, and the magnetic coil through the centre of the detector

air gap is also present, leading to an average separation between successive steel planes of 5.94 cm. The front planes of each supermodule are uninstrumented.

A full detector plane is too large to fit down the narrow elevator shaft at the Soudan Mine. Instead, they were assembled underground, with each $8\text{ m} \times 8\text{ m} \times 2.54\text{ cm}$ steel sheet constructed from eight smaller ones, 2 m wide and 1.27 cm thick, plug welded together. The mean density of the steel was measured to be 7.85 g cm^{-3} [66].

3.2.1 Scintillator planes

The scintillator planes comprise 192 separate strips, each 4.1 cm across and between 3.4 m and 8 m long depending on its position within the plane. Alternating scintillator planes are aligned with the strips at $\pm 45^\circ$ to the vertical. This orthogonal arrangement allows the 2D position of a track crossing several successive planes to be determined, and the 45° angle avoids the need to place connectors and electronics on the underside of the detector. It also means that successive planes are symmetric about the beam-axis, and the vertical axis about which the cosmic ray and neutrino spectrum is also approximately symmetric.

The arrangement of the scintillator in alternating planes defines the U-V coordinate system, where:

$$\hat{\mathbf{U}} = \frac{1}{\sqrt{2}} (\hat{\mathbf{x}} + \hat{\mathbf{y}}) \quad (3.1)$$

$$\hat{\mathbf{V}} = \frac{1}{\sqrt{2}} (-\hat{\mathbf{x}} + \hat{\mathbf{y}}) \quad (3.2)$$

Here $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ represent unit vectors parallel to the x- and y-axes.

The physical mechanism underpinning this design is that the steel plates provide a heavy target with which the neutrinos can interact, and the charged particles produced then pass through the layers of plastic, exciting molecules in the scintillator which emit photons as they relax. The scintillation photons are collected using lengths of wavelength shifting (WLS) fibres which transport them to the end of the scintillator strip, and direct them via a length of clear fibre into a multi-anode photomultiplier tube (PMT).

The WLS fibres, manufactured by Kuraray, are 1.2 mm in diameter and are embedded in grooves carved into the top of scintillator strips. The tops of the grooves are covered with a reflective aluminised Mylar tape to increase light collection. The WLS fibre serves several purposes: it has a smaller absorption length than the scintillator allowing longer strips to be used, it increases the wavelength of the scintillation light from ~ 420 nm at emission to ~ 520 nm where the quantum efficiency of the PMT is higher, and it directs the light from the scintillator down the clear fibres and into the readout system.

Within each plane, scintillator strips are organised into submodules containing 20 or 28 strips placed side by side. These are glued together to improve their rigidity and covered with a 0.25 mm thick coating containing a mixture of polystyrene and 15% TiO_2 which provides a reflective surface so that more of the scintillation light reflects into the WLS fibre. The complete assembly of one scintillator strip with its fibre and coating is shown in Figure 3.2.




Figure 3.2: Diagram of a single scintillator strip in the MINOS FD, showing the plastic scintillator with the wavelength shifting fibre running through its channel and covered with reflective Mylar tape. Inset: The polystyrene and TiO_2 coating designed to prevent scintillation light from escaping before it enters the WLS. Diagram taken from [66].

3.3 Instrumentation

After exiting the detector the WLS fibres couple to approximately 4 m of clear fibre, which directs the collected light onto Hamamatsu R5900-00-M16 multi-anode photomultiplier tubes (PMTs). Each tube contains 16 $4\text{ mm} \times 4\text{ mm}$ pixels, and each pixel is shared (multiplexed) between eight different fibre ends. These eight fibres are typically separated by $\sim 1\text{ m}$ within a plane, and the mapping of which strips share the same pixel is different at each end of the fibre. In the case of a single hit, this allows the eightfold ambiguity to be exactly resolved; with multiple hits timing and event reconstruction information is also used for demultiplexing. The PMTs are operated with gains of approximately 10^6 , have a quantum efficiency of 13% at 520 nm, and a dark noise rate of around 25Hz per pixel.

Groups of 3 PMTs are contained together in a single light-tight steel container, referred to as a MUX box, which also directs the clear fibres for multiplexing. A pair of U and V planes is read out by three PMTs for each side of the detector, requiring 1452 PMTs in total. The PMTs are triggered using the output from the last dynode, which is a sum over all the PMT's anode pixels. The trigger threshold corresponds

to roughly 0.3 times that of the single photoelectron level. For comparison, a single minimum ionising particle (MIP) gives rise to around 5 p.e. per strip end.

MINOS is also notable for being the first large underground neutrino detector to be magnetised. The field is provided by a coil consisting of 190 turns of copper wire with water cooling that runs through the centre of each detector plane. Each supermodule has its own coil, carrying a current of 80 A which gives rise to a toroidal field whose average strength over the two supermodules is 1.3 T. During normal running (associated with the NuMI beam running in the forward horn current configuration that produces a neutrino-enhanced beam) the field focuses negative muons travelling in the beam direction. For high energy muons whose momentum cannot be determined from range, the field allows this to be measured using the track curvature instead. Finally, the field allows the charge of muons to be determined based on the direction of their curvature. This allowed MINOS to make the first oscillation measurement [68] to directly compare atmospheric neutrinos and antineutrinos on an event-by-event basis.

3.4 Readout electronics

The design of the readout electronics used at the Far Detector is motivated by the expected event rate. Interactions of cosmic rays and atmospheric neutrinos with the FD occur at a rate of approximately 0.5 Hz with a noise rate of around 10 kHz at each plane side due mainly to the dark noise of the PMTs and fibre noise from the scintillator. The front-end electronics were designed to cope with this rate whilst remaining sensitive to minimum ionising particles and without encountering significant dead time. Specifically, the readout was designed to operate with a threshold of 0.3 p.e., PMT gains of order 10^6 , and with a timing resolution of 3-5 ns to enable upward-going neutrino tracks to be distinguished from downward-going cosmics [66] [67].

3.4.1 Overview

The electronics used at the Far Detector operate in an environment where the low data rate is dominated by the detector noise. At this level, commercially available




Figure 3.3: Schematic showing the readout electronics of the Far Detector, taken from [66].

10MHz digitisers can be shared across multiple channels, saving on cost while still providing minimal dead time. Each ADC serves three PMTs, and a schematic of the full chain is shown in Figure 3.3.

Each PMT is read out by a single ASIC: a model VA32-HDR11 chip henceforth referred to as a VA chip. Three VA chips are mounted on a single VA Front-end Board (VFB) which is located outside the PMTs' MUX box. The VFB provides power and voltage biasing to the VA and also houses two PIN diodes that monitor the light injection system used to calibrate the PMT gain curves.

The analogue signals from the VA are sent through an ASDLite discriminator into the VA Readout Controller (VARC). A VARC contains six VARC Mezzanine Modules (VMMs), each of which houses a single ADC which digitises the output from two VFBs. Under this scheme, each ADC is responsible for 6 PMTs, each with

16 channels. Each VARC is implemented as a single VME card, and groups of three VARCs, a timing card, and a Motorola VME processor share a single VME crate. Sixteen of these crates have a capacity of 27648 channels and are used to read out the 22000 channels available to MINOS.

3.4.2 The VA chip and VA Front-end Board

The VA chip is a model VA32-HDR11 Viking chip [69] manufactured by IDE AS, Norway. It has 32 channels of preamplifiers and sample-and-hold circuitry, and a multiplexing output switch. Of these channels, sixteen are used; one for each of the PMT anodes.

Of the sixteen remaining channels, four are used to identify noise effects that occur simultaneously in all channels of the chip. Additionally, a Light Injection (LI) system is used to calibrate the PMTs for linearity and stability. The LI system is monitored using two PIN diodes per VFB whose signal is read out on a further VA channel in coincidence with the PMTs.

Three VA chips are mounted on each VFB, which is used to distribute power to the chips and control their voltage bias. The VFB also houses the discriminator used to compare the analogue output from the PMT dynode with a programmable threshold, and electronics for temperature and voltage monitoring. The VFB is used for readout only, and operates in slave mode directly controlled by the VARC upstream.

3.4.3 The VMM and VARC

Pairs of VFBs are digitised by a single ADC, with a typical response of 70 ADC/p.e., which resides in the VARC on a VARC Mezzanine Module. The VARC is responsible for the digitisation, triggering, and timing of the VA chips it controls, as well as for setting their bias voltage.

The VARC receives the PMT dynode signal, affixes a timestamp using a 640MHz FPGA TDC, and generates a hold signal for the VA. Digitisation operates on a ‘two-out-of-36’ method, where the readout only takes place if the VARC receives two discriminated signals from its 36 PMTs within 400 ns of one another.

If this condition is met, the VARC reads out each VA chip one after the other. Any dynode hits that occur during this process will be ignored, so a dead time is incurred during this phase: at $5\ \mu\text{s}$ per chip this is between 5 and $30\ \mu\text{s}$ depending on how many of the six chips were above threshold. The output from the ADC is passed to a sparsifier which subtracts the PMT pedestal from the signal and removes any signals below 20 ADC (electronic noise is typically ~ 2.5 ADC [66]). Two buffers operate alternately to collect data from the VARC and transfer it out, allowing this part of the process to run without dead time.

3.4.4 Data acquisition and triggering

Each crate contains a computer known as the Read Out Processor (ROP) which is used to extract the digitised data from the front end electronics, and these are synchronised with one another using timing cards present in each crate. Each buffer readout constitutes a single time block, and the job of the ROP is to assemble these blocks into 1 s long timeframes, and to append monitoring and calibration information as appropriate.

The ROPs are connected to further DAQ computers known as Branch Readout Processors (BRPs), of which the Far Detector has six. One master BRP instructs the others to send a given timeframe to one of a batch of Trigger Processing (TP) machines. In this manner, each TP receives one full timeframe for the entirety of the detector. Each ROP, BRP, and TP machine can buffer and queue multiple timeframes, enabling the data rate of approximately 8 MB/s to be accommodated comfortably.

On the Trigger Processor, software-based triggering is applied to each of the timeframes, using a variety of triggers. Zero bias triggers are used for data in time with the beam spill, extracting all in-spill digitisations into a single event without further processing. A number of special triggers are also available for calibration and debugging etc.

For signals outside of the spill window, the TP first sorts hits within a single timeframe into groups separated by at least 156 ns. The trigger option used for atmospheric and cosmic events is the so-called ‘4/5 plane trigger’ which requires 4 out of any 5 contiguous planes in the detector to register a hit. Other triggers are

in place searching for high energy deposition in contiguous planes, and the total number of planes hit throughout the detector.

The output from a single successful trigger can in theory contain more than one event, particularly in the ND during beam spills. It is commonly referred to as a ‘snarl’.

3.5 Calibration

The readout of the Far Detector provides a timestamp and pulse height reading for each strip end, but both of these quantities may be subject to offsets and variations with respect to one another. These differences may also vary with time, and fluctuate depending on external factors such as temperature. Because of this, a comprehensive calibration scheme has been developed to convert these raw readings into measurements that are directly comparable between different channels and detectors, and to convert them into a measurement of the calorimetric energy deposited.

The timing calibration is explained in detail in Chapter 7; this Section will briefly describe the procedure and equipment used to calibrate the charge response of the Far Detector. This can vary from one channel to the next for a variety of reasons, including:

- Random variation in the light level produced by different scintillator strips. Further differences may occur due to scintillator ageing, as different mechanical and temperature stresses will be experienced by strips in different locations.
- The response of the WLS fibre that collects and transports the scintillation light. The length of fibre used for each strip also affects the attenuation.
- Signal attenuation in cables of varying lengths.
- Variation in the transmission of light through optical connections.
- The response of the PMTs, and the electronics used to read out the signal.

A dedicated Calibration Detector (CalDet) [70] was constructed to develop the calibration procedure for the MINOS detectors and to enable a determination of their

absolute energy response. Typically this is determined for a detector by exposing it to a beam of known composition and energy, but this is not possible with the large underground MINOS Near and Far Detectors. Instead, the CalDet, which was smaller than but otherwise similar to the Far Detector, was exposed to test beams at CERN’s Proton Synchrotron (PS), then reassembled in front of the Near and Far Detectors to cross-calibrate them.

The CalDet consisted of 60 $1\text{ m} \times 1\text{ m}$ square steel-scintillator planes, each of which contained 24 scintillator strips, with successive planes oriented at 90° to one another. It employed the same readout electronics as the FD, and a scaled-down but otherwise identical DAQ system. The CalDet was operated in the T11 and T7 test beams at the CERN PS to characterise the calorimeter’s response to known beams of hadrons (pions and protons), muons and electrons with momenta between 0.2 and 10 GeV/c. The full calibration of the CalDet is described in References [70] and [71].

The calibration procedure for the Far Detector is described below. It uses the CalDet calibration as a reference point, corrects the FD channels for consistency with this reference, and hence determines the overall energy scale.

3.5.1 FD Calibration Stages

The aim of the calibration chain is to start with the responses of individual channels in one of the MINOS detectors and convert these numbers into an absolute energy measurement. This requires a number of stages. The final calibrated charge, Q_{corr} is related to the raw signal Q_{raw} by a series of multiplicative factors that depend on the detector d , strip s , position x along the strip, and time t :

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

Here, L is a linearity correction that accounts for the response of the combined PMT and readout system as a function of the number of photoelectrons at the PMT photocathode. For the M16 PMTs of the FD this becomes nonlinear at around 100 p.e. [72]. It is evaluated using the light injection system described in Section 3.5.2.

The D term is a drift correction that accounts for the change in whole-detector performance over time caused by effects such as scintillator ageing and varying temperature. The median pulse height per detector plane is calculated daily for throughgoing cosmic muons, and the drift correction scales the pulse height so that this median matches the value on a reference date, 1st December 2005 [66].

The attenuation term A corrects for attenuation of the scintillator light within a strip, primarily due to passing along the WLS fibre. It is implemented as a double exponential (i.e. the sum of two exponentials with different decay lengths) fitted for each strip [66]. The attenuation corrected response for each strip is then scaled by a strip-to-strip correction, S , which normalises the response of all strips to that of the average strip. This is described further in Section 3.5.3.

At this stage, the channels within a detector behave consistently, but the response of the Near and Far Detectors is not necessarily the same. A final scaling, M , is applied to convert the signal into “Muon Energy Units” (MEU), an energy unit which is consistent across the MINOS detectors and can be related to absolute energy using the CalDet results. This is described in Section 3.5.3.

3.5.2 Measuring linearity: LI and CI systems

MINOS uses a light injection (LI) and charge injection (CI) system to measure the response of each channel’s PMT and readout electronics. The LI system uses pulsed LEDs to measure the PMTs’ gain curves and to characterise their nonlinearity.

The LI system was designed to use pulsed blue LEDs to inject light directly into the scintillator strips. The LEDs are located in rack-mounted pulser boxes with 20 LEDs per box, and each LED is fanned out to illuminate several optical fibres. These carry light to the back of the pulser box where they connect to fibres that lead into Light Injection Modules, highly reflective cavities located at the edges of the scintillator planes. Inside these cavities the LED light illuminates the WLS fibres that connect to the PMTs. Additional fibres from the pulser box fanout lead to PIN diodes that measure the LED light intensity.

During CalDet testing it was discovered that the wavelength of the blue LEDs shifted slightly at low applied voltages. This occurred in a region in which the absorption of the WLS fibre varies sharply with wavelength, giving rise to nonlinear

behaviour, so ultraviolet LEDs are used at the FD instead [71]. The shape of the LED pulse is chosen to mimic that of scintillation light and the wavelength spectrum provides a reasonable approximation to that of scintillation light after passing through the WLS.

The LED brightness is measured independently of the PMT system using Hamamatsu S5971 PIN diodes, with two such diodes located on each VFB. The currents from these diodes are integrated and amplified, and read out using the VA chips in the same way as the PMTs. Because the VA triggers using the signal from a PMT dynode, the fibres are arranged such that each diode and the PMT that triggers its readout are both illuminated by the same LED.

One complication is that both the PMTs and the VA electronics of the FD readout are nonlinear in their response to large signals. Measuring the PMT's response to the LI incorporates the combined nonlinearity from the two sources, which is the information needed for the calibration. But the measurement of the LED brightness with the PIN diode uses only the VA chips, and so their nonlinearity must be calibrated out of the LED measurements. For this, the VA's CI system is used: charge is injected across the entire VA dynamic range, and the output of this procedure is used to linearise the PIN diode response, which can in turn be used to linearise the PMT \otimes electronics response. The PIN diode itself is known to be linear to within 1-2% [73].

The LI system runs in two main modes. A long run to determine the gain curves is performed at intervals of approximately one month (it can also be initiated manually during periods of scheduled beam downtime). This illuminates every channel 1000 times at forty intervals between ~ 10 and several-hundred p.e. The second mode is interspersed with normal data-taking and pulses each FD strip 300 times per hour at a single constant light level equivalent to approximately 50 p.e. The short-term variation in these single gain points is used to correct the linearity between full gain curve runs.

3.5.3 Strip-to-strip calibration with cosmic muons

The strip-to-strip calibration equalises the response of each strip with the mean response across the entire detector. This requires a source of events that are

the same throughout the detector, and so uses a sample of throughgoing cosmic muons. At the Far Detector, the expected rate of these muons is approximately $530 \text{ strip}^{-1} \text{ month}^{-1}$ [73] and so approximately one month of cosmic data is required before this calibration can be updated.

The linearised light level at each strip end is taken from the sample of muons and corrected for known angular and statistical effects [71]. A correction must also be made for the higher Poisson probability of producing zero photoelectrons in strips with lower scintillator yields; these are not triggered and go undetected.

The strip-to-strip correction factor was originally defined as:

$$S_i = \frac{\langle R \rangle}{R_i} \quad (3.4)$$

where R_i is the response of strip i corrected for drift and linearity, and $\langle R \rangle$ is the average across the whole detector. As the scintillator ages, the variation between strips increases, and this broadening increases the average correction factor required. This simply reflects the fact that, for a Gaussian distribution of strip responses centred at 1, a strip with a response of 0.5 requires a correction factor of two, whilst an equally-likely strip at 1.5 requires a correction factor of $2/3$.

Because the correction factor is a ratio this broadening increases the mean correction factor. This effect is absorbed by a corresponding change in the MEU scale, but to avoid biasing the MEU as the detector ages we instead normalise to a fixed mean calibration constant.

3.5.4 Energy scale calibration: Muon Energy Units

The Far (and Near) Detector strip-to-strip calibration is performed using throughgoing cosmic muons, whose spectrum is assumed to be constant throughout an individual detector. But between the two detectors the cosmic ray spectrum is very different due to the different overburdens. To compare results between the two detectors and with the CalDet, a ‘standard candle’ is required; we need to define and measure an energy that is the same in each detector. This motivates the Muon Energy Unit, or MEU.

Because the MINOS detectors are so similar in design, the stopping power of a single steel-scintillator plane can be used to define a common energy scale. For muons which stop in the detector, the muon range provides a measurement of its energy which is independent of the calorimetric response of the detector.

The energy change of charged particles as they travel relativistically through a material and lose energy to ionisation is described by the Bethe-Bloch equation [45]:

$$-\left\langle \frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \left(\frac{2m_e c^2 \beta^2 \gamma^2 T_{max}}{I^2} \right) - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \quad (3.5)$$

Here, z is the charge of the particle, Z and A are the atomic and mass numbers of the material through which the particle propagates, while m_e is the electron mass. The particle is moving at a speed βc with relativistic gamma factor $\gamma = (1 - \beta^2)^{-1/2}$. The mean excitation energy is I and T_{max} is the maximum kinetic energy available to a free electron in a collision. For a particle with mass M :

$$T_{max} = \frac{2m_e c^2 \beta^2 \gamma^2}{1 + 2\gamma m_e/M + (m_e/M)^2} \quad (3.6)$$

The energy loss for muons in iron and carbon is plotted in Figure 3.4. It can be seen that the rate of energy loss for muons varies by only a small amount for muon energies between 0.5 and 1.1 GeV but rapidly increases the closer to stopping the muon becomes. For a typical stopping muon away from the very end point of its track, this translates to about 30 MeV lost per plane of steel traversed and 2 MeV per scintillator plane.

The large increase in energy loss in the final 10% of the stopping muon's track length combined with the segmentation of the MINOS detectors, and uncertainties in the uniformity of the magnetic fields, steel, etc. yields an uncertainty of 2% in the muon range. This is already at the level of the desired overall detector-to-detector energy uncertainty, so a more precise scale is required.

Instead, we define a 'track window', illustrated in Figure 3.5, corresponding to part of the muon's trajectory for which $\frac{dE}{dx}$ varies slowly as a function of momentum. This window begins 16 plane widths (95 cm) back from the muon stopping point and extends back a further 14 plane widths (83 cm). Now the uncertainty in the stopping point becomes an uncertainty in the placement of this window, in a region




Figure 3.4: Bethe-Bloch equation showing the rate of energy loss of charged particles to ionisation as they pass through various materials, including iron. Taken from [45].




Figure 3.5: Definition of the track window used to calculate the MEU factor. The muon energy loss is considered in a window 83 cm (14 planes) long, beginning 95 cm (16 planes) from the stopping point. This is a region where the muon energy loss as a function of distance from the track end varies slowly, and so is less affected by uncertainty in the track end position. Figure taken from [74].

where the effect of this uncertainty is small. Consequently the 2% uncertainty in the muon stopping point translates to a 0.2% error in the energy deposition over the track window.

We can then define a quantity:

$$\frac{1}{N_p} \sum_1^{N_p} \frac{S_i}{L_i} \quad (3.7)$$

Where S_i is the strip-to-strip corrected response in the i th plane traversed in the window, L_i is the muon path length through the plane ($L_i = \frac{1}{\cos \theta_z}$ in units of plane-widths, where θ_z is the angle of the muon track to the z-axis), and $N_p \leq 14$ is the number of planes in the track window and depends on the steepness of the track.

This quantity is evaluated for a large sample of n muons, and the final MEU scale defined as the median of this measure over all events in the sample:

$$\text{MEU} = \text{Median} \left(\left(\frac{1}{N_p} \sum_1^{N_p} \frac{S_i}{L_i} \right)_1, \dots, \left(\frac{1}{N_p} \sum_1^{N_p} \frac{S_i}{L_i} \right)_n \right) \quad (3.8)$$

A rough rule of thumb is that one MEU corresponds to the response of the detector to a 1 GeV muon crossing a single steel-scintillator plane at right-angles.

With this MEU factor determined for each detector, the relative energy scales of the Near and Far detectors are calibrated to within 2.1% of one another. In combination with the CalDet data [70] the absolute hadronic energy scale can be determined to 5.1%.

Chapter 4

Monte Carlo simulation

In order to predict the neutrino flux as a function of the oscillation parameters, and to develop an analysis scheme, a detailed model of the incoming neutrino flux and the detector response is required. For this reason, a comprehensive Monte Carlo (MC) simulation has been developed, which is summarised in this Chapter.

4.1 Cosmic rays

Knowledge of the cosmic ray spectrum is important for two reasons. Firstly, cosmic ray interactions are responsible for the production of atmospheric neutrinos through decay of the secondary hadrons produced when they collide with molecules in the atmosphere. Hence the energies and directions of the primary cosmic ray flux feed directly into those of the neutrinos. Secondly, the muons produced in these interaction chains are themselves incident on the detector. These are used for calibration of the timing and charge response, but also form the main background to the atmospheric neutrino selection.

Cosmic ray primaries are predominantly protons (95%) with a smaller component of helium nuclei (4%) and heavy nuclei (1%) [45] [75]. At lower energies they can be detected directly using airborne and space-borne experiments, while at higher energies they are detected via the extensive air showers they produce. Figure 4.1 shows the flux of cosmic ray primaries as a function of their kinetic energy, with the components due to different primaries plotted separately and scaled to separate the lines. A power law dependence of the flux on the energy is apparent. This is




Figure 4.1: Cosmic ray flux data from a range of experiments separated out by primary particle type, as a function of kinetic energy. A power law dependence is evident for each of the different components. Figure taken from [45].

also visible in Figure 4.2, which shows the total cosmic ray primary flux over seven orders of magnitude in energy. This spectrum is shown on a logarithmic scale, and has been multiplied by $E^{2.6}$ for display purposes.

A power law distribution is evident in the all particle spectrum too, with a number of features. Below about 5×10^{15} eV the spectrum follows an approximate $E^{-2.7}$ distribution, until the ‘knee’ where the spectrum steepens. One possible explanation for this is that it represents a maximum energy for some accelerating process in our galaxy, and the beginning of the transition from galactic to extragalactic sources. A possible ‘second knee’ observed by Kascade-Grande experiment [76] may be indicative of a transition from lighter to heavier primaries. The flux beyond the knee goes as approximately $E^{-3.0}$ up to 10^{18} eV before flattening at the ‘ankle’. This softening of the spectrum may reflect extragalactic cosmic rays beginning to dominate, or may be a sign of protons interacting with the CMB via $p\gamma \rightarrow e^+e^- + p$.




Figure 4.2: The ‘all-particle’ combined cosmic ray flux as a function of energy, with key features highlighted. The y-axis has been scaled by $E^{2.6}$ in order to make the features more apparent. Figure taken from [45].

At lower energies ($E < 10^{10}$ GeV) the flux of cosmic rays is modulated by the solar wind. The Sun is a large ball of plasma and as such produces a magnetic field, which varies with solar activity following an 11-year solar cycle. A strong solar wind shields the Earth against lower-energy cosmic rays, reducing the overall flux. Further magnetic shielding is provided by the Earth's own field, which can deflect lower energy cosmic rays depending on their energy, incident angle, and the local field strength.

4.2 Atmospheric neutrino MC

To generate the atmospheric neutrino flux at the MINOS FD, it is necessary to simulate cosmic ray interactions in the upper atmosphere that produce showers of hadrons, and to model the decay chains and interactions of these hadrons to produce the incident flux of neutrinos. For this, MINOS uses the model of Barr et al. [77]. Uncertainties related to this model are discussed in Section 9.1, where comparisons to other models are also considered.

In the model, the flux of primary cosmic rays is determined by using power law fits to the fluxes of different incident nuclei. The primaries are introduced to the atmosphere, the secondary hadrons produced from their interactions are tracked in three dimensions as they propagate, and the flux is recorded for a detector centred on the MINOS FD. This 3D model updates an earlier 1D version in which only the shower component directed towards the detector was simulated.

The model accounts for the effects of the solar cycle on the incoming cosmic ray flux. The solar cycle can also be monitored using the cosmic neutron flux, and MINOS uses a sinusoidal parametrisation [78] of neutron data taken by the CLIMAX experiment [79] to describe the modulation of the neutrino flux. Two MC samples are generated, at the solar maximum and minimum, and this parametrisation is used to weight the total exposure according to the solar cycle, combining the two samples in the appropriate ratios. The data used in this analysis span a period from 2003 to early 2014, and therefore contain almost an entire solar cycle.

The ability of a charged particle to reach the Earth without being deflected by its magnetic field depends on the ratio of its momentum to its charge. This produces

a cutoff momentum below which particles are deflected away, which depends on the magnetic latitude and the incident angle of the cosmic ray. At high latitudes, the cosmic rays are deflected weakly because they are incident almost parallel to the magnetic field, whereas at lower latitudes they are almost perpendicular to it and so are deflected quite strongly. Northern Minnesota is at a relatively high magnetic latitude, so there is a rather large upward/downward asymmetry.

An additional asymmetry occurs because the Earth's magnetic field is such that it will focus positive particles incident from the east, and negative ones incident from the west. The cosmic ray flux is dominated by positively-charged primaries, and so this leads to an east-west asymmetry in the number of atmospheric neutrinos. However, this effect is weaker at higher latitudes because the deflection becomes approximately parallel to the Earth's surface.

These various cutoff effects are imposed by propagating any cosmic rays that produce neutrinos which interact in the detector backwards through the atmosphere until they reach a height of $30R_E$, where R_E is the radius of the Earth. Any which travel a path length $> 300R_E$ when performing this backwards propagation are deemed to be so strongly affected by the magnetic field that it is likely they would be cut off, and are rejected.

The secondary particle cascades are simulated using the package called *Target* [80]. Several particles (e.g. pions) are able both to interact and decay, and the competition between these two processes affects the neutrino flux. For such particles, the MC draws a random sample from probability density functions for the decay and interaction lengths, and applies whichever process drew the shortest length. *Target's* hadron production model parametrises results from accelerator experiments [81–85], but these leave a large region of phase space unexplored and so it relies on extrapolation or interpolation to describe these regions. For this reason, uncertainties relating to hadron production make a large contribution to the flux uncertainty. The Bartol model divides this phase space into regions, assigns uncertainties based on the amount of data available in each region, and propagates these into the overall atmospheric neutrino flux uncertainties.

The atmospheric neutrino flux at the Soudan Mine as a function of zenith angle and assuming no oscillations is shown in Figure 4.3. The peak for horizontal neutrinos in the 3D model arises because primaries not directed towards the detector can

descend to some height and decay, producing a secondary that is directed towards the detector. For almost horizontal neutrinos, the same fraction of solid angle encompasses a larger area of sky at a given distance from the surface than for vertical ones, hence the enhancement [86]. At higher energies the peak flattens because of the increased likelihood that a muon survives long enough to impact the Earth (this also reduces the ν_e flux arising from $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ decays).

Neutrino interaction cross-sections are modelled using Neugen3 [87], which includes models of quasielastic interactions, resonance scattering, and deep inelastic scattering. The QE and RES modes model the nucleus as a Fermi gas, where each nucleon carries ~ 230 MeV of energy, and has a binding energy of ~ 30 MeV, and the resonant mode can produce one of 17 possible resonances. The DIS mode involves scattering off a single quark instead of a whole nucleon. It is parametrised using structure functions derived from parton distributions, and employs Intranuke [88] to determine whether the resulting hadrons are scattered or reabsorbed by other nuclei.

The MC sample of neutrino induced muon (NIM) events is generated using the NUANCE event generator [89] and uses the earlier 1D Bartol flux model [90]. The NIM events derive from higher energy neutrinos ($\gtrsim 10$ GeV) so the difference between the two models is small, as can be seen in Figure 4.3. Nevertheless, this sample is reweighted by the ratio of the newer 3D calculation to the 1D one, as a function of neutrino energy and zenith angle. A fuller description of this MC is provided in [91].

4.3 Cosmic muon MC

A sample of cosmic ray muon events was used to tune the selection criteria and to study the level of agreement between cosmic data and MC. It also allows the cosmic muon background to be studied.

The flux of muons from cosmic rays at the surface can be parametrised [92] as:

$$N_0(E_0, \cos \theta) \approx \frac{0.14 E_0^{-2.7}}{\text{cm}^2 \text{ sr GeV s}} \left(\frac{1}{1 + \frac{1.1 E_0 \cos \theta}{115 \text{ GeV}}} + \frac{0.054}{1 + \frac{1.1 E_0 \cos \theta}{850 \text{ GeV}}} \right) \quad (4.1)$$

where the muon has energy E_0 and zenith angle θ , and the subscript ‘0’ indicates that the flux is on the surface, at a depth of 0 m. This approximation is valid for




Figure 4.3: The atmospheric neutrino flux at the Soudan Underground Laboratory, as a function of zenith angle. The solid lines correspond to a 3D shower simulation model, and the dashed lines to the earlier 1D model. Plot taken from [86], which is used as the basis for the atmospheric neutrino flux in the MC.

large energies $E_0 > 100 \text{ GeV}$ and $\theta < 60^\circ$, where the muon can be assumed not to decay before it hits the ground and the Earth's magnetic field can be neglected.

