Atmospheric Neutrino Observations In The MINOS Far Detector

John Derek Chapman
of
Churchill College, University of Cambridge

A dissertation submitted to the University of Cambridge
for the degree of Doctor of Philosophy

Abstract

This thesis presents the results of atmospheric neutrino observations from a 12.23 ktyr exposure of the 5.42 kt MINOS Far Detector between 1st August 2003 until 1st March 2006. The separation of atmospheric neutrino events from the large background of cosmic muon events is discussed. A total of 277 candidate contained vertex $\nu/\bar{\nu}_\mu$ CC data events are observed, with an expectation of 354.4 ± 47.4 events in the absence of neutrino oscillations. A total of 182 events have clearly identified directions, 77 data events are identified as upward going, 105 data events are identified as downward going. The ratio between the measured and expected up/down ratio is: $R_{u/d}^{\text{data}}/R_{u/d}^{\text{MC}} = 0.72_{-0.11}^{+0.13}(\text{stat.}) \pm 0.04(\text{sys.})$. This is 2.1σ away from the expectation for no oscillations. A total of 167 data events have clearly identified charge, 112 are identified as ν_μ events, 55 are identified as $\bar{\nu}_\mu$ events. This is the largest sample of charge-separated contained-vertex atmospheric neutrino interactions so far observed. The ratio between the measured and expected $\bar{\nu}_\mu/\nu_\mu$ ratio is: $R_{\bar{\nu}/\nu}^{\text{data}}/R_{\bar{\nu}/\nu}^{\text{MC}} = 0.93_{-0.15}^{+0.19}(\text{stat.}) \pm 0.12(\text{sys.})$. This is consistent with ν_μ and $\bar{\nu}_\mu$ having the same oscillation parameters. Bayesian methods were used to generate a $\log(L/E)$ PDF for each selected event with a clearly identified direction. This information was used to give a measure of the uncertainty in the reconstructed $\log(L/E)$ value for each event. A maximum likelihood analysis is used to determine the allowed regions for the oscillation parameters Δm_{32}^2 and $\sin^2 2\theta_{23}$. The likelihood function uses the uncertainty in $\log(L/E)$ to bin events in order to extract as much information from the data as possible. This fit rejects the null oscillations hypothesis at the 98% confidence level. A fit to independent ν_μ and $\bar{\nu}_\mu$ oscillation assuming maximal mixing for both is also performed. The projected sensitivity after an exposure of 25 ktyr is also discussed.

Declaration

This dissertation is the result of my own work, except where explicit reference is made to the work of others, and has not been submitted for another qualification to this or any other university. This dissertation does not exceed the word limit for the respective Degree Committee.

John Chapman

Acknowledgements

Firstly, I would like to thank my parents, for their help and support over the past 26 years. I would also like to thank my supervisor, Mark Thomson, for his help and guidance throughout my PhD. I am also indebted to Andy Blake for taking the time to answer so many of the questions I have had about MINOS over the years. Thanks to Caius Howcroft and Andy Buckley for teaching me C++, without which 99% of my PhD would have been impossible. I would also like to mention my office mates Andy Culling and John Marshall, who have tolerated me over the last few years and for making our office a fun place to be. Thanks to Martin White, Chris Jones, Alan Phillips, Mark Slater and Chris White for making days at work much more enjoyable than they would have been otherwise. Thanks also to Pat and David Ward for their help and advice over the years. Thanks to the MINOS DAQ group for their help and support in my work on MINOS Run Control and not laughing at some of the stupid questions I asked. Last, but not least, thanks to my family and friends for all their support and encouragement.

Contents

1	Introduction	1
1.1	A Brief History of Neutrinos	1
2	Neutrino Physics	6
2.1	Theory	6
2.1.1	Theory of Neutrino Oscillations	6
2.1.2	Matter Effects	12
2.1.3	Neutrino Decoherence	13
2.1.4	Neutrino Decay	13
2.1.5	Massive Neutrinos	13
2.2	Neutrino Mass Measurements	16
2.2.1	Mass Measurements Using ν_e	16
2.2.2	Mass Measurements Using ν_μ	17
2.2.3	Mass Measurements Using ν_τ	17
2.2.4	Constraints From Astrophysics	17
2.3	An Overview Neutrino Oscillation Experiments	17
2.4	Solar Neutrino Experiments	18
2.4.1	Solar Neutrino Production	18
2.4.2	Matter Effects For Solar Neutrinos	19

2.4.3	Radiochemical Experiments	21
2.4.4	Water Čerenkov Experiments	23
2.5	Reactor Neutrino Experiments	27
2.5.1	KamLAND	27
2.5.2	CHOOZ	30
2.5.3	Future Experiments	32
2.6	Atmospheric Neutrino Experiments	32
2.6.1	Atmospheric Neutrino Production	34
2.6.2	Soudan 2	35
2.6.3	Super-Kamiokande	35
2.6.4	MINOS	38
2.7	Long-baseline Accelerator Neutrino Experiments	39
2.7.1	K2K	39
2.7.2	MINOS	40
2.7.3	Future Experiments	45
2.8	Short-baseline Accelerator Neutrino Experiments	46
2.8.1	LSND	46
2.8.2	MiniBooNE	46
2.9	Summary	48
3	The MINOS Experiment	50
3.1	Overview	50
3.2	NuMI Beam	51
3.3	The MINOS Detectors	53
3.3.1	MINOS Near Detector	53

3.3.2	MINOS Far Detector	55
3.3.3	MINOS Calibration Detector	56
3.4	MINOS Physics Analyses	56
3.4.1	Beam ν_μ disappearance Analyses	56
3.4.2	Contained Vertex Atmospheric Neutrino Analyses	56
3.4.3	Neutrino Induced Muon Analyses	57
3.4.4	Other Analyses	57
4	The MINOS Far Detector	59
4.1	Overview	59
4.2	Detector Structure	60
4.3	Detector Technologies	61
4.3.1	Plastic Scintillator	61
4.3.2	M16 Photo-multiplier Tubes	63
4.3.3	Front-End Electronics	67
4.4	Data Acquisition and Trigger Processing	70
4.4.1	Far Detector Data Acquisition System Operation	70
4.4.2	The Trigger Farm	71
4.4.3	Data Collection Process	73
4.4.4	Run Control	73
4.5	Detector Control System	77
4.6	Data Quality	78
4.6.1	Coil Currents	79
4.6.2	Dead Electronics	79
4.6.3	Busy Electronics	83

4.7	The Veto Shield	83
5	Calibrating The MINOS Far Detector	89
5.1	Overview	89
5.2	Timing Calibration	90
5.2.1	Calibration Software	90
5.2.2	Time Walk Corrections	91
5.2.3	Electronics Swaps	92
5.2.4	Timing Drift	97
5.2.5	Strip-to-Strip Calibration	98
5.2.6	Validating The Calibration	99
5.3	Spill Time Calibration	104
5.4	Pulse Height Calibration	105
5.4.1	Gain Calibration	105
5.4.2	Linearity Calibration	106
5.4.3	Cosmic-ray Muon Drift Calibration	106
5.4.4	Strip-to-Strip Calibration	107
5.4.5	Attenuation Calibration	108
5.4.6	MIP Calibration	108
5.4.7	Energy Calibration	109
6	Monte Carlo Simulation	111
6.1	Overview	111
6.2	Cosmic-Ray Muon Simulation	112
6.2.1	Cosmic-Ray Muon Flux	112
6.2.2	Simulating the Underground Cosmic Muon Flux	115

6.3	Atmospheric Neutrino Simulation	117
6.3.1	Normalising Atmospheric Neutrino Monte Carlo over the Solar Cycle	121
6.3.2	Simulation of Atmospheric Neutrino Cross-Sections	123
6.4	Detector Simulation	125
6.4.1	Active Detector Simulation	125
6.4.2	Detector Electronics Simulation	126
6.4.3	Electronics Timing Smearing	127
6.5	Evaluating The Simulation	130
6.5.1	Evaluating The Detector Simulation	130
6.5.2	Data-Monte Carlo Comparison of Cosmic Muon Flux	131
6.6	Monte Carlo Samples	133
7	Reconstructing Events	136
7.1	Reconstruction Overview	136
7.2	From Raw Data to Digits	136
7.3	Demultiplexing	137
7.4	Strip making	139
7.5	Track Finding	139
7.6	Track Fitting	143
7.7	Shower Finding	147
7.8	Veto Shield Reconstruction	148
7.9	Event Reconstruction	149
7.10	Shower Energy Calibration	150
8	Atmospheric Neutrino Event Selection	153
8.1	Selection Overview	153

8.2	Data Quality	154
8.3	Removal of Beam Events	155
8.4	Pre-selection	156
8.4.1	Upward-going Pre-selection	157
8.4.2	Contained-Highest-End Pre-selection	161
8.5	Upward-Going Events	166
8.5.1	Reconstruction Quality Cuts	167
8.5.2	Topology Cuts	168
8.5.3	Timing Cuts	169
8.5.4	Vertex Containment Cuts	173
8.5.5	Fully Contained/Partially Contained Separation	173
8.5.6	Summary of Upward-Going Selection	174
8.6	Contained-Highest-End Events	174
8.6.1	Highest End Containment Cuts	174
8.6.2	Topology Cuts	176
8.6.3	Fully Contained/Partially Contained Separation	180
8.6.4	Downward-Going Partially Contained Event Selection	183
8.6.5	Veto Shield Cuts	184
8.6.6	Summary of Contained Highest-End Selection	187
8.7	Combining the Two Branches of The Selection	187
8.7.1	Selection Efficiencies	192
8.8	Identifying A High Resolution Sample	192
8.9	Charge Identification	196
8.9.1	Splitting Tracks	198
8.10	Summary	199

9	Searching For Atmospheric Neutrino Oscillations	204
9.1	Overview	204
9.2	Double Ratios	205
9.2.1	Determination of Up/Down Ratio	205
9.2.2	Determination of $\nu_\mu/\bar{\nu}_\mu$ Ratio	208
9.3	L/E Distribution	211
9.4	Bayesian Analysis	211
9.4.1	Bayes's Theorem	211
9.4.2	Applying Bayes's Theorem To Atmospheric Neutrinos	213
9.4.3	Implementation of The Bayesian Method	214
9.5	Fits to Neutrino Oscillations	222
9.5.1	Likelihood Function	222
9.5.2	Systematic Uncertainties	223
9.5.3	Fit Mechanics	224
9.5.4	Fit To The Combined $\nu_\mu/\bar{\nu}_\mu$ Sample	225
9.5.5	Fits To ν_μ Dominated And $\bar{\nu}_\mu$ Dominated Samples	229
9.5.6	Separated $\nu_\mu/\bar{\nu}_\mu$ Fit	232
9.6	Projected Sensitivities	236
10	Conclusions	239
A	Selected Events	1
	Bibliography	13
	List of Figures	22
	List of Tables	29

“It is a bad plan that admits of no modification. – Maximus”

— Publilius Syrus 100 BC

Chapter 1

Introduction

“Success is the ability to go from one failure to another with no loss in enthusiasm.”