The flux is extrapolated downwards from the surface through the rock towards the FD. For the Soudan Laboratory, the rock overburden is predominantly composed of Lake Vermillion Greenstone with a density of 2.8 g cm^{-3} although, being an iron mine, there are also pockets of iron ore. A density map of the rock above the MINOS cavern was computed by the Soudan 2 experiment [93] and is used to determine the profile of the material through which the muons travel.

The energy lost by a muon travelling through the rock can be parametrised by [92]:

$$-\frac{dE}{dx} = A + \frac{E(x)}{L} \quad (4.2)$$

The term in A ($\approx 1.9 \text{ MeV g}^{-1} \text{ cm}^2$) corresponds to energy lost by ionisation, and the term in E/L (with $L \approx 2.5 \times 10^4 \text{ g cm}^{-2}$) corresponds to radiative losses. In reality the loss is stochastic and both A and L are themselves functions of energy, but they vary slowly enough that they can be treated as approximately constant.

Integrating, we find:

$$E(x) = (E_0 + AL)e^{-x/L} - AL \quad (4.3)$$

where x is the depth below the surface and $E(x)$ is the energy of the muon at that depth.

In this model, the number of muons remains the same but their energies asymptotically decrease as they get deeper underground. We can track a fixed flux of muons, $N_0(E_0, \cos \theta) dE_0 d(\cos \theta)$, from the surface to the detector by adjusting the energy according to Equation 4.3 and transforming the range dE_0 correspondingly. Because Equation 4.3 has no $\cos \theta$ dependence we omit that term and find:

$$N(E_x, \cos \theta) dE_x = N(E_0, \cos \theta) dE_0 \quad (4.4)$$

$$= N(E_0, \cos \theta) \frac{dE_0}{dE_x} dE_x \quad (4.5)$$

$$= N(E_0, \cos \theta) e^{x/L} dE_x \quad (4.6)$$

Hence the flux of muons in a given energy range at a given depth can be expressed in terms of the flux at the surface. This equation defines the probability distribution function used in the MC for the neutrino energy and direction at the FD. Muons are generated according to this PDF on the surface of a cuboidal box that encompasses the detector, and those that would not intersect the detector are discarded while the rest are retained.

4.4 Detector simulation

Once the particles produced in the detector from the various event types have been determined, it is necessary to simulate the response of the detector. This involves propagating the particles, simulating their interactions with the detector material, determining the number of photons produced, and modelling their detection by the PMTs and the response of the electronics and DAQ.

The MINOS detector simulation is part of the *gminos* software package. It uses a GEANT3 [94] simulation to propagate particles through the material, using GCALOR [95] to simulate hadronic interactions, as this was found to best reproduce the CalDet data [96]. This portion of the simulation reads in the four-momenta of the incident particles, inserts them into the GEANT3 detector geometry, and eventually outputs the energy deposited into the scintillator.

The response of the scintillator to these deposits is modelled using the *Photon-Transport* package. This generates photons according to Birks' law [97]:

$$N = C \frac{L_0}{C_0} \frac{dE}{1 + k_B \frac{dE}{dx}} \quad (4.7)$$

Here, dE is the energy deposited in the scintillator, L_0 is the light output of the scintillator, C_0 is the calibration constant for this strip, C is an overall normalisation tuning parameter, and k_B is Birks' constant and is taken to be 0.133 m GeV^{-1} .

After generating the photons, their capture by and propagation along the WLS fibre is simulated using probability density functions derived from simulations, which relate the position and timing of the incoming blue photons to that of the output green ones. At the PMT, the incoming photons are converted to photoelectrons

using the measured quantum efficiency of 13%, and single photoelectron noise is also introduced at this stage. Finally, the *DetSim* package is used to perform the simulation of the PMT, accounting for the amplification along the dynode chain, the digitisation process, and triggering. The final output is a series of Monte Carlo files that have the same format as regular MINOS data files, but include additional MC truth information.

Chapter 5

Data Quality

To ensure the integrity of data used for the analysis it is required to pass a series of quality checks. These are intended to remove data taken outside of normal physics run conditions, or when the FD magnet, high voltage, or GPS timing hardware were not functioning correctly.

Validation software runs automatically on a daily basis, processing all the data taken on the previous day. This daily processing did not occur between 10th October 2012 and 3rd June 2013 as the automatic processing code was migrated from computer systems at Soudan to Fermilab during the extended NuMI shutdown. The data for this period were retrospectively processed in September 2013 in preparation for this analysis. The results of this processing are used to fill a database table, which can then be queried with timestamps to produce a list of good data files. The overall data quality is determined by combining four separate checks, each of which must be passed if the data is to be declared good.

The following sections summarise these standard cuts.

5.1 Run quality

The run quality cuts are designed to select data taken during normal physics running with the full detector reading out, and to reject runs with anomalously high or low rates. MINOS data is labelled with a ‘run’ and ‘subrun’. When the FD is taking normal physics data a run consists of 24 subruns, each of which is one hour long. However, shorter subruns or runs with fewer than 23 subruns can occur, for example

when the run sequence is interrupted in order to perform maintenance and then restarted, or if a gain calibration run is initiated early to take advantage of beam downtime.

The run quality checks require:

1. The run type is a physics run and is not flagged as test running.
2. All sixteen electronics crates are powered on and functioning.
3. The subrun is at least two minutes long, unless it is the first subrun of a run, in which case it must be five minutes long.
4. The subrun contains at least 100 “good” snarls, defined as one recorded under a physics trigger, with fewer than 1000 digitised hits.
5. Good snarls occur at a rate lower than 30 per one-second timeframe.
6. The median snarl rate for the run is lower than 100 Hz and the maximum rate is lower than 300 Hz.

5.2 High Voltage status

The status of the high voltage supply to the FD is determined by counting the number of ‘cold’ VA chips in each of the two supermodules, i.e. those PMTs with anomalously low rates. A PMT is classified cold if it has a mean singles rate below 50Hz (this should be compared to a nominal dark rate of ~ 250 Hz). A supermodule fails the HV check if it contains 20 or more cold VA chips. This is rare during normal operations but occasionally the high voltage supply trips, and this is associated with an increased number of cold channels. The data quality cuts require that both supermodules pass the cold chip cut.

5.3 Magnet coil current

This check requires that the FD electromagnet be switched on and operating within normal current ranges, as this governs the magnetic field strength inside the detector which is vital for the reconstruction of muon momentum and charge sign. During

normal operations the current in the FD coil is 80.0 A, but the actual value is monitored and recorded at two-minute intervals. A check is made against the database table containing information on the coil state, and data passes the check if the coil status is ‘on’, and the coil current is between 79.0 and 81.0 A, inclusive.

5.4 GPS timing uncertainty

MINOS uses a GPS system to synchronise the timing between the ND and FD, and to match data taken with NuMI beam spills. For the atmospheric analysis only events occurring outside of beam spills are considered. The worst-case timing uncertainty is the sum of the ND and FD uncertainties and has a nominal value of 400 ns, but this increases if the GPS unit loses satellite coverage or suffers some other malfunction. Data fails the GPS timing cut if the nearest spill was reported and within 5 minutes of the data in question but the worst-case GPS timing uncertainty is greater than 1 μ s.

5.5 Data quality results

The quality checks are applied by querying the data quality database for each one-second timeframe of data, and selecting all the one-hour subruns containing only good data. Figure 5.1 shows a summary of the data quality for all of the new data taken since the 2012 analysis, with the green region indicating the fraction of data passing all the quality cuts. The majority of data failing the cuts comes from the changeover between Runs VII and VIII in the first months of 2013, and the long NuMI shutdown and upgrade in preparation for NO ν A-era running, when the FD magnet was deliberately switched off to reduce power consumption.

5.6 Summary of data used for the analysis

Table 5.1 summarises the data used for this analysis broken down into data already used for the 2012 analysis, and data added to this updated sample. The runs and timestamps shown for each period are not continuous, but run from the first second

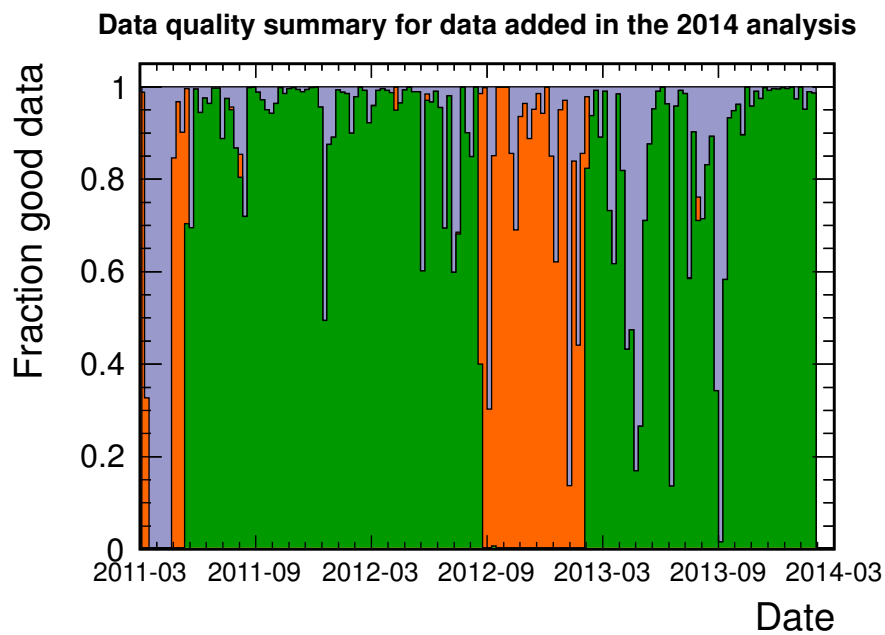


Figure 5.1: Stacked histogram summarising the data quality checks on the new data taken since the 2012 analysis, shown in bins of one week. The blue region denotes all new data, with the fraction of this data that passes all the data quality cuts shown in green. The fraction of data passing run quality checks but failing due to the detector checks is shown in orange; the majority of this is from the long NuMI shutdown when the FD coil was intentionally switched off.

of the first good run to the last second of the last good run in each period (the full period was tested, but the missing time period contains commissioning runs that fail one or more of the quality cuts).

		Date	Time (UTC)	Run/subrun	Total kton-yrs
2012 analysis	Start	2003-08-01	02:48	18143/00	37.90
	End	2011-03-07	11:45	47397/20	
New data	Start	2011-03-08	17:17	47435/00	10.79
	End	2014-01-31	23:51	53474/00	
Total	Start	2003-08-01	02:48	18143/00	48.67
	End	2014-01-31	23:51	53474/00	

Table 5.1: Summary of good data used for the atmospheric neutrino analysis

The new data adds an additional 726 live days, equivalent to 10.79 kton-yrs, yielding a total atmospheric neutrino sample to date of 3279 live days, or 48.67 kton-yrs. This represents an increase of 28% over the previous sample.

Chapter 6

Veto shield

A significant challenge in selecting atmospheric neutrinos lies in distinguishing muons and electrons produced by charged current neutrino interactions from those originating from cosmic rays. This is particularly difficult in the case of the MINOS FD as the planes of steel and scintillator are placed vertically so as to align the face of the detector perpendicular to the incoming NuMI beam. This means that cosmic ray muons, whose directions are strongly peaked about the downward vertical direction, typically travel almost parallel to the face of the detector planes.

With respect to particles travelling vertically downwards, the detector has only 20% active coverage: the arrangement of planes means the top of the detector is approximately 40% steel, 40% air gap, and 20% scintillator. A steep muon track may therefore deposit energy in only a few detector planes before exiting, and may propagate through the steel or between the instrumented planes before producing light in any of the scintillator. Such muons may be mistaken for neutrinos interacting within the detector rather than particles entering from outside.

To combat this effect, a veto shield was constructed whose planes are made from the same scintillator material as in the body of the detector, but are placed across the top perpendicular to the rest of the planes. Due to the shape of the FD hall, the shield was constructed in separate sections as shown in Figure 6.1. This configuration provides $> 99\%$ coverage of the top of the detector, the majority of which is with two layers of scintillator.

Unlike the main body of the detector, the shield planes have no light injection system. Instead, the PMT gains were estimated by setting a low dynode threshold and

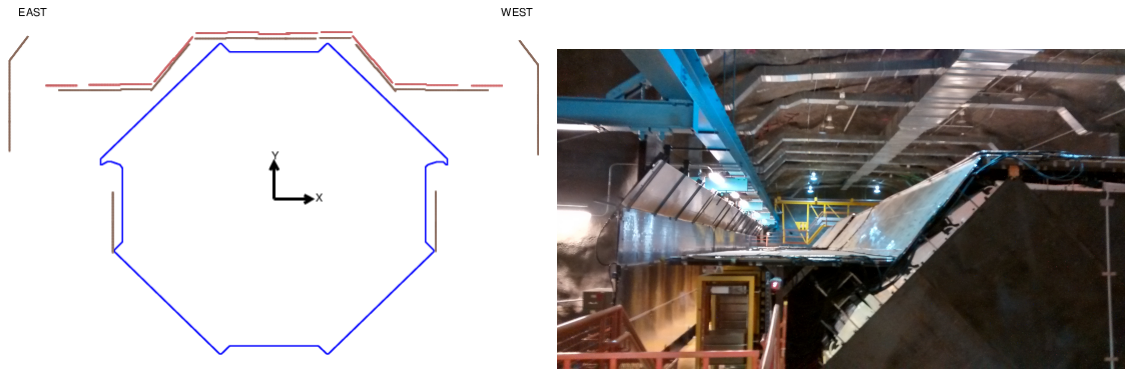


Figure 6.1: Left: Schematic showing cross-section of the FD veto shield [66]. Right: Photograph of the shield installed above the Detector.

fitting the single photoelectron peak. For normal data taking, the dynode threshold is set to $\sim 1 - 2$ p.e. to reduce inefficiencies due to single photoelectron noise that would cause spurious veto hits.

6.0.1 Veto criteria

Events are rejected as cosmic rays if they can be associated with activity in the veto shield, matched in both space and time. The highest point in an event is identified: the vertex of a reconstructed shower or either end of a reconstructed track, whichever is higher. Shield activity in the section above this point is then examined, and events are vetoed if there is any shield activity within $|\Delta t| < 50$ ns and $\Delta Y > -10$ cm (i.e. the shield hit is not more than 10 cm vertically lower than the high point of the event). If the highest point in the event is located less than 50 cm from the boundary between two shield sections then both sections are considered. Such a broad time window is necessary because cosmic events tend to be poorly-reconstructed.

A sample of stopping muons is used to provide a data-driven measurement of the veto shield efficiency. These are defined as events containing a track spanning at least eight planes, with at least six track-like planes (planes where at least 50% of the charge is deposited within one strip of the reconstructed track), and where only one end of the track is contained within the fiducial volume [98]. Figure 6.2 shows this sample and indicates the 50 ns time window in which shield hits lead to an event being vetoed.

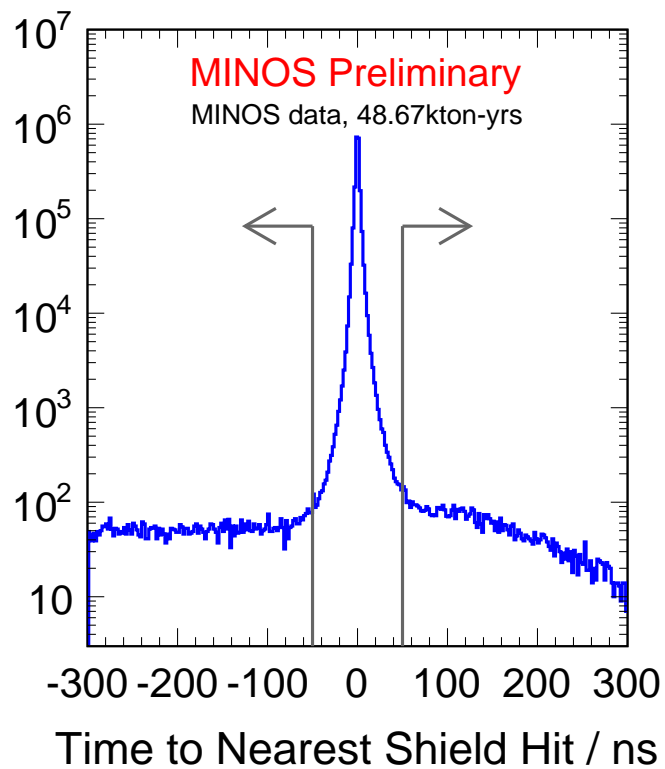


Figure 6.2: The time between the event vertex and the nearest-timed hit in the section of the veto shield above the highest point in the reconstructed event, for a sample of stopping muons selected from the data. The grey arrows indicate the cut applied to select events without an in-time shield hit: events with shield activity within ± 50 ns are removed.




Figure 6.3: The average efficiency of the FD Veto Shield over time, from 1st August 2003 to 31st January 2014. The efficiency is defined as the fraction of stopping muon events failing the veto shield cuts, and decreases linearly from 97.1% to 95.7%, with a mean value of 96.4%.

The veto criteria are applied to this sample and the fraction of events rejected is calculated. This is taken to be the efficiency of the veto shield and is shown in Figure 6.3. The data show a slow linear decrease in shield efficiency over time, from 97.1% in 2003 to 95.7% as of January 2014 (0.14 percentage points per year) with no evidence of a change in this pattern since the 2012 analysis. The mean efficiency over the entire run period is 96.4%.

Chapter 7

Timing calibration

A good timing calibration is particularly important to the atmospheric neutrino analysis as timing information is used both to distinguish upward- and downward-going tracks and for vetoing cosmic events. The timing calibration uses a sample of through-going muons, and relies on the fact that scintillator strips in the FD are read out by PMTs at both ends of the fibre.

The calibration uses an iterative procedure to correct for three effects:

- Time-walk: the rise time of PMT signal is dependent on the number of photoelectrons being detected; statistically, the first photon arrives on average earlier for larger signals.
- Hardware swaps: occasionally PMTs or a component of their readout electronics are replaced as part of detector maintenance. This introduces slightly different hardware delays for that channel and causes sudden jumps in the average difference between times measured at the East and West strip ends which must be calibrated out.
- Time constants: each individual channel has its own timing offset, so a full strip-to-strip calibration is performed to calculate calibration constants for each channel to account for these offsets.

7.1 Timing fits to through-going cosmics

The timing calibration is performed using a sample of high energy through-going cosmic muons, with events required to span at least twenty scintillator planes. Straight line fits to their trajectories are then performed separately using U and V planes, and only tracks with an RMS deviation of < 1 cm from a straight line are accepted. This criterion ensures that the muons are travelling in a straight line, which also requires that their energy be large, and hence that they can be assumed to be travelling at c . The U and V view fits are then combined to give a 3-dimensional straight line fit for each track. This produces approximately 750000 muons per month, or 400-500 hits per strip end per month.

Using this fit, the point at which the track crosses each strip is evaluated, and from these points the hit times recorded at the East and West strip ends can be used to calculate the muon's vertex time. This is achieved by projecting back from the PMT hit time to the time at which the muon track crossed the strip, via:

$$T_{\mu}^{E(W)} = T_{rec}^{E(W)} - \frac{n}{c} L^{E(W)} \quad (7.1)$$

Here, $T_{\mu}^{E(W)}$ is the predicted time at which the muon crossed the strip, using information from the East (West) end of the strip, $T_{rec}^{E(W)}$ is the hit time recorded by the PMT at the East (West) end of the strip, c is the speed of light in a vacuum, $n = 1.77$ is the refractive index of the wavelength-shifting fibre, and $L^{E(W)}$ is the distance along the fibre from the point at which the muon crossed it to the East (West) strip end. This is illustrated in Figure 7.1

Given our assumption that the muon travels at c , the time at which a muon has travelled a distance S is:

$$T_{fit}^{E(W)}(S) = T_{fit}^{E(W)}(0) + \frac{S}{c} \quad (7.2)$$

Using this parametrisation and weighting hits by charge, timing fits are carried out separately for the East and West strip ends, and calibration corrections are calculated for each strip from the mean residuals to these fits. These corrections are tuned by applying them and iterating the fit.




Figure 7.1: Diagram illustrating the variables used to perform the timing fit

7.2 Time-walk

The timing correction as a function of recorded ADC is calculated as in [99]: timing fits are applied to a sample of through-going muons to find straight-line tracks. These fits are performed using only hits with > 5000 ADCs for which the difference in time-walk from one hit to another is considered small. The residual to this fit is then evaluated for all hits, and parametrised as:

$$\frac{\Delta t_{\text{time-walk}}}{\text{ns}} = C_0 + C_1 \ln Q + C_2 (\ln Q)^2 + C_3 (\ln Q)^3 \quad (7.3)$$

where Q is the recorded charge in ADC and is on the order of 300 ADC (for a minimum-ionising particle) to several thousand ADC. The constants C_0 , C_1 , C_2 , and C_3 are approximately 20, 2, 0.2 and 0.02 respectively [99].

7.3 Hardware swaps

It is occasionally necessary to replace some components of a channel's readout system. When this occurs there is insufficient data to immediately recalibrate the channel. Instead, the hardware replacement introduces a constant time offset for this channel relative to the previous calibration. This can be identified by comparing the muon




Figure 7.2: Example of a jump in the East-West time difference caused by a hardware replacement. The x-axis shows days since 1st August 2003. A step was identified by the calibration software at around 3500 days, and is indicated by the dotted line.

time predicted at the East and West ends of the same strip (which should find the same time up to the level of the timing resolution if the detector is fully-calibrated).

At the point where the channel is replaced, a step change in the mean East-West time difference occurs. These are identified semi-automatically using straight line fits to the mean East-West difference for a give strip over time. Figure 7.2 shows an example of a timing jump identified by the calibration software. The ambiguity as to which of the two channels has been replaced can be resolved by comparing the mean residuals for each channel individually.

The size of the correction required is calculated by fitting a constant value to the time difference either side of the step, and adjusting the appropriate channel's overall correction by this amount.

7.4 Strip to strip calibration

The final stage of the timing calibration is to produce a full set of calibration corrections for each individual strip, using the iterated fits to cosmic muons described in Section 7.1. To encourage the fits to converge more rapidly, this is first carried

out at the VARC level by combining all the strips read out by the same VARC. The same procedure is then applied at the PMT level, combining all strips read out by each PMT before the full strip to strip calibration is performed using individual strip ends. The final calibration constant for each channel is given by the sum of the corrections for the VARC, PMT and individual strip end.

Once the constants have been calculated for each strip end, a final overall East-West time difference is calculated for the entire detector, and the constants are then corrected so as to set this overall time difference to zero, by adding half the difference to the West channels, and subtracting half the difference from the East ones.

7.5 Validation

The two years of additional data added to the analysis were calibrated in batches of six months each. Because the East and West sides of the detector are calibrated separately, the calibration can be tested by comparing the differences between times measured using only the East or West strip ends.

Figure 7.3 shows a single run from March 2013. The upper plot with red points shows the mean time difference between East and West strip ends before calibration, for every plane and with the differences averaged over all strips in the plane. The points show large excursions away from zero, and a repeating pattern due to the sixteen crates (eight pairs) reading out the planes is evident. The lower plot with green points shows the same quantities after calibration, and is approximately flat and centred on zero.

The mean time differences also provide a method to estimate the timing resolution of the calibrated detector. Figure 7.4 shows the distribution of East-West time differences for each of the four calibration periods. Each entry corresponds to the mean time difference for a single strip averaged over the six month period. Each entry is the difference between two timing measurements, so the RMS of the distribution can be converted into the timing resolution at a single strip end using:

$$\sigma_{\text{strip}} = \frac{1}{\sqrt{2}} \sigma_{\text{E-W}}^{\text{RMS}} \quad (7.4)$$

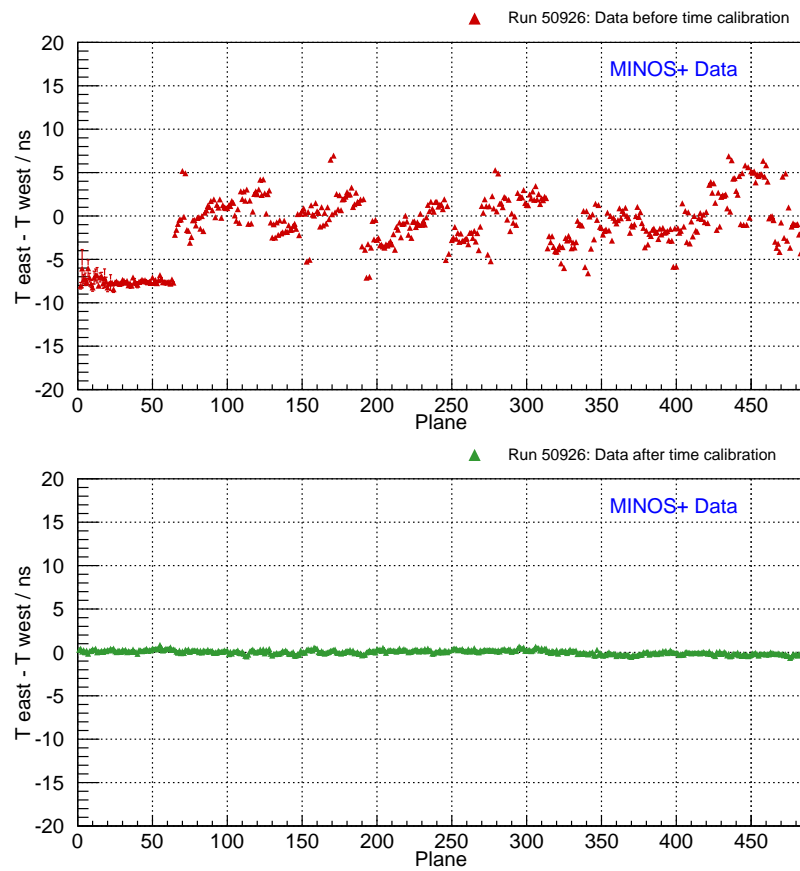


Figure 7.3: Uncalibrated (top) and calibrated (bottom) average East-West time differences for each plane in the detector. The time differences are for a single run from March 2013 and for each plane they are averaged across all that plane's strips.

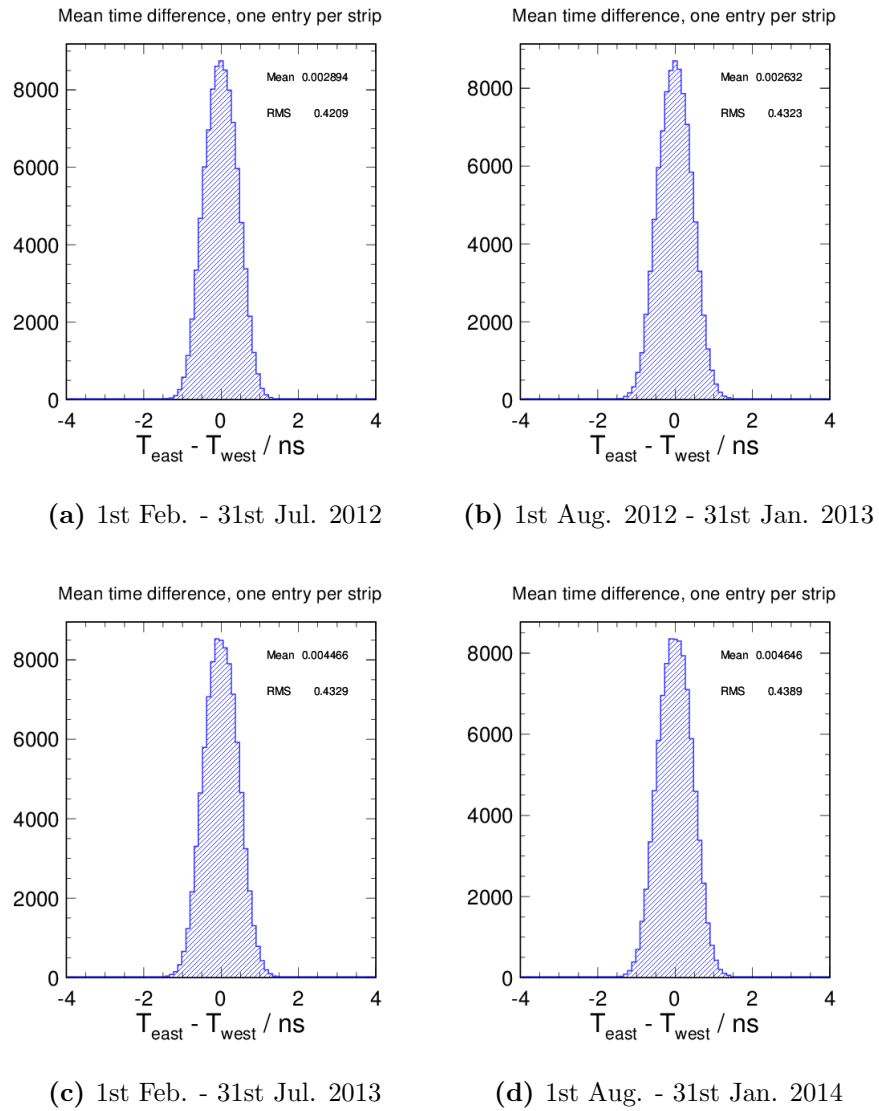


Figure 7.4: The mean timing difference, after applying the new calibration constants, between the East and West ends of each strip in the detector. The average is taken over the full six months of data used for each calibration period.

Using this method, the single strip end timing resolution is between 0.297 and 0.311 ns for each period.

Chapter 8

Event selection

8.1 Event classification

The atmospheric neutrino analysis considers three different event classifications:

1. Contained vertex (CV) muons
2. Neutrino induced muons
3. Contained showers

The CV muon sample contains muon-like tracks whose vertex is contained within the fiducial volume of the detector, and corresponds to atmospheric neutrinos interacting directly inside the detector volume. The NIM sample contains horizontal or upwards-going muon tracks whose vertices may or may not be contained within the detector. Because the rock overburden is so large for upward tracks, these correspond to neutrinos that have traversed the Earth and interacted, either in the detector itself or in the rock below it, producing a muon. The final sample consists of contained shower-like events, which contain a combination of neutral current and $\nu_e/\bar{\nu}_e$ events.

The main challenge in selecting events for the analysis is to remove the comparatively enormous cosmic ray background (~ 50000 events daily, compared to ~ 1 atmospheric neutrino interaction), requiring a cosmic rejection factor on the order of 10^6 without also cutting out the atmospheric signal. This is achieved via a series of cuts based on timing, the event topology, vertex location, and track direction.

8.2 Preselection

Preselection cuts are applied to distinguish track-like and shower-like events, and to build a sample of ‘good’ events.

8.2.1 Definition of good events:

Good events are defined as those events that pass the data quality cuts described in Chapter 5 which also satisfy the following:

- Contain at least 4 good planes overall
- Contain at least 2 good planes in both the U-Z and V-Z views
- At least 50% of the total pulse height in each view is double-ended

Here, a good plane is defined as a plane containing double-ended strips (i.e. strips that recorded a charge at both ends) with a pulse height corresponding to at least 2 photoelectrons.

8.2.2 Track-like and shower-like events

The good events are then separated into track-like and shower-like events by considering the longest track and longest shower present. This is achieved by calculating the number of track-like planes (which are planes through which tracks pass that contain ≤ 80 p.e., where at least 80% of the charge was deposited within ± 2 strips of the track) and track-only planes (which are planes containing strips associated with a track but no strips associated with a shower).

Events are declared track-like if they satisfy all of:

- At least 8 planes traversed by the longest track
- At least 5 track-like planes associated with the longest track
- At least 3 track-only planes associated with the longest track

These cuts were chosen because the muon reconstruction improves for tracks that span a larger number of planes, and because the number of MC events reconstructed with a single muon track that were true ν_μ CC interactions likewise increases [99].

Similarly, shower-like events are all events satisfying:

- At least 4 planes spanned by the longest shower
- No more than 20 track planes associated with the longest track
- No more than 12 track-like planes associated with the longest track
- No more than 8 track-only planes associated with the longest track

At this point, it is possible for events to pass both the track-like and the shower-like event preselection (for example, events containing both a shower and a short track). These events are included in both samples for now. The track-like sample is used to select CV muon and NIM events, and the shower-like sample is used to select CV shower events. If, after these further selection cuts have been applied, the same event is still present in the two samples, the event is removed from the CV shower sample to prevent double-counting.

8.2.3 Refining track-like and shower-like events

Once the events have been classified as track-like or shower-like, a series of further cuts are applied to select well-reconstructed events. For track-like events we require:

- Only 1 reconstructed track
- At least 8 planes in total, including at least 3 in each of the U and V views
- The difference between the number of planes in each view is no greater than 9
- More than 33% of the pulse height in the event is due to the track
- The centre of the track lies within the fiducial volume defined in Section 8.3
- The difference between the directions at the beginning and end of the track is less than 90°

The reconstruction software estimates the track's curvature in terms of the ratio of the muon's charge sign, $Q = \pm 1$, to the magnitude of its momentum, $p = |\mathbf{p}|$. Beginning with:

$$\frac{d\mathbf{p}}{ds} = cQ\hat{\mathbf{p}}(s) \times \mathbf{B}(s) + \hat{\mathbf{p}}(s) \frac{dp}{ds} \quad (8.1)$$

Rearranging and taking the inner product with $\hat{\mathbf{p}} \times \mathbf{B}$ gives:

$$\frac{Q}{p(s)} = \frac{\frac{d\hat{\mathbf{p}}}{ds} \cdot \hat{\mathbf{p}}(s) \times \mathbf{B}(s)}{c|\hat{\mathbf{p}}(s) \times \mathbf{B}(s)|^2} \quad (8.2)$$

where s is the distance along the track's trajectory, $\mathbf{B}(s)$ is the magnetic flux density at that distance, and c is the speed of light (in m ns^{-1} in the MINOS software). The reconstruction divides the track into 15-plane segments and fits each of these with a quadratic in s in the U and V views, before combining them into a 3D trajectory. The sign of the average of $\frac{Q}{p}$ from each segment indicates the muon's charge, and the standard deviation $\sigma_{Q/p}$ represents the uncertainty on the charge measurement.

We therefore require:

- The track fitter converges, and reports $\sigma_{Q/p} < 10^{-4}$

Two further cuts are also applied that remove events associated with a high rate of cosmic muon backgrounds:

- The reconstructed vertex must lie within 1 m of the track
- No significant activity after the last plane of the track

The equivalent quality cuts to select events where the primary shower is well-reconstructed are:

- Only 1 reconstructed shower
- At least 2 shower planes in each view
- More than 66% of pulse height in the event is due to the shower
- More than 50% of the pulse height is located within the fiducial volume defined in Section 8.3

A further cut to remove shower events associated with a high cosmic background is applied:

- The RMS width of the shower strips must be < 0.5 m in both the U and V views

A final cut is applied to increase the fraction of selected shower events that are due to ν_e or neutral current events, instead of background ν_μ CC events: for events with short showers (≤ 8 planes) ν_μ CC events typically have fewer than strips per plane than ν_e or NC events, so a cut requiring $(\# \text{ shower strips} - \# \text{ shower planes}) \geq 3$ is applied.

8.3 Containment

A fiducial volume cut is applied at the level of individual digit hits, primarily to remove events due to through-going cosmic ray muons. Further containment cuts are later applied at the track and shower level. For these digit-level cuts, the fiducial volume is defined as the region of the detector greater than 30cm from each of the edges of an octagonal plane, and > 5 planes away from either end of each supermodule. The detector volume outside of this fiducial region is then divided into regions corresponding to each of the plane edges, and each end of the supermodules, and the hits inside and outside these regions are compared.

Events are rejected if they contain ≤ 10 p.e. inside the fiducial volume. The remaining events are then classified based on how many edge regions contain > 6 p.e. as follows:

- Events with 0 such edge regions are declared fully-contained (FC)
- Events with 1 such edge region are partially-contained (PC), as are events with two adjacent edge regions, where the mean positions in each edge differ by less than 1.33 m
- Remaining events are tagged as through-going

8.4 Selecting Contained Vertex Muons

When a CV ν_μ interaction occurs, the neutrino enters without leaving a signal until it interacts in the main body of the detector. Thus the signature of these events is a muon track appearing inside the detector volume without activity near the edges of the detector. However, downward cosmic ray muons can enter through a steel or air gap section so can also penetrate into the detector before leaving a signal in any scintillator. The main aim of the CV ν_μ selection is to reduce this background by applying a series of cuts based on timing and topology. The following section describes these cuts in detail.

8.4.1 Stage 1: Fiducial volume

The first stage is to reject events where the track lies outside of a fiducial volume. In addition to the containment cuts of Section 8.3, this requires that the vertex of the track lie at least 20 cm from the edges of the detector planes, at least 5 planes away from the ends of either supermodule in the z-direction, and at least 40 cm away from the centre of the (uninstrumented) magnetic coil hole. A loose cut requiring that the vertex z-trace be greater than 15 cm is also applied. The trace variable is described more fully in Section 8.4.2.

Tracks travelling into the detector relative to the beam axis (the z-axis) whose vertices are within 21 planes of the front or back of the detector are subjected to a more stringent coil hole cut and must have vertices further than 1 m away from the centre of the coil hole; this is to prevent steep muons entering through the coil hole and penetrating into the detector while travelling between two instrumented planes.

A fiducial cut is also applied by combining reconstructed strips in the event into a 3D hit collection. The track vertex determined by this method is required to lie at least 30 cm from the edge of the detector and 5 planes from the end of a supermodule.

8.4.2 Stage 2: Trace cut

The z-trace, Δz , is a geometric quantity used to require that more vertical tracks originate further from the edge of the detector than shallower ones. The z-trace

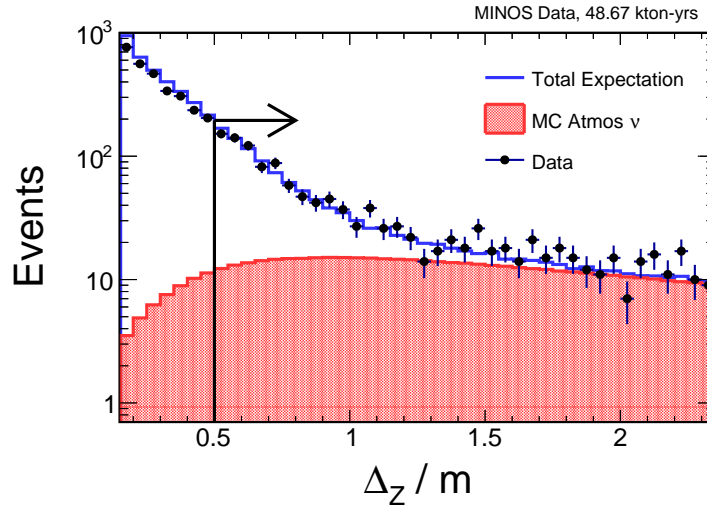


Figure 8.1: Z-trace variable used to select contained events for the CV ν_μ sample. The data points are shown in black, with the expectation from Monte Carlo atmospheric neutrinos (two-flavour oscillations with $\Delta m^2 = 2.32 \times 10^{-3} \text{ eV}^2$ and $\sin^2 \theta = 1$) shown in red, and scaled cosmic ray data shown in blue.

is calculated by extrapolating backwards from the track vertex along the direction of the track until the edge of the detector is reached, and Δz is defined as the z-component of this extrapolation. The CV ν_μ selection requires that Δz for the highest end of the track be greater than 50 cm. The effect of this cut is shown in Figure 8.1

A further digit trace cut is applied by considering reconstructed strips that occurred beyond the highest end of the track but were not associated with it. If more than five such strips are present, the distance d_z between the highest end of the track and the furthest additional strip in z is calculated, and we require $\Delta z_{digit} = \Delta z - d_z > 40 \text{ cm}$.

8.4.3 Stage 3: Topology cuts

Cosmic muons that survive the trace cut tend to be steep tracks that travel a long distance in a single steel plane and then change direction, either because they scatter or because they bend in the magnetic field. In these cases there is often activity above the vertex (e.g. for downward muons that entered high up a steel plane, reversed lateral direction in the field and exited lower down), or large deposits close

to the start of the track. A set of topology cuts are made to identify events with a large number of hits, or that deposit a large amount of energy close to their vertex. Cuts are applied in succession on the following six variables:

- $\langle \Delta_{UV} \rangle$ and $\langle \Delta_{UV}^2 \rangle^{1/2}$: Every hit strip located within ± 4 planes of the highest end of the track is considered, and their charge-weighted average and RMS separation from the track end is calculated separately for the U-view and V-view planes. A high $\langle \Delta_{UV} \rangle$ indicates a large energy deposition at this end of the track, and a high $\langle \Delta_{UV}^2 \rangle^{1/2}$ indicates a large scatter. Events are rejected if $\langle \Delta_{UV} \rangle > 0.25$ m or $\langle \Delta_{UV}^2 \rangle^{1/2} > 0.5$ m in either view.
- ΔR_{max} : Every hit strip within ± 4 planes of the highest end of the track is considered and its 3D distance ΔR from the track end is calculated, with ΔR_{max} defined as the maximum of these distances. Events are rejected if $\Delta R_{max} > 1.25$ m.
- Q_{vtx} : This is the total pulse height deposited in strips within ± 4 planes of the highest end of the track. A high Q_{vtx} implies significant activity around this track end, so events with $Q_{vtx} > 300$ p.e. are rejected.

The distribution of these variables for tracks passing the Stage 2 cuts are shown in Figures 8.2 and 8.3.

8.4.4 Stage 4: Futher topology requirements

Because the cosmic muon spectrum is peaked in the downward direction, a more stringent topology constraint is placed on events with shorter tracks by altering the Q_{vtx} cut for tracks shorter than 25 planes if $Q_{vtx} > 75$ p.e. and $\cos \theta_y > -0.7$ or $|\cos \theta_z| < 0.5$, where θ_y and θ_z are the angles between the reconstructed track direction at the highest end of the track, and the y - and z - axes respectively. In the FD coordinate system, the z -axis is the NuMI beam axis, and the y -axis points vertically upwards. The regions allowed by this cut and its effect on cosmic ray data and atmospheric neutrino MC is shown in Figure 8.4

Events passing all four stages of cuts are included in the CV ν_μ sample.




Figure 8.2: Distribution of $\langle \Delta_{UV} \rangle$ (left) and $\langle \Delta_{UV}^2 \rangle^{1/2}$ (right) for events passing the CV ν_μ trace cuts. Events passing the cuts are indicated by the arrows.




Figure 8.3: Distribution of ΔR_{max} (left) and Q_{vtx} (right) for events passing the CV ν_μ trace cuts. Events passing the cuts are indicated by the arrows.

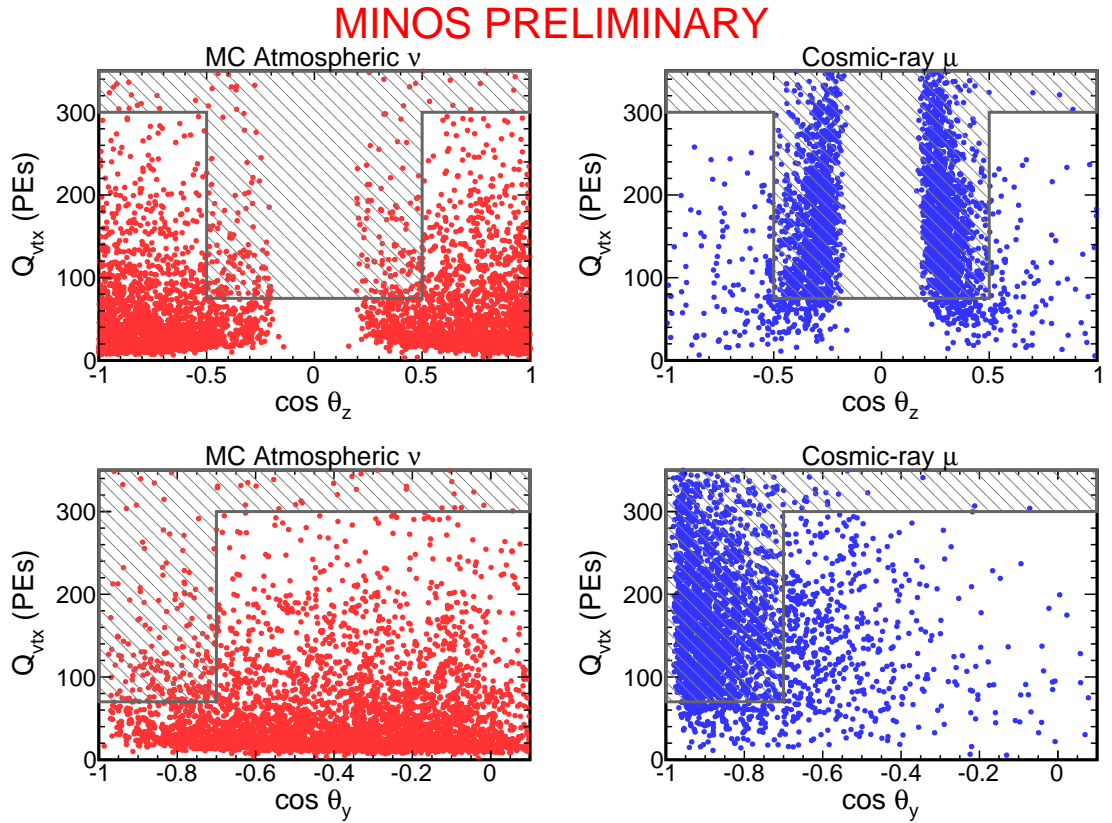


Figure 8.4: Distribution of Q_{vtx} , the total charge in all strips within ± 4 planes of the highest track end, against $\cos \theta_z$ (top) $\cos \theta_y$ (bottom) or for samples of MC atmospheric neutrinos (left) and cosmic ray data (right). The events shown pass the Stage 2 selection cuts and additionally have leading tracks shorter than 25 planes. The grey hatched region shows events rejected by the Q_{vtx} cuts, demonstrating the rejection of downward cosmic events while retaining the atmospheric neutrinos.

8.5 Selecting CV showers

A similar series of cuts are applied to select CV showers, with the focal point of the selection becoming the shower vertex instead of the highest track end. However, many shower-like events also contain one or more tracks, so the containment cuts are imposed on both the primary shower and the primary track in these events. These cuts are similar to those developed in [100].

8.5.1 Stage 1: Fiducial Volume

The first stage is to apply a fiducial volume cut on the shower vertex. This is tighter than the equivalent track cut, requiring that the shower vertex lie no more than 40 cm from the edge of the detector. The requirement to be at least 5 planes from either end of a supermodule is retained, as is the 40 cm coil hole cut, and the 1 m coil hole cut for shower vertices in the first and last 21 planes.

If present, both ends of the leading track are also required to be 20 cm from the edge of the detector, 5 planes from the end of a supermodule, and 40 cm from the centre of the coil hole, with a z-trace of 15 cm

8.5.2 Stage 2: Trace Cut

The z-trace variable for showers is defined in a similar way to tracks: a linear fit for the shower direction is performed using strips associated with the shower, and a projection is made from the vertex to the edge of the detector along this direction. The z-component of this projection is $\Delta^{shw} z$ for the shower. The shower trace cut requires $\Delta^{shw} z > 80$ cm. If present, the leading track must also satisfy $\Delta z > 50$ cm at both ends.

The shower trace variable is shown Figure 8.5 for shower-like events passing the fiducial cuts.

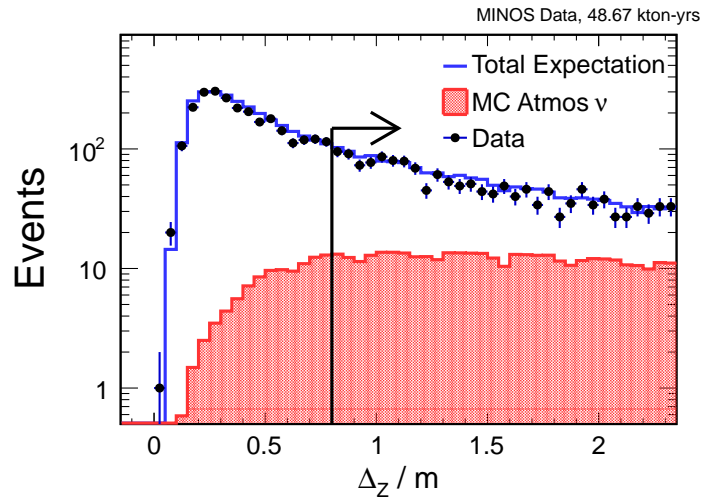


Figure 8.5: Shower z-trace variable used to select contained shower events for the CV ν_e sample. The data points are shown in black, with the expectation from Monte Carlo atmospheric neutrinos (two-flavour oscillations with $\Delta m^2 = 2.32 \times 10^{-3} \text{ eV}^2$ and $\sin^2 \theta = 1$) shown in red, and scaled cosmic ray data shown in blue.

8.5.3 Stage 3: Topology Cuts

These cuts are primarily focused on distinguishing showers associated with cosmic ray muon events from electromagnetic showers produced in ν_e interactions, and fall into two categories. The first uses the shape of the shower to reject showers associated with cosmic muons that deposit a large amount of energy in one plane, or have a pattern of energy deposition that shows strong fluctuations. Short and long showers are defined, where a short shower is ≤ 8 planes in length and a long one contains > 8 planes. Short showers must satisfy:

- Each plane contains a mean number of strips < 4 with an RMS < 3
- And has a mean pulse height per plane < 100 p.e. with an RMS < 100 p.e.

The equivalent cuts for long showers are:

- Each plane contains a mean number of strips < 5 with an RMS < 4
- And has a mean pulse height per plane < 150 p.e. with an RMS < 150 p.e.

The second set of cuts uses the eigenvalues of the shower's moment of inertia tensor to reject muon-like showers associated with a long, straight deposition of energy. The motivation for these cuts is described in greater detail in [100].

8.5.3.1 Moment of inertia

In the classical mechanics of rigid bodies, the moment of inertia tensor is used to parametrise rotation about three axes, and satisfies $\mathbf{L} = \mathbf{I}\omega$ where \mathbf{L} is the angular momentum, ω is the angular velocity, and \mathbf{I} is the moment of inertia tensor. For a collection of points, this tensor can be constructed by [101]:

$$\mathbf{I}_{ij} = \sum_{\text{particles}} m(|\mathbf{x}|\delta_{ij} - x_i x_j) \quad (8.3)$$

Where \mathbf{x} is the particle's position, $i, j = 1, 2, 3$, x_i is the i th component of the particle's position vector, and δ_{ij} is the Kronecker delta .

MINOS defines a generalised moment of inertia tensor, replacing individual particles with hit strips. Hits in U and V planes are treated separately, creating two sets of points in (U, Z) and (V, Z) space. The masses of the particles are substituted for the energy recorded by the strips. The ‘‘centre of mass’’ of the hit strips in each view can be calculated by finding the energy-weighted mean (U or V, Z) coordinates of this system, and the moment of inertia about this point calculated from Equation 8.3. This is then divided by the total strip energy to factor out the dependence on absolute energy, giving two energy-averaged moment of inertia tensors.

It can be seen from Equation 8.3 that \mathbf{I} is real and symmetric. Hence it has real eigenvalues and can be diagonalised by transforming into the orthonormal basis of its eigenvectors. The largest eigenvalue represents the largest energy-averaged moment of inertia in this view. The eigenvectors can be thought of as defining the major and minor axes of an ellipse, with the eigenvalues corresponding to the length along this axis. Hence a large value of the eigenvalue corresponds to a long, narrow ellipse, i.e. a long straight energy deposition, whilst a small value corresponds to a shorter, more diffuse topology.

The largest moment of inertia I_{UV}^{max} is determined by considering both the U-Z and V-Z views, and the following cuts are applied:




Figure 8.6: Moment of inertia variables used to select showers with an electron-like topology by rejecting long, narrow showers. Events to the left of the arrows are selected.

- For short showers events are selected if $I_{UV}^{max} < 0.05 \text{ m}^2$
- For long showers events are selected if $I_{UV}^{max} < 0.15 \text{ m}^2$

The distribution of I_{UV}^{max} for long and short showers is shown in Figure 8.6

Surviving events are classified as CV shower events.

8.6 Neutrino Induced Muons

The requirement that CV ν_μ have contained vertices is imposed because of the large cosmic ray background; requiring the vertex to lie inside the detector helps to veto cosmic ray muons that have penetrated down to the physics level of the Soudan Mine. Here the overburden of 2070 m.w.e. is still sufficiently shallow that downward-going muons dominate over the neutrino signal. As zenith angle increases, so too does the effective overburden through which the muons must penetrate. Measurements from Soudan II [102], also located at Soudan, demonstrate that at overburdens greater than 14000 m.w.e. the cosmic ray background becomes negligible and the

flatter neutrino-induced muon flux becomes dominant. In the MINOS cavern, this occurs for $\theta_{\text{zenith}} \gtrsim 82^\circ$. Consequently, horizontal or upward-going muon events with $\cos(\theta_{\text{zenith}}) < 0.14$ indicate an atmospheric neutrino interaction, regardless of whether their vertex was contained or originated in the rock outside the detector.

An additional set of selection criteria is applied to identify these neutrino induced muons. For this sample, the primary background is due to downward cosmic ray events misreconstructed as upward-going, so the cuts are designed to select events with well-determined directions.

8.6.1 Classifying event directions

Events passing the track-like preselection criteria may be classified as having an upward, downward, or horizontal muon. Tracks are categorised according to the y-component of the track direction, $\cos\theta_y$. In the detector coordinate system, θ_y is the angle between the y-axis, which points vertically upwards, and the trajectory. Consequently, a track travelling vertically downwards (zenith angle = 0°) has $\theta_y = 180^\circ$ and $\cos\theta_y = -1$.

Track directions are classified as follows:

- $\cos\theta_y \leq -0.05$ is a downward track
- $-0.05 < \cos\theta_y < 0.05$ is a horizontal track
- $0.05 \leq \cos\theta_y$ is an upward track

A further set of cuts is applied to filter out events whose reconstructed direction is inconsistent with their start and end points. Events are removed from the upwards and horizontal samples if the track vertex is more than 0.5 m higher than the track end, and from the downward sample if it is more than 0.5 m lower. More stringent directional cuts are applied later in the chain.

Finally, two fits are applied using the strip hit times associated with the track. The track is assumed to be travelling in a straight line at c , either forwards (i.e. upwards for an upward-going track and downwards for a downward-going track) from the reconstructed vertex to the track end, or backwards from track end to vertex.

The RMS deviation of the measured hit times from the best-fit is calculated, and events are removed if $rms_{\text{forward}} \geq rms_{\text{back}}$.

8.6.2 Selecting neutrino induced muons: event sample

The same preselection criteria used to select good track-like events for the CV ν_μ sample in Section 8.2.3 are used for the NIM sample. Successive cuts are then applied to select well-reconstructed tracks that are upward-going.

All fully-contained and partially-contained preselected muon events are initially eligible for selection as NIM events if they are upward-going or horizontal. In addition, upwards and horizontal rock and anti-fiducial (RAF) events are also included: these are events whose vertex does not pass the CV containment cuts, but does pass the coil hole cut.

8.6.3 Stage 1: Track Length

The first stage is to select good quality tracks, which are sufficiently long and cross enough planes that timing information can be used to help determine their direction. This stage requires that:

- The track spans > 15 planes
- The track contains > 5 track-like planes
- The track has a range > 1.5 m

8.6.4 Stage 2: Direction from Timing

These cuts use the timing fits described in Section 8.6.1, where fits are applied to the strip times assuming the muon is travelling at c forwards (upwards) or backwards (downwards) between the reconstructed vertex and end of the track. An additional fit is performed where the muon speed is not fixed at c but is free to vary as a fit parameter.

The first series of cuts are designed to ensure reliable fits by requiring sufficient degrees of freedom (NDoF) for the timing fit:

- Events require $\text{NDoF} \geq 20$ for both the forward and backward fits (tracks can cross multiple strips per plane, so this is not the same as a cut of 20 planes)
- Forward NDoF - backward NDoF must be > -5 (the fit has an iterative procedure that removes outlying hits and repeats the fit, so this indicates that the backward fit is less-consistent with the data than the forward one)
- The timing fit with direction free must have $\text{NDoF} \geq 20$

Then a series of cuts are applied using the RMS of these fits to select events that agree well with the upward (i.e. forward) hypothesis and disagree with the downward hypothesis:

- The upward fit is required to have $rms_{\text{upward}} \leq 4.66$ ns
- While the downward fit must satisfy $rms_{\text{downward}} \geq 3.66$ ns
- The fit with unconstrained direction must satisfy $\chi^2/\text{NDoF} \leq 2.25$

A further cut is placed to compare the goodness of fit between the upward and downward hypotheses:

- $rms_{\text{upward}} - rms_{\text{downward}} < -1.66$

Finally, the ratio of the timing fit RMS to the track length can also help to discriminate between true upward-going events and events with misreconstructed directions. For a track reconstructed in the correct direction, the RMS as a fraction of the total track length is likely to be small, whereas for tracks with misreconstructed direction it tends to be broadly-distributed around a central value of $\sim \sqrt{1/3} = 0.577$. Events are therefore rejected if $rms_{\text{upward}}/n_{\text{planes}} \geq 0.577$.

In the absence of the zenith angle cuts of Section 8.6.1, the track length and direction cuts prove effective at selecting clean, separable samples of tracks with well-reconstructed direction. This is illustrated in Figures 8.7 and 8.8. These contain all the upward, horizontal and downward events passing the track length and direction cuts.

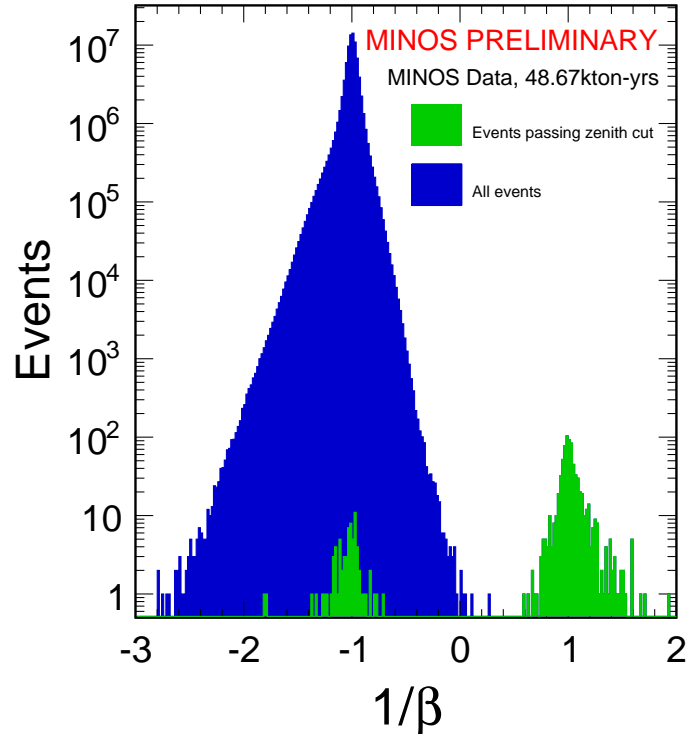


Figure 8.7: Distribution of $\frac{1}{\beta}$, the inverse of the muon velocity, obtained by fitting the muon trajectory as a straight line between the track vertex and end point, for events passing the track length and timing selection for well-reconstructed directions. Events passing the zenith angle selection cut are shown in green: a clear peak of downward neutrino induced muons is evident at +1 separated from the large background of cosmic muon events at -1 . A small subset of horizontal neutrino induced muons with slightly downward trajectories is also apparent near -1 .

The timing fit with the muon velocity free fits the recorded time from a strip as a function of the distance s , using:

$$t = a_0 + a_1 s \quad (8.4)$$

Hence the gradient is the inverse of the muon's velocity: $a_1 = \frac{1}{\beta}$. Figure 8.7 shows the distribution of $\frac{1}{\beta}$ for these events, and is strongly peaked at ± 1 . It also demonstrates a clear separation between the downward ($\beta \approx -1$) cosmic muons, and the upward ($\beta \approx +1$) neutrino induced muons.

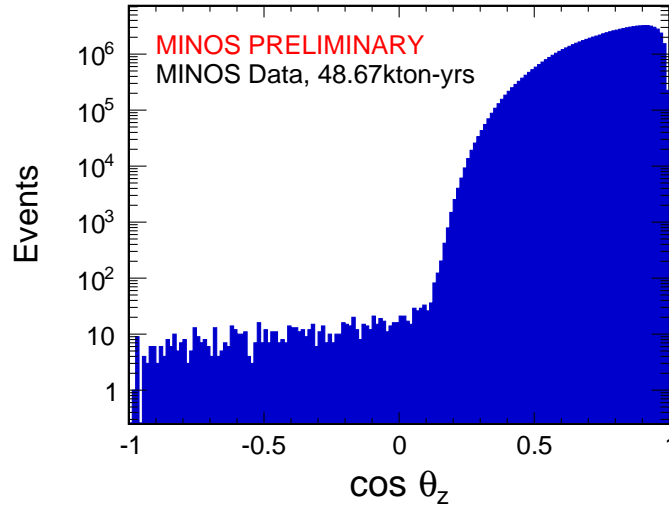


Figure 8.8: Distribution of the reconstructed zenith angle for upward, horizontal, and downward muon events passing the timing and topology cuts. The background from cosmic ray muons dominates at high values of $\cos \theta_{\text{zenith}}$ but falls sharply in the region $0.10 < \cos \theta_{\text{zenith}} < 0.30$ as the effective overburden increases until the flatter neutrino-induced muon flux dominates. The arrow indicates the cut placed at $\cos \theta_{\text{zenith}} < 0.05$ to reduce the cosmic ray background.

The zenith angle cut is illustrated in Figure 8.8. The cosmic ray background falls by six orders of magnitude as $\cos(\theta_{\text{zenith}})$ decreases from 1 to ~ 0.10 , where the NIM spectrum begins to dominate.

8.7 Avoiding double-counting

It is possible for events to pass more than one set of selection criteria, for example a contained vertex event with both a shower and a track object, or an upward-going muon with a contained vertex. Events for which this is the case are assigned to the sample with the greatest power for the oscillation analysis (the CV ν_μ sample, followed by the NIM and finally the CV showers).