— Sir Winston Churchill, 1874–1965

1.1 A Brief History of Neutrinos

The story of our awareness of neutrinos starts in 1914 when James Chadwick observed a continuous energy spectrum from β -decay [1]. If β -decay was a two body process, as was believed at the time, then Chadwick’s results indicated a violation of energy conservation. This problem remained a thorn in the side of physics until Wolfgang Pauli proposed a possible solution in a letter to a conference on radioactivity at Tübingen on the 14th December 1930 [2]. Pauli proposed that an electrically neutral, low mass, spin-1/2 *neutron* “was emitted along with the electron in such a way that the sum of the energies of the *neutron* and electron is constant”, thus conserving energy while maintaining conservation of charge and angular momentum.

In 1931 Chadwick discovered a heavy neutron [3] (for which he was awarded the Nobel Prize in 1935). This was clearly not Pauli’s *neutron* though, as its mass was much too high. Pauli’s particle was then renamed the neutrino by Enrico Fermi later that year. In 1934 Fermi published a quantitative theory of β -decay [4], including the neutrino, giving it a sound theoretical grounding.

Later in 1934, Hideki Yukawa published his theory of mesons (referring to what we

now call pions), giving an explanation of the interaction between protons and neutrons in the nucleus [5]. He was awarded the Nobel Prize “for his prediction of mesons on the basis of theoretical work on nuclear forces” in 1949. In 1943 Shoichi Sakata and Takesi Inoue proposed that the new particle, the *meson*, in the hard component of cosmic rays, was not the meson (Yukawa particle, see below) of Yukawa theory, rather it was “an elementary particle which has close correlations to the Yukawa particle, but it should be considered as an elementary particle of a different sort”. They proposed that the cosmic-ray *meson* underwent the following interaction:

$$m^{\pm} \rightleftharpoons n + Y^{\pm} \quad (1.1)$$

similar to the electron-neutrino case:

$$e^{-} \rightleftharpoons \nu + Y^{-} \quad (1.2)$$

where $Y^{\pm}$ is the charged Yukawa particle, $m^{\pm}$ is the charged *meson* and n is “a neutral *meson* ... assumed to have negligible mass, and consequently may be regarded as equivalent to the neutrino”. The cosmic ray *meson* mentioned here subsequently became known as the muon (μ). This was the first suggestion of the existence of the muon neutrino (ν_{μ}). The work was published in 1946 [6].

In 1951 Frederick Reines and Clyde Cowan, both from Los Alamos Laboratory, attempted to detect the high flux of neutrinos emitted by nuclear reactors at Hanford [7]. This research was dubbed “Project Poltergeist”. After early attempts were dogged by high backgrounds, they finally observed conclusive evidence of neutrino interactions in 1956 using an underground detector constructed near the Savannah River nuclear reactor in South Carolina [8]. Frederick Reines was awarded the Nobel prize in 1995 for this work.

In 1957 Pontecorvo suggested that the neutrino might be a combination of two neutral Majorana particles. The difference between ν and $\bar{\nu}$ being whether the combination of was symmetric or antisymmetric. If $m_{\nu} > 0$ then there was the possibility of $\nu \leftrightarrow \bar{\nu}$ oscillations [9]. This was the first mention of ‘neutrino oscillations’; albeit not of the kind that have become familiar in recent years though.

Goldhaber *et al.* measured the helicity of the neutrino in 1958 [10]. They observed that it was a left handed particle - in agreement with the V-A (Vector-Axial Vector) theory of the weak interaction proposed by Feynman and Gell-Mann [11] and in dis-

agreement with the S-T (Scalar-Tensor) theory proposed by Lee and Yang [12], which had been popular previously.

One year later, Ray Davis *et al.* had put a limit on the neutrino flux from the sun and showed that $\bar{\nu}$ could be distinguished from ν [13]. Soon after, Melvin Schwartz proposed creating a neutrino beam by making a beam of π^+/π^- and allowing them decay into $\nu/\bar{\nu}$ [14]. This offered a powerful new technique for studying the interactions of neutrinos. Around the same time a problem had been observed in muon decay data. Muons had been observed to decay thus: $\mu^- \rightarrow e^- \nu \bar{\nu}$, however no $\nu \bar{\nu}$ annihilation had been observed. The two neutrinos being of different flavours would explain this. In 1962, using the AGS accelerator at Brookhaven, Ledermann *et al.* showed that the ν produced from π^+/π^- decay only produced μ mesons and so were different from the ν produced in β -decay [15]. This was the first confirmed observation of ν_μ .

In 1962 Maki, Nakagawa and Sakata published “Remarks on the Unified Model of Elementary Particles” [16]. This article proposes that the *weak neutrinos* (referring to ν_e and ν_μ) may be linear combinations of the *true neutrinos*. They go on to propose that if these *true neutrinos* are the eigenstates of some new interaction X then the “*weak neutrinos* are not stable due to the occurrence of a virtual transmutation $\nu_e \leftrightarrow \nu_\mu$ induced by the interaction”. They call the self-energies of the *true neutrinos* under the interaction X “masses”. They go on to discuss the level of mass difference that could be detected by the experiments of the time. This is the first discussion of oscillations between neutrino flavour states.

The first atmospheric neutrinos (see Section 2.6.1) were detected by experiments performed in the early 1960s. These experiments were located in Kolar Gold Field in Southern India [17] and at the East Rand Proprietary mine in South Africa [18]. Both experiments were located extremely deep underground; the former at 7000 mwe (metres of water equivalent) and the latter at 8800 mwe, to reduce backgrounds from cosmic ray muons. These experiments detected muons produced by atmospheric neutrino interactions in the surrounding rock.

Throughout the 1960s Davis *et al.* continued to perform experiments to measure the solar neutrino flux [19] [20]. Then in 1971 Davis *et al.* measured the solar ν_e flux and found that it was much lower than had been predicted [21]. This came to be known as the “solar neutrino problem”. In early 1968 Pontecorvo suggested that the possibility of $\nu_\mu \leftrightarrow \nu_e$ oscillations were not inconsistent with the experimental evidence [22]. Then after the 1968 results of Davis *et al.* [20] Pontecorvo put forward $\nu_\mu \leftrightarrow \nu_e$ oscillations

as another possible explanation for the results in addition to the various nuclear and astrophysical uncertainties that still remained [23].

In 1973 the Gargamelle experiment at CERN discovered the weak neutral current (NC) interaction, mediated by the Z^0 boson, by observing events induced by $\nu/\bar{\nu}$ that produced hadrons, but no muon or electron [24].

The next advance in neutrino physics came from the water Čerenkov detectors. These experiments were originally set up to look for proton decay, but neutrinos created by the interactions of high energy cosmic rays in the atmosphere were a major source of background. Proton decay was not observed, but the US experiment IMB [25] and the Japanese experiment Kamiokande [26] both reported a deficit of ν_μ created in the atmosphere with respect to the number of ν_e . This deficit became known as the *atmospheric neutrino problem*. In addition to observing atmospheric neutrino events in 1989 the Kamiokande collaboration published results of observation of solar neutrinos from the ${}^8\text{B} \rightarrow {}^8\text{Be}^* + e^+ + \nu_e$ decay chain [27]. They observed that the solar neutrino flux was around half the value predicted by solar models. This was consistent with the measurements made by Davis *et al.* using their ${}^{37}\text{Cl}$ radiochemical detector, located in the Homestake mine, South Dakota [21].

Two other neutrino related major discoveries were made in the 1980s. Neutrinos from outside our solar system were observed for the first time, when observations of neutrino events were correlated with astronomical observations of the SN 1987A supernova. The LEP experiments at CERN precisely measured the width of the Z resonance. This measurement was highly significant for neutrino physics as it provided very strong evidence there were only 3 light ($m_\nu < 45 \text{ GeV}$), active neutrino flavours [28].

The goals of this thesis are:

- To identify a sample of charged current atmospheric muon neutrino events in the MINOS Far Detector.
- Then to further identify a sub-sample of events where the charge of the muon produced can be clearly identified.
- To use the identified muon neutrino events to look for evidence of atmospheric neutrino oscillations and attempt to measure the oscillation parameters, assuming identical $\nu_\mu \rightarrow \nu_\tau$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ oscillations.
- Finally to test for the possibility of different oscillation parameters for $\nu_\mu \rightarrow \nu_\tau$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ oscillations, using the events with clearly identified muon charges.

In Chapter 2 more recent results in neutrino physics are discussed along with neutrino oscillation theory. An overview of the *The MINOS Experiment* is given in Chapter 3. Followed by a more detailed discussion of *The MINOS Far Detector* in Chapter 4. Chapter 5 discusses *Calibrating The MINOS Far Detector*. Details of the *Monte Carlo Simulation* are discussed in Chapter 6. Chapter 7 contains a discussion of the process of *Reconstructing Events* in the MINOS Far Detector. Chapter 8 details the *Atmospheric Neutrino Event Selection* used in this thesis. In Chapter 9, “*Searching For Atmospheric Neutrino Oscillations*”, the process of measuring the oscillation parameters using the selected data events is described. Lastly, Chapter 10 gives the final *Conclusions*.

The author contributed to this effort by working on the data acquisition software for both MINOS detectors (see Section 4.4.4); in addition to being on-call 24 hours a day for many months providing technical support for both detectors. The author rewrote the existing Far Detector timing calibration software more efficiently and wrote additional programs to automate finding jumps in readout timing caused by electronics swaps and searching for long term timing drift (see Section 5.2). Software was also written to select atmospheric neutrino events (see Chapter 8). Using ideas from [29] software was written by the author to estimate the $\log(L/E)$ resolution of each event using Bayesian methods (as described in Section 9.4). The author wrote software to fit the data to both combined and separate $\nu_\mu/\bar{\nu}_\mu$ oscillations (see Section 9.5).

Chapter 2

Neutrino Physics

“In the beginning the Universe was created. This has made a lot of people very angry and been widely regarded as a bad move.