The CV ν_μ sample contains all event passing the CV ν_μ selection, plus all upward-going and horizontal NIM events with contained vertices, even if they had associated veto hits or did not pass the remainder of the CV ν_μ selection. Remaining events that pass the NIM selection are placed into the NIM sample, and then events passing the

	Data	Expectation (no oscillations)					
		cosmic	atmos $\nu_e/\bar{\nu}_e$ & $\nu_\mu/\bar{\nu}_\mu$ CC	$\nu_\tau/\bar{\nu}_\tau$ CC	NC	ν -induced μ	Total
CV $_\mu$	1134	44 ± 4	1327 ± 196	0 ± 0	32 ± 8	11 ± 3	1414 ± 204
NIM	590	0 ± 0	33 ± 5	0 ± 0	0 ± 0	699 ± 175	732 ± 175
CV $_e$	899	110 ± 11	661 ± 83	0 ± 0	159 ± 40	1 ± 0	932 ± 124
Total	2623	3078 ± 296					

Table 8.1: Selected events from 48.67 kton-yrs of MINOS atmospheric neutrino exposure. The left column shows the number from data, while the right columns show the expectation from Monte Carlo, assuming null oscillations. A total of 2623 events are selected, compared to the null oscillation prediction of 3078 ± 296 .

	Data	Expectation ($\Delta m_{32}^2 = 2.1 \times 10^{-3}$, $\sin^2 \theta_{23} = 0.5$)					
		cosmic	atmos $\nu_e/\bar{\nu}_e$ & $\nu_\mu/\bar{\nu}_\mu$ CC	$\nu_\tau/\bar{\nu}_\tau$ CC	NC	ν -induced μ	Total
CV $_\mu$	1134	44 ± 4	1023 ± 150	3 ± 1	32 ± 8	7 ± 2	1109 ± 158
NIM	590	0 ± 0	20 ± 3	0 ± 0	0 ± 0	571 ± 143	591 ± 143
CV $_e$	899	110 ± 11	636 ± 79	5 ± 1	159 ± 40	1 ± 0	911 ± 120
Total	2623	2611 ± 244					

Table 8.2: Selected events from 48.67 kton-yrs of MINOS atmospheric neutrino exposure. The left column shows the number from data, while the right columns show the expectation from Monte Carlo, assuming nominal oscillation parameters of $\Delta m_{32}^2 = 2.1 \times 10^{-3} \text{eV}^2$ and $\sin^2(\theta_{23}) = 1$.

CV shower selection that have not been assigned to either of the other two samples are placed in the CV ν_e sample.

8.8 Summary of selected events

In total, 2623 candidate atmospheric neutrino events are selected, of which 1134 are CV ν_μ candidates, 590 are NIM, and 899 are CV showers. These are summarised in Tables 8.1 and 8.2, where Table 8.1 shows the Monte Carlo prediction for the number of selected events assuming null oscillations, and Table 8.2 shows the predicted number of selected events assuming nominal oscillation parameters of $\Delta m_{32}^2 = 2.1 \times 10^{-3} \text{eV}^2$ and $\sin^2(\theta_{23}) = 1$.

The rate at which events are selected into any of the three samples is shown in Figure 8.9 as a function of time. Events passing the selection cuts occur at a mean rate of 0.80 events per live day of running, and the rate is seen to be stable over time.




Figure 8.9: Rate at which candidate atmospheric neutrino events are selected, plotted over time. New data for this analysis are shown in red. Each point corresponds to approximately two months (61 days) and shows the number of events selected per live day of running during that 61 day period. The mean rate is 0.80 events per live day, and is shown by the grey dashed line.

Chapter 9

Oscillation Fit

The aim of the fitting process is to constrain the mixing angle $\sin^2 \theta_{23}$ and mass splitting Δm_{32}^2 using the Far Detector atmospheric neutrino samples. The selected events are fitted using the standard MINOS three-flavour fitting package, GhostFitter [103]. GhostFitter is designed to perform a simultaneous fit to a variety of different samples. It compares the energy spectra recorded in data with predictions at different values of the oscillation parameters, and uses a binned log-likelihood method to identify the oscillation parameters that provide the best agreement with data once statistical and systematic errors have been accounted for. The same package is used for MINOS beam analyses [104] and for combined analyses incorporating the beam ν_μ disappearance, ν_e appearance, and atmospheric samples [42] [105] [106].

9.0.1 Fit method

The input to GhostFitter is a series of ‘samples’ (e.g. the atmospheric CV ν_μ sample) which may optionally be broken down further into ‘vessels’ (e.g. the grouping of events into subsamples based on their energy resolution employed for the beam CV ν_μ sample [104]). Each vessel consists of a set of unoscillated events, including backgrounds, in MC and data where appropriate. The fitting proceeds by applying oscillations at a given set of parameters and any relevant systematics to each vessel, and comparing the oscillated prediction with FD data.

The fitter then uses a maximum likelihood method to identify the parameters \mathbf{x} that minimise the negative log-likelihood (multiplied by a conventional factor of 2

for comparison with a χ^2 test statistic):

$$-2 \ln \mathcal{L} = 2 \sum_i^{N_v} \sum_j^{N_b} \left(\mu_{ij}(\mathbf{x}) - N_{ij} + N_{ij} \ln \frac{N_{ij}}{\mu_{ij}(\mathbf{x})} \right) \quad (9.1)$$

Here, N_v counts the total number of vessels across all samples, and N_b represents the number of energy bins in a given vessel. The number of entries in each bin is taken to be described by a Poisson distribution whose expectation μ_{ij} is calculated by applying oscillations at a point in parameter space described by \mathbf{x} [45].

Systematic errors are included by adding nuisance parameters to the expression for $-2 \ln \mathcal{L}$. After including a total of N_s systematics, the final expression becomes:

$$-2 \ln \mathcal{L} = 2 \left[\sum_i^{N_v} \sum_j^{N_b} \left(\mu_{ij}(\mathbf{x}) - N_{ij} + N_{ij} \ln \frac{N_{ij}}{\mu_{ij}(\mathbf{x})} \right) + \sum_k^{N_s} \frac{\Delta_k^2}{2\sigma_k^2} \right] \quad (9.2)$$

The Δ_k terms represent the shift of the k th systematic from its nominal value, and the σ_k represent the uncertainty ascribed to that systematic.

9.0.2 Oscillation templates

The most computationally expensive part of the fit procedure is applying the combination of oscillations and systematic shifts to each of the samples. To mitigate this the oscillations are precomputed, producing a series of ‘templates’. For each subset of events we have a set of unoscillated Monte Carlo histograms consisting of reconstructed vs. true energy and zenith angle distributions for nominal systematics, and $\pm 1\sigma, \pm 2\sigma$ shifts to each of the N_s systematics.

Sets of oscillated templates are then produced by weighting the unoscillated templates according to the oscillation probabilities. In an update since [103], this process is now performed using the exact three-flavour probabilities, including matter effects, calculated using fast matrix methods derived from [107]. For the purpose of the atmospheric neutrino sample, MINOS employs a model Earth with four layers of constant density derived from the Preliminary Reference Earth Model (PREM) [108], illustrated in Figure 9.1. Tests with a more-complex model composed of 42 layers yielded a considerably slower fit but a negligible difference in the final result.

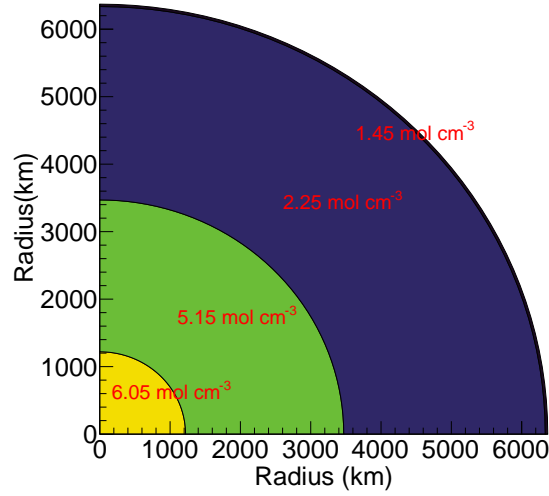


Figure 9.1: The four-layer model of the Earth used to calculate the oscillation probabilities of atmospheric neutrinos. This is a simplified version of the Preliminary Reference Earth Model [108], with electron densities n_e given as a function of the distance, r , from the Earth’s core. The model describes four regions: the inner core ($r < 1220$ km), outer core ($1220 \leq r < 3470$ km), mantle ($3470 \leq r < 6336$ km), and crust ($r \geq 6336$ km).

Templates are calculated over a large grid of points in oscillation parameter space. For the atmospheric-only analysis (the combined beam and atmospheric analysis used a finer, narrower grid [106]), the grid points are defined by:

Parameter	Number of bins	Minimum	Maximum
$ \Delta m_{32}^2 $	60	0.001 eV ²	0.003 eV ²
$\sin^2 \theta_{23}$	21	0.3	0.7
δ_{CP}	5	0	2π
$\sin^2 \theta_{13}$	3	0.0151	0.0333

Table 9.1: Grid points used to generate oscillated templates for the atmospheric neutrino analysis. Templates are generated for both the normal and inverted hierarchies, giving a total of 37800 points.

The oscillations use fixed values for the terms related to solar mixing: $\Delta m_{21}^2 = 7.54 \times 10^{-5}$ eV² and $\sin^2 \theta_{12} = 0.307$ [109]. The allowed range of $\sin^2 \theta_{13}$ is based on a combined average of results from the Daya Bay [110], RENO [54], and Double

CHOOZ [111] reactor experiments. A penalty term is also added to the nuisance parameters and penalises the likelihood if $\sin^2 \theta_{13}$ moves away from this average of $\sin^2 \theta_{13} = 0.0242 \pm 0.002$ [112]. In addition to oscillations, predicted atmospheric neutrino fluxes are also weighted to account for the solar cycle.

To enable the use of pre-calculated oscillation templates, the negative log likelihood function is minimised using a grid search method. It is also not practical to generate templates for every possible value of every systematic shift. Instead, templates are calculated at 5 separate values for each systematic: the nominal value, $\pm 1\sigma$, and $\pm 2\sigma$ shifts. The effect of an arbitrary shift is calculated by interpolation, using a weighting scheme that exactly reproduces the appropriate histograms for $0, \pm 1\sigma, \pm 2\sigma$ and varies smoothly in between [103]. At each grid point, the oscillated predictions are taken, and Minuit [113] is used to determine the vector of systematic shifts that minimise the total systematic nuisance parameter. It is possible to only apply certain systematics to certain vessels, or to fully correlate a systematic across several vessels.

9.1 Systematic errors for the atmospheric sample

A total of 14 sources of systematic error are accounted for in the analysis. These can be broken down into four main categories:

- **Normalisations** - a parameter controls the overall normalisation for each of the CV ν_μ , CV ν_e , upward NIM and neutral current background samples, with two further parameters describing the ratio of ν_μ to $\bar{\nu}_\mu$ in the CV ν_μ and NIM samples.
- **Spectral shapes** - a parameter is used to represent the uncertainty in the shape of the energy spectrum for the CV ν_μ and $\bar{\nu}_\mu$ samples (separately), and the upward NIM from ν_μ and $\bar{\nu}_\mu$ (also separately).
- **Charge identification** - the CV ν_μ and upward NIM samples are divided into neutrinos and antineutrinos based on the reconstructed muon charge. A systematic uncertainty is associated with each of these.
- **Resolutions** - separate systematics are used for the energy estimation of contained and escaping tracks, shower energy, and the reconstructed zenith angle.




Figure 9.2: Muon neutrino (upper red points) and antineutrino (lower red points) fluxes and their percentage uncertainties (upper and lower green lines respectively). The ratio of ν_μ to $\bar{\nu}_\mu$ is shown by the blue points, and its uncertainty by the hollow black points. Plot taken from [114].

The dominant uncertainties arise from the normalisations of the CV ν_μ and $\bar{\nu}_\mu$, and NIM samples, and these are dominated by cross-section and flux modelling uncertainties.

9.1.1 Normalisation uncertainties

The uncertainty used for the CV ν_μ normalisation follows the treatment of Barr et al [114]. As shown in Figure 9.2, the flux uncertainty is approximately 15% for energies in the range 1-10 GeV, and increases above 25% as the energy increases past 100 GeV. Hence a 15% uncertainty is assigned to the normalisation of the CV ν_μ sample, whose energies are typically below 10 GeV. The neutrino induced muons originate from higher energy cosmic rays (up to 1 TeV, with a median energy of 50 GeV) and so a larger 25% uncertainty is assigned.

The uncertainty in the normalisation of the CV shower sample is largely attributable to the modelling of hadronic processes. This was analysed using model spread in [96], replacing the GCALOR [95] hadron propagation model used in the MINOS MC with an alternative GHEISHA model [115] [116]. The difference in the rate of CV showers between the two models is 7%. In the oscillation fit these are

separated into ν_e and NC components, and normalised relative to the CV ν_μ sample. The ν_e and NC components are assigned uncertainties of 5% and 20% respectively, with the NC uncertainties dominated by the cross-section [117].

There are three main sources of uncertainty that contribute to the systematic error on the $\nu_\mu/\bar{\nu}_\mu$ ratio: the fluxes, cross-sections, and reconstruction effects. As evidenced by Figure 9.2, at energies below ~ 10 GeV many of the overall flux uncertainties cancel in the ratio, so the total uncertainty is small. This is because cosmic rays with these energies that interact with the atmosphere mostly produce charged pions, and the decay chain:

$$\pi^+ \rightarrow \nu_\mu + \mu^+ \rightarrow \nu_\mu + \bar{\nu}_\mu + e^+ + \nu_e \quad (9.3)$$

gives rise to a ratio that approaches unity with low uncertainty even if the overall rate is poorly-known. For the CV ν_μ , a conservative uncertainty of 4% is used. As energy increases the uncertainty in the charge ratio becomes higher due to the increased production of kaons, and the potential for muons to survive until they hit the ground. For this reason, the flux contribution to the charge ratio uncertainty is set to 10% for the NIM sample [118].

The total MINOS cross-section uncertainty for CC events is highest at 8% for 1-5 GeV due to difficulties modelling the region where neutrino interactions change from quasielastic to the DIS regime. At higher energies, the uncertainty decreases to 2% due to constraints from DIS experiments. The uncertainty is larger for $\bar{\nu}_\mu$ than ν_μ due to the smaller data sample available [117]. To evaluate the contribution to the charge ratio uncertainty, the input parameters to the NEUGEN3 event generator are adjusted according to their systematic uncertainties. The shifts in the neutrino energy spectrum obtained when each parameter is adjusted are added in quadrature to produce an error band as a function of neutrino energy. The final contribution to the uncertainty is determined by integrating this error band, weighted according to the MC energy spectrum, and yields an uncertainty of 8.5% for the CV ν_μ sample and 4% for the NIM sample, whose events fall mainly into the DIS regime.

An additional contribution to the charge ratio uncertainty due to detector and reconstruction effects was included. This employed a scanning study of 2262 events [119,120], comparing the curvature of events scanned by eye to the charge determined by the reconstruction in order to estimate the level of discrepancy in the purity due

to any reconstruction differences between data and MC. An additional uncertainty of 3% is allocated to the CV ν_μ sample, and 6% to the NIM charge separation purity. This covers the discrepancy between reconstructed and MC truth purities, and the consistency between MC and data identified by this study.

Adding these three components in quadrature yields final charge ratio uncertainties of:

$$\text{CV } \nu_\mu : \text{flux} \oplus \text{cross-section} \oplus \text{reco.} = 4.0\% \oplus 8.5\% \oplus 3.0\% = 10\% \quad (9.4)$$

$$\text{NIM} : \text{flux} \oplus \text{cross-section} \oplus \text{reco.} = 10.0\% \oplus 4.0\% \oplus 6.0\% = 12.5\% \quad (9.5)$$

9.1.2 Spectral shape uncertainties

In addition to scaling effects, uncertainties in the flux and cross-section can also affect the shape of the atmospheric neutrino energy spectrum. Systematic errors are included to account for both of these effects by allowing an energy-dependent scaling of the Monte Carlo CV ν_μ sample, described by a single free parameter, α :

$$f(E_\nu) = \begin{cases} 1 + \alpha(E_\nu - E_0) & \text{if } E_\nu \leq E_0 \\ 1 + \alpha \ln(E_\nu/E_0) & \text{if } E_\nu > E_0 \end{cases} \quad (9.6)$$

Here the high energy form is motivated by the approximate power-law distribution of high energy atmospheric neutrinos. The threshold energy E_0 is set to 3 GeV as this is approximately the median neutrino energy for events whose direction is well-measured. The systematic nuisance term in the final oscillation fit assumes the α parameters are distributed according to a Gaussian, with mean zero and width σ_α .

For the flux component, σ_α is determined by comparing three different MC schemes: the 3D flux model produced by the Bartol group [121] (used in the default MINOS atmospheric MC) and alternative 3D models from the FLUKA [122] and Honda *et al.* [123] groups. The ratios of FLUKA/Bartol and Honda/Bartol are plotted separately for neutrinos and antineutrinos, normalised so that the ratios are set to unity at $E_0 = 3$ GeV. The value of α required to generate an error band that covers the spread between the different models is determined, and σ_α for the flux is set to this value. In the case of the cross-section contribution, instead of comparing different models the quasielastic axial vector mass term in the cross-section is varied

by $\pm 15\%$ and the same procedure is applied to construct an error band that covers the distortion in the spectrum caused by these changes [124].

The final uncertainties using this procedure are 5% (flux) and 3% (cross-section) for the ν_μ , and 6% (flux) and 2% (cross-section) for the $\bar{\nu}_\mu$. Combining in quadrature yields a 6% shape uncertainty for neutrinos and antineutrinos [118]. A shape uncertainty is assigned to the CV ν_μ , CV $\bar{\nu}_\mu$, NIM μ and NIM $\bar{\mu}$ samples.

9.1.3 Reconstruction uncertainties

The energy resolution for muons whose energy is measured by range is 3%. The resolution of events whose energy is measured by curvature was estimated by comparing estimates using range and curvature for a sample of stopping cosmic muons. The two techniques agree at the 3% level, and a conservative 5% error is assigned to muon energies measured from curvature. The shower energy scale uncertainty is set to 15% and reflects a combination of a 10% uncertainty due to the hadronisation model, a 5% calibration error, and a 10% error in the energy reconstruction itself, estimated by considering the shift in the mean reconstructed energy of vertex showers accompanying ν_μ events when their length increases by one plane [124].

Uncertainties in the predicted ratio of upward- to downward-going neutrinos also give rise to a systematic error. Assuming that oscillations occur, one expects a deficit of upward-going muon neutrinos compared to downward-going ones, due to the large L/E . However, even in the absence of oscillations this ratio would not be unity due to deflections of the primary cosmic rays caused by the Earth's magnetic field.

Figure 9.3 shows the uncertainty on the up/down and up/horizontal ratios of the atmospheric neutrino flux, as a function of true neutrino energy. The uncertainty on the up/down ratio cancels almost exactly above ~ 3 GeV because the energy of the parent cosmic rays is high enough that the effects of the Earth's magnetic field are negligible. Away from this region, the uncertainties reflect hadron production uncertainties. Selection effects due to the Earth's magnetic field and the effect of the density of the atmosphere on secondary interactions mean that these uncertainties do not cancel exactly for the up/down ratio at low energies, and the up/horizontal ratio at higher energies [114]. An overall 2% uncertainty is attributed to the flux model.




Figure 9.3: Uncertainty in the up/down and up/horizontal ratio, taken from [114]. For the purposes of this plot, ‘up’ means a zenith angle satisfying $\cos(\theta_{\text{zenith}}) < -0.6$, ‘down’ means $\cos(\theta_{\text{zenith}}) > 0.6$, and ‘horizontal’ means $|\cos(\theta_{\text{zenith}})| < 0.3$.

For a reconstructed track with a given inclination, the overall direction (i.e. whether the particle was travelling upwards or downwards along that trajectory) is determined using timing information as described in Chapter 8. Purity measurements made using MC have a systematic error related to the level of agreement between simulation and reality, and the MC parameters related to the timing are an effective refractive index that controls the propagation speed of light and a smearing term that accounts for the timing resolution. These were tuned to match a sample of cosmic ray data. A conservative systematic was assigned by comparing the nominal MC with this tuned version, which results in a 1.1% discrepancy.

A total uncertainty of 3% is assigned to the up-down ratio.

9.1.4 Summary of systematic errors

The complete set of systematic errors considered is summarised in Table 13.2. The normalisations of the CV $\bar{\nu}_\mu$ and shower samples are implemented by the fitter through their ratio to the CV ν_μ sample, and the normalisation of the NIM $\bar{\mu}$ relative to the NIM μ sample.

Systematic name	Description	Size (%)
AtmosNormCV	Normalisation of CV ν_μ sample	15.0
AtmosNormRock	Normalisation of NIM μ sample	25.0
AtmosNormNue	Normalisation of ν_e shower sample	5.0
AtmosNCBkg	Normalisation of NC shower background	20.0
AtmosChgCV	Normalisation of CV $\bar{\nu}_\mu$	10.0
AtmosChgRock	Normalisation of NIM $\bar{\mu}$ sample	12.5
AtmosSpecNumuCV	Spectral shape of CV ν_μ sample	6.0
AtmosSpecNumuBarCV	Spectral shape of CV $\bar{\nu}_\mu$ sample	6.0
AtmosSpecNuRock	Spectral shape of NIM μ sample	6.0
AtmosSpecNubarRock	Spectral shape of NIM $\bar{\mu}$ sample	6.0
AtmosZenith	Upward/downward ratio	3.0
AtmosNumuTrkEn	Track energy resolution (range)	3.0
AtmosNumuTrkEnExit	Track energy resolution (curvature)	5.0
AtmosNumuShwEn	Shower energy resolution	15.0

Table 9.2: Summary of the type and size of systematic errors used in the atmospheric neutrino oscillation fit.

Chapter 10

Fit results

The output from the fit is a two-dimensional surface containing the value of $-2 \ln \mathcal{L}$ in bins of $\sin^2(\theta_{23})$ and Δm_{32}^2 in both the normal and inverted hierarchies. This follows the binning scheme used to generate the template grid. The best-fit oscillation parameters are the ones that generate \mathcal{L}_{\min} , corresponding to the minimum of $-2 \ln \mathcal{L}$, and two dimensional contours can be constructed through constant values of $-2 \ln \mathcal{L}$.

The n-sigma confidence intervals are represented by the contours through points satisfying:

$$n^2 = -2(\ln \mathcal{L} - \ln \mathcal{L}_{\min}) \quad (10.1)$$

$$= -2\Delta \ln \mathcal{L} \quad (10.2)$$

These are converted into 68% and 90% confidence intervals by constructing contours through the points at which $-2\Delta \ln \mathcal{L} = 2.30(68\%)$ and $4.61(90\%)$, using the standard p-values for a two-dimensional χ^2 distribution.

The contour obtained from fitting the atmospheric neutrino sample is shown in Figure 10.1. The best fit point is indicated by a star, and occurs at $\Delta m_{32}^2 = 2.03 \times 10^{-3} \text{ eV}^2$, $\sin^2(\theta_{23}) = 1.0$. The preference is for the normal hierarchy and maximal $\sin^2(\theta_{23})$, although the very similar size and shape of the contours (allowing for the sign inversion when switching between the NH and IH) indicates that the hierarchy preference is small.

One dimensional limits can be defined for Δm_{32}^2 and $\sin^2(\theta_{23})$ separately by profiling the likelihood surface to find the minimum value of $-2\Delta \ln \mathcal{L}$ for each value




Figure 10.1: The 68% and 90% confidence limits for the $(\sin^2 \theta_{23}, \Delta m_{32}^2)$, parameter space, resulting from a fit to 48.67 kt-yrs of MINOS atmospheric neutrino data. The best fit is indicated by the star.




Figure 10.2: The 1D likelihood profile for the $|\Delta m_{32}^2|$ parameter, calculated separately for the normal hierarchy (red curve) and the inverted hierarchy (blue curve) by profiling the 2D likelihood surface in $(\sin^2(\theta_{23}), \Delta m_{32}^2)$.

of the parameter in question. The profile for Δm_{32}^2 is shown in Figure 10.2, and for $\sin^2(\theta_{23})$ in Figure 10.3. Dotted lines are drawn at the values of $-2\Delta \ln \mathcal{L}$ corresponding to the 68% and 90% confidence levels for a one-dimensional χ^2 distribution ($\Delta\chi^2 = 1$ and 2.71 respectively).

Table 10.1 summarises the best-fit points in each hierarchy. The best-fit point in the inverted hierarchy has $\Delta m_{32}^2 = -2.13 \times 10^{-3} \text{ eV}^2$, and is disfavoured by 0.559 units of $-2\Delta \ln \mathcal{L}$ compared to the overall best-fit. This preference is not statistically significant. For a one dimensional χ^2 distribution this corresponds to disfavoured at the 55% level, although care should be taken when interpreting the significance with which the mass hierarchy can be determined [125] in this fashion.




Figure 10.3: The 1D likelihood profile for the $\sin^2 \theta_{23}$ parameter, calculated by profiling the 2D likelihood surface in $(\Delta m_{32}^2, \sin^2 \theta_{23})$

Hierarchy	Best fit oscillation parameters				$-2\Delta \ln \mathcal{L}$
	Δm_{32}^2 ($\times 10^{-3}$ eV 2)	$\sin^2 \theta_{23}$	$\sin^2 \theta_{13}$	δ_{CP}	
Normal	2.03	0.50	0.0242	0	-
Inverted	2.13	0.50	0.0242	1.57	0.559

Table 10.1: Best fit oscillation parameters and $-2\Delta \ln \mathcal{L}$ relative to the overall best fit, for each mass hierarchy, using the MINOS and MINOS+ atmospheric neutrino data combined. Note that the fit includes an external constraint of $\sin^2 \theta_{13} = 0.0242 \pm 0.0025$, based on reactor neutrino experiments.

The fitted systematic errors at the best-fit point in oscillation parameter space are displayed in Table 10.2. The shifts are consistent with the nominal values, with

Systematic Parameter	Best fit	Systematic parameter	Best fit
AtmosNormCV	+0.50	AtmosSpecNumuBarCV	+1.15
AtmosNormRock	+0.10	AtmosSpecNuRock	+0.35
AtmosChgCV	-0.67	AtmosSpecNuBarRock	+0.36
AtmosChgRock	+1.42	AtmosZenith	-0.17
AtmosNueNorm	-0.32	AtmosNumuTrkEn	-0.004
AtmosNCBkg	-0.52	AtmosNumuTrkEnExit	+0.17
AtmosSpecNumuCV	-0.59	AtmosNumuShwEn	+0.40

Table 10.2: Fitted systematic nuisance parameters at the best-fit oscillation point.

the largest ones being a 1.42σ pull on the NIM charge ratio, and a 1.15σ shift in the CV $\bar{\nu}_\mu$ spectral parameter.

10.0.5 Atmospheric neutrino spectra

A selection of spectra from the atmospheric samples are shown below. Figure 10.4 contains the measured CV ν_μ and $\bar{\nu}_\mu$ samples as a function of zenith angle, in three bins of energy. This is the sample with the most events and the greatest sensitivity to the oscillations, particularly in the lower energy bins. The cosmic ray background, shown in blue, is determined from data by weighting the spectrum of vetoed cosmic ray events that would otherwise pass the CV ν_μ selection by the measured veto inefficiency.

Similar spectra for the neutrino induced muons and antimuons are shown in Figure 10.5. The event count for this sample is smaller and so its effect on the fit result is less powerful than the CV ν_μ sample. These different samples, each with their own statistical fluctuations, give rise to the non-elliptical contour when a combined fit is performed using all of them. The CV shower sample has a lower sensitivity to the oscillation parameters than the ν_μ and NIM samples due to poorer energy resolution and the large NC component, but its inclusion in the fit helps to constrain the normalisation parameters.




Figure 10.4: Spectra for the contained vertex ν_μ (top) and $\bar{\nu}_\mu$ (bottom) samples, showing the number of selected and predicted events as a function of zenith angle. For display purposes, three plots are shown for different energy bins, with the energy increasing from left to right. The blue shaded region represents the cosmic ray background. The three lines represent the predicted spectra in the case of no oscillations (grey line), the overall best-fit oscillations (red dashed line) which occur in the normal hierarchy, and the best-fit point assuming the inverted hierarchy (blue dashed line).




Figure 10.5: Zenith angle spectra for the neutrino induced muon samples, divided into ν_μ events (top) and $\bar{\nu}_\mu$ events (bottom). For display purposes, two plots are shown for different energy bins, with the energy increasing from left to right. The three lines represent the predicted spectra in the case of no oscillations (grey solid line), the overall best-fit oscillations (red dashed line) which occur in the normal hierarchy, and the best-fit point assuming the inverted hierarchy (blue dashed line).

Chapter 11

CHIPS

In order to measure three flavour oscillation effects with the precision required to determine the mass hierarchy or to measure CP violation, precise detectors capable of identifying electrons from ν_e appearance are required. At the same time, these detectors must be massive to provide a suitably large sample of events. Unless the oscillation parameters, in particular δ_{CP} , lie at particularly favourable values, currently-running experiments may be unable to settle these questions definitively. Several future experiments are proposed [126] [127], with typical costs on the order of a billion dollars and long lead times, with data-taking beginning around 2025.

The CHIPS project aims to reduce the cost of large neutrino detectors to \sim \$200 – \$300k /kton by constructing water Cherenkov detectors in deep bodies of water on the Earth’s surface instead of in underground caverns [128]. Such bodies include lakes, reservoirs, and flooded mine pits. The detector would be sunk to the bottom of the body of water so that the water above it would provide a modest overburden, and purified water would be separated from the lake using a watertight and light-tight flexible membrane. The lake water surrounding the detector would provide support to a lightweight structure, avoiding the need for costly excavation of underground rock caverns. Detectors would be constructed from smaller modules, beginning with CHIPS-10, a 10 kton R&D phase.

11.0.6 Water Cherenkov detectors

The operating principle of a water Cherenkov detector is that charged particles travelling faster than the local speed of light in a dielectric medium will emit photons

due to the Cherenkov effect. A water Cherenkov detector consists of a large body of water typically instrumented with PMTs.

An incoming neutrino interacts with one of the nuclei in the water, emitting a charged particle such as an electron or muon. The speed of light in water is approximately $\frac{3}{4}c$, and so the charged particle can exceed this speed if it is sufficiently relativistic. The induced polarisation of the water molecules in the wake of the charged particle causes a coherent shockwave of light to be emitted. The Cherenkov angle at which the light is emitted is given by:

$$\cos \theta_{Ch} = \frac{1}{n\beta} \quad (11.1)$$

Where β is the particle's speed as a fraction of c , and n is the refractive index of the medium (and is in general a function of wavelength). When the particle is travelling with $n\beta < 1$, the previous equation gives $\cos \theta_{Ch} > 1$ and so there can be no emission. This speed requirement can be translated into a minimum energy for a particle of mass m :

$$E_{\text{thresh}} = \gamma m = \frac{m}{\sqrt{(1 - (1/n)^2)}} \quad (11.2)$$

Using the approximate value of $n = \frac{4}{3}$ gives a Cherenkov angle of $\sim 41^\circ$ and a threshold energy of $\sim 1.5m$.

The Cherenkov emission per unit wavelength per centimetre travelled by the particle is given by [45]:

$$\frac{d^2 N}{d\lambda dx} = \frac{2\pi\alpha^2}{\lambda^2} \left(1 - \frac{1}{\beta^2 n^2(\lambda)} \right) \quad (11.3)$$

Where α is the fine structure constant. Setting $\beta = 1$ and $n = \frac{4}{3}$ and integrating between 350 nm and 650 nm, the approximate optical wavelength range over which a PMT is sensitive, gives ~ 240 photons cm^{-1} before absorption, scattering and PMT efficiencies are taken into account.