– (The Restaurant at the End of the Universe, 1980) ”

— Douglas Adams, 1952–2001

2.1 Theory

2.1.1 Theory of Neutrino Oscillations

Neutrinos can only be observed through their couplings to the weak force. The eigenstates of these interactions are known as the neutrino flavour states $|\nu_\alpha\rangle$ ($\alpha = e, \mu, \tau$) (electron, muon and tau). By observing the weak charged-current (CC) interactions of neutrinos the neutrino flavour can be determined by observing the flavour of the charged lepton produced. The theory of neutrino oscillations predicts that a neutrino produced as one flavour state may interact as another flavour state, as shown in Figure 2.1.

If neutrinos have mass then there will exist a set of neutrino mass eigenstates $|\nu_i\rangle$ ($i = 1, 2, 3$). In general the flavour eigenstates need not be commensurate with these

Figure 2.1: Feynman diagram of a neutrino produced in flavour state α propagating as a superposition of mass eigenstates and being detected in flavour state β .

mass eigenstates, but can be written as a linear combination of them as below:

$$|\nu_\alpha\rangle = \sum_{i=1,2,3} U_{\alpha i} |\nu_i\rangle \quad (2.1)$$

where U is a unitary lepton mixing matrix.

Each mass eigenstate propagates as a free particle in a vacuum, so a state $|\nu_i\rangle$ at position $\mathbf{x}$ with four-momentum $\mathbf{p}_i$ evolves as follows:

$$|\nu_i(\mathbf{x})\rangle = e^{-i\mathbf{p}_i \cdot \mathbf{x}} |\nu_i\rangle \quad (2.2)$$

Thus, a neutrino produced in a flavour eigenstate $|\nu_\alpha\rangle$ will evolve as below:

$$|\nu_\alpha(\mathbf{x})\rangle = \sum_{i=1,2,3} e^{-i\mathbf{p}_i \cdot \mathbf{x}} U_{\alpha i} |\nu_i\rangle \quad (2.3)$$

Inverting equation (2.1) using the unitarity of U and inserting the result into equation (2.3) it can be seen that

$$|\nu_\alpha(\mathbf{x})\rangle = \sum_{\beta} \left[\sum_i U_{\alpha i} e^{-i\mathbf{p}_i \cdot \mathbf{x}} U_{\beta i}^* \right] |\nu_\beta\rangle \quad (2.4)$$

It is important to note that quantum theory requires all the phases $(\mathbf{p}_i \cdot \mathbf{x})$ are taken at the *same space-time point*. If it is assumed that all neutrino mass eigenstates can be

analysed at the same spatial position in the lab frame (*i.e.* at a distance L from the where they were produced), making the substitution $t_i = E_i L / p_i$ gives:

$$\mathbf{p}_i \cdot \mathbf{x} = E_i t_i - p_i L = \left(\frac{E_i^2 - p_i^2}{p_i} \right) \times L = \frac{m_i^2 L}{p_i}. \quad (2.5)$$

In general this is *invalid*, as the values of t_i used are effectively different for each mass eigenstate. It should be emphasised that a more complete understanding of the oscillation phase needs a full quantum field-theoretical approach [30]. If we consider the two neutrino oscillations case, as in [31], the phase difference of the two components of the neutrino wave function is:

$$\begin{aligned} \Phi(x, t) &= (E_1 - E_2)t - (p_1 - p_2)x \\ &= \frac{m_1^2 - m_2^2}{E_1 + E_2} + (p_1 - p_2) \left[\frac{p_1 + p_2}{E_1 + E_2} t - x \right]. \end{aligned} \quad (2.6)$$

The first term in equation (2.6) is the ‘standard’ one. The second term is not usually considered. It vanishes for points related by

$$x = v_o t = \frac{p_1 + p_2}{E_1 + E_2} t, \quad (2.7)$$

where v_o is the average velocity of the two components. The components of the neutrino wave function can retain their coherence and one can observe neutrino oscillations only for space-time points obeying equation (2.7).

At the energies relevant for most neutrino oscillation experiments the neutrinos can be considered to be ultra-relativistic, so the approximation $\mathbf{x} = (t, x, 0, 0) \simeq (L, L, 0, 0)$ is valid.

$$\mathbf{p}_i \cdot \mathbf{x} = E_i t - p_i x \simeq (E_i - p_i)L \quad (2.8)$$

Substituting:

$$p_i \simeq E_i - \frac{m_i^2}{2E_i}, \quad (2.9)$$

into equation (2.8) gives:

$$\mathbf{p}_i \cdot \mathbf{x} \simeq \frac{m_i^2 L}{2E_i} = \frac{m_i^2 L}{2E} \quad (2.10)$$

Where all mass eigenstates are taken to have the same energy. Substituting equation (2.10) into equation (2.4) gives the following:

$$|\nu_\alpha(\mathbf{x})\rangle = \sum_\beta \left[\sum_i U_{\alpha i} e^{-im_i^2 L/2E} U_{\beta i}^* \right] |\nu_\beta\rangle \quad (2.11)$$

The probability that a neutrino produced at the origin in flavour state ν_α is observed at $\mathbf{x}$ in a flavour state ν_β , as shown in Figure 2.1, is given by:

$$P(\nu_\alpha \rightarrow \nu_\beta) = |\langle \nu_\beta | \nu_\alpha(\mathbf{x}) \rangle|^2 \quad (2.12)$$

Expanding this gives:

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) = & \delta_{\alpha\beta} - 4 \sum_{i < j} \mathbf{Re}(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin^2 \left(\frac{\Delta m_{ij}^2 L}{4E} \right) \\ & + 2 \sum_{i < j} \mathbf{Im}(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin \left(\frac{\Delta m_{ij}^2 L}{2E} \right), \end{aligned} \quad (2.13)$$

where $\Delta m_{ij}^2 (= m_i^2 - m_j^2)$ is the squared mass splitting between the i^{th} and j^{th} mass eigenstates.

The probabilities for anti-neutrino flavours can be obtained from the transition probabilities for neutrinos as follows. Assuming $\mathcal{CP}\mathcal{T}$ invariance holds then,

$$P(\bar{\nu}_\beta \rightarrow \bar{\nu}_\alpha) = P(\nu_\alpha \rightarrow \nu_\beta). \quad (2.14)$$

Equation (2.13) shows that

$$P(\nu_\alpha \rightarrow \nu_\beta; U^*) = P(\nu_\beta \rightarrow \nu_\alpha; U), \quad (2.15)$$

where $P(\nu_\alpha \rightarrow \nu_\beta; U^*)$ means equation (2.13) with U replaced by U^* . Together these two results imply

$$P(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta; U) = P(\nu_\alpha \rightarrow \nu_\beta; U^*). \quad (2.16)$$

This indicates that if U is complex then the probabilities for neutrino and anti-neutrino transitions between the same flavour states will differ. This is due to the sign change of the last term in equation (2.13) for anti-neutrinos compared to the neutrino case. This

difference would be a source of $\mathcal{CP}$ violation.

In general the survival probability for neutrinos is given by equation (2.13) is:

$$P(\nu_\alpha \rightarrow \nu_\alpha) = 1 - 4 \sum_{i < j} |U_{\alpha i}|^2 |U_{\alpha j}|^2 \sin^2 \left(\frac{\Delta m_{ij}^2 L}{4E} \right). \quad (2.17)$$

Without assuming $\mathcal{CP}\mathcal{T}$ invariance survival probability for anti-neutrinos is then:

$$P(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\alpha) = 1 - 4 \sum_{i < j} |V_{\alpha i}|^2 |V_{\alpha j}|^2 \sin^2 \left(\frac{\overline{\Delta m_{ij}^2} L}{4E} \right), \quad (2.18)$$

where V is the unitary mixing matrix for anti-neutrinos. In the case that $\mathcal{CP}\mathcal{T}$ is violated *i.e*

$$P(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\alpha) \neq P(\nu_\alpha \rightarrow \nu_\alpha), \quad (2.19)$$

then $V_{\alpha i} \neq U_{\alpha i}$ and/or $\Delta m_{ij}^2 \neq \overline{\Delta m_{ij}^2}$.

Three Neutrino Flavour Mixing

The idea of massive neutrinos mixing was first discussed by Pontecorvo [9] [22] and Maki, Nakagawa and Sakata [16]. In recognition of the pioneering contributions of these scientists to the physics of oscillation and mixing; the unitary matrix U is known as the *Pontecorvo-Maki-Nakagawa-Sakata* or *PMNS* matrix. As mentioned in Chapter 1, measurements made at LEP in the 1980s provided very strong evidence there were only 3 light ($m_\nu < 45$ GeV), active neutrino flavours [28] [32]. For this reason the PMNS matrix is commonly shown as a 3×3 matrix. It is possible that there are other neutrino mass and flavour eigenstates, but they would either have to have masses greater than 45 GeV or be weak singlets (*sterile* neutrinos). See Section 2.8 for further discussion of this topic. A common parameterisation for the PMNS matrix is:

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}s_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \begin{pmatrix} e^{i\alpha} & 0 & 0 \\ 0 & e^{i\beta} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (2.20)$$

Here $c_{ij} = \cos \theta_{ij}$ and $s_{ij} = \sin \theta_{ij}$, where θ_{ij} is the mixing angle between the i^{th} and j^{th} mass eigenstates. The complex phase δ is known the *Dirac phase* and is the leptonic analogue of the single complex phase in the CKM matrix of the quark sector[33]. The size of δ determines the extent of $\mathcal{CP}$ violation in the neutrino sector. The complex phases α and β represent the *Majorana phases* and only have physical consequences if neutrinos are Majorana particles (identical to their anti-particles). The Majorana phases have been separated out as they have no effect on the neutrino oscillation probabilities.

Two Flavour Mixing Approximation

The current empirical evidence (detailed in this chapter) suggests a regime in which $|\Delta m_{32}^2|$ (and $|\Delta m_{31}^2|$) is much greater than $|\Delta m_{21}^2|$. For an oscillation experiment with L/E such that $\Delta M^2 L/E = \mathcal{O}(1)$ equation (2.13) is simplified considerably to the following forms:

$$P(\nu_\alpha \rightarrow \nu_\beta) \simeq S_{\alpha\beta} \sin^2 \left(\frac{\Delta M^2 L}{4E} \right) \quad (2.21)$$

for $\beta \neq \alpha$, and

$$P(\nu_\alpha \rightarrow \nu_\alpha) \simeq 1 - 4 T_\alpha (1 - T_\alpha) \sin^2 \left(\frac{\Delta M^2 L}{4E} \right) \quad (2.22)$$

Here,

$$S_{\alpha\beta} \equiv 4 \left| \sum_{i \text{ Up}} U_{\alpha i}^* U_{\beta i} \right|^2 \quad (2.23)$$

and

$$T_\alpha \equiv \sum_{i \text{ Up}} |U_{\alpha i}|^2 \quad (2.24)$$

where “ $i \text{ Up}$ ” denotes a sum over only those neutrino mass eigenstates that lie above ΔM^2 .

Empirical evidence further suggests that θ_{12} and θ_{23} are large and that θ_{13} is small. Further simplifications are possible in this regime and the three flavour oscillations of equation (2.13) effectively decouple into the following sets of two flavour oscillations:

Long Range Oscillations: Over long ranges where $\Delta m_{21}^2 L/E = \mathcal{O}(1)$, the dominant

mode of oscillations is $\nu_e \leftrightarrow \nu_\mu/\nu_\tau$, controlled by the squared mass difference Δm_{21}^2 and the mixing angle θ_{12} . This mode of oscillations is associated with measurements of ν_e disappearance in solar neutrinos and long baseline reactor experiments. The oscillation probability is given by:

$$P(\nu_e \rightarrow \nu_\mu/\nu_\tau) \simeq \sin^2 2\theta_{12} \sin^2 \left(\frac{\Delta m_{21}^2 L}{4E} \right) \quad (2.25)$$

Short Range Oscillations: Over short ranges where $\Delta m_{32}^2 L/E = \mathcal{O}(1)$, the dominant mode of oscillations is $\nu_\mu \leftrightarrow \nu_\tau$, controlled by the squared mass difference Δm_{32}^2 and the mixing angle θ_{23} . This short range mode of oscillations is associated with measurements of ν_μ disappearance in atmospheric neutrinos and long baseline neutrino experiments. The oscillation probability is given by:

$$P(\nu_\mu \rightarrow \nu_\tau) \simeq \sin^2 2\theta_{23} \sin^2 \left(\frac{\Delta m_{32}^2 L}{4E} \right) \quad (2.26)$$

Sub-Dominant Short Range Oscillations: Over short ranges, a sub-dominant mode of oscillations $\nu_e \leftrightarrow \nu_\mu/\nu_\tau$, is permitted at a level controlled by the small mixing angle θ_{13} . This short range mode of oscillations is associated with searches for ν_e disappearance in short baseline reactor experiments. The oscillation probability is given by:

$$P(\nu_e \rightarrow \nu_\mu/\nu_\tau) \simeq \sin^2 2\theta_{13} \sin^2 \left(\frac{\Delta m_{32}^2 L}{4E} \right) \quad (2.27)$$

2.1.2 Matter Effects

The preceding discussion of neutrino oscillations has assumed the neutrinos are propagating through a vacuum, but often a significant part of their path from creation to interaction will be through matter. The effect of matter must therefore be considered. When neutrinos travel through matter they undergo coherent elastic forward scattering. The amplitude of this scattering is enhanced for ν_e , as they can undergo CC and NC interactions with the electrons in matter, while ν_μ and ν_τ can only undergo neutral current interactions. The effect of this coherent forward scattering is to endow the ν_e with an “effective” mass relative to the ν_μ and ν_τ .

The effect of these interactions on neutrino oscillations is known as the MSW (Mikheev-Smirnov-Wolfenstein) effect after Wolfenstein who first proposed the concept in the late 1970s [34] and Mikheev and Smirnov who expanded his ideas [35]. The effect of

the medium on the neutrinos traversing it is described by an effective potential that depends on the composition and density of matter. For ν_e the potential is given by $V = \sqrt{2G_F N_e}$, where G_F is the Fermi constant and N_e is the electron density in the medium. This gives rise to an effective mixing and mass matrix. The MSW effect has a significant impact on the oscillations of solar neutrinos due to the high electron density in the sun. For oscillations of atmospheric neutrinos the consequences of a matter effect depends on whether the mode of oscillations is $\nu_\mu \rightarrow \nu_\tau$ or $\nu_\mu \rightarrow \nu_s$, where ν_s is a sterile neutrino. The signature for matter effects in the sterile case is a suppression of the oscillation probability and a higher mass squared splitting, which increases with neutrino energy.

2.1.3 Neutrino Decoherence

Neutrino Quantum Decoherence is a generalised term used to explain neutrino flavour change due to loss of coherence of the neutrinos' quantum mechanical phases. This process can occur in the Standard Model due to the mass eigenstates of neutrinos travelling at different velocities and hence separating out over large distances (for example neutrinos travelling from supernovae). Some coherency models invoke mechanisms using beyond Standard Model physics to increase the coherence effects and shorten the range over which they would be observable [36].

2.1.4 Neutrino Decay

Neutrino Decay attempts to explain the differences between observed and predicted neutrino fluxes by allowing at least one neutrino to decay into another active neutrino flavour, a sterile neutrino or some new particle. Decay models do not necessarily preclude the occurrence of oscillations [37]. Recent results from neutrino experiments which appear to favour oscillations over neutrino decay are discussed later on in this chapter.

2.1.5 Massive Neutrinos

Traditionally in the Standard Model (SM) there is no need to have massive neutrinos. Baryon number B and lepton number L automatically emerge as approximately conserved charges. Baryon number and $\mathcal{CP}$ are only violated in a very specific way,

described by the CKM matrix. Observations of neutrino oscillations call for an extension of the SM to include massive neutrinos. The new physics required can either be at energies lighter or heavier than 100 GeV. When neutrinos are assumed to be massless, the SM only contains the chirally left-handed field ν_L that couples to the $W^\pm$ and Z^0 bosons. The chirally right-handed field ν_R is not required. If the new physics occurs at energies less than 100 GeV, then the exclusion by LEP of new particles coupled to the Z boson and lighter than $M_Z/2$, still allows right-handed neutrino singlets ν_R to be added to the SM. These ν_R must be neutral under all SM gauge interactions.

If neutrino masses are accommodated in the same manner as quark masses, then right-handed neutrino singlets are added to the SM. This enables a “Dirac mass term” for neutrinos to be added to the SM Lagrangian:

$$\mathcal{L}_D = -m_D \bar{\nu}_L \nu_R + h.c. , \quad (2.28)$$

where m_D is a constant. This term, as with the mass terms of quarks and charged leptons, conserves the lepton number L . Since the rest of the SM conserves L already, then each neutrino mass eigenstate ν_i is distinct from its antiparticle $\bar{\nu}_i$ (as $L(\bar{\nu}_i) = -L(\nu_i)$). In this $\bar{\nu}_i \neq \nu_i$ world, ν_i and $\bar{\nu}_i$ are called “Dirac neutrinos”. Masses for Dirac neutrinos can be generated via the Higgs mechanism. Naively, this would suggest that neutrinos have masses similar to the other particles in the Standard Model. This is not the case, so in order to avoid this problem, the strength of the neutrino interactions with the Higgs boson must be made less than 10^{-12} that of the top quark interactions with the Higgs. Some models explain this weakness by not constraining ν_R to stay on the three-brane where the rest of the SM particles are constrained to lie. This means their interactions with other SM particles would be extremely weak [38].

A Majorana particle has only one 2 component spinor degree of freedom available. This is often stated as: “a Majorana particle is its own antiparticle”. As ν_R must be neutral under all SM gauge interactions, gauge invariance allows a Majorana mass term for right-handed neutrinos (such a term is not possible for other fermions as it would lead to violation of charge conservation).

$$\mathcal{L}_M = -m_R \bar{\nu}_R^c \nu_R + h.c. , \quad (2.29)$$

Here ν_R^c is the charge conjugate of ν_R . Since both ν_R and $\bar{\nu}_R^c$ absorb ν and create $\bar{\nu}$, $\mathcal{L}_M$ mixes ν and $\bar{\nu}$. Thus a Majorana mass does not conserve L . If there is no conserved lepton number to distinguish each neutrino mass eigenstate ν_i from its antiparticle $\bar{\nu}_i$. It

Figure 2.2: Possible neutrino mass hierarchies. The compositions of each mass eigenstate in terms of the flavour eigenstates are indicated by the fractions of the bars in each colour. Figures taken from [39].

is possible that neutrino masses may have contributions from both Dirac and Majorana terms. This is because the Dirac mass term, while it does not violate L conservation, does not require L conservation either. If L conservation is imposed though, then Majorana mass terms are ruled out.

Irrespective of how neutrinos gain their masses, experimentalists hope to measure the hierarchy of the neutrino mass eigenstates. Figure 2.2 shows the two hierarchies (based on three neutrinos) which are compatible with current experimental evidence. The fractions of the bars in each colour represent the current best fit compositions of each mass eigenstate in terms of the flavour eigenstates. The ordering of ν_1 and ν_2 has been determined by Solar neutrino experiments (see Section 2.4.4) and reactor neutrino experiments (see Section 2.5.1). Physicists hope to be able to decide between the normal and inverted hierarchies, due to Δm_{32}^2 , using future off-axis long-baseline accelerator neutrino experiments (see Section 2.7.3).

In addition to the possible hierarchies of the neutrino mass eigenstates, there are two types of model for the absolute values of neutrino masses. In hierarchical models the magnitude of each mass splitting is comparable to the mass of the larger mass eigenstate involved. In the case of a quasi-degenerate neutrino mass spectrum, neutrino masses are all of the same order of magnitude and so much larger than the mass splittings [40].

2.2 Neutrino Mass Measurements

The SM makes no predictions for the values of the neutrino masses and in the absence of any observations of right-handed neutrino states they are assumed to be massless. Direct searches for ν masses have only ever yielded upper limits. Neutrinos take part in weak interactions in their flavour eigenstates, so any direct measurement of neutrino mass will be tied to a particular flavour eigenstate. In general the mass eigenstates of neutrinos need not be commensurate with the flavour eigenstates. (*c.f. The Strong and Weak eigenstates of quarks.*) Unlike quarks no single mass eigenstate can be associated as being the dominant component of a flavour eigenstate due to the large mixing angles so far observed (see Section 2.3). Therefore flavour eigenstates must always be considered as

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i} |\nu_i\rangle \quad (2.30)$$

Many experiments have attempted to constrain the neutrino masses. Results are quoted according to the flavour of the neutrinos detected in the experiment, but the limits can no longer be thought of as limits on the masses of the weak neutrino eigenstates themselves. The current limits are summarised below. In each case the quantity constrained is

$$m_{\nu_\alpha}^{2(\text{eff})} = \sum_i |U_{\alpha i}|^2 m_{\nu_i}^2, \quad (2.31)$$

where ν_α is the neutrino flavour state observed. Here the sum is over all $m_{\nu_i}^2$ which are too close together to be resolved experimentally.