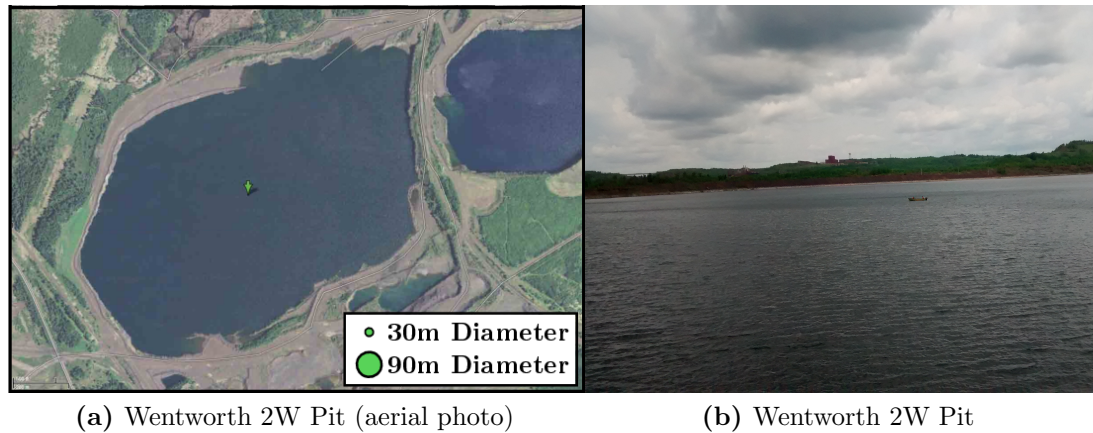


Figure 11.1: Photographs of the Wentworth 2W pit, the proposed site of CHIPS. Indicated on the aerial photograph (left) are markers 30 m and 90 m in size; a 10 kton detector 20 m high requires a diameter of approximately 25 m.

11.0.7 Location

The location identified for CHIPS is the Wentworth 2W Pit, a disused surface iron mine near to Hoyte Lakes in Northern Minnesota. The Wentworth Pit lies 7 mrad off-axis from the NuMI beam at a baseline of 712 km and has a large flat region at its deepest point, approximately 60 m below the surface. It has road access and nearby power lines associated with its history as a mine.

Figure 11.1 shows a satellite image of the Wentworth Pit with two markers indicating a 30 m and 90 m diameter shown for scale. Alongside is a photograph of the Pit taken in June 2015, during the recovery of the CHIPS-M prototype.

11.1 CHIPS-M

CHIPS-M (for ‘Model’), a small prototype detector shown in Figure 11.2, was constructed and deployed into the Wentworth 2W pit during the summer of 2014 in order to test the suitability of materials (especially the geomembrane liner) and the performance of the water filtration system. CHIPS-M was an octagonal prism, with an inner volume 131 in (3.318 m) tall and 125 in (3.186 m) across, kept isolated from the main water of the lake by a layer of Seaman 8130 XR-5 geomembrane, a reinforced flexible polymer membrane.



Figure 11.2: Photograph of the CHIPS-M prototype detector at the end of July 2014, shortly before deployment. Also pictured is the large umbilical cable carrying pipes from the detector to the small building housing the water filtration system, and the smaller black cables that provide power and communication to the DOMs.

The end caps were constructed from aluminium stage struss, held apart by eight stainless steel columns and braced with stainless steel cables to prevent the two caps from twisting relative to one another. The detector was initially instrumented with five 10" IceCube DOMs [129] (digital optical modules housing a PMT and readout board), but was recovered in 2015 and fitted with additional instrumentation. This consisted of a plane of 3" PMTs constructed by unfolding a spherical KM3NeT optical module [130] into a flat panel, and a second plane constructed by potting 3" tubes formerly used for the NEMO-3 experiment [131] and protecting the electronics using off-the-shelf plumbing equipment.

Analysis of the prototype data has been discussed elsewhere [132, 133], but the process of constructing and operating the detector proved invaluable in developing a water filtration test stand [134] and for informing the design of a larger module.




Figure 11.3: Early design drawing of CHIPS-10 showing a lightweight structure (red and green lines) supporting panels of photomultiplier tubes (dark grey rectangles) and separated from the main body of the lake water by a flexible liner (light grey region). Image taken from [135].

11.2 CHIPS-10

CHIPS-10 is intended to be a 10 kton R&D module with a cylindrical fiducial volume approximately 25 m in diameter and 20 m tall. Preliminary designs call for construction from tiling rectangular panels consisting of a rigid frame supporting a rectangular portion of flexible watertight membrane. An internal support structure with sliding rails onto which PMT planes can be mounted will be attached approximately 1 m inside the outer structure. An early schematic is shown in Figure 11.3.

The complete structure for the bottom endcap would be constructed first. A first layer of the cylindrical barrel, approximately 3.5 m high would then be attached to the bottom cap, with the whole structure floating in a shallow area of the Pit, controlled using flotation bags. This would be filled with purified water and covered, with the top endcap constructed underneath this cover, floating on the purified water. It would then be joined to the barrel wall to complete a small ‘slice’ of the full detector. Additional height can be incrementally added to the barrel by detaching the top endcap, and adding one layer of side panels at a time. In this way, construction could be staged to take place during the summer months before submerging the

detector to take data throughout the winter when the water on the surface of the lake freezes over, and retrieving it for later expansion as resources allow.

11.2.1 Novel analysis strategies for water Cherenkov detectors

The physics goals of CHIPS differ from those of previous water Cherenkov experiments in a number of ways, and present their own design challenges and opportunities. Because CHIPS is concerned with accelerator neutrinos the timing of beam spills can be used to veto out-of-spill cosmic events. Furthermore, CHIPS does not also intend to search for supernova neutrinos or processes such as proton decay so a larger cosmic ray background rate is acceptable, enabling a shallower overburden.

Studies using models of the cosmic ray flux as a function of depth and detector geometry [136] [137] indicate a cosmic ray background rate of 50.5 kHz (0.51 events per $10 \mu\text{s}$ beam spill) given a 40 m overburden. Assuming a conservative 500 ns dead time is incurred by the entire detector, with no capability to reconstruct a simultaneous cosmic muon and beam interaction, this translates to a 2.5% total dead time. Those calculations applied to a 40 kton detector; with a 10 kton detector of the same height (which is fixed to 20 m because of the 60 m depth of the lake and the need for sufficient overburden) this can be expected to scale approximately with the area of the top cap, by a factor of 1/4.

Traditional design strategies for water Cherenkov detectors such as Super-Kamiokande [138] utilise large (20" diameter) PMTs and endeavour to cover a high fraction of the internal surface area of the detector walls with PMT photocathode. This enables the detection of low energy events but requires many thousands of PMTs, a substantial expense. In the case of CHIPS observing the NuMI beam in its medium energy tune from 7 mrad off-axis, the typical energies of the leading muon or electron from ν_μ or ν_e CC events is greater than 500 MeV (c.f. thresholds \sim MeV for solar neutrino experiments) as shown in Figure 11.4. This suggests that a lower photocathode coverage may be permissible because of the higher light yield from each event. Furthermore, being a dedicated beam experiment CHIPS is concerned with neutrino events that are primarily forward-going, so instrumenting all walls with the




Figure 11.4: Energies of the leading electrons and muons produced from NuMI CCQE ν_e and ν_μ interactions at the CHIPS baseline, 7 mrad off-axis.

same coverage may not be optimal. Instead photodetectors could be concentrated more densely on the downstream walls of the detector.

Finally, the use of smaller and more ubiquitous 3" PMTs in underwater neutrino detectors has been demonstrated, notably by KM3NeT [130]. This allows alternative manufacturers to be considered, and the greater granularity compared to larger tubes may assist event reconstruction. Such tubes, with timing resolutions on the order of 1-2 ns may also allow timing information as well as recorded charge to contribute significantly to the reconstruction.

A 10 kton CHIPS R&D module would enable these opportunities to be tested in a real neutrino beam environment. Initially, however, a Monte Carlo simulation has been written to construct generic water Cherenkov detector geometries with PMTs placed according to customisable patterns. Alongside this, a full track reconstruction program has been written to evaluate the effect of the design choices described above on the physics performance of the detector. The simulation and reconstruction are described in Chapter 12, and the studies and their results in Chapter 13.

Chapter 12

Reconstruction for CHIPS

An event reconstruction algorithm has been written for CHIPS, which can be used both for Monte Carlo design studies and as the base of a framework for reconstructing real data. It can be used to guide the design of CHIPS-10, for instance by proposing a PMT layout, simulating neutrino events in such a detector, and using the reconstruction to examine the effect on physics performance such as correctly identifying electron neutrinos.

12.1 Reconstruction considerations

The principal challenge of reconstruction in CHIPS is to identify electron neutrinos against a large background of muon neutrinos and neutral current interactions. Particularly of concern are those neutral current interactions that produce a π^0 .

Neutral pions decay to a pair of photons with a branching ratio of 98.82% [45] and both of these will pair-produce, forming one e^+e^- pair per photon. Each of these produces an electromagnetic shower, leading to two electron-like rings of PMT hits. For a boosted π^0 the separation angle, θ_{AB} , between the two photons is given by:

$$(1 - \cos \theta_{AB}) = \frac{m_\pi^2}{2E_A E_B} \quad (12.1)$$

where E_A and E_B are the energies of the two photons, and m_π is the invariant mass of the π^0 . A π^0 decaying to two 500 MeV photons gives rise to a separation angle of only 15° . This yields two overlapping rings which could be misreconstructed

as a single electron, providing a background to ν_e appearance searches. The π^0 production is irreducible, so this background can only be reduced to the extent that the reconstruction can resolve the two separate rings.

Traditionally, water Cherenkov detectors have used a Hough transform method to identify rings of struck PMTs [139]. Information relating to the shape of the ring can then be used for particle identification: muon rings tend to be sharper than electron rings owing to the fact that the light is released by a single particle instead of a diffuse shower. This method is effective for reconstructing single or well-separated rings, but degrades considerably as the rings overlap.

12.1.1 Hough transformation for ring reconstruction

The Hough transform is an established method in image processing [140] which is used to identify shapes such as lines and rings that can be described with a small number of parameters. For example, a straight line can be parametrised:

$$y = mx + c \quad (12.2)$$

where m is the gradient and c is the y-intercept. However this formulation has the undesirable property that as the line becomes vertical, $m \rightarrow \pm\infty$.

An alternative parametrisation can be deduced by considering a circle of radius r centred on the origin, as shown in Figure 12.1. Any straight line whose distance of closest approach to the origin is r must lie tangent to this circle, and the point at which the line and the circle intersect can be specified by its polar angle θ . The locations x_0 and y_0 at which the straight line crosses the x- and y-axes respectively therefore satisfy:

$$\cos(\theta) = \frac{r}{x_0} \quad (12.3)$$

$$\cos(90^\circ - \theta) = \sin(\theta) = \frac{r}{y_0} \quad (12.4)$$




Figure 12.1: Illustration of an alternative parametrisation of a straight line which uses the distance, r , between the line and the origin at its point of closest approach, and the angle θ of the line joining the origin to this point. This parametrisation is commonly used for a Hough transformation because both r and θ remain finite even for vertical lines.

Converting into m and c , we find:

$$c = y_0 = \frac{r}{\sin \theta} \quad (12.5)$$

$$m = \frac{y_0}{-x_0} = -\cot \theta \quad (12.6)$$

And so, substituting into $y = mx + c$:

$$y = -x \frac{\cos \theta}{\sin \theta} + \frac{r}{\sin \theta} \quad (12.7)$$

$$\Rightarrow r = x \cos \theta + y \sin \theta \quad (12.8)$$

Hence the straight line can alternatively be parametrised as $r = x \cos \theta + y \sin \theta$ which removes the difficulty associated with vertical lines.

The Hough transform is essentially a voting procedure: the Hough space (r, θ) is divided into bins, and for each data point (e.g. a dark pixel in an image) every (r, θ) combination that would yield a straight line passing through this data point has a count added to its bin. If several of the data points lie in a straight line, they will have one (r, θ) point in common and so will produce a peak at this point in Hough space. Thus the problem of finding a complex shape is reduced to one of peak-finding. An additional benefit is that the method is reasonably robust when some information is missing, e.g. a partially obscured straight line in an image, or when a low light level means that not every PMT that lies in the Cherenkov ring registers a hit.

In the more relevant case of fitting a Cherenkov ring the method remains the same but the problem becomes three dimensional, with two parameters corresponding to the direction vector from the event vertex to the centre of the ring, and a third corresponding to the Cherenkov cone's half-angle.

First, the event vertex is estimated based on averaging the location of PMT hits and simple parametrisation of the charge and timing distributions of a conical pattern of hits. Given this vertex, a Hough transform is used to estimate the direction and opening angle of the ring more precisely. The basic Hough transform method is illustrated in Figure 12.2.

Ignoring the geometry of the detector for the time being, the task is to reconstruct a circular ring shown in magenta in Figure 12.2a, with the black circles corresponding to hit PMTs in the ring. Assuming that the diameter of the ring is already known, we can consider a single PMT (highlighted in cyan) and construct all the possible rings that would intersect this PMT (the grey circles are examples of these). The locus of centres for these circles is shown by the green ring, i.e. a Cherenkov ring of the correct radius centred anywhere on the green ring would intersect the highlighted PMT. In Figure 12.2b equivalent green rings have been constructed for two additional PMTs. These all intersect at exactly one point, and this intersection point is the centre of the true Cherenkov ring. So in this case, the Hough space could consist of the (x, y) coordinates of the ring centre, and the circle's radius, r . For each radius the algorithm would construct a green ring around each PMT, step around that ring, and add a count to the 3D histogram bin containing (r, x, y) .




Figure 12.2: Illustration of a Hough transformation method for finding rings. Top left: The green ring shows the locus of points where the grey potential Cherenkov rings could be centred in order to intersect the highlighted PMT. Top right: Equivalent circles constructed for different PMTs intersect at the centre of the true (magenta) ring. Bottom: The same method extended to identifying a cone using the direction of the cone's axis instead of the location of a circle's centre.

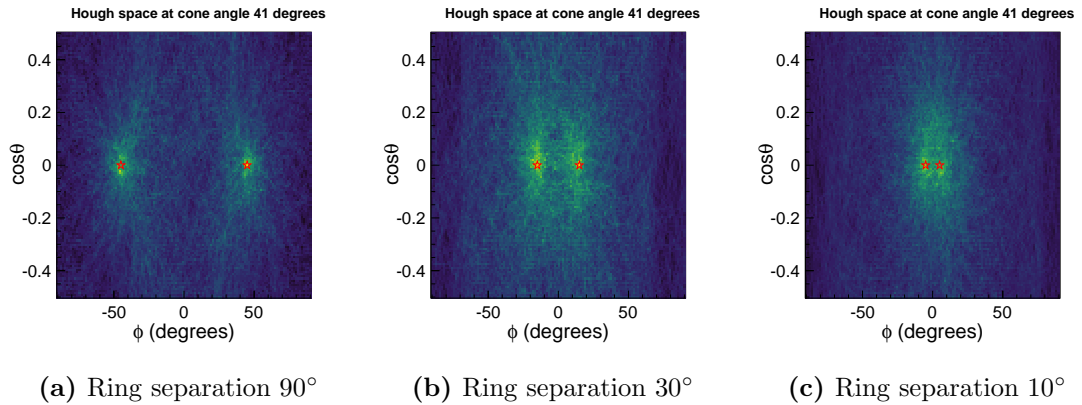


Figure 12.3: The output from a Hough transform with a cone angle of 41° for two 1 GeV electrons produced at the origin with separation angles of 90° (left), 30° (centre) and 10° (right). The hollow stars represent the directions of the two true tracks. The peaks align well with the true tracks but become harder to resolve as the separation decreases.

To reconstruct a ring in a detector where it may extend over several walls the same approach can be used, but the coordinates must be adapted. Instead of the (x, y) coordinate of the circle's centre the algorithm uses the polar angle θ and azimuthal angle ϕ that describe the direction of the cone's axis (which is the same as the direction of the particle). The diameter of the circle is replaced by the opening angle of the cone. This is illustrated in Figure 12.2c; instead of constructing circles and finding a common point at which they intersect the algorithm now constructs cones and identifies a common axis. This yields the track direction and cone angle.

Identifying a single peak in Hough space is straightforward as the coordinates of the maximum value can easily be found. However, when multiple peaks are present the viability of this method depends on their separation. Figure 12.3 shows the Hough transform in (ϕ, θ) space at the Cherenkov cone angle of 41° for three events, each featuring a pair of 1 GeV electrons sharing a common vertex and separated by 90° , 30° , and 10° . As the tracks become closer together the peaks in Hough space become increasingly difficult to resolve. This becomes more problematic for photons from π^0 decay where one may be considerably more energetic than the other, giving peaks of different heights. Masking hit PMTs in the first reconstructed ring can help with identifying further rings, but generally as the separation between the two tracks decreases so too does the efficiency with which π^0 events can be rejected.

Instead of this geometrical method, a more modern approach employs a likelihood based reconstruction which considers the expected charge and timing response of the detector's PMTs to each of the charged particle tracks. The CHIPS reconstruction uses a Hough transform to provide the seed for an algorithm of this kind, which is described in the following Section.

12.2 Reconstruction method

The reconstruction uses a series of C++ algorithms which run over the files output by the detector simulation. It comprises two separate components: one concerned with the charge recorded by each PMT, and one concerned with the timing of PMT hits. The time component was devised specifically for CHIPS, whilst the charge component is an implementation from scratch of an algorithm originally devised for the MiniBooNE experiment [141] [142]. It has been adapted for a larger detector which is cylindrical as opposed to spherical, and contains pure water rather than scintillating mineral oil.

The procedure is conceptually simple: for a hypothesised track, or combination of tracks, the number of photoelectrons recorded by each PMT is predicted, as well as the time at which the first photoelectron is detected. This prediction is then compared to the measured hit charges and times, and the likelihood that the measured signal would be produced by the proposed track combination is calculated. The track parameters are then varied until the negative logarithm of the likelihood is minimised. This identifies the most-likely track configuration. The value of the log-likelihood serves as an indicator of the quality of the fit, and can be used for particle identification by comparing, for example, the best-fit log-likelihood for muon and electron track hypotheses.

The main challenge of the method is to predict the charge and time distribution of PMT hits for arbitrary track combinations. To achieve this, a large Monte Carlo sample is used to produce lookup tables that describe the pattern of Cherenkov emission from different particle types as a function of their kinetic energy. These are combined with parametrisations of detector properties such as the water clarity and PMT efficiencies.

An advantage of this method, particularly when designing a detector, is its modularity: changes to the design (such as using a different size or type of PMT) only require changes to the corresponding part of the fitter and the remainder of the reconstruction can be left intact.

12.3 Calculating the likelihood

Let the parameters used to describe the track hypothesis be written as a single vector \mathbf{x} . Then the likelihood function to be minimised can be written as:

$$\mathcal{L} = \prod_{i,\text{unhit}} P_i(\text{unhit}|\mathbf{x}) \times \prod_{j,\text{hit}} P(Q_j, t_j|\mathbf{x}) \quad (12.9)$$

Here, i and j are indices corresponding to the unhit and hit PMTs respectively, and $P(Q_j, t_j)$ is the probability that PMT j recorded a hit with charge Q_j and time t_j . Unlike a Hough transform, the algorithm uses information from all PMTs in the detector, rather than only those that were hit.

The fit attempts to find the track configuration that maximises the likelihood, which is equivalent to maximising the logarithm of the likelihood. This simplifies the computation by converting the large product into a large sum of logarithms. Likelihoods lie between 0 and 1 by definition so the log-likelihood will be negative, meaning that minimising the negative log-likelihood is equivalent and more intuitive. By convention a factor of two is included (this makes the test statistic equal to the chi-squared in the case of Gaussian errors). Hence the function to be minimised is:

$$-2 \ln \mathcal{L} = -2 \sum_{i,\text{unhit}} \ln (P_i(\text{unhit}|\mathbf{x})) - 2 \sum_{j,\text{hit}} \ln (P(Q_j, t_j|\mathbf{x})) \quad (12.10)$$

12.3.1 Track parameters

A single electron or muon track can be described with seven independent parameters:

- Three parameters corresponding to the track's vertex x , y , and z coordinates
- The vertex time

- Two parameters describing the track's direction: the polar angle θ and azimuthal angle ϕ
- The track's kinetic energy

The fitter stores an additional (fixed) flag that identifies the particle type (electron, muon, or photon). In the case of a photon track, one further parameter is defined:

- The conversion distance, i.e. the distance the photon travels before it converts into an electron-positron pair. This causes the beginning of the electromagnetic shower to be offset from the vertex, and is fixed to zero for electron and muon tracks.

In principal, adding additional tracks to the fit adds a complete set of extra parameters. However, for π^0 hypotheses the two photons are constrained to share the same vertex position and time, and the invariant mass of the π^0 may be used to constrain the energy of the second photon. Hence the π^0 hypothesis involves eight free parameters from the first photon track, and the two direction angles and the conversion distance (and optionally the energy) from the second.

12.4 The charge component of the likelihood

The contribution to the likelihood due to the charge measured by the PMT is calculated in three steps. First, the mean predicted number of photons incident on the PMT photocathode is calculated. Secondly, this is converted into the mean number of photoelectrons liberated from the photocathode. Finally, the probability that this number of photoelectrons would lead to the final digitised charge recorded by the PMT is evaluated. This probability becomes the charge contribution to the likelihood for this PMT.

12.4.1 Predicting the number of photons at the PMT: isotropic point-like source

Before considering the more complex case of an extended track emitting a cone of light, we will first consider a simpler isotropic point-like source. If the source emits a

total number Φ of photons, then the number of photons incident on the PMT can be written:

$$\mu = \Phi T(R) \epsilon(\psi) \frac{\Omega(R)}{4\pi} \quad (12.11)$$

Here, R is the distance from the point source to the PMT, $T(R)$ is the probability for a photon to be transmitted a distance R through the water, $\epsilon(\psi)$ is the angular acceptance of the PMT as a function of the angle of the incident photon relative to the PMT normal (with $\psi = 0$ corresponding to a photon arriving head-on), and $\Omega(R)$ is the solid angle of the PMT viewed from the point source.

The functions $\epsilon(\psi)$, $T(R)$, and $\Omega(R)$ are independent of the source particle and depend only on the detector and PMT. Conversely, Φ is independent of the detector properties, and depends only on the particle's type and energy.

12.4.2 Predicting photons for an extended, non-isotropic source

With Equation 12.11, one can calculate the predicted number of photons arriving at the PMT if they were produced by an isotropic point source. However, real particle tracks differ from this in two ways: the emission of photons occurs along an extended trajectory, and it is not isotropic as emission is most likely close to the Cherenkov angle.

Ideally we could model the extended track as a sum of several point sources: the track is divided into steps, Equation 12.11 is calculated for each step, and the contribution from each step is summed to produce the final predicted number of photons:

$$\mu = \sum_{\text{step } i} \Phi_i T(R_i) \epsilon(\psi_i) \frac{\Omega(R_i)}{4\pi} \quad (12.12)$$

To evaluate this contribution we must calculate Φ_i , the total number of photons emitted at a given step. The extended and anisotropic nature of the source must be considered when calculating Φ_i , and this is achieved using functions termed 'emission profiles'.




Figure 12.4: Illustration demonstrating the calculation of the emission profile ρ to describe the number of photons a particle emits as a function of the distance along the track. Here, the (fictional) track has been divided into three steps, and ρ is the fraction of photons released in each step.

12.4.2.1 Definition of the emission profiles

The normalisation Φ_i is expressed in terms of the total number of photons Φ via two multiplicative factors:

$$\Phi_i = \Phi(E) \times \rho(E, s_i) \times g(E, s_i, \cos \theta_i) \quad (12.13)$$

Here, ρ and g are the linear and angular emission profiles respectively; ρ accounts for dividing the track up into steps, and g accounts for the photons not being emitted isotropically. To simplify the notation, the energy dependence of the emission profiles will be treated as implicit in this Section.

For the sake of illustration, consider a fictional particle that survives for 3 m and emits a thousand photons at angles close to (but spread around) the Cherenkov angle. We wish to divide the track into 1 m steps, and we observe that the particle emits 600 photons in the first metre, 300 photons in the second, and 100 photons in the third. Thus, labelling by the middle of each step, the fraction of photons emitted is given by $\rho(0.5 \text{ m}) = 0.6$, $\rho(1.5 \text{ m}) = 0.3$, and $\rho(2.5 \text{ m}) = 0.1$. This defines $\rho(s)$ and is illustrated in Figure 12.4.

Using ρ alone we can only describe a sum over isotropic sources, whereas the actual emission is peaked around the Cherenkov angle. Figure 12.5 illustrates the first step of the track: most of the 600 photons are emitted at close to the Cherenkov angle, with very few of them emitted backwards, but for an isotropic source we would expect the same number in each bin. So using ρ alone, for PMTs located close to the Cherenkov angle we would underpredict the number of photons, and for PMTs far from the Cherenkov angle we would overpredict. This is accommodated using g , which can be thought of as defining a different point source for each bin of emission angle, as illustrated in Figure 12.6.




Figure 12.5: Illustration of a section of a particle track where 600 photons are released, divided into six bins of $|\theta|$. The real emission would be peaked around the Cherenkov angle, with very few photons emitted backwards (left). Scaling by ρ alone would produce the situation shown on the right, with the same density of photons emitted into each bin.

In practice the fitter uses bins of $\cos \theta$ rather than $|\theta|$. It may also be desirable to use bins of variable width, so instead of multiplying by the number of bins we multiply by the range of possible values ($\cos \theta$ varies from -1 to +1, so this is 2) and divide by the width of the angular bin in question. The distribution of emission angles changes as the particle propagates (e.g. as the electromagnetic shower develops, or the particle slows down and the Cherenkov angle increases), and so the procedure shown in Figure 12.6 is repeated for each step along the track.

Both ρ and g depend on the type and energy of the particle. They are evaluated from large Monte Carlo samples by producing histograms of the distance and angle at which particles of a given energy emit optical photons. For a discretely-binned histogram, ρ and g are formally defined using:

$$\rho(s_i) = \frac{\# \text{ photons emitted in } i\text{th distance bin}}{\text{total number of photons emitted by particle}} \quad (12.14)$$

$$g(s_i, \cos \theta_j) = \left(\frac{\# \text{ photons emitted in } i\text{th distance bin and } j\text{th } \cos \theta \text{ bin}}{\# \text{ photons emitted in } i\text{th distance bin}} \right) \div \left(\frac{\text{width of } j\text{th } \cos \theta \text{ bin}}{\int_{-1}^{+1} d(\cos \theta)} \right) \quad (12.15)$$

Under this formalism, the sum over all bins of the ρ histogram is 1, as is the sum of $g(s, \cos \theta) \times (\text{width of } \cos \theta \text{ bin})$ when taken over all $\cos \theta$ bins at a given value of




Figure 12.6: Demonstration of the angular emission profile g for three of the six illustrated bins of $|\theta|$. The left circle shows the true angular distribution of photons, and the three circles on the right show the calculation of g for three different angle bins. As a function of the angle at which a photon would need to be released in order to hit the PMT, g determines the factor by which the right-hand image from Figure 12.5 should be scaled in order to emit the correct number of photons in the relevant direction. In the fitter, the binning scheme is in terms of $\cos \theta$, not $|\theta|$.




Figure 12.7: Sample emission profiles for electrons and muons. The linear profile $\rho(s)$ is shown in the top plot for 1 GeV electrons and muons, and the angular profile $g(s, \cos\theta)$ is shown for electrons in the bottom left plot, and muons in the bottom right. The long, sharp emission near to the Cherenkov angle is evident for the muon, compared to the shorter and more diffuse shower from the electron. The narrowing of the muon profile and its peak being slightly displaced from the start of the track are due to the normalisation scheme's requirement that each row integrates to one.

s . Note also that for a given track vertex, track direction, and PMT position, the value of $\cos\theta$ required to hit the PMT is a function of s_i only. Emission profiles for 1 GeV electrons and muons are shown in Figure 12.14

With these emission profiles in hand, the predicted number of photons at a given PMT due to a given track can be written as:

$$\mu = \Phi(E) \sum_{\text{step } i} T(R_i) \epsilon(\psi_i) \frac{\Omega(R_i)}{4\pi} \rho(s_i) g(s_i, \cos \theta(s_i)) \quad (12.16)$$

Because the vertex and direction of the track, and the position of the PMT are fixed, R_i and ψ_i depend only on the distance s_i that the particle has travelled, so Equation 12.16 can be recast as:

$$\mu = \Phi(E) \sum_{\text{step } i} T(s_i) \epsilon(s_i) \frac{\Omega(s_i)}{4\pi} \rho(s_i) g(s_i, \cos \theta(s_i)) \quad (12.17)$$

12.5 Calculating the components of the predicted charge

Each component of the charge prediction is evaluated separately, typically from large Monte Carlo samples configured to isolate the effect of interest. Salient points for these components are listed below.

12.5.1 Number of photons $\Phi(E)$

The number of photons produced by a particle is determined simply by simulating a large Monte Carlo sample at a range of energies and averaging the number of photons produced at each energy. A linear fit to the kinetic energy is used in the reconstruction, which reproduces the number of photons well at NuMI energies, and is shown in Figure 12.8.

12.5.2 Transmission function, $T(R)$

The transmission function is evaluated using Monte Carlo. An isotropic point-like source of Cherenkov photons is produced by generating large numbers of isotropic 3 MeV electrons, and the distance travelled by each photon is calculated. A histogram is produced containing, for each distance bin shown on the x-axis, the fraction of




Figure 12.8: Number of Cherenkov photons produced by Monte Carlo electrons, in 50 MeV increments of kinetic energy (blue points) and the linear fit (red line) used in the reconstruction.

photons that survive this far or further. This histogram is fitted with a function:

$$T(R) = \sum_{i=0}^n A_i e^{-R/\lambda_i} \quad (12.18)$$

Subject to the constraint that $\sum_{i=0}^n A_i = 1$.

In the reconstruction, $R(s)$ can be calculated using $R = |\mathbf{P} - \mathbf{V} + s\hat{\mathbf{d}}|$, where \mathbf{P} is the PMT position, \mathbf{V} is the track vertex, and $\hat{\mathbf{d}}$ is a unit vector in the direction of the track's momentum. The fitter models the transmission using a sum over three absorption lengths ($n = 3$), and this fit is shown in Figure 12.9. It also uses a lookup table for each of the terms in order to avoid repeatedly performing time-consuming exponential calculations.

The default model of the absorption length used in the simulation is based on measurements of pure water [143–145]. To allow for impurities in the water, the entire absorption length curve was scaled by a constant factor so that the absorption




Figure 12.9: Transmission length used in the fitter. Histogram entries (pink) are the fraction of Monte Carlo Cherenkov photons surviving at least as far as the distance on the x-axis. This is fitted with a sum of three exponentials, shown in blue.