2.2.1 Mass Measurements Using ν_e

The strongest limits on the ' ν_e mass' are obtained by fitting the shape of the β -decay spectrum of tritium (^3H). The two current best limits from experiments of this type are from the Troitsk [41] and Mainz [42] groups who obtain limits of $m_{\nu_e}^{(\text{eff})} < 2.5 \text{ eV}$ and $m_{\nu_e}^{(\text{eff})} < 2.2 \text{ eV}$ respectively at the 95% C.L..

Interestingly, if the experimental resolution was better than $\Delta m_{ij}^2 = m_{\nu_i}^2 - m_{\nu_j}^2$, then m_{ν_i} and Δm_{ij}^2 could be determined from fitting the spectral anomaly which would result. Such precision is probably some way off however [43]. The KATRIN (*Karlsruhe Tritium Neutrino Mass Experiment*) is a next generation tritium β -decay experiment

designed to reach a sensitivity of 0.2 eV. It will be able to discriminate between so-called hierarchical and quasi-degenerate neutrino mass models (see Section 2.1.5). The experiment is scheduled to begin taking data in late 2008 [44].

2.2.2 Mass Measurements Using ν_μ

Measuring the momentum of the muon produced in pion decay ($\pi^+ \rightarrow \mu^+ \nu_\mu$) and fitting the observed spectrum in a similar way to that done for β -decay allows us to constrain $m_{\nu_\mu}^{2(\text{eff})}$. The most sensitive measurement made so far was made at the Paul Scherrer Institute in Zurich and obtained a limit of $m_{\nu_\mu}^{(\text{eff})} < 170 \text{ keV}$ at 90% C.L. [45].

2.2.3 Mass Measurements Using ν_τ

In the case of ν_τ the decays of the τ lepton can be used to determine $m_{\nu_\tau}^{(\text{eff})}$. The τ lepton can decay via a number of leptonic and semi-leptonic modes. Once again a fit to the shape of the spectra of the decay products can be used. The current best upper limit on $m_{\nu_\tau}^{(\text{eff})}$ of $m_{\nu_\tau}^{(\text{eff})} < 18.2 \text{ MeV}$ at 95% C.L. was obtained by the ALEPH collaboration by considering 3-prong and 5-prong decays of the τ lepton (where the term ‘prong’ refers to a charged particle detected in the tracking chambers of the experiment) [46].

2.2.4 Constraints From Astrophysics

The WMAP (*Wilkinson Microwave Anisotropy Probe*) makes detailed measurements of the cosmic microwave background. By fitting the parameters of the standard cosmological model to these data along with the 2dFGRS and Ly α forest data, a limit on sum of neutrino masses of $\sum_i m_{\nu_i} < 0.68 \text{ eV}$ (95% C.L.) [47] can be obtained.

2.3 An Overview Neutrino Oscillation Experiments

As mentioned above, neutrino detection is claimed on the basis of observations of particles produced in charge-current or neutral-current neutrino interactions. The fact that neutrino cross-sections are so small means that intense neutrino fluxes and large detectors are required in order to obtain useful event rates. Steps must be taken to minimise

experimental backgrounds from cosmic radiation. The best way to shield neutrino detectors from cosmic rays is to build them deep underground.

The non-Majorana part of the PMNS matrix (equation (2.20)) can be factorised into a product of three matrices as below:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{-i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} e^{i\alpha} & 0 & 0 \\ 0 & e^{i\beta} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (2.32)$$

The details of each of these matrices are determined by different experiments. The first and third parts have already been determined to some extent. The elements of the first matrix can be determined by atmospheric neutrino experiments (see Section 2.6) and experiments looking at neutrinos from accelerators (see Section 2.7). The elements of the third matrix can be determined from solar neutrino experiments (see Section 2.4) and reactor neutrino experiments (Section 2.5). As for the second matrix; at the time of writing no experiment has yet observed values of $\theta_{13} \neq 0$. As is clear from the second matrix a non-zero value of θ_{13} is required in order for observation of $\delta \neq 0$. The best limits on θ_{13} so far were obtained by the CHOOZ experiment (see Section 2.5.2). Two possible strategies for future experiments in this area are to:

1. Look for ν_e disappearance in ν_e flux (reactor neutrino experiments, discussed in Section 2.5.3)
2. Look for ν_e appearance in ν_μ flux (accelerator neutrino experiments, discussed in Section 2.7.2 and Section 2.7.3)

2.4 Solar Neutrino Experiments

2.4.1 Solar Neutrino Production

The Sun is a main-sequence star at a stage of stable hydrogen burning. The dominant chain of fusion reactions is the “ pp chain”, in which four protons fuse together to form

a helium nucleus. The combined effect of these reactions is:

Our best picture of the structure and dynamics of the solar core is given by the Solar Standard Model (SSM). Currently, the preferred SSM is BS05(OP) developed by Bahcall and Serenelli [48] [49]. The SSM is used to determine the rate of each fusion reaction and calculate the emitted flux of solar neutrinos. Figure 2.3 shows the reactions that form the pp chain, and the energy spectrum of solar neutrinos predicted by the SSM. According to the model, the pp chain accounts for over 90% of the neutrinos emitted by the sun. However, all these neutrinos have energies below 0.42 MeV. Other processes in the sun produce neutrinos of higher energies. For example, the interactions ${}^7\text{Be} + e^- \rightarrow {}^7\text{Li}^* + \nu_e$ (7% of flux) and $p + e^- + p \rightarrow d + \nu_e$ (0.2% of flux) produce mono-energetic neutrinos with energies of 0.86 MeV and 1.44 MeV respectively. Despite only making up 0.01% of the flux the decay ${}^8\text{B} \rightarrow {}^8\text{Be}^* + e^+ + \nu_e$ produces neutrinos with energies up to 14.1 MeV, making it experimentally advantageous to detect them.

As mentioned in Chapter 1, since the first measurements of solar neutrinos by the Homestake experiment [21], a deficit in the observed solar neutrino flux compared to that predicted by the SSM caused scepticism about the SSM particularly in the neutrino physics community. Satellite based measurements of the speed of sound in the sun (helioseismology) from the mid-1990s onward have verified many of the predictions made by the SSM [50]. Now it is generally accepted that new physics, such as neutrino oscillations, must be invoked in order to explain the neutrino deficit.

2.4.2 Matter Effects For Solar Neutrinos

The propagation of solar neutrinos through solar matter provides a good illustration of the matter effects mentioned in Section 2.1.2. When combined, the information on atmospheric neutrino oscillations (see Section 2.6) and results from short baseline reactor neutrino experiments (see Section 2.5) show that, in the absence of sterile neutrinos (ν_s), only two neutrino eigenstates, ν_1 and ν_2 , are significantly involved in the evolution of solar neutrinos. Correspondingly, only two flavours are involved: ν_e and ν_x (some

Figure 2.3: The pp chain of solar fusion reactions and the solar neutrino energy spectrum as predicted by the Standard Solar Model [48].

combination of ν_μ and ν_τ). The Hamiltonian is then a 2×2 matrix in $\nu_e - \nu_x$ space [43].

$$\mathbf{H} = \mathbf{H}_V + \mathbf{H}_M(r) = \frac{\Delta m_\odot^2}{4E} \begin{pmatrix} -\cos(2\theta_\odot) & \sin(2\theta_\odot) \\ \sin(2\theta_\odot) & \cos(2\theta_\odot) \end{pmatrix} + \begin{pmatrix} V(r) & 0 \\ 0 & 0 \end{pmatrix} \quad (2.33)$$

Here ν_2 is *defined* to be heavier than ν_1 , so $\Delta m_\odot^2 (= m_2^2 - m_1^2)$ is positive. $\mathbf{H}_V$ is the Hamiltonian in a vacuum and $\mathbf{H}_M$ is the modification due to matter. In $\mathbf{H}_M$, $V(r) = \sqrt{2}G_F N_e(r)$, where G_F is the Fermi constant and $N_e(r)$ is the electron density of the solar matter as a function of the radial distance from the centre of the sun.

As equation (2.25) shows, in a vacuum two neutrino oscillations cannot distinguish between $\theta_\odot$ and $\theta'_\odot = \pi/2 - \theta_\odot$. The two situations are physically distinct however. If $0 < \theta_\odot < \pi/4$ then ν_1 is mostly composed of ν_e , while if $\pi/4 < \theta_\odot < \pi/2$ then ν_1 is more ν_x than ν_e . The degeneracy is broken by matter effects, as can be seen from equation (2.33). The sign of $\cos(2\theta_\odot)$ relative to $V(r)$ differs in the two cases. This means the matter effects differ also. Experiments such as SNO (Section 2.4.4) fit $\Delta m_\odot^2$ versus $\tan^2 \theta_\odot$ in order to distinguish the two regimes.

2.4.3 Radiochemical Experiments

Radiochemical experiments exploit the absorption of ν_e by nuclei followed by the subsequent decay of those nuclei via orbital electron capture. The Auger electrons produced in this process are then counted.

Homestake

The Homestake chlorine experiment in the USA was the first to attempt to observe solar neutrinos. Homestake relied on the reaction:

which has an energy threshold of 814 keV. This means the solar neutrinos captured by the experiment are predominantly from the ${}^8\text{B}$ decay chain. The SSM predicts Homestake should observe 8.5 ± 1.8 Solar Neutrino Units ($1 \text{ SNU} \equiv 1 \text{ interaction per } 10^{36} \text{ atoms per second}$) [51]. The experiment was located 1478 m underground in the Homestake Gold Mine in order to shield it from cosmic muons. Homestake used 615 t of liquid

tetrachloroethylene (C_2Cl_4) as its source of chlorine. This amount of chlorine results in 1 interaction per day on average. The ^{37}Ar was extracted from the tetrachloroethylene by periodically flushing the experiment with helium. Typically this was done after an exposure time of two to three times the half-life of ^{37}Ar (34.8 days). This helium was then passed through a charcoal trap to separate out any ^{37}Ar atoms. Proportional counters were then used to measure the 2.8 keV Auger electrons which are emitted when ^{37}Ar undergoes K-capture.

Initial results from 1968 showed no events above background, but set an upper limit on the solar neutrino flux of 3 SNU [20]. An improved electronics system, which discriminated signal from background using the rise time of the pulses from proportional counters, was introduced in 1970. From then onward a finite solar-neutrino flux was observed. Homestake took 108 runs between 1970 and 1994. Combining all these results, Homestake observed a solar neutrino capture rate of $2.56 \pm 0.16 \pm 0.16$ SNU [21]; only around 30% of the rate predicted by the SSM.

Gallium Experiments

In addition to the Homestake experiment three other radiochemical experiments have published results on the solar neutrino flux. These experiments, GALLEX (succeeded by GNO from 1998 onward) at Gran Sasso in Italy and SAGE at Baksan in Russia, use the following reaction of gallium nuclei:

$$\nu_e + {}^{71}\text{Ga} \rightarrow {}^{71}\text{Ge} + e^- \quad (2.35)$$

which has an energy of 233 keV. This allows observations of solar neutrinos emitted by the primary pp interaction. Consequently the flux of neutrinos available to interact is much higher and thus the gallium experiments can be much smaller (30-60 t). GALLEX/GNO uses gallium in a hydrochloric solution, while SAGE uses liquid gallium. The germanium is extracted and counted in a similar fashion to the argon produced in Homestake. The expected flux from the SSM for these experiments is 131^{+12}_{-10} SNU [51]. GALLEX observed a solar neutrino flux of $77.5 \pm 6.2^{+4.3}_{-4.7}$ SNU [52]. The results from SAGE are similar. SAGE observed a flux of $70.8^{+5.3+3.7}_{-5.2-3.2}$ SNU [53]. Combined results of GALLEX and GNO were published in 2005; they observed a flux of $69.3 \pm 4.1 \pm 3.6$ SNU [54]. All three observations again show a significant deficit in the solar neutrino flux.

2.4.4 Water Čerenkov Experiments

The big difference between the water Čerenkov experiments and the radiochemical experiments is that they can measure neutrino interactions in real-time. These experiments detect the Čerenkov light from electrons recoiling as a consequence of elastic scattering of neutrinos:

$$\nu_x + e^- \rightarrow \nu_x + e^- \quad (2.36)$$

All three (known) neutrino flavours contribute to this process via NC scattering, but for ν_e there is a large additional contribution from CC scattering. This results in the ν_e cross-section being around six times that of ν_μ and ν_τ . The close correlation between the direction of the recoil electrons and the direction of the incident neutrino assists the separation of solar neutrinos from the natural background radioactivity.

Kamiokande and Super-Kamiokande

Kamiokande-II experiment ran from 1987-1995 and was located 1000 m underground in the Kamioka mine in the Japanese Alps. The 2400 mwe (metres of water equivalent) minimum overburden over the experiment helped to reduce backgrounds. A total of 948 photo-multiplier tubes (PMTs) were used to collect the Čerenkov light and covered 20% of the total surface area of the 2140 t fiducial volume. The threshold for detecting the recoil electrons in Kamiokande-II was 7 MeV. Consequently the experiment was only sensitive to solar neutrinos from the ^8B solar interaction, plus a small contribution from the hep fusion reaction. Kamiokande measured a ^8B ν flux of $2.80 \pm 0.19 \pm 0.33 \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$ [55], significantly less than the SSM prediction of $5.69(1 \pm 0.16) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$ [48].

The Super-Kamiokande experiment (SK) is the large scale successor to Kamiokande and has been collecting data since 1996. SK uses 50 kt of ultra-pure water (22.5 kt fiducial) to detect neutrino interactions. The large mass of the SK detector provides a high neutrino interaction rate. The fiducial volume is instrumented with 11000 PMTs. The outer detector, containing 27.5 kt of ultra-pure water and instrumented with 2000 PMTs, is used to veto through-going cosmic muons. Significant improvements in water purity over that achieved for the Kamiokande experiment have allowed SK to identify solar neutrinos down to energies of 5 MeV.

In November 2001, SK suffered an accident in which a large number of PMTs were destroyed in chain reaction as the shock wave from the concussion of each imploding tube cracked its neighbours. Within a year the detector had been rebuilt by redistributing the intact PMTs, but only half the original number of PMTs remained. Between June 2005 and July 2006 the detector was restored to its original state. Before November 2001 the experiment is known as Super-Kamiokande-I (SK1). After the accident it is known as Super-Kamiokande-II. The fully restored experiment is known as Super-Kamiokande-III.

The final results for the 1496 day sample of SK1 have been published [56]. The flux of ^8B solar neutrinos observed by SK1 is $2.35 \pm 0.02(\text{stat.}) \pm 0.08(\text{sys.}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$. Again this is significantly less than the SSM prediction of $5.69(1 \pm 0.16) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$ [48].

Sudbury Neutrino Observatory

The Sudbury Neutrino Observatory (SNO) began data taking in 1999. It is located deep underground (2092 m) in Sudbury, Canada. SNO uses 1 kt of ultra-pure heavy water (D_2O) contained in a spherical acrylic vessel and surrounded by an ultra-pure H_2O shield. Using D_2O gives the experiment a unique sensitivity to all three neutrino flavours. Three different neutrino interactions are detectable at SNO:

- Charged current (CC) interactions ($\nu_e + d \rightarrow e^- + p + p$), sensitive only to ν_e .
- Neutral current (NC) interactions ($\nu_\ell + d \rightarrow \nu_\ell + p + n$), equally sensitive to all three neutrino flavours.
- Elastic scattering (ES) interactions ($\nu_\ell + e^- \rightarrow \nu_\ell + e^-$), sensitive to all three neutrino flavours, but with a reduced cross-section for ν_μ and ν_τ .

By combining the measurements of each interaction type, separated measurements of the ν_e and ν_μ/ν_τ fluxes can be obtained. The energy threshold of SNO is currently 5 MeV, making it predominantly sensitive to ^8B neutrinos.

The SNO experiment is being performed in three phases. Each phase enhances the experiment's sensitivity to NC interactions. In phase I NC events were detected by observing the emission of 6.25 MeV of energy as γ -rays following the capture of an emitted neutron by a deuteron. In phase II, 2 t of salt (NaCl) was added to the heavy-water. This increased the neutron capture efficiency, the γ -ray multiplicity and total energy emitted in NC interactions. Neutron capture by a chloride ion emits 8.6 MeV

of energy as γ -rays. The third phase is currently underway. Phase III uses an array of ^3He proportional counters deployed in the ultra-pure D_2O to detect neutrons directly ($^3\text{He} + n \rightarrow ^1\text{H} + ^3\text{H}$).

The measured fluxes for each interaction from phase II of the experiment [57] are as follows:

$$\begin{aligned}\phi_{CC} &= 1.68^{+0.06}_{-0.06}(\text{stat.})^{+0.08}_{-0.09}(\text{sys.}) \times 10^6 \text{cm}^{-2}\text{s}^{-1} \\ \phi_{ES} &= 2.35^{+0.22}_{-0.22}(\text{stat.})^{+0.15}_{-0.15}(\text{sys.}) \times 10^6 \text{cm}^{-2}\text{s}^{-1} \\ \phi_{NC} &= 4.94^{+0.21}_{-0.21}(\text{stat.})^{+0.38}_{-0.34}(\text{sys.}) \times 10^6 \text{cm}^{-2}\text{s}^{-1}\end{aligned}$$

The measurement of ϕ_{NC} is a mixing-independent result and therefore tests solar models and shows reasonable agreement with the SSM(BP05) predictions [48]. The phase II result is shown in Figure 2.4. The non- ν_e component of the flux, $\phi_{\mu\tau}$, is plotted against the ν_e component, ϕ_e . The intersection of the three bands allows the ν_μ/ν_τ flux to be resolved and shows the consistency of the three measurements. This provides compelling evidence that a non- ν_e component of the flux exists, and provides direct evidence that neutrino oscillations, rather than neutrino decoherence or neutrino decay are the cause of the observed deficits in the solar ν_e flux.

The parameters from equation (2.25) that control long range $\nu_e \rightarrow \nu_\mu/\nu_\tau$ oscillations, (θ_{12} and Δm_{21}^2) can be determined using the SNO measurements. Any fit to the SNO data must take account of the matter effects which are predicted to be particularly strong while the neutrinos propagate through the sun from the solar core (see Section 2.1.2). Figure 2.5 shows the confidence limits for $\tan^2 \theta_{12}$ and Δm_{21}^2 calculated from the SNO phase II data, assuming two-flavour oscillations. The fit is consistent with more than one region of parameter space, but favours “Large Mixing Angle” (LMA) oscillations. The best fit point is $\Delta m_{21}^2 = 5.0^{+6.2}_{-1.8} \times 10^{-5} \text{eV}^2$ ($m_2 > m_1$), $\tan^2 \theta_{12} = 0.45^{+0.11}_{-0.10}$ [57]. The SNO collaboration has also carried out a global analysis of all solar neutrino experiments as shown in Figure 2.6. This strongly favours LMA oscillations, excluding all other regions of parameter space at 99% confidence. The best fit point is: $\Delta m_{21}^2 = 6.5^{+4.4}_{-2.3} \times 10^{-5} \text{eV}^2$ ($m_2 > m_1$), $\tan^2 \theta_{12} = 0.45^{+0.09}_{-0.08}$ [57]. As discussed in Section 2.4.2, the range of allowed $\tan^2 \theta_{12}$ values indicate that the ν_1 mass eigenstate is mostly composed of ν_e .

Figure 2.4: The results of phase II of the SNO experiment [57]. Flux of $\mu + \tau$ neutrinos versus flux of electron neutrinos. CC, NC and ES flux measurements are indicated by the filled bands. The total ${}^8\text{B}$ solar neutrino flux predicted by the Standard Solar Model [48] is shown as dashed lines, and that measured with the NC channel is shown as the solid band parallel to the model prediction. The narrow band parallel to the SNO ES result corresponds to the Super-Kamiokande result in [58]. The intercepts of these bands with the axes represent the $\pm 1\sigma$ uncertainties. The non-zero value of $\phi_{\mu\tau}$ provides strong evidence for neutrino flavour transformation. The point represents ϕ_e from the CC flux and $\phi_{\mu\tau}$ from the NC-CC difference with 68%, 95%, and 99% C.L. contours included.

Figure 2.5: The confidence limits for the oscillation parameters Δm_{21}^2 and θ_{12} obtained from an analysis of the results from the phase II of the SNO experiment [57]. The best fit point is indicated by a star.