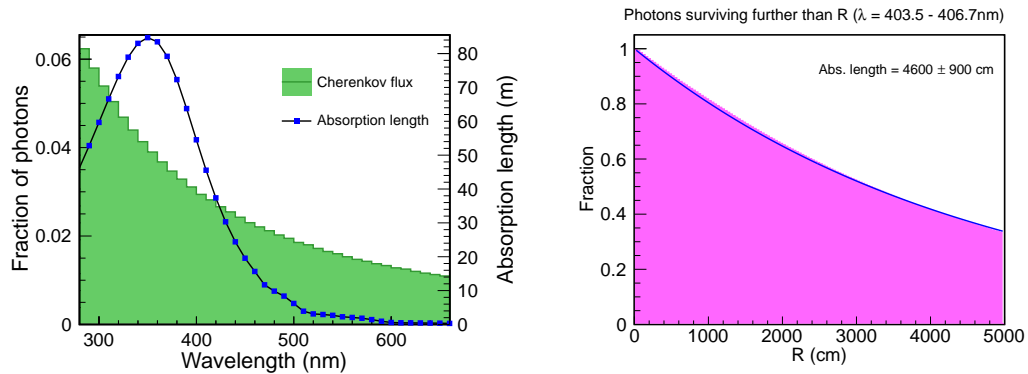


Figure 12.10: The absorption lengths as a function of wavelength used in the fitter (which have been scaled to 50 m at 405 nm) and the wavelength spectrum of the Cherenkov photons produced (left). The right plot shows a single exponential decay fitted to the transmission length considering only photons with wavelengths close to 405 nm, and is consistent with 50 m.

length of the water was 50 m for 405 nm photons. This is conservatively based on measurements taken with a test stand that filtered water from the Wentworth pit down to $0.1 \mu\text{m}$, passed it through a UV steriliser, and measured the attenuation length using a laser and PIN diode [134]. The fitted transmission lengths represent an average over the wavelength spectrum of Cherenkov light and the absorption length of the water as a function of wavelength and so should not be expected to reproduce this 50 m value.

This is illustrated in Figure 12.10 (left) which shows the wavelength spectrum of Cherenkov light produced in the simulation, and the absorption length as a function of wavelength. In Figure 12.10 (right), only the photons between 403.5 and 406.7 nm in wavelength have been fitted with a single exponential decay, resulting in an absorption length of 46 ± 9 m. The large error and the difference from 50 m reflects the fact that other processes beside absorption are present. In particular, the distance between a photon's start and end point may be shorter than its total path length if the photon has scattered, which will lower the apparent absorption length.

12.5.3 PMT Solid Angle, $\Omega(R)$

The simulation models the PMT as a sphere partially protruding through an opaque inner liner of the detector, by less than one full PMT radius (see Figure 12.11).




Figure 12.11: Diagram of PMT protruding through the opaque inner liner

Calculating the solid angle from an arbitrary vantage point is, in general, not straightforward, so the approximation that the PMT is viewed head-on is used.

In this approximation the solid angle subtended by the PMT becomes that subtended by a cone, namely:

$$\Omega = 2\pi(1 - \cos \alpha) \quad (12.19)$$

Where α is the cone half-angle. In our case, for a PMT of radius ρ , exposed height h , and whose centre is a distance R from the source, we have:

$$\Omega(R) = 2\pi \left(1 - \frac{R + h}{\sqrt{(R + h)^2 + \rho^2}} \right) \quad (12.20)$$

12.5.4 PMT efficiency

The fitter considers two contributions to the efficiency of the photomultiplier: the acceptance as a function of the incident photon angle, and the quantum efficiency of the photocathode. These are both taken directly from the properties assigned to the PMTs when they are constructed in the simulation, but are simple to replace once measurements are made using real PMTs.

The angular efficiency is calculated by interpolating a simple array of efficiencies at different angles, whilst the mean quantum efficiency is calculated by taking the wavelength spectrum of Cherenkov photons produced in Monte Carlo, weighting it by the PMT quantum efficiency at each wavelength, and dividing by the total number of Cherenkov photons.

Because the absorption and scattering of the photons are also functions of wavelength, the wavelength spectrum seen by the PMT is different depending on how far from the PMT the photons were emitted. The average quantum efficiency is therefore calculated as a function of this distance when used in the fitter, although the impact of this effect is small.

12.5.5 Evaluating the sums

The fitter must evaluate the sums in Equation 12.17 once per track, per PMT, per evaluation of the likelihood. In suitably small distance steps along the track (the reconstruction uses 5 cm steps) this quickly becomes impractically slow, especially if a large number of small PMTs is used. To counteract this the sums are manipulated so that lookup tables can be used instead of evaluating them each time, again employing the MiniBooNE method [141].

This formalism notes that we can recast the sum in Equation 12.17 as:

$$\Phi(E) \sum_i J(s_i) \rho(s_i) g(s_i, \theta(s_i)) \quad (12.21)$$

with $J = T(s) \epsilon(s) \frac{\Omega(s)}{4\pi}$. We have seen that each of the component functions of $J(s)$ varies smoothly and therefore assume that we can approximate $J(s)$ using a quadratic expansion in s :

$$J(s) \approx j_0 + j_1 s + j_2 s^2 \quad (12.22)$$

Rather than evaluating J for each step and performing a fit, we evaluate it at three distances only, chosen to cover a large fraction of the track's range. These are:

1. 10 cm
2. The distance, D , at which 75% of the track's photons have been emitted, ($\sum_i \rho(s_i) = 0.75$)
3. $D/2$

Or, if this would require using a point at which the track had already exited the detector:

1. 10 cm,
2. 80% of the distance at which the particle exits the detector
3. 40% of the distance at which the particle exits the detector

These points yield three simultaneous equations which can be solved by matrix inversion to yield the j_i coefficients.

After making this expansion the quantities to evaluate become:

$$\mu = \Phi(E) \left(\begin{aligned} & j_0 \sum_i \rho(s_i) g(s_i, \theta(s_i)) \\ & + j_1 \sum_i s_i \rho(s_i) g(s_i, \theta(s_i)) \\ & + j_2 \sum_i s_i^2 \rho(s_i) g(s_i, \theta(s_i)) \end{aligned} \right) \quad (12.23)$$

This approach serves to decouple properties of the detector and track geometry (J) from quantities derived solely from the emission profiles (the summations of ρ and g). The PMT position and track parameters now only determine j coefficients, and the path through $(s, \cos \theta)$ space over which the sum must be carried out (i.e. the mapping of s to $\cos \theta(s)$).

The summations can therefore be calculated in advance, in bins of track energy. The path through $(s, \cos \theta)$ space can be specified by R_0 , the distance from the track vertex to PMT, and $\cos \theta_0$, determined by constructing a vector from the vertex to the PMT, and another heading away from the vertex in the direction of the track, and finding the angle θ_0 between the two. The sums are stored in three dimensional tables which the fitter loads into memory and can rapidly look-up. To ensure that these integrals vary smoothly as a function of track energy, cubic splines (in the form of native ROOT TSpline3 objects) are stored. When an integral is required, the correct spline for the desired $(R_0, \cos \theta_0)$ bin is retrieved and the spline is evaluated at the desired energy.

12.5.6 Converting the charge prediction into a likelihood

The final step required to compute the contribution to the likelihood made by charge is to calculate the likelihood that a PMT whose photocathode produces μ photoelectrons registers a reading of n_γ digitised photoelectrons. A number of different models for this process are available in both the simulation and reconstruction.

The original version of the MC defined a PDF corresponding to the PMT's response to a single photoelectron. It drew repeated samples distributed according to this PDF, taking one sample per photoelectron and then summed the digitised outputs.

An improved PMT simulation was introduced, which uses a more physical model. The total gain, number of dynodes, and relative gain of each dynode compared to the first one are defined, and the digitiser simulates the amplification at each stage along the chain, sampling randomly from Poisson distributions to determine the number of photons after each stage. An additional probability to miss the cathode or the first dynode is also included. Nonlinearity is also simulated, based on the results of [129]. This reduces the digitised response when the input current is $\gtrsim 30$ p.e. ns⁻¹, although with 3" PMTs where the charge is typically ~ 1 p.e. this is a rare occurrence.

Both these methods are modelled in the fitter. The full PMT simulation is modelled using a toy Monte Carlo which, in bins of predicted mean photoelectrons, samples a Poisson distribution to determine the number of actual photoelectrons, and passes this through a replica of the PMT simulation. This process is repeated to yield 10^6 samples per bin using a scheme whose bins become increasingly coarse up to 1000 p.e. and the result is shown in Figure 12.12.

For the repeated sampling method, the same procedure is applied for undigitised charges below 10 p.e.: the 1 p.e. distribution is repeatedly sampled and this is used to construct a histogram from which the fitter can look up the digitisation probability. For predicted charges greater than 10 p.e., a fit to a Gaussian smeared with an exponential is used, and for any predicted charges above 200 p.e., a simple Gaussian is sufficient to model the response.

The fitter also features an option to calculate the likelihood using a smoothed Poisson distribution. With small PMTs and light absorption and absorption lengths




Figure 12.12: Results of a Monte Carlo model of the PMT simulation which, as a function of the predicted mean number of photoelectrons, samples a Poisson distribution to determine the actual number of photoelectrons, then passes this through the full PMT simulation to determine the probability distribution of the digitised charge. Because the bin widths vary, the probability is shown per 0.05 p.e. of digitised charge. Saturation effects can be seen in the flattening at high charges.

on the order of the size of the detector, the number of photoelectrons per PMT is generally small, and the three methods behave similarly.

12.6 The time component of the likelihood

An important concept for CHIPS is to use smaller PMTs in order to benefit from increased spatial granularity. The timing resolution of modern PMTs, particularly smaller tubes which have a physically shorter dynode chain and are less vulnerable to the effects of magnetic fields, is sufficient that considerable information could be gleaned about the track using timing information. For instance, the timing resolution for 3" PMTs used in KM3NeT is required to be < 2 ns [146], whilst a particle travelling at c would take 106 ns to traverse the diagonal of a 25 m diameter by 20 m high 10 kton detector.

The potential value of timing is evident in Figure 12.13, which shows an event display for a 2 GeV electron travelling in the beam direction from (-2 m, 0 m, 7 m). The colour of the markers indicates the time at which the PMT was hit, with each colour corresponding to a 5 ns window, and the size corresponds to the hit charge. The trajectory of the electron can clearly be traced using timing, first hitting PMTs on the top cap of the detector before depositing a ring of charge on the back wall of the cylinder.

The reconstruction aims to harness this timing information to improve track fitting. The timing method described in [141] is not used; instead a new method has been devised. Similarly to the charge, this new method aims to predict the time at which a phototube would be hit, and the expected error on this prediction, and then compare this to the recorded hit time to construct a likelihood.

12.6.1 Predicting the PMT hit time

Consider a charged particle originating at time t_0 at a vertex position specified by \mathbf{V} . It travels a distance s at speed v in a direction given by the unit vector of its momentum, $\hat{\mathbf{p}}$, and then emits a photon which strikes a PMT located at \mathbf{P} . If the water has refractive index n then we can predict the time at which the photon reaches




Figure 12.13: Event display showing a simulated 2 GeV electron (cyan line) travelling along the positive x direction from (-2 m, 0 m, 7 m). The markers correspond to hit PMTs, where the marker size is the hit charge and the marker colour is the recorded time in 5ns bins with purple corresponding to the earliest hits and red to the latest. The front edge of the detector has been cut away to show the PMTs more clearly.

the PMT:

$$t_{PMT} = t_0 + \frac{s}{v} + \left(\frac{n}{c} \times |\mathbf{P} - (\mathbf{V} + s\hat{\mathbf{p}})| \right) \quad (12.24)$$

If photons were emitted only at the Cherenkov angle, there would be at most one possible positive value of s at which a photon could be emitted and go on to hit the PMT. However, in reality there is a spread of angles at which photons can be emitted which varies as a function of s . The charge component of the likelihood accounts for this by way of the emission profiles, and the time likelihood takes a similar approach.

This leads to a spread of possible PMT hit times which should be accounted for when calculating the time likelihood, in addition to smearing due to the PMT time resolution. It should also be noted that t_{PMT} in general does not map onto a unique value of s due to the different propagation speeds of the photon and charged particle.

12.6.2 Method for predicting PMT hit times

The basic method is to start at the track vertex and move along the track direction in discrete steps. For each step, emission profiles similar to those involved in the charge prediction are used to calculate the relative number of photons emitted by the charged particle in a direction that would go on to hit the PMT. The arrival time at the PMT of photons from each step is calculated according to Equation 12.24.

The predicted mean time at which photons hit the PMT is found by taking the weighted average of the arrival times from all steps, where the weights correspond to that step's number of photons. The smearing of times due to the track being an extended source, σ_{source} , is taken to be the weighted RMS of the predicted arrival times.

12.6.3 Emission profiles for timing

These differ from the profiles used for the charge likelihood in two main ways, but are constructed using the raw photon information from the same Monte Carlo sample. The first difference is simply a matter of binning: the PMT hit time depends linearly on the distance between the PMT and the point at which the photon was emitted. This is related to the PMT and vertex positions via trigonometric functions, making it difficult to pre-calculate integrals. Instead we simply take coarser steps so the time profiles have 25 cm bins in s and perform the summation manually.

The second difference is more subtle. The charge emission profiles are normalised such that the sum over all s bins in $\rho(s)$ is 1, and the sum over all θ bins at fixed s in $g(s, \cos \theta(s))$ is also 1. This is so as to mimic a point source at each step along the track.

For the timing we require that for each step Δs , the value

$$\rho_t(s)g_t(s, \cos \theta(s))\Delta s \Delta \cos \theta(s) \quad (12.25)$$

gives the fraction of the total number of photons emitted towards the PMT.

Under this scheme, the timing histograms contain the following:

$$\rho_t(s_i) = \left(\frac{\# \text{ photons emitted in } i\text{th distance bin}}{\text{total number of photons emitted by particle}} \right) \times \frac{1}{\text{width of } i\text{th distance bin}} \quad (12.26)$$

$$g_t(s_i, \cos \theta_j) = \left(\frac{\# \text{ photons emitted in } i\text{th distance bin and } j\text{th } \cos \theta \text{ bin}}{\# \text{ photons emitted in } i\text{th distance bin}} \right) \times \frac{1}{\text{width of } j\text{th } \cos \theta \text{ bin}} \quad (12.27)$$

Further processing is applied to these profiles to ensure that the time likelihood remains smooth as a function of energy. These are discussed in Section 12.6.5 after the main method for predicting arrival times is described. The final timing emission profiles are shown in Figure ??

12.6.4 Predicting the arrival time

The quantities used to calculate the predicted time for the i th step along the track are shown in Figure 12.15. For each step, the angle θ_i and distance d_i to the PMT is calculated. For a particle travelling at speed v through a material with refractive index n , the predicted hit time from this step is simply:

$$t_{PMT,i} = t_0 + \frac{s_i}{v} + \frac{n}{c}d_i \quad (12.28)$$

To calculate the probability weight, the angle θ_{i+1} to the PMT at the beginning of the next step is calculated. The weight can then be calculated using the timing profiles:

$$w_i = \rho_t(s_i)g_t(s_i, \cos \theta_i)\Delta s_i\Delta \cos \theta_i \quad (12.29)$$

Where $\Delta s_i = s_{i+1} - s_i = 25 \text{ cm}$ and $\Delta \cos \theta_i = |\cos \theta_{i+1} - \cos \theta_i|$

The algorithm continues stepping along the track until $\rho_t(s_i) = 0$, producing an array of points such as those shown in Figure 12.16. Each point represents a single step along the track, with the x-axis showing the arrival time of photons from that step and the y-axis showing the emission profile weighting assigned to the step. Note





Figure 12.14: Timing emission profiles for electrons and muons showing the coarser binning scheme in the distance along the track than is used in the emission profiles for charge and the different normalisation scheme.




Figure 12.15: Diagram showing the calculation of the predicted hit time for a single step along the track. The time at which the PMT is hit is the sum of the terms in green, and this time is assigned a probability weight using the values of the time emission profiles for s_i and $\cos \theta(s_i)$.

that different steps along the track can result in the same arrival time because of the different propagation speeds of the particle and photon.

This can be seen in the left part of Figure 12.16: at the start of the track (point A) the angle to the PMT is less than the Cherenkov angle, so the probability is low. The entire distance to the PMT is travelled by the photon, which is slower than the charged particle track, and so the arrival time is also quite late. At B the probability is maximised suggesting that here the angle to the PMT is approximately the Cherenkov angle. The arrival time is earlier than A even though the total distance travelled by the particle and photon is larger because the particle travels at $\sim c$ and the photon travels at c/n . As the particle continues the arrival time passes through a minimum. At C, the particle has moved past the PMT and this extra distance means the arrival time is late. The photons must be emitted at very wide angles, and so the probability is also very low.

The predicted mean photon arrival time for the PMT is the weighted average of this array:

$$\bar{t}_{PMT} = \left(\sum_i t_i w_i \right) / \left(\sum_i w_i \right) \quad (12.30)$$

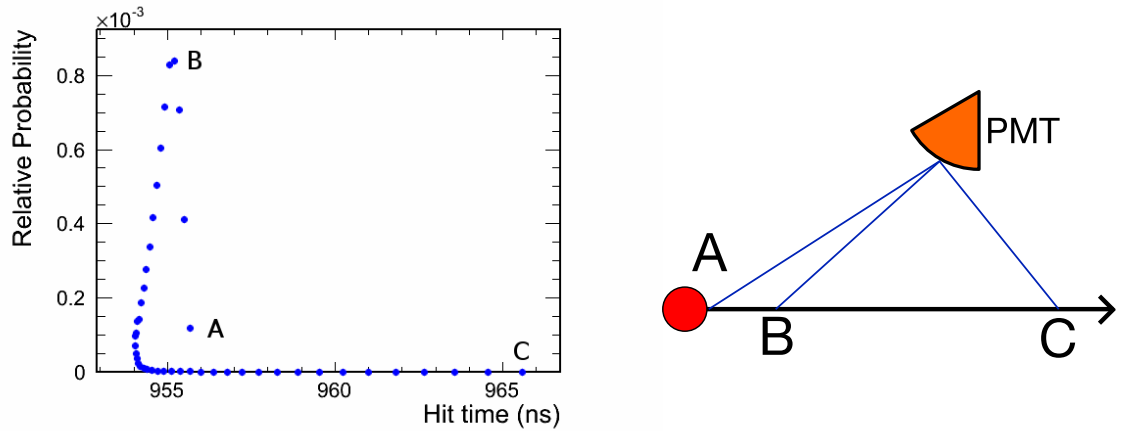


Figure 12.16: Predicted arrival times at a single PMT from each step along a track (left). Each point represents a step, with the x-axis showing the arrival time from that step and the y-axis showing the probability weighting. An arrangement of the track and PMT which could lead to such a pattern is shown on the right, with the labels A, B and C showing the beginning of the track, emission at close to the Cherenkov angle, and the end of the track.

Similarly, the predicted error is the weighted RMS of the array:

$$\sigma_{source} = \left(\sum_i (t_i - \bar{t}_{PMT})^2 w_i \right) / \left(\sum_i w_i \right) \quad (12.31)$$

Occasionally PMTs (typically ones outside of the main Cherenkov ring) can only be hit from one filled bin of the emission profile, and in this case the RMS is zero. To enable a likelihood to be calculated later, a constant additional width of 0.1 ns is added to σ_{source} for all PMTs, which allows them to smoothly ‘turn on’ in the calculation.

Typical predicted times have a weighted RMS ~ 1 ns for PMTs located in the main Cherenkov ring, where the weights are dominated by a small number of points with $\theta_i \approx \theta_{Ch}$, and an RMS of ~ 5 ns for PMTs away from the ring, which require photons to be emitted at wide angles and are similarly unlikely to be hit for a wide range of s .

12.6.5 Smoothly parametrising the arrival times

The use of the emission profiles directly in order to calculate the expected arrival time and RMS can give rise to problematic discontinuities in the likelihood surface as a function of energy. This is due to the binned nature of the profiles, and the fact that they are only evaluated at discrete energies.

The effect is particularly significant in the regions of $(\cos \theta, s)$ space where emission is unlikely. Here, the number of photons emitted into each bin of the profile is more affected by statistical fluctuations in the MC sample used to produce the profiles than by a small change in the energy of the charged particle from one sample to the next. In these cases, the predicted arrival time (or the simple question of whether a PMT can be hit at all) can change suddenly when moving from the emission profile at one energy to the next, as unoccupied bins may become occupied and vice-versa.

This effect is counteracted in three main ways. Firstly, the emission profiles are smoothed by applying a smoothing algorithm to the raw histograms containing the number of photons as a function of distance and angle.

Secondly, any $(s, \cos \theta)$ bins in the raw histograms with very few photons are removed. Under Poisson statistics there may be a large number of bins which would be expected to contain a fraction of a photon, given the size of the MC sample. Most of these will contain zero photons, and a few will contain one or more, but which bins they are will fluctuate randomly from sample to sample. For bins with a larger number of photons, these fluctuations are smoother and differences from one MC sample energy to the next are more likely to be a consequence of the change in energy. Fluctuations between zero (the PMT cannot be hit from this step) and almost-zero (the PMT could be hit from this step, but very improbably) also have a more pronounced effect on the likelihood than fluctuations around larger numbers. For this reason, bins containing fewer than 10 photons are removed from the histogram.

When simulating muons, occasionally one of them will undergo a hard scatter as it passes through the water, resulting in a large change of direction. It then continues to travel and emit photons at $\theta \approx \theta_{Ch}$ in this new direction. Relative to the initial direction of the muon, however, it appears that the photons are being emitted at very wide angles. This converts an effect that occurs for one in every several thousand

muons into an effect that occurs for one in every several thousand photons. These events are removed by requiring, for muons, that all $(s, \cos \theta)$ bins contain at least 0.1% of all photons emitted at that distance along the track.

All empty bins for distances at which photons are still emitted, including bins which have failed the previous two checks, are filled with 10 photons (for comparison, 2500 events simulated for a 1 GeV muon yield around 10^9 photons in total). These histograms are then normalised as described previously to produce the timing emission profiles.

Finally, during the fitting process the predicted mean and RMS hit times, and their probability weights, are also smoothly interpolated between emission profiles as a function of energy. A cubic spline method is used rather than simple linear interpolation so that both the predictions and their derivatives vary smoothly as a function of energy. Specifically, a Catmull-Rom spline method [147] is chosen:




Figure 12.17: Example of Catmull-Rom spline. We have pre-calculated $f(x)$ at four equidistant values of x , and wish to estimate $f(x)$ at a value of x lying between the central two (red) points by interpolating using a cubic spline. The black line is the cubic that performs the interpolation: its gradient at x_1 is equal to that of the (pink dashed) straight line joining the first and third points, and at x_2 is that of the (green dashed) line joining the second and fourth.

Given a set of four points (x_i, y_i) where $0 \leq i \leq 3$ and the x_i are evenly-spaced, we wish to estimate $y(x)$ at a point where $x_1 < x < x_2$. A Catmull-Rom spline constructs a cubic polynomial that passes through the middle two points, subject to

the requirement that the gradient at x_1 is that of the straight line joining (x_0, y_0) with (x_2, y_2) , and the gradient at x_2 is that of the straight line joining (x_1, y_1) with (x_3, y_3) , as illustrated in Figure 12.17. Transforming coordinates such that $x_1 \rightarrow 0$ and $x_2 \rightarrow 1$ greatly simplifies the matrix inversion that yields the coefficients of the cubic, and so the interpolated point (x, y) can be quickly evaluated.

In this way a spline can be constructed which is continuous in E and $\frac{d}{dE}$ across the full range of emission profile energies, but which only requires the summation of steps through the profiles to be performed for the four energies nearest to that of the track in question. This yields our final predicted arrival time, width, and weighting.

12.6.6 Accounting for hits with multiple photoelectrons

In the previous section we arrived at a prediction of the mean time at which a photon hits the PMT, but this does not predict the time that the PMT will record. Instead, a hit is typically registered as soon the output voltage from the tube exceeds some threshold, and so the time that is required is the expected time at which the *first photoelectron* is produced. Converting from photons to photoelectrons merely requires the application of the quantum efficiency, which scales each step's weight by a constant amount and has no effect on the mean or RMS arrival time.

Converting from the mean to the first photoelectron time is more complicated. Consider a PMT that detects two photoelectrons. On average, one expects that one will hit before the mean time, and one will hit after it. The more photoelectrons there are, the earlier the expected time of the first one becomes.

If $P(t)$ is the probability density function (PDF) for the time of each photoelectron, the probability that $t \leq t'$ is:

$$P(t \leq t') = \int_{-\infty}^{t'} P(t) dt = C(t) \quad (12.32)$$

where $C(t)$ is the cumulative distribution function (CDF) for t . If we draw a sample of n random arrival times from $P(t)$ the probability that all of them are greater than

t' is:

$$P(\text{all } t_i > t') = (1 - C(t))^n \quad (12.33)$$

Hence the probability that at least one of the $t_i \leq t'$ is:

$$P(t_{min} \leq t') = 1 - (1 - C(t))^n \quad (12.34)$$

This is nothing more than the cumulative distribution function of the minimum time, so the probability density function for the minimum time can be found by differentiating:

$$P(t_{min}) = \frac{d}{dt} (1 - (1 - C(t))^n) \quad (12.35)$$

In the algorithm, we make the approximation that the probability distribution for the individual photoelectron times is Gaussian, with mean and width given by the weighted mean and RMS previously calculated. Given this assumption:

$$P(t) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(t-\mu)^2}{2\sigma^2}\right) \quad (12.36)$$

$$C(t) = \frac{1}{2} \left(1 + \operatorname{erf}\left(\frac{t-\mu}{\sqrt{2}\sigma}\right)\right) \quad (12.37)$$

$$P(t_{min}) = \frac{n}{2^n \sigma \sqrt{2}} \left(1 - \operatorname{erf}\left(\frac{t-\mu}{\sqrt{2}\sigma}\right)\right)^{n-1} \exp\left(-\frac{(t-\mu)^2}{2\sigma^2}\right) \quad (12.38)$$

Here, μ and σ are the weighted mean and RMS of the array of arrival times calculated from the emission profiles. The number of photoelectrons, n , is taken from the digitised output of the PMT. The predicted distribution of first arrival times using this formula is shown in Figure 12.18, and is compared to a distribution produced using a toy Monte Carlo to repeatedly sample three arrival times from a Gaussian distribution and record the earliest one.




Figure 12.18: The probability density function describing the arrival time of the first of three photoelectrons (red line) at the PMT and its integral (purple line). The time of each photoelectron is distributed according to a Gaussian, with mean 100 ns and width 1.0 ns. The blue filled region shows the distribution of times obtained using a toy Monte Carlo to repeatedly sample three times from a Gaussian and record the earliest one.

12.6.7 Propagation time of particles

The speed of a relativistic particle is related to its kinetic energy, E_k according to:

$$\beta = \frac{v}{c} = \sqrt{1 - \frac{1}{(1 + E_k/m)^2}} \quad (12.39)$$

This means that a muon whose momentum is > 651 MeV, and essentially all electrons in the detector are travelling at $v > 0.99c$. For the fitter it is necessary to define an effective propagation speed because the emission profiles are used relative to the direction in which the particle was travelling at the start of the track. Multiple scattering and showering effects will lower the speed at which the emission of photons propagates in this direction.

Monte Carlo is used to determine the propagation speed that is used in the fitter: tracks of a given particle type are generated travelling vertically upwards from the bottom of a tall, narrow (10 m diameter) detector with a 40% coverage of 3" PMTs. The time at which PMTs on the barrel are hit should therefore vary linearly as a function of their height, with the hit time t_{hit} as a function of height z given by:

$$t_{hit} = vz + \text{const.} \quad (12.40)$$

Fitting this distribution yielded an effective muon speed of $0.94c$ and an electron shower propagation speed of $0.81c$.

12.6.8 Calculating the time contribution to the likelihood

The contribution to the total likelihood due to timing is calculated by testing the compatibility of the hit time measured by the PMT with the predicted first arrival time. The distribution associated with the PMT is taken to be a normalised Gaussian distribution centred on the measured time with a width given by the time resolution of the PMT. In simulation, the resolution in nanoseconds is described using an approximate parametrisation of the resolutions in [138]:

$$\sigma_{PMT} = 0.33 + \sqrt{a/Q} \quad (12.41)$$




Figure 12.19: The time likelihood is calculated by finding the area of overlap between the predicted PDF for the arrival time of the first photon, and a Gaussian centred on the time measured by the PMT whose width is the PMT time resolution.

where Q is recorded charge and a is a tunable constant that depends on the model of PMT. For the PMTs simulated, a is set to 2, giving a resolution of 1.74ns at 1 p.e.

The likelihood contribution is then the area of overlap between this Gaussian, f_{PMT} , and the predicted PDF for the first arrival time, $g_{pred} = P(t_{min})$ from Equation 12.38, as illustrated in Figure 12.19. This area is equivalent to

$$\int_{-\infty}^{\infty} \min(f_{PMT}(t), g_{pred}(t)) dt \quad (12.42)$$

and is simple to evaluate numerically. However, this requires many hundreds of calls to the two functions, and becomes impractical with 3^n PMTs when several thousand can be hit. Instead, the points at which the two functions overlap are determined numerically and then the exact forms of the integrals of $f_{PMT}(t)$ and $g_{pred}(t)$ are used to calculate the overlap: $f_{PMT}(t)$ is simply a Gaussian and can be integrated using the error function, while $g_{pred}(t)$ was determined by differentiating $1 - (1 - C(t))^n$, so its integral can easily be evaluated using Equation 12.37.

12.7 Minimising the log-likelihood

Identifying the combination of track parameters that maximises the likelihood for a single track requires performing a minimisation over a seven dimensional space. This is made more challenging by the presence of false minima and the discretisation imposed by the lattice of PMT positions and the binning of the emission profile energies and the lookup table coordinates. The smoothing procedures for the charge and time components go some way towards mitigating this but the minimiser is still unable to use gradient-based minimisation methods, so instead the Simplex method of Nelder and Mead is used [148].

In addition to the large phase space there are also nontrivial correlations between several variables that can give rise to local minima, in particular between the vertex time, energy, and position in the direction parallel to the beam. The minimisation is performed in several stages to accommodate this, with certain parameters kept fixed at different stages. First, PMT hits are clustered into ‘slices’ based on their hit times and location in the detector. A vertexing algorithm is applied to all hit PMTs in each slice to determine a seed vertex, and this is used to perform a Hough transformation for that slice. The transformation with the largest Hough peak is used to seed the position and direction of the main fit.

The first two stages of the minimisation routine provide the seed track with a time and energy:

1. The position and direction are fixed, the energy is set to 1 GeV, and the vertex time is fitted using the time component only.
2. The vertex, time and direction are fixed, and the energy is fitted using charge and time.

The next stage addresses the correlated region of phase space. A special minimisation routine is used which has three parameters: a displacement Δx of the vertex along the seeded track direction, a shift Δt of the vertex time, and a shift ΔE in the energy. When the vertex is moved the time is adjusted by $\Delta x/v$ and Δt is applied in addition to this shift. The aim is for the mimiser to explore this region of phase space while taking proper account of the correlations.

3. Adjust the energy, vertex and time along the track direction, then re-fit the track direction. This is performed twice.

This refined seed is then used for the full fit:

4. Perform a completely free fit, using charge and time.

The correlation problems affect the time component more strongly than the charge, so a final set of refinements is performed:

5. The energy and vertex position are adjusted along the track direction, using charge only.
6. The energy is re-fitted using charge and time, with all other components fixed.
7. The time is re-fitted, with all other components fixed.

The output from the minimisation routine is the best-fit track, and the charge and time likelihood components for this track.

12.8 Extending to multiple tracks

Fits to a multiple-track hypothesis are desirable for reclaiming beam events that coincide with a cosmic interaction, and for when interactions produce multiple Cherenkov rings (in particular, NC π^0 production). In these cases the charge and time predictions are evaluated separately for each track, and combined in the following way:

The total predicted number of photoelectrons at the PMT μ is the sum of the predicted charges from all the tracks

$$\mu = \sum_i \mu_i \quad (12.43)$$

where μ_i is the predicted charge from the i th track.

For the overall timing contribution, the time likelihood \mathcal{L}_i^t from each track is calculated, and the overall contribution to the likelihood is taken to be the sum of

these components over all the tracks, weighted by their predicted charge:

$$\mathcal{L}^t = \frac{\sum_i \mu_i \mathcal{L}_i^t}{\sum_i \mu_i} \quad (12.44)$$

Seeding multiple-track events is also challenging. For candidate events with a simultaneous beam and cosmic interaction, the locations of hit veto PMTs are used to seed the cosmic direction. PMTs within the cosmic ring are then masked before seeding the other track. For other multi-track events, a clustering algorithm is used which attempts to slice the detector into well-separated hit regions which can be treated separately. A more comprehensive seeding method is employed for candidate π^0 events.

12.8.1 Seeding π^0 events

A more involved seeding procedure is used for the π^0 fitter. Two common failure modes both involve fitting the two rings as one single large ring. In the first, the two photon tracks overlap completely and both contribute predicted hits to the same ring, and in the second the large ring is entirely due to one track and the second track energy is negligible. The seeding procedure aims to avoid these cases.

A schematic of a situation in which the second case can arise is shown in Figure 12.20: the purple ring overlaps the PMT hits quite well, while the green ring covers some of the PMT hits but performs poorly on the left-hand side of the plot. In order to find the correct configuration, the minimiser must adjust the two tracks in the right directions simultaneously. However, the likelihood might improve if the energy of the purple ring were increased (allowing more of the overlapping PMTs to be hit) and the energy of the green track decreased (to improve the disagreement at the left of the Figure). This can result in the minimiser becoming trapped in a false minimum, increasing the energy of one track until the entire pattern is fitted with a single ring, and decreasing the energy of the second until it hits no PMTs and makes no contribution to the likelihood. This eventually results in a single track fit, with the second track oriented randomly in whatever direction it happened to be pointing when the best single track fit was obtained.




Figure 12.20: Example of a situation in which the two track minimiser find a false minimum where one track is responsible for the entire double ring and the second track energy is driven to zero. The PMT hits from the two true tracks are the fuzzy red and blue rings, and the current tracks being tried by the minimiser are the solid purple and green ones. The purple track accounts quite well for most of the charge, while the green track contributes to the right edge of the red ring, but is a poor fit to the left edge.

The first stage of the seeding procedure is to fit the event using a single electron track. This produces a likelihood against which multi-track seeds can be compared, and a track which can be perturbed to generate such seeds.

A Hough transform method is used next: first, the detector hits are separated into a number of slices by clustering together PMT hits based on their separation in space and time. Each slice is seeded individually as an electron, a Hough transform is applied, and multiple peaks are searched-for using ROOT's TSpectrum libraries. The candidate rings are sorted by the number of votes assigned to their Hough transform, and the two leading rings are picked. The π^0 vertex is constructed by propagating the two tracks backwards and finding their point of closest approach. The two photon energies are constructed so that their ratio is that of the Hough peaks, and the invariant mass of the two photons is that of a π^0 . This vertex, and the two track directions and energies define the π^0 seed. The likelihood for this combination is calculated and compared to the single electron fit.

If the two-track Hough seed has a better log-likelihood than the single electron track, it is used as the seed and no further processing takes place. Otherwise the single electron best-fit is adjusted to explore the parameter space that two nearly-

overlapping photons could occupy. This procedure is similar to the method described for MiniBooNE in [141].

In brief, the PMT hit positions are projected into a plane that contains the single electron vertex, and is perpendicular to its direction. Their covariance ellipse in this plane is found, and the major and minor axes determined. The single electron track is then rotated about the major axis by $0, \pm\frac{\pi}{20}, \pm 0.45, \pm\frac{\pi}{5}$ radians, and by $\pm\frac{\pi}{20}$ about the minor axis. This defines the first photon's direction, and the second photon's is found by performing a grid search. Because the Hough transform method is suited to larger separation angles, the grid search is performed between $\pm\frac{\pi}{8}$ in θ and ϕ . The second track energy is chosen so that the two tracks have the invariant mass of a π^0 .

The two overall configurations producing the best likelihood when the smaller track energy is greater than or less than the half the larger track energy are retained and passed to the fitter. Both photons are currently assigned a conversion length of 50 cm, although the procedure can be altered to use different combinations, returning two seeds for each one.

Chapter 13

CHIPS Reconstruction Results

The reconstruction framework was used to perform a study comparing detector designs. The event generator Genie [149] was used to generate a large sample of events using the NuMI beam spectrum at the CHIPS baseline, 715 km downstream and 7 mrad off-axis. Separate samples of 5000 pure ν_e CCQE, ν_μ CCQE, and ν_μ NC events were generated, in addition to samples of combined ν_e and combined ν_μ events of all types.

These events were passed through two different detector geometry simulations. Each featured 10 kton detector in the form of a 20-sided prism, 25 m in diameter and 20 m tall with a uniform coverage of PMTs on the walls and endcaps. One geometry had a 10% coverage of 10" diameter PMTs, and the second had a 10% coverage of 3" PMTs. The reconstruction was run twice over each data set, once with a muon hypothesis and once with an electron hypothesis. The performance of the various detector designs can be assessed by comparing the difference between the best-fit reconstructed tracks and the Monte Carlo truth.

13.1 Comparison of 3" and 10" tubes

Figure 13.1 shows a comparison between the reconstructed-minus-true vertex coordinates for the geometries with a 10% coverage of 3" and 10" tubes with ν_e CCQE events. The beam direction is along the positive x-axis, which gives rise to the poorer resolution in the x-direction: moving the vertex in $\pm x$ causes the projection of the Cherenkov ring onto the wall to expand or contract, which mainly affects the PMTs

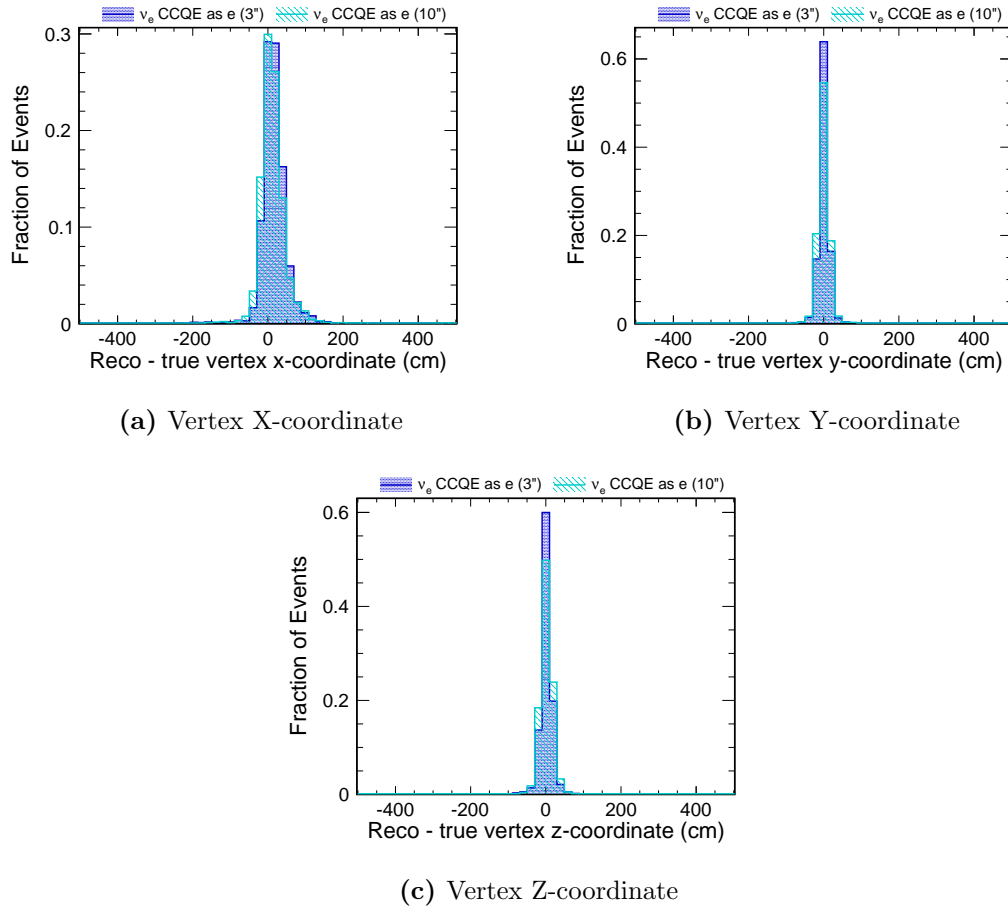


Figure 13.1: (Reco - true) distributions for the coordinates of the vertex when ν_e CCQE interactions are fitted with single electrons, for 10'' (light blue, striped) and 3'' (dark blue, dotted) PMTs.

on the outer and inner edges of the ring. This can often be accommodated by slightly adjusting the track energy (and hence its length, or the thickness of the ring) or relying on emission at slightly wider or narrower angles. For these reasons, the effect on the likelihood is smaller compared to a translation parallel to the y-axis or z-axis which moves the entire ring.

Equivalent plots for ν_μ CCQE events are shown in Figure 13.2 with a similar pattern of worsened resolution in the direction parallel to the beam. Both plots show only contained events: it is assumed that the full detector would have a cosmic veto region consisting of an outer volume instrumented more sparsely with PMTs that would efficiently tag escaping tracks, so containment is determined using MC truth

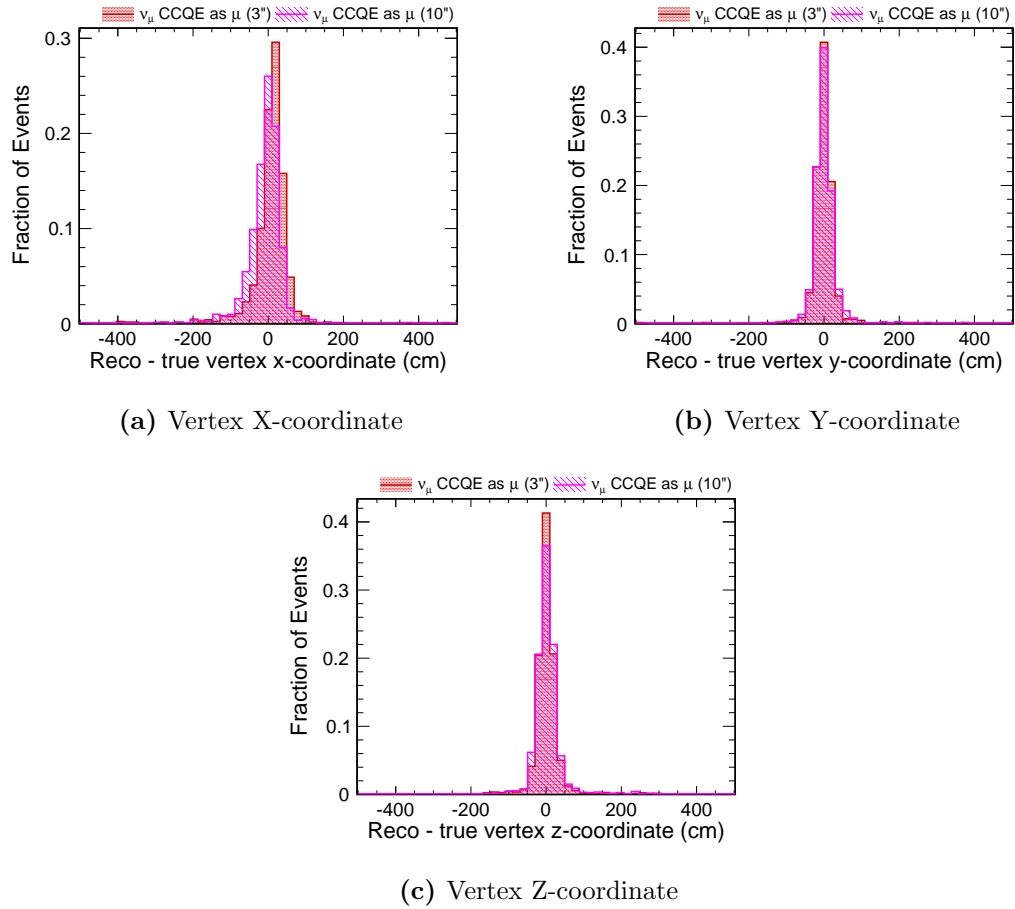


Figure 13.2: (Reco - true) distributions for the coordinates of the vertex when ν_μ CCQE interactions are fitted with single muons, for 10'' (magenta, striped) and 3'' (dark red, dotted) PMTs.

to exclude any escaping tracks. The vertex resolution with the two different PMT types is comparable but slightly better with the smaller tubes.

Figure 13.3 compares the direction, energy and vertex time resolution performance for ν_e CCQE events reconstructed as electrons, and Figure 13.4 shows the same comparisons for ν_μ CCQE events. Again the performance between the two types of tube is similar. The long tail on the muon energy distribution is associated with muon tracks for which the reconstruction fails and returns an energy approximately equal to that of the highest energy emission profile, which can easily be removed. A slight bias is evident in the timing and energy plots, and investigations into this are ongoing.

In general the minimum of the likelihood surface tends to be shallow, with a three-dimensional ‘valley’ of points with similar likelihoods. Moving along this valley consists of moving the reconstructed vertex closer to the wall, the vertex time slightly later, and decreasing the track energy, or vice versa. These changes largely compensate for one another: moving the vertex away from the wall means the vertex time has to be earlier in order to hit the PMTs at the same time, and the energy has to increase to offset the slightly higher absorption. Consequently the likelihood surface is almost flat along this valley.

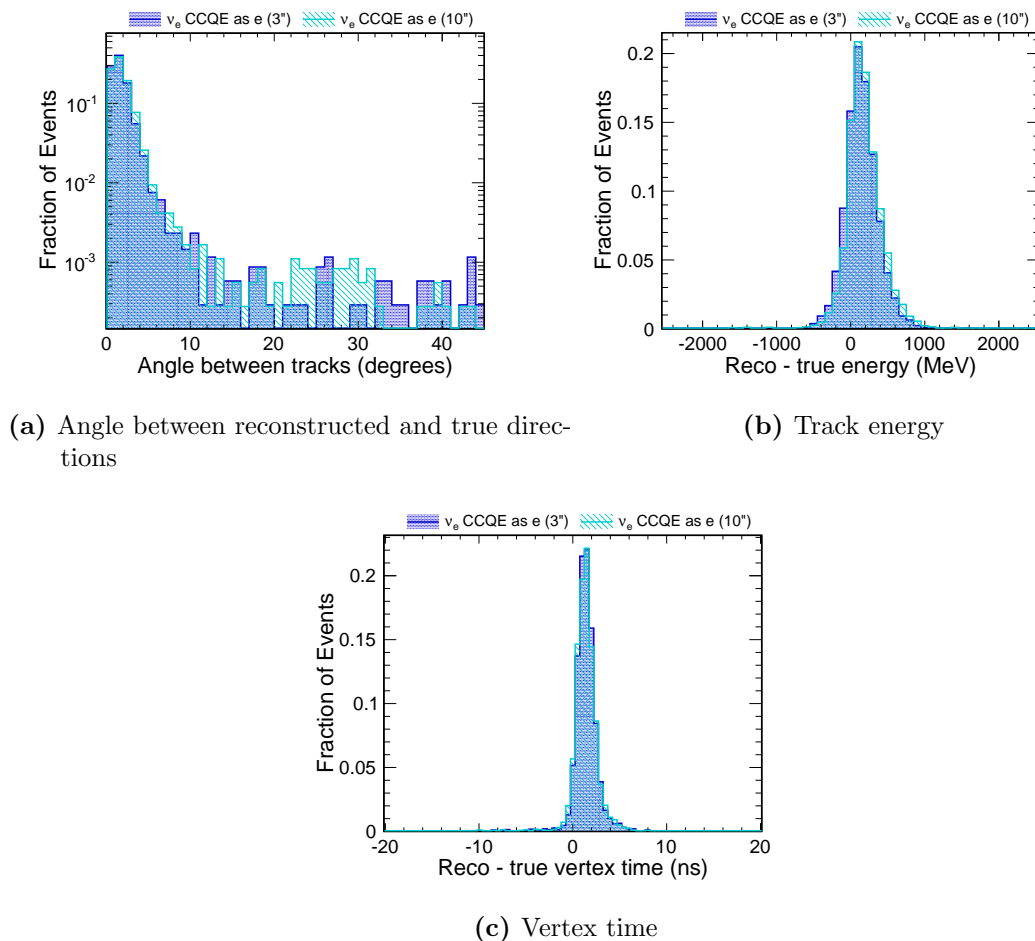


Figure 13.3: Distributions of the angle between the reconstructed and true track, the vertex time resolution, and the track energy resolution when ν_e CCQE events are fitted with an electron track hypothesis, for layouts with 10'' (light blue, striped) and 3'' (dark blue, dotted) PMTs.




Figure 13.4: Distributions of the angle between the reconstructed and true track, the vertex time resolution, and the track energy resolution when ν_μ CCQE events are fitted with a muon track hypothesis, for layouts with 10" (magenta, striped) and 3" (dark red, dotted) PMTs.

	ν_μ CCQE events		ν_e CCQE events	
	3" PMTs	10" PMTs	3" PMTs	10" PMTs
Position (cm)	44	47	35	35
Angle ($^\circ$)	2.7	2.6	1.9	2.1
Time (ns)	1.14	1.35	0.84	0.9
Energy (MeV)	112	110	210	208

Table 13.1: Comparison of the reconstruction resolutions for a 10kton detector with a 10% coverage of 3" or 10" PMTs.

Table 13.1 shows a comparison of the reconstruction performance with the two different geometries. The overall distance from the reconstructed to true vertex is calculated by taking the magnitude of the vector joining these two vertices. For the distance and angle between tracks, which are always positive by construction, the resolution is defined as the width of the region containing 68% of events. For the energy and vertex time, the resolution is defined as the width, σ , of a Gaussian fit to the appropriate (reco - true) variable.

For the energy and time components, the uncertainty on the resolution can be estimated using the uncertainty on σ reported by the Gaussian fit. For the energies this is at the 2-3 MeV level, and for the times it is between 0.01 and 0.03 ns. This information is not available for the distance and angle resolutions due to the way in which they are calculated, but typical uncertainties on σ are 0.3 cm when the individual vertex components are fit with Gaussians, and 0.1° when the polar and azimuthal angle components are fitted.

In general, the two geometries perform similarly, with the 3" tubes offering a slight improvement overall compared to the 10" ones. This demonstrates that the use of 3" tubes is viable.

13.2 Comparison of 10% and 6% coverage

Given that 3" tubes can reasonably be used, a second study was performed to assess the effect of decreasing the coverage from 10% to 6%, as removing 40% of PMTs would result in a considerable cost saving. The same set of neutrino interactions as used in Section 13.1 were also simulated using a geometry with a 6% coverage of 3" tubes. Figures 13.5 and 13.6 show comparisons of the distance from reconstructed to true vertex, the angle between the reconstructed and true tracks, and the (reco - true) energy and vertex time plots for the sample of ν_e CCQE and ν_μ CCQE events, respectively.

As expected, decreasing the coverage degrades the performance of the reconstruction, but the extent to which the results worsen is small: a 3 cm increase in the vertex resolution and a 0.2° worsening of the angular resolution for ν_e events, with a very small effect on the energy resolution. This performance indicates that it is possible

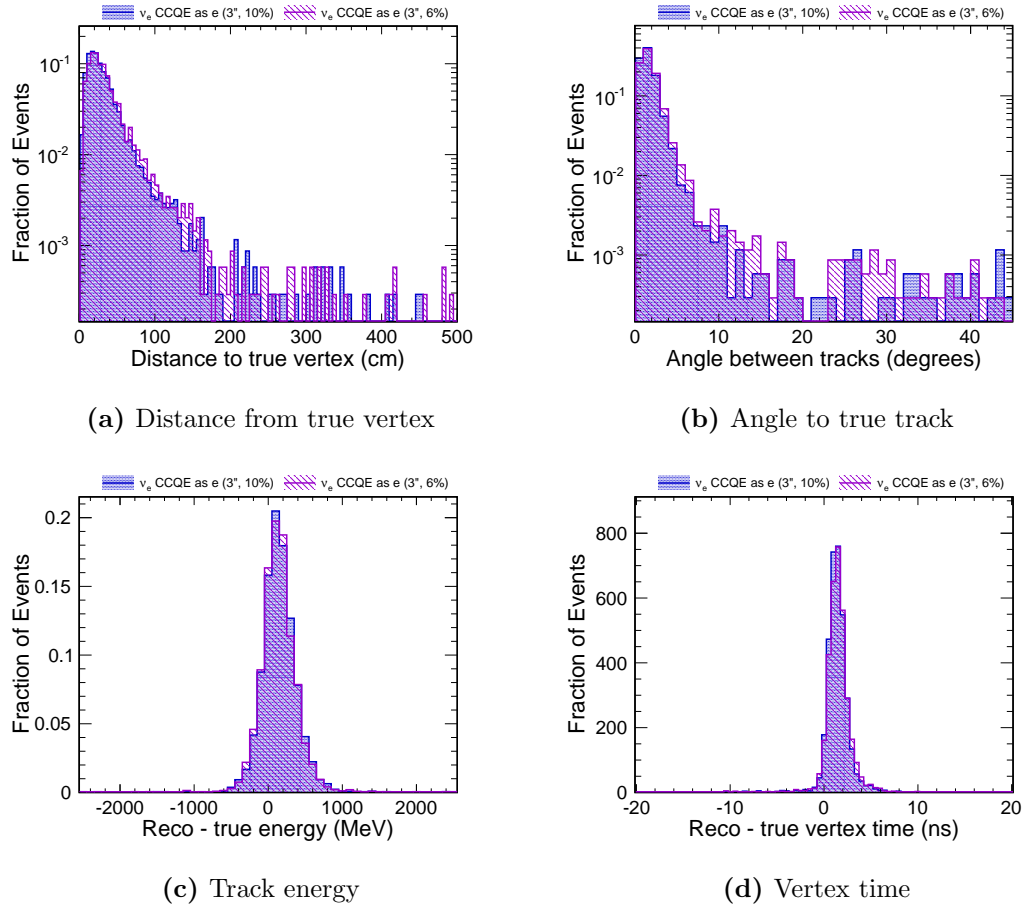


Figure 13.5: Distributions of the distance and angle between the reconstructed and true track, the track energy resolution, and the vertex time resolution when ν_e CCQE events are fitted with an electron track hypothesis, for layouts of 3" PMTs with a 6% (purple, striped) and 10% (dark blue, dotted) PMTs

	ν_μ CCQE events		ν_e CCQE events	
	10% coverage	6% coverage	10% coverage	6% coverage
Position (cm)	44	51	35	38
Angle ($^\circ$)	2.7	3.0	1.9	2.1
Time (ns)	1.14	1.28	0.84	0.89
Energy (MeV)	112	113	210	211

Table 13.2: Comparison of the reconstruction resolutions for a 10kton detector with a 10% or 6% coverage of 3" PMTs

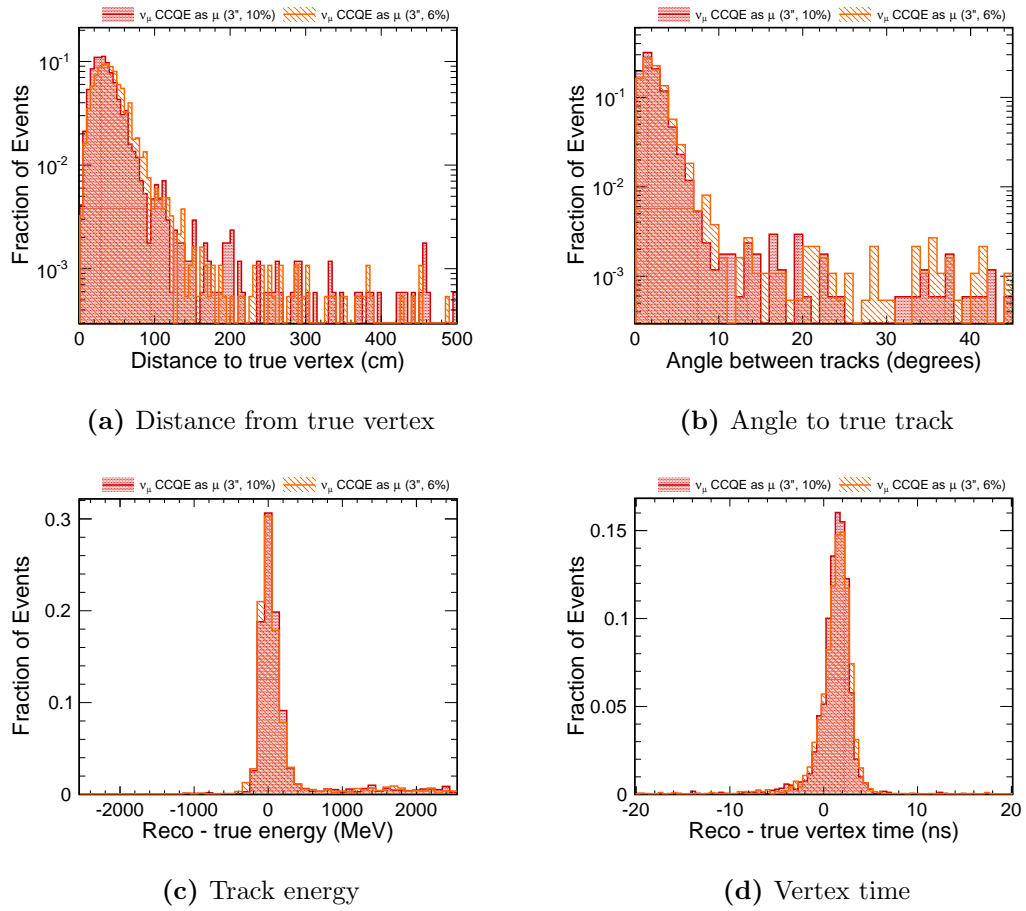


Figure 13.6: Distributions of the distance and angle between the reconstructed and true track, the track energy resolution, and the vertex time resolution when ν_μ CCQE events are fitted with a muon track hypothesis, for layouts of 3" PMTs with a 6% (orange, striped) and 10% (dark red, dotted) PMTs

to reduce the coverage to 6% without compromising the experiment. The simulation and reconstruction are capable of operating with different PMT layouts in different detector regions, so further studies could assess whether some of the decrease in performance could be offset by placing a higher coverage of PMTs in the downstream part of the detector while instrumenting the upstream walls more sparsely, a feature which exists already in the simulation and reconstructoin software.

13.3 Particle identification: Simple Cuts

A major aim of CHIPS is to measure the rate of ν_e appearance in a ν_μ beam in order to help constrain δ_{CP} . For this reason, a method of identifying ν_e events against a large background of ν_μ and neutral current interactions is required. A simple method would rely only on the comparisons between the time and charge components of the best-fit likelihood for fits using muon and electron hypotheses. For the following plots, events are required to pass a preselection cut which requires at least 50 PMTs were hit. The efficiencies quoted are the combined efficiency of this preselection and any further cuts unless otherwise stated.

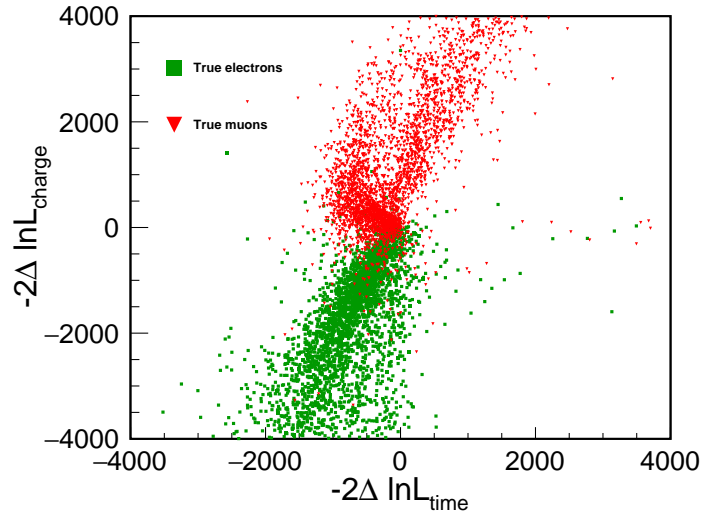


Figure 13.7: Comparison of the charge and time components of $-2\Delta \ln \mathcal{L}$ between electron and muon hypotheses, for ν_e CCQE events (green squares) and ν_μ CCQE events (red triangles). In each case $-2\Delta \ln \mathcal{L} = (-2 \ln \mathcal{L}_e) - (-2 \ln \mathcal{L}_\mu)$, so that more negative values correspond to a more electron-like result.

Figure 13.7 shows such a comparison for the samples of ν_e and ν_μ CCQE events using a 6% coverage of 3" PMTs. While these two variables alone show a reasonable separation between the electron and muon fits, there is a sizeable region of overlap near to (0, 0), and it is clear that the time component tends to favour the electron hypothesis, even when a true muon is present. This is because the time component considers only hit PMTs, and the greater variety of emission angles from an electron shower gives greater flexibility in the time at which a PMT can be hit. Even if the wrong charge is predicted, the time will be approximately correct and so most of

the penalty is absorbed by the charge component of the fit. The effect of placing different cut values in the time and charge likelihood components is shown in Figure 13.8.

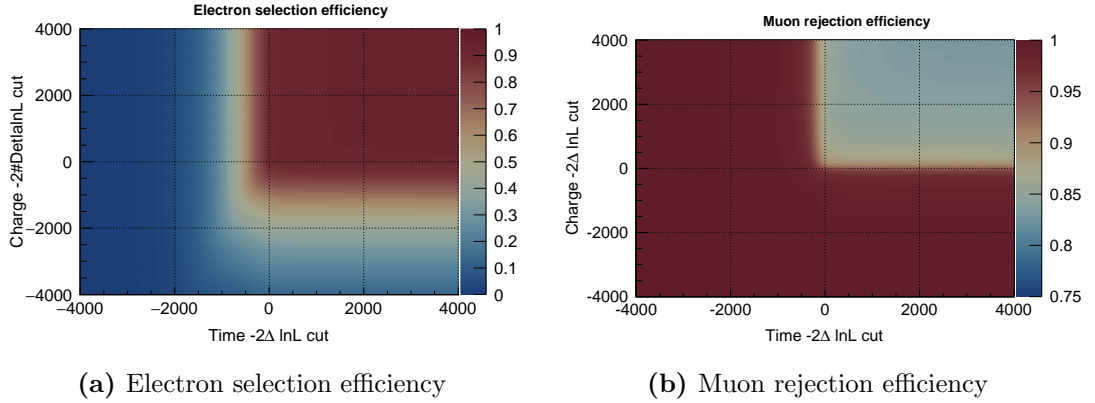


Figure 13.8: Efficiency with which electrons can be selected (left) and muons can be rejected (right) using rectangular cuts on the two likelihood components. The x-axis and y-axis show the values of the cut made in each component of the likelihood, with events selected as electrons if both $-2\Delta \ln \mathcal{L}$ components are below their respective cut values.

Using only cuts on these two variables, and assuming a ratio of $\frac{\nu_e \text{ CCQE}}{\nu_\mu \text{ CCQE}} = \frac{5\%}{95\%}$ yields the efficiencies and purities for selecting ν_e events out of a combined sample of CCQE interactions that are shown in Figure 13.9. A common metric for optimising selectors is (efficiency \times purity), and this is maximised for cuts of $-2\Delta \ln \mathcal{L}_{time} < 208$ and $-2\Delta \ln \mathcal{L}_{charge} < -592$, yielding a selection efficiency for ν_e CCQE events of 83.3% and a purity of 63.0%.