2.5 Reactor Neutrino Experiments

2.5.1 KamLAND

KamLAND (*Kamioka Liquid scintillator Anti-Neutrino Detector*) [59] is a 1 kt ultra-pure liquid scintillator detector located in the Kamioka mine, in the cavern previously occupied by the Kamiokande experiment. Figure 2.7 shows the KamLAND experimental setup. An inner detector holds 1 kt of ultra-pure liquid scintillator. An outer detector holding 3200 t of water surrounds the inner detector. This acts to shield it from radioactivity in the rock, as well as allowing cosmic muons to be vetoed.

The ultimate goal of the experiment is to observe ${}^7\text{Be}$ solar neutrinos with a much lower energy threshold. This will require a reduction in the background from radioactivity in the detector. It is currently in its initial phase concentrating on observations of $\bar{\nu}_e$ emitted from 52 reactor cores located at 16 nuclear power reactors, with a flux-weighted average baseline of ~ 180 km. KamLAND uses the inverse β -decay reaction $\bar{\nu}_e + p \rightarrow e^+ + n$ to detect reactor $\bar{\nu}_e$. The prompt scintillation light produced by the e^+ , combined with a delayed coincidence of a 2.2 MeV γ -ray from neutron capture

Figure 2.6: (a) Global neutrino oscillation analysis using only solar neutrino data, and (b) including KamLAND 766 t – year data. The stars indicate the best-fit parameters for the two datasets. Taken from [57].

Figure 2.7: A cut-away view of the KamLAND experiment. Taken from [59].

on a proton, reduces backgrounds considerably.

The expected $\bar{\nu}_e$ flux and energy spectrum from each reactor can be calculated using the information on the initial components of the fuel rods and hourly thermal power measurements along with existing data on β -decays. This removes the need for a series of ‘near detectors’ next to the reactors.

The KamLAND Collaboration has published an analysis of data taken in the period March 2003 - January 2004 [60]. The expectation is 365.2 ± 23.7 events in the absence of oscillations, however only 258 events are observed. This corresponds to a measured reactor $\bar{\nu}_e$ flux which is $0.658 \pm 0.044(\text{stat.}) \pm 0.047(\text{sys.})$ of the expected unoscillated value. Figure 2.8 shows strong evidence for spectral distortion of the $\bar{\nu}_e$ flux. There is a clear energy dependent disappearance of $\bar{\nu}_e$. The best fit to the data shown in Figure 2.8 was obtained by fitting the data to a two flavour oscillation model. It should be noted that the MSW effect has no significant influence on the propagation of reactor $\bar{\nu}_e$ through the Earth. This simplifies the determination of oscillation parameters by reducing systematic errors. Figure 2.9(a) shows the confidence limits obtained by KamLAND for Δm_{21}^2 and $\tan^2 \theta_{12}$. The confidence limits obtained from the global analysis of solar neutrino data are also shown. The KamLAND confidence limits are consistent with the LMA oscillation regime and there is clear overlap between the regions of pa-

Figure 2.8: Observed energy spectrum of prompt e^+ from $\bar{\nu}_e$ interactions detected by KamLAND. Also shown are the predicted spectrum in the absence of oscillations and the best-fit to a two flavour oscillations model. Taken from [60].

parameter space allowed by the two analyses. KamLAND also published a joint analysis of KamLAND and solar neutrino data. Figure 2.9(b) shows the confidence limits obtained by combining the two datasets. The best fit point is: $\Delta m_{21}^2 = 7.9_{-0.5}^{+0.6} \times 10^{-5} \text{ eV}^2$, $\tan^2 \theta_{12} = 0.40_{-0.07}^{+0.10}$ [60].

KamLAND also considered neutrino decay (see Section 2.1.4) and neutrino decoherence (see Section 2.1.3) hypotheses. These hypotheses have different L_0/E dependencies, as shown in Figure 2.10. The best-fit points were found to be $(\sin^2 \theta, m/c\tau) = (1.0, 0.011 \text{ MeV/km})$ for decay and $(\sin^2 2\theta, \gamma^0) = (1.0, 0.030 \text{ MeV/km})$ for decoherence. It was found that the data was not well described by either hypothesis, as is clear from Figure 2.10.

2.5.2 CHOOZ

The CHOOZ experiment was a short baseline reactor neutrino experiment located approximately 1 km from the CHOOZ nuclear power plant in northern France. CHOOZ measured the flux of $\bar{\nu}_e$ emitted from the reactor using a 5 t liquid scintillator detector. As with KamLAND, CHOOZ used the inverse β -decay reaction $\bar{\nu}_e + p \rightarrow e^+ + n$ to detect reactor $\bar{\nu}_e$. The signal was prompt scintillation from the e^+ , followed by a delayed

(a) Allowed region of oscillation parameter space from analysis of KamLAND data (indicated by shaded regions) and global analysis of solar neutrino data (indicated by lines)

(b) Allowed region of oscillation parameter space from combined analysis of KamLAND and solar neutrino data.

Figure 2.9: KamLAND neutrino oscillation confidence limits. Figures taken from [60].

Figure 2.10: Ratio of the observed $\bar{\nu}_e$ spectrum to the expectation for no-oscillation versus L_0/E [60]. The curves show the expectation for the best-fit oscillation, best-fit decay and best-fit decoherence models taking into account the individual time-dependent flux variations of all reactors and detector effects. The data points and models are plotted with $L_0=180$ km, as if all anti-neutrinos detected in KamLAND were due to a single reactor at this distance.

coincidence of a 2.2 MeV γ -ray. The CHOOZ detector was loaded with gadolinium to enhance both the efficiency of neutron capture and the energy of the resulting photon emissions. CHOOZ sought to probe the sub-dominant mode of $\nu_e \leftrightarrow \nu_\mu/\nu_\tau$ oscillations described by equation (2.27). This mode depends on a non-zero value of θ_{13} , which is a crucial ingredient for potentially observing $\mathcal{CP}$ violation in the lepton sector. The CHOOZ experiment has established the tightest bound set so far on the value of θ_{13} . The signal of oscillations at CHOOZ is a deficit of $\bar{\nu}_e$ events and a distorted $\bar{\nu}_e$ spectrum with respect to the unoscillated predictions. As can be seen in Figure 2.11(a), no significant distortion or deficit in the $\bar{\nu}_e$ spectrum was observed [61]. Figure 2.11(b) shows the confidence limits obtained from a two flavour oscillation analysis of the CHOOZ data. The 90% confidence limit is that $\sin^2 2\theta_{13} < 0.17$ [61].

2.5.3 Future Experiments

A proposed successor to the CHOOZ experiment called Double CHOOZ is currently under development. The Double CHOOZ proposal consists of an upgrade/replacement of the CHOOZ detector (the ‘far detector’), followed by the construction of another identical ‘near detector’ 280 m from the nuclear cores at the Chooz site. The experiment will run with only the far detector for 1.5 years. This will enable the experiment to reach a limit of $\sin^2 2\theta_{13} < 0.08$ if no oscillation signal is detected. The near detector will then be installed in a 45 m deep shaft. This second phase will considerably reduce the overall systematic uncertainties and could limit $\sin^2 2\theta_{13}$ to between 0.022 and 0.030 (for $\Delta m_{31}^2 = 3.5 - 2.5 \times 10^{-3} \text{ eV}^2$).

2.6 Atmospheric Neutrino Experiments

As mentioned in Chapter 1, the *atmospheric neutrino problem* was discovered due to the importance of atmospheric neutrino interactions as a background for proton decay searches. The proton decay search experiments discovered a deficit in the number of observed atmospheric ν_μ interactions compared to expectations, but no discrepancy in the atmospheric ν_e flux.

(a) CHOOZ observed prompt e^+ spectrum (points) and Monte-Carlo prediction in the absence of oscillations (solid line).

(b) CHOOZ confidence limits on $\sin^2 2\theta_{13}$ and Δm^2_{32} . The region of parameter space to the right of the line is excluded at 90% confidence. The two lines correspond to the *Feldman and Cousins* (FC) and *strong Confidence Level* (sCL) methods of obtaining the 90% confidence level.

Figure 2.11: CHOOZ results, taken from [61].

2.6.1 Atmospheric Neutrino Production

Atmospheric neutrinos are produced when cosmic-rays strike the Earth's atmosphere. Protons dominate the primary cosmic-ray flux (95%), with helium nuclei comprising (4.5%) at energies above 2 GeV/nucleon. C, N and O plus heavier elements make up the remaining 0.5% of the flux. Helium and the heavier nuclei carry $\sim 20\%$ of the incident nucleons and thus contribute about 20% of the neutrino flux. When high energy cosmic-rays strike the atmosphere they produce a cascade of secondary pions and kaons. The subsequent decays of which result in the emission of a high flux of neutrinos, as described in the equations below:

$$\begin{aligned}
 p, He + N &\rightarrow X + \pi^\pm/\kappa^\pm \\
 \pi^\pm/\kappa^\pm &\rightarrow \mu^\pm + \nu_\mu(\bar{\nu}_\mu) \\
 \mu^\pm &\rightarrow e^\pm + \nu_e(\bar{\nu}_e) + \bar{\nu}_\mu(\nu_\mu)
 \end{aligned}$$

For secondaries $\lesssim 2$ GeV the ratio

$$R = \frac{N(\nu_\mu + \bar{\nu}_\mu)}{N(\nu_e + \bar{\nu}_e)} \quad (2.37)$$

is expected to be ~ 2 . Higher energy secondaries, particularly μ , are likely to strike the Earth's surface before they decay thus R will increase with energy.

The calculations required to predict the atmospheric neutrino flux are complex with many factors to consider [62][63][64]. The experimentally determined primary cosmic-ray compositions and energy spectra are used as a starting point. Primaries are propagated through a 3-dimensional model of the Earth's magnetic field. This produces significant variations in the neutrino flux depending on latitude and longitude (*i.e.* the atmospheric neutrino fluxes at Kamioka, Japan and Soudan, Minnesota have significant differences). Where and in what quantities the primaries interact is determined using a model of the density structure of Earth's atmosphere. The absolute flux of atmospheric neutrinos can only be predicted to an accuracy of $\sim 20\%$ due to uncertainties in the production of secondaries, the charged-pion multiplicities and the resulting pion momentum spectrum. However, because the fluxes of ν_μ and ν_e are correlated the error in R is only $\sim 5\%$ between 0.1 GeV and 10 GeV. The simulation of the atmospheric neutrino flux is discussed in more detail in Chapter 6.

2.6.2 Soudan 2

Soudan 2 used a 1 kt calorimeter detector composed of active drift tubes within a passive steel structure to sample the particle tracks and showers produced by neutrino interactions [65]. An active shield containing layers of proportional tubes was constructed around the detector to tag particles entering or exiting the detector. The signature that an atmospheric neutrino interaction had occurred in the detector was an event with a contained interaction vertex. The ν_μ CC and ν_e CC samples were separated by analysing the topology of the events. Muon tracks were used to tag ν_μ CC events and electromagnetic showers tagged ν_e CC events. Soudan 2 observed a deficit in the ν_μ flux relative to expectation. The ratio in equation (2.38) was calculated for data and for Monte Carlo simulations. A ratio of ratios was then calculated. Using data from a fiducial exposure of 5.9 ktyr this ratio was found to be:

$$\mathfrak{R} = \frac{R_{Data}}{R_{MC}} = 0.69 \pm 0.10(\text{stat.}) \pm 0.06(\text{sys.}) [66]. \quad (2.38)$$

The probability of the no oscillation hypothesis was found to be 5.8×10^{-4} from an unbinned analysis of the neutrino L/E distribution.

2.6.3 Super-Kamiokande

As well as studying solar neutrinos (see Section 2.4.4), the Super-Kamiokande (SK) experiment has also studied atmospheric neutrinos. The 22.5 kt fiducial volume of the experiment has enabled it to record the largest sample of atmospheric neutrinos of any experiment so far undertaken. SK analysed the neutrino events as a function of both energy and direction. This was possible as for the energies considered the emitted lepton closely follows the path of the incident neutrino, giving an accurate measurement of neutrino propagation distance.

In 2004 SK published an L/E analysis of the SK-I dataset [67]. Only regions of the data with a mean L/E resolution of $< 70\%$ were used. This vetoed events pointing close to the horizon, where the propagation distance varies rapidly with zenith angle, and low energy events, where the mean angle between the neutrino and muon is large. Figure 2.12 shows a clear distortion in the ratio of the observed L/E distribution to the unoscillated expectation. For low values of L/E there is good agreement with expectation, but for larger values of L/E the ratio drops to ~ 0.5 . A dip is observed at $L/E \simeq 500 \text{ km/GeV}$,

Figure 2.12: The ratio of the data taken by Super-Kamiokande to their MC expectation without neutrino oscillation (points) as a function of the reconstructed L/E together with the best-fit expectation for 2-flavor $\nu_\mu \leftrightarrow \nu_\tau$ oscillations (solid line). The error bars are statistical only. Also shown are the best-fit expectation for neutrino decay (dashed line) and neutrino decoherence (dotted line) [67].

the second oscillation maximum is not visible due to the intrinsic L/E resolution of the experiment. The ν_e flux was in good agreement with expectations. The data were interpreted as atmospheric $\nu_\mu \leftrightarrow \nu_\tau$ oscillations. Figure 2.13 shows the contour plot of the allowed oscillation parameter regions. The 90% C.L. allowed parameter region was obtained as $1.9 \times 10^{-3} \text{ eV}^2 < \Delta m_{32}^2 < 3.0 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta > 0.90$. The best fit in the physical-region was at $(\sin^2 2\theta, \Delta m_{32}^2) = (1.00, 2.4 \times 10^{-3} \text{ eV}^2)$.

In 2005 Super-Kamiokande published a combined analysis of Fully Contained (FC), Partially Contained (PC) and upward-going muon (UPMU) atmospheric neutrino data [68]. In FC events the neutrino interaction vertex and emitted particles stay within the fiducial volume. In PC events the interaction vertex is contained, but one or more of the emitted particles leaves the fiducial volume. UPMU events are caused by neutrinos which interact in the rock below the detector, producing a muon which enters the detector. Figure 2.14 shows the confidence limits obtained by this analysis and how the contours compare to those obtained by the L/E analysis. For the combined analysis the 90% confidence limits are: $1.5 \times 10^{-3} < \Delta m_{32}^2 < 3.4 \times 10^{-3} \text{ eV}^2$, $\sin^2 2\theta_{23} > 0.92$. Both results are consistent. However, the combined analysis gives stronger limits on $\sin^2 2\theta$, while the L/E analysis allows a smaller region of Δm_{32}^2 . The best fit point for the

Figure 2.13: Super-Kamiokande L/E analysis. 68, 90 and 99% C.L. allowed oscillation parameter regions for 2-flavor $\nu_\mu \leftrightarrow \nu_\tau$ oscillations obtained by the present analysis [67].

combined analysis is: $\Delta m_{32}^2 = 2.1 \times 10^{-3} \text{ eV}^2$, $\sin^2 2\theta_{23} = 1.0$.

A three flavour neutrino oscillation analysis of atmospheric neutrinos was published by Super-Kamiokande in 2006 [69]. The analysis assumed “one mass scale dominance” (*i.e.* $|m_2^2 - m_1^2| \ll |m_3^2 - m_{1,2}^2|$), so they take $\Delta m_{21}^2 = 0$. This reduces the number of parameters to three: two mixing angles (θ_{23} , θ_{13}) and one mass squared difference Δm^2 . The same FC, PC and UPMU samples were used as in the 2005 publication, along with an additional multi-ring e^- -like sample, selected using a likelihood method. This was added to increase e^- statistics and improve sensitivity to θ_{13} . The best fit point was found to be $(\Delta m^2, \sin^2 \theta_{23}, \sin^2 \theta_{13}) = (2.5 \times 10^{-3}, 0.5, 0.0)$. Both normal and inverted mass hierarchy hypotheses (see Section 2.1.5) agree with the data. For a normal mass hierarchy the region of $\sin^2 \theta_{13} < 0.14$ and $0.37 < \sin^2 \theta_{23} < 0.65$ is allowed at the 90% confidence level. For an inverted mass hierarchy the allowed region is slightly wider; $\sin^2 \theta_{13} < 0.27$ and $0.37 < \sin^2 \theta_{23} < 0.69$ is allowed at the 90% confidence level. The upper limit on θ_{13} is consistent with that obtained by CHOOZ.

The most recent analysis published on the SK-I dataset is a search for ν_τ appearance in the atmospheric neutrino flux [70]. This is a difficult analysis to perform for many reasons. Firstly, the neutrino energy threshold for τ production is 3.5 GeV. The atmospheric flux above this level is relatively low. Over the 1489 day live-time of the SK-I experiment, the MC expectation is for $78 \pm 26(\text{sys.})$ events to be observed (for

Figure 2.14: Comparison of confidence levels for SK combined and L/E analyses. 68, 90 and 99% C.L. allowed oscillation parameter regions are shown [68].

$\Delta m_{32}^2 = 2.4 \times 10^{-3} \text{ eV}^2$ and maximal mixing). Secondly, the τ lifetime is only 290 fs. Any τ produced decays immediately into one of many possible final states involving e , μ , or one of more pions. Water Čerenkov detectors are not suited for identifying individual CC ν_τ interactions with their multiple Čerenkov rings with no easily identified leading lepton. Likelihood and neural network techniques were used to identify ν_τ CC interactions in the data. A best fit ν_τ appearance signal of $138 \pm 48(\text{stat.})_{-32}^{+15}(\text{sys.})$ events was obtained. The hypothesis of no ν_τ appearance is disfavoured by 2.4 sigma.

2.6.4 MINOS

The MINOS experiment, discussed in more detail in the following chapters, has made observations of atmospheric neutrinos using its 5.42 kt Far Detector, located in the Soudan Mine, Minnesota. The Far Detector is magnetised to an average value of 1.3 T by a 15 kA current loop. This enables the separate observation of ν_μ and $\bar{\nu}_\mu$ using the curvature of the observed μ . In 2006 MINOS published an atmospheric neutrino analysis based on an exposure of 418 days [29]. A total of 107 candidate contained-vertex neutrino

interactions were observed compared to an expectation of 127 ± 13 events in the absence of oscillations.

The ratio of upward-going to downward-going events in the data is compared to the Monte Carlo expectation in the absence of oscillations:

$$\frac{R_{\text{up/down}}^{\text{data}}}{R_{\text{up/down}}^{\text{MC}}} = 0.62_{-0.14}^{+0.19}(\text{stat.}) \pm 0.02(\text{sys.}).$$

An extended maximum likelihood analysis of the observed L/E distributions excluded the no oscillations hypothesis at the 98% C.L., with a best fit value of $(\Delta m_{32}^2, \sin^2 2\theta_{23}) = (1.3 \times 10^{-3} \text{ eV}^2, 0.90)$. The 90% confidence limits were $7.0 \times 10^{-5} \text{ eV}^2 < \Delta m_{32}^2 < 5.0 \times 10^{-2} \text{ eV}^2$, $\sin^2 2\theta_{23} > 0.2$. MINOS found that it was possible to identify the charge of 52 of the 107 selected events; 18 were identified as $\bar{\nu}_\mu$ candidates and 34 as ν_μ candidates. This gave an observed $\bar{\nu}_\mu$ to ν_μ ratio of $0.53_{-0.15}^{+0.21}(\text{stat.}) \pm 0.15(\text{sys.})$. The fraction of $\bar{\nu}_\mu$ events in the data was compared to the MC expectation assuming neutrino and anti-neutrinos oscillate in the same manner, giving:

$$\frac{R_{\bar{\nu}_\mu/\nu_\mu}^{\text{data}}}{R_{\bar{\nu}_\mu/\nu_\mu}^{\text{MC}}} = 0.96_{-0.27}^{+0.38}(\text{stat.}) \pm 0.15(\text{sys.}).$$

2.7 Long-baseline Accelerator Neutrino Experiments

2.7.1 K2K

The K2K (KEK-to-Kamioka) long-baseline accelerator experiment [71], uses a ν_μ beam to study neutrino oscillations. The neutrino beam is manufactured by the KEK accelerator facility in Tsukuba, Japan. A beam of 12 GeV protons from the synchrotron at KEK is fired onto a fixed aluminium target. This produces an intense flux of secondary pions and kaons. The positively charged secondaries are focused by a pair of magnetic horns and then directed into a 200 m long decay pipe where they decay to produce a 98% pure ν_μ beam with a mean energy of 1.3 GeV. The initial unoscillated ν_μ spectrum is measured at a distance of 300 m from the proton target using a 1 kt water Čerenkov detector and a system of fine-grained detectors. The beam travels 250 km through the Earth and is then sampled by the Super