13.4 ANN-based particle identification

In reality, the neutrino interactions occurring inside CHIPS will not be exclusively CCQE. There will also be a large background of NC interactions, as well as non-quasielastic CC events such as resonance or deep inelastic scattering interactions. Examining the events generated by Genie, the composition was approximately 70% CC and 30% NC. Within the CC events, approximately 20% were quasielastic, 30% were resonant, and 50% were DIS. These are the compositions which were assumed in the following study, and are summarised in Table 13.3.




Figure 13.9: The purity of the selected ν_e CCQE sample as a function of the selection efficiency, assuming an initial sample composed of 95% ν_μ CCQE and 5% ν_e CCQE events.

Event type	Fraction of Events (%)
ν_e CCQE	0.70
ν_e CC Res.	1.05
ν_e CC DIS	1.75
ν_e NC	1.50
ν_μ CCQE	13.30
ν_μ CC Res.	19.95
ν_μ CC DIS	33.25
ν_μ NC	28.50

Table 13.3: Assumed contribution of different event classifications to the events detected in CHIPS, for the purposes of training a particle identification routine.

Figure 13.9 demonstrates that ν_e and ν_μ CCQE events can be distinguished from one another with around 80% efficiency, producing a sample that is over 60% pure assuming a 95% : 5% breakdown of ν_μ to ν_e CCQE events. However Table 13.3 indicates that this is not sufficient for CHIPS: there is a large NC background that also needs to be removed, and 80% of CC events are not quasielastic. To address these challenges, a more sophisticated particle ID routine was devised, making use of two artificial neural networks (ANNs). Each ANN is a multilayer perceptron (MLP),

with one trained to separate ν_e CCQE from ν_μ CCQE events, and the other used to separate the ν_e CCQE events from NC interactions. Using this method the ν_μ CCQE events, which should be the easiest to reject, can be removed before the more difficult task of rejecting more complicated topologies. Training using the simpler CCQE subsamples provides a conservative estimate of how well a selector could perform, as future work could potentially improve the ability to identify CC non-QE events.

In addition to the $\Delta \ln \mathcal{L}$ components, a number of other variables are made available to the ANN. These were chosen by examining the distribution of the variables and their correlation matrices for ν_e CC, ν_μ CC and NC events and selecting a set of variables that appeared to have some discriminating power and were not trivially correlated.

In the absence of a two-ring fit, a large amount of activity outside the reconstructed ring serves as a good proxy indicator that a second ring was present, so the fraction of hits inside and outside the reconstructed ring (assuming a Cherenkov angle of 41°) are made available. The fraction of hits inside the ring is further broken down into those in the main body of the ring and those in the ‘hole’ in the middle of the ring that exists if the particle comes to a stop before reaching the edge of the detector. The fraction of *predicted* charge outside the ring is also used - this reflects a tendency for the fitter to produce a larger single ring if there were two tracks present so that it overlaps both of them.

Two ‘goodness-of-fit’ variables are also used: the ratio of the charge component of the likelihood to the number of hit PMTs, and the ratio of the total predicted charge to the total detected charge. Energy effects are accommodated by including the total number of hits (which also helps to indicate the presence of a second ring), and the ratio of reconstructed energy to the total recorded charge.

Two additional variables were found to be useful for rejecting neutral current events, and are made available to the NC selector only; these are the reconstructed electron’s direction (as events with a lot of hadronic activity tend to produce fits that are less strongly directed along the beam direction) and the fraction of hits in the downstream ($x > 0$) half of the detector.

The variables are summarised in Table 13.4.

Name	Type	Description
deltaCharge2LnL	Fit result	Difference in $-2 \ln \mathcal{L}_{\text{charge}}$ between electron and muon fits
deltaTime2LnL	Fit result	Difference in $-2 \ln \mathcal{L}_{\text{time}}$ between electron and muon fits
chargeLikelihoodRatio	Fit quality	Ratio of $-2 \ln \mathcal{L}_{\text{charge}}$ to number of hits (e & μ)
predictedCharge/totalQ	Fit quality	Ratio of the total predicted charge to the total measured charge (e & μ)
fracPredQOutsideRing	Topology	Fraction of predicted charge outside the reconstructed ring (e & μ)
fracQInRing	Topology	Fraction of measured charge inside the reconstructed ring (e & μ)
fracQOutsideRing	Topology	Fraction of measured charge outside the reconstructed ring (e & μ)
nHits	Energy	Total number of detected hits
recoE/totalQ	Energy	Ratio of reconstructed energy to total charge (e & μ)
dirX	Topology	The x-component of the reconstructed direction unit vector (e only)
upstreamHitFrac	Topology	Fraction of total measured charge located in the upstream half of the detector

Table 13.4: Summary of variables made available to the two artificial neural networks used for identifying ν_e interactions. The bottom two variables are only used by the ANN that separates CC and NC events.

13.4.1 ANN to separate CCQE ν_e and ν_μ events

The ANN used to separate CCQE ν_e events from CCQE ν_μ events was trained using the ROOT TMVA package [150] and the pure CCQE event samples for the geometry with 6% coverage of 3" tubes. To avoid overtraining the samples were split in two, with half of the events used to train the MLP and half used to test its performance. Before training, a preselection is applied to remove events which escape the detector, have fewer than 50 hit PMTs, have a reconstructed vertex outside of a fiducial volume 1 m from the edge of the detector, or have energies within 50 MeV of the maximum and minimum allowed by the fitter.

The performance of this network is shown in Figure 13.10. The output of the network is a single discriminating variable whose value indicates the favoured particle hypothesis. A low value corresponds to a likely ν_μ CCQE interaction, and a high value corresponds to a ν_e CCQE-like interaction. Efficiencies are displayed relative to the sample of preselected events. For a cut value q , the ν_e efficiency is the fraction of ν_e CCQE events selected if we identify all events for which the ANN output is $\geq q$ as electrons. Similarly, the ν_μ selection efficiency is the same number for ν_μ CCQE events, and the rejection efficiency shown in the plot is $(1 - \text{selection efficiency})$.

13.4.2 ANN to separate CCQE ν_e events from NC events

This second ANN was trained using the sample of ν_e CCQE events, and one of pure ν_μ NC events. Events were required to pass the preselection cuts, and to have an output ≥ 0.8 from the previous network in order to discourage the two networks from identifying similar features. The results from this network are shown in Figure 13.11, again relative to the sample of preselected events. The performance appears poorer compared to the network for CCQE events, which occurs for two main reasons. Firstly, the preselection cut is already very effective for removing a large fraction of the NC events, in particular the $nHits > 50$ component. Secondly, CCQE ν_e and ν_μ interactions are particularly simple channels, producing single charged particles with markedly different track lengths and topologies, so should be the easiest to separate. Nevertheless, this network demonstrates that the NC background can be substantially reduced without adversely affecting the ν_e selection efficiency, and before performing any time-consuming two-track fits.




Figure 13.10: Output of the ANN designed to separate ν_e CCQE events from ν_μ CCQE ones. The filled histograms show the distribution of the ANN output variable for samples of pure ν_μ CCQE events (red) and ν_e CCQE events (green). The markers correspond to the efficiency with which ν_e CCQE events are selected (green) or ν_μ CCQE events are rejected (red) by placing a cut at that value.

13.4.3 Combining PID networks

To evaluate the expected performance of the detector, the two networks were combined and used to estimate the efficiency and purity with which ν_e CC interactions (both QE and non-QE) could be identified using a 10 kton CHIPS detector. Samples of ν_μ and ν_e events of all types were used, and the CCQE-only and NC-only samples were also included to increase the statistics, with events of each type reweighted to reflect the ratios of Table 13.3.

A two-dimensional scan was performed to optimise the ANN cuts applied. Figure 13.12 shows a plot of the figure of merit (efficiency \times purity) as a function of these two cuts. This figure of merit is maximised for cuts of ANN > 0.90 for the ν_e CCQE vs ν_μ CCQE network and ANN > 0.75 for the CC vs NC network. Figure 13.13 shows the distribution of the ANN output variables broken down by event category, with each event type scaled to the appropriate size. For each plot, all of the other selection criteria have already been applied. Although the neural networks were




Figure 13.11: Output of the ANN designed to separate ν_e CCQE events from NC ones. The filled histograms show the distribution of ANN output variable for samples of pure NC events (red) and ν_e CCQE events (green). The markers correspond to the efficiency with which ν_e CCQE events are selected (green) or NC events are rejected (red) by placing a cut at that value.

trained using CCQE events, the selection performs well when applied to CC non-QE events as well.

Figure 13.14 shows the final breakdown of the selection efficiencies for different signal (ν_e CCQE and all-CC) and background (ν_μ CC and combined ν_e and ν_μ NC) event categories, as a function of the energy of the reconstructed electron track. Also shown is the purity of the selected sample. The statistical errors on the purity are due to the small number of accepted background events, due to the efficiency with which these cuts reject them. The overall flux-averaged efficiency for selecting ν_e CC events is 30.4%, producing a sample with 58.2% purity, and compares favourably with the performance assumed in the Letter of Intent [128] ($\sim 34\%$ for CCQE and $\sim 8\%$ for CC non-QE events with a reconstructed neutrino energy of 2 GeV) when assessing the potential physics reach even though this assumed a larger (20%) PMT coverage.




Figure 13.12: Efficiency \times purity figure of merit for the selected ν_e sample shown as a function of the cut values applied to the two ANN output variables, with the maximum indicated by the green star. This is used to determine the cut values applied for the ν_e selection.




Figure 13.13: Distributions of the ANN selection variables used to identify ν_e events, shown for different event categories and normalised according to Table 13.3. Each plot shows events which have passed preselection and the cut based on the other ANN, and events to the right of the arrow are selected as ν_e events. All CC ν_e events are shown in pale green, with the CCQE component highlighted in dark green.




Figure 13.14: Efficiency with which signal ν_e events and background NC and ν_μ CC events are selected, and the purity of the resulting ν_e sample as a function of the energy of the reconstructed electron track.

13.5 Fitting pions

Figures 13.13 and 13.14 demonstrate that NC interactions make a large contribution to the selected background events, so improving their rejection is a priority for the fitter. One method of doing so would be to perform a dedicated π^0 fit to candidate electron events.

A preliminary test of the π^0 fitter was performed as follows: for any event in the sample of ν_μ NC events produced by GENIE that contained at least one π^0 , the highest energy one was isolated and run separately through the Geant4 detector simulation. This produced a clean sample of neutral pions with appropriate energy and direction distributions for the NuMI beam but with any other potentially confounding tracks removed, against which the fitter could be tested. The π^0 fitter was then used to reconstruct these events, without applying the π^0 mass constraint to the two photons.

Figure 13.15 shows the spectrum obtained when these events were fitted, and the invariant masses of each pair of photons were plotted. A clear peak is visible in the region of the true mass of the π^0 ($m_{\pi^0} = 134.98$ MeV), and fitting this with a Gaussian yields $m_{\gamma\gamma} = 130 \pm 38$ MeV.




Figure 13.15: Invariant mass, $m_{\gamma\gamma}$, of the two photons when a sample of clean neutral pions isolated from ν_{μ} NC events is reconstructed using the π^0 fitter, without applying the π^0 mass constraint. The pion mass peak has clearly been reconstructed, and a Gaussian fit returns $m_{\gamma\gamma} = 130 \pm 38$ MeV.

This demonstrates that the fitter is clearly capable of reconstructing neutral pions. Going forward, the challenge is to integrate this with the rest of the reconstruction and particle ID scheme. The rejection of NC events may also benefit from the fitter being extended with additional emission profiles to identify charged pions.

Chapter 14

Summary and conclusion

MINOS has a long history of measuring the neutrino mixing parameters relevant to ν_μ disappearance. In addition to its usual beam data, MINOS has also collected a rich sample of atmospheric neutrinos, spanning almost a complete solar cycle. This thesis documents the first dedicated MINOS atmospheric neutrino analysis to employ a full three-flavour description of the mixing, and contains the first atmospheric data analysed in the MINOS+ era. Samples of contained vertex ν_μ interactions, neutrino induced muons, and contained vertex ν_e and NC showers have been isolated and used to perform a fit to the oscillation parameters. Contours corresponding to 68% and 90% confidence intervals in the mass splitting and mixing angle were constructed, with the data favouring $\Delta m_{32}^2 = 2.03 \times 10^{-3} \text{ eV}^2$ and $\sin^2(\theta_{23}) = 1.0$.

While the matter effect provides the atmospheric sample with some sensitivity to the mass hierarchy, it is clear that new large detectors are required in order to constrain three-flavour effects with precision. In the second part of this thesis, design studies were presented for CHIPS, a low-cost proposed water Cherenkov detector for measuring ν_e appearance in the NuMI beam with high statistics. A reconstruction framework which uses both charge and timing information has been written to benchmark the performance of a variety of designs for a 10 kton module. It was demonstrated that acceptable performance can be achieved using a sparsely-instrumented detector with a 6% coverage of small 3" PMTs. Finally, a particle identification routine to select ν_e CC events against a large background of ν_μ and neutral current interactions was described.

Although MINOS+ ceased running at the end of June 2016, the coming decade remains an exciting time for neutrino physics. The full MINOS+ dataset, including

two more years of atmospheric data, has yet to be analysed and will continue to set competitive limits on the mixing parameters, while the higher beam energy also enables strong limits to be set on sterile neutrinos and other exotic phenomena. Meanwhile the worldwide neutrino oscillation programme continues apace, with exciting results and hints of CP-violation emerging from NO ν A and T2K, reactor experiments measuring θ_{13} and Δm_{31}^2 with ever-increasing precision, and a suite of short-baseline liquid argon experiments paving the way for the ambitious DUNE experiment. Prospects for atmospheric neutrinos are also promising, with neutrino telescopes collecting enormous samples, and proposals for Hyper Kamiokande and ICAL@INO moving forward.

In this sense, at least, neutrinos will not be disappearing any time soon.

Bibliography

- [1] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, and T. Schwetz, Updated fit to three neutrino mixing: exploring the accelerator-reactor complementarity, 2016, arXiv:1611.01514.
- [2] NuFIT group, Nufit 3.0, www.nu-fit.org/, 2016.
- [3] B. Pontecorvo, Zhur. Eksptl'. i Teoret. Fiz. **33** (1957).
- [4] M. Gell-Mann and A. Pais, Physical Review **97**, 1387 (1955).
- [5] W. Pauli, Letter to the 'Radioactive Ladies and Gentlemen', in *Tübingen Meeting (4 December 1930)*, 1930.
- [6] C. L. Cowan, Jr., F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, Science **124**, 103 (1956).
- [7] M. Goldhaber, L. Grodzins, and A. W. Sunyar, Phys. Rev. **109**, 1015 (1958).
- [8] G. Danby *et al.*, Phys. Rev. Lett. **9**, 36 (1962).
- [9] K. Kodama *et al.*, Physics Letters B **504**, 218 (2001).
- [10] ALEPH, DELPHI, L3, OPAL and SLD Collaborations, LEP Electroweak Working Group and SLD Electroweak and Heavy Flavour Groups, Physics Reports **427**, 257 (2006).
- [11] Planck Collaboration, P. A. R. Ade *et al.*, Astron. Astrophys. **571**, A16 (2014), 1303.5076.
- [12] R. Davis, D. S. Harmer, and K. C. Hoffman, Phys. Rev. Lett. **20**, 1205 (1968).
- [13] J. N. Bahcall, N. A. Bahcall, and G. Shaviv, Phys. Rev. Lett. **20**, 1209 (1968).
- [14] J. N. Bahcall, M. Pinsonneault, and G. Wasserburg, Reviews of Modern

- Physics **67**, 781 (1995).
- [15] S. Turck-Chieze and I. Lopes, *The Astrophysical Journal* **408**, 347 (1993).
- [16] I. Sackmann *et al.*, *The Astrophysical Journal* **360**, 727 (1990).
- [17] K. S. Hirata *et al.*, *Physical Review Letters* **63**, 16 (1989).
- [18] A. I. Abazov *et al.*, *Phys. Rev. Lett.* **67**, 3332 (1991).
- [19] P. Anselmann *et al.*, *Physics Letters B* **285**, 376 (1992).
- [20] J. N. Bahcall and R. K. Ulrich, *Rev. Mod. Phys.* **60**, 297 (1988).
- [21] SAGE Collaboration, J. N. Abdurashitov *et al.*, *Phys. Rev. C* **80**, 015807 (2009).
- [22] W. Hampel *et al.*, *Physics Letters B* **447**, 127 (1999).
- [23] C. V. Achar *et al.*, *Phys. Lett.* **18**, 196 (1965).
- [24] F. Reines *et al.*, *Phys. Rev. Lett.* **15**, 429 (1965).
- [25] T. Haines *et al.*, *Physical review letters* **57** (1986).
- [26] K. Hirata *et al.*, *Physics Letters B* **205**, 416 (1988).
- [27] Super-Kamiokande Collaboration, Y. Fukuda *et al.*, *Phys. Rev. Lett.* **81**, 1562 (1998).
- [28] M. Ambrosio *et al.*, *Physics Letters B* **566**, 35 (2003).
- [29] M. Sanchez *et al.*, *Phys. Rev. D* **68**, 113004 (2003).
- [30] B. P. Pontekorvo, *Zhur. Eksptl'. i Teoret. Fiz.* **37** (1959).
- [31] Z. Maki, M. Nakagawa, and S. Sakata, *Prog. Theor. Phys.* **28**, 870 (1962).
- [32] CKMfitter Group, J. Charles *et al.*, *Phys. Rev. D* **91**, 073007 (2015).
- [33] I. J. Aitchison and A. J. Hey *Gauge Theories in Particle Physics* Vol. 2 (CRC Press, 2012).
- [34] B. Kayser and F. Gibrat-Debu, *The Physics of Massive Neutrinos* (World Scientific, 1989).

- [35] S. Nussinov, Physics Letters B **63**, 201 (1976).
- [36] B. Kayser, Phys. Rev. D **24**, 110 (1981).
- [37] C. Giunti and C. W. Kim, Phys. Rev. D **58**, 017301 (1998).
- [38] S. Mikheev and A. Y. Smirnov, Sov. J. Nucl. Phys.(Engl. Transl.);(United States) **42** (1985).
- [39] L. Wolfenstein, Phys. Rev. D **17**, 2369 (1978).
- [40] B. Aharmim *et al.*, Physical Review C **72**, 055502 (2005).
- [41] T. Araki *et al.*, Physical Review Letters **94**, 081801 (2005).
- [42] MINOS Collaboration, P. Adamson *et al.*, Phys. Rev. Lett. **110**, 251801 (2013).
- [43] F. P. An *et al.*, Phys. Rev. Lett. **108**, 171803 (2012).
- [44] SNO Collaboration, B. Aharmim *et al.*, Phys. Rev. C **88**, 025501 (2013).
- [45] Particle Data Group, K. A. Olive *et al.*, Chin. Phys. **C38**, 090001 (2014).
- [46] KamLAND Collaboration, A. Gando *et al.*, Phys. Rev. D **88**, 033001 (2013).
- [47] Super-Kamiokande Collaboration, K. Abe *et al.*, Phys. Rev. Lett. **107**, 241801 (2011).
- [48] IceCube Collaboration, M. G. Aartsen *et al.*, Phys. Rev. Lett. **111**, 081801 (2013).
- [49] S. Adrin-Martnez *et al.*, Physics Letters B **714**, 224 (2012).
- [50] K2K Collaboration, M. H. Ahn *et al.*, Phys. Rev. D **74**, 072003 (2006).
- [51] P. Adamson *et al.*, Phys. Rev. Lett. **112**, 191801 (2014).
- [52] T2K Collaboration, K. Abe *et al.*, Phys. Rev. Lett. **112**, 181801 (2014).
- [53] Daya Bay Collaboration, F. P. An *et al.*, Phys. Rev. Lett. **115**, 111802 (2015).
- [54] RENO Collaboration, J. K. Ahn *et al.*, Phys. Rev. Lett. **108**, 191802 (2012).
- [55] Y. Abe *et al.*, Journal of High Energy Physics **2014** (2014).
- [56] T2K Collaboration, K. Abe *et al.*, Phys. Rev. Lett. **112**, 061802 (2014).

-
- [57] MINOS Collaboration, P. Adamson *et al.*, Phys. Rev. Lett. **110**, 171801 (2013).
- [58] NOvA Collaboration, P. Adamson *et al.*, Phys. Rev. Lett. **116**, 151806 (2016).
- [59] H. Nunokawa, S. Parke, and R. Z. Funchal, Phys. Rev. D **72**, 013009 (2005).
- [60] C. Giunti, C. Kim, and M. Monteno, Nuclear Physics B **521**, 3 (1998).
- [61] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, Phys. Rev. Lett. **13**, 138 (1964).
- [62] A. D. Sakharov, Soviet Physics Uspekhi **34**, 392 (1991).
- [63] F. R. Klinkhamer and N. S. Manton, Phys. Rev. D **30**, 2212 (1984).
- [64] M. Fukugita and T. Yanagida, Physics Letters B **174**, 45 (1986).
- [65] H. Nunokawa, S. Parke, and J. W. Valle, Progress in Particle and Nuclear Physics **60**, 338 (2008).
- [66] D. Michael *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **596**, 190 (2008).
- [67] MINOS Collaboration and others, NuMI-L-337 (1998).
- [68] MINOS Collaboration, P. Adamson *et al.*, Phys. Rev. D **73**, 072002 (2006).
- [69] ASA Ideas, PO Box 1, 1330 Fornebu.
- [70] P. Adamson *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **556**, 119 (2006).
- [71] R. Nichol, *Calibration of the MINOS Detectors*, PhD thesis, University College London, 2003.
- [72] K. Lang *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **545**, 852 (2005).
- [73] P. Adamson *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment

- 492, 325 (2002).
- [74] J. J. Hartnell, *Measurement of the Calorimetric Energy Scale in MINOS*, PhD thesis, Oxford University, 2005.
- [75] T. Sanuki *et al.*, The Astrophysical Journal **545**, 1135 (2000).
- [76] W. Apel *et al.*, Physical Review Letters **107**, 171104 (2011).
- [77] G. D. Barr, T. K. Gaisser, P. Lipari, S. Robbins, and T. Stanev, Phys. Rev. D **70**, 023006 (2004).
- [78] J. D. Chapman, *Atmospheric neutrino observations in the MINOS Far Detector*, PhD thesis, University of Cambridge, 2007.
- [79] Neutron Monitor Datasets, National Science Foundation Grant ATM- 9912341.
- [80] R. Engel, T. Gaisser, P. Lipari, and T. Stanev, Hadronic interactions for atmospheric cascades, in *Proceedings of ICRC Vol. 2001*, 2001.
- [81] Y. Cho *et al.*, Phys. Rev. D **4**, 1967 (1971).
- [82] T. Eichten *et al.*, Nuclear Physics B **44**, 333 (1972).
- [83] E-802 Collaboration, T. Abbott *et al.*, Phys. Rev. D **45**, 3906 (1992).
- [84] G. Prior, Nuclear Physics A **752**, 24 (2005).
- [85] NA49, G. Barr, Hadroproduction in proton carbon collisions at the NA49 experiment, in *Proceedings, 28th International Cosmic Ray Conference (ICRC 2003): Tsukuba, Japan, July 31-August 7, 2003*, pp. 1551–1554, 2003.
- [86] J. Dumarchez, T. Patzak, F. Vannucci, and G. Barr, Nuclear Physics B - Proceedings Supplements **143**, 89 (2005).
- [87] H. Gallagher, Nuclear Physics B - Proceedings Supplements **112**, 188 (2002).
- [88] R. Merenyi *et al.*, Phys. Rev. D **45**, 743 (1992).
- [89] D. Casper, Nuclear Physics B-Proceedings Supplements **112**, 161 (2002).
- [90] V. Agrawal, T. K. Gaisser, P. Lipari, and T. Stanev, Phys. Rev. D **53**, 1314 (1996).

- [91] B. J. Rebel, *Neutrino - Induced Muons in the MINOS Far Detector*, PhD thesis, Indiana U., 2004.
- [92] T. K. Gaisser, *Cosmic Rays and Particle Physics* (Cambridge University Press, 1990).
- [93] S. M. Kasahara *et al.*, Physical Review D **55**, 5282 (1997).
- [94] R. Brun *et al.*, *GEANT: Detector Description and Simulation Tool; Oct 1994* CERN Program Library (CERN, Geneva, 1993), Long Writeup W5013.
- [95] C. Zeitnitz and T. Gabriel, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **349**, 106 (1994).
- [96] M. A. Kordosky, *Hadronic interactions in the MINOS detectors*, PhD thesis, Texas U., 2004.
- [97] J. B. Birks *The Theory and Practice of Scintillation Counting: International Series of Monographs in Electronics and Instrumentation* Vol. 27 (Elsevier, 2013).
- [98] C. L. F. Howcroft, *Atmospheric Neutrinos In The MINOS Far Detector*, PhD thesis, University of Cambridge, 2004.
- [99] A. Blake, *A Study of Atmospheric Neutrino Oscillations in the MINOS Far Detector*, PhD thesis, University of Cambridge, 2005.
- [100] B. P. Speakman, *Atmospheric Electron Neutrinos in the MINOS Far Detector*, PhD thesis, University of Minnesota, 2007.
- [101] E. Lifshitz, L. Pitaevskii, and V. Berestetskii, Statistical physics **5** (1980).
- [102] D. DeMuth *et al.*, Astroparticle Physics **20**, 533 (2004).
- [103] J. Mitchell, *Measuring ν_μ Disappearance with the MINOS Experiment*, PhD thesis, University of Cambridge, 2011.
- [104] MINOS Collaboration, P. Adamson *et al.*, Phys. Rev. Lett. **106**, 181801 (2011).
- [105] MINOS, P. Adamson *et al.*, Phys. Rev. Lett. **112**, 191801 (2014), 1403.0867.
- [106] M. M. d. Medeiros *et al.*, *Estudo da oscilação de neutrinos muônicos usando*

- dados atmosféricos e de acelerador nos experimentos MINOS e MINOS+*, PhD thesis, Universidade Federal de Goiás, 2015.
- [107] J. Kopp, International Journal of Modern Physics C **19**, 523 (2008).
- [108] A. M. Dziewonski and D. L. Anderson, Physics of the Earth and Planetary Interiors **25**, 297 (1981).
- [109] G. L. Fogli *et al.*, Phys. Rev. D **86**, 013012 (2012).
- [110] F. P. An *et al.*, Chinese Physics C **37**, 011001 (2013).
- [111] Double Chooz Collaboration, Y. Abe *et al.*, Phys. Rev. D **86**, 052008 (2012).
- [112] A. Blake, Impact of θ_{13} and δ_{CP} on three-flavour analysis, Internal presentation, MINOS-doc-9810.
- [113] F. James and M. Roos, Computer Physics Communications **10**, 343 (1975).
- [114] G. D. Barr, T. K. Gaisser, S. Robbins, and T. Stanev, Phys. Rev. **D74**, 094009 (2006), astro-ph/0611266.
- [115] H. S. Fesefeldt, Aachen TU 3. Inst. Phys. Report No. PITHA-85-02, 1985 (unpublished).
- [116] G. Bower and R. Cassell, SLAC-GHEISHA.
- [117] MINOS Collaboration, P. Adamson *et al.*, Phys. Rev. D **81**, 072002 (2010).
- [118] The MINOS Collaboration, P. Adamson *et al.*, Phys. Rev. D **86**, 052007 (2012).
- [119] L. Corwin, Results of hand scan to assess systematic uncertainties in Atmospheric Neutrino Analysis, Minos doc-6925.
- [120] A. Blake, Atmospheric Neutrino Cross-Checks - Charge Separation, MINOS doc-6840.
- [121] G. D. Barr, T. K. Gaisser, P. Lipari, S. Robbins, and T. Stanev, Phys. Rev. **D70**, 023006 (2004), astro-ph/0403630.
- [122] G. Battistoni, A. Ferrari, T. Montaruli, and P. Sala, Astroparticle Physics **19**, 269 (2003).
- [123] M. Honda, T. Kajita, K. Kasahara, S. Midorikawa, and T. Sanuki, Phys. Rev.

- D 75, 043006 (2007).
- [124] A. Blake, Results of 2011 Atmospheric Neutrino Analysis Part II: Systematic Uncertainties, Internal presentation, MINOS-doc-8477.
- [125] X. Qian *et al.*, Phys. Rev. D **86**, 113011 (2012).
- [126] K. Abe *et al.*, (2011), hep-ex/1109.3262.
- [127] DUNE, R. Acciarri *et al.*, (2016), 1601.05471.
- [128] P. Adamson *et al.*, arXiv preprint arXiv:1307.5918 (2013).
- [129] R. Abbasi *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **618**, 139 (2010).
- [130] S. Adrián-Martínez *et al.*, The European Physical Journal C **74**, 1 (2014).
- [131] R. Arnold *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **536**, 79 (2005).
- [132] A. Perch, Construction of the CHIPS-M prototype and simulations of a 10 kiloton module, in *Proceedings, Topical Research Meeting on Prospects in Neutrino Physics (NuPhys2014)*, edited by F. Di Lodovico and S. Pascoli, 2014.
- [133] M. Pfützner, Prototype detection unit for the CHIPS experiment, in *Proceedings, XXVII International Conference on Neutrino Physics and Astrophysics: Neutrino 2016 (forthcoming)*.
- [134] F. Amat *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **844**, 108 (2017).
- [135] K. Lang, Prospects for CHIPS, in *Proceedings, Topical Research Meeting on Prospects in Neutrino Physics (NuPhys2014)*, edited by F. Di Lodovico and S. Pascoli, 2014.
- [136] E. V. Bugaev *et al.*, Physical Review D **58**, 054001 (1998).
- [137] L. Bogdanova, M. Gavrilov, V. Kornoukhov, and A. Starostin, Physics of

- Atomic Nuclei **69**, 1293 (2006).
- [138] S. Fukuda *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **501**, 418 (2003).
- [139] M. Shiozawa, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **433**, 240 (1999).
- [140] E. R. Davies, *Computer and machine vision: theory, algorithms, practicalities* (Academic Press, 2012).
- [141] R. B. Patterson, *A Search for Muon Neutrino to Electron Neutrino Oscillations at $\delta(m^2) > 0.1 eV^2$* , PhD thesis, Princeton U., 2007.
- [142] R. Patterson *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **608**, 206 (2009).
- [143] T. I. Quickenden and J. A. Irvin, The Journal of Chemical Physics **72** (1980).
- [144] R. M. Pope and E. S. Fry, Appl. Opt. **36**, 8710 (1997).
- [145] H. Buiteveld, J. H. M. Hakvoort, and M. Donze, Optical properties of pure water Vol. 2258, pp. 174–183, 1994.
- [146] KM3NeT, P. Bagley *et al.*, KM3NeT: Technical Design Report for a Deep-Sea Research Infrastructure in the Mediterranean Sea Incorporating a Very Large Volume Neutrino Telescope, 2009.
- [147] E. Catmull and R. Rom, Computer aided geometric design **74**, 317 (1974).
- [148] J. A. Nelder and R. Mead, The Computer Journal **7**, 308 (1965), <http://comjnl.oxfordjournals.org/content/7/4/308.full.pdf+html>.
- [149] C. Andreopoulos *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **614**, 87 (2010).
- [150] A. Hoecker *et al.*, PoS ACAT, 040 (2007), physics/0703039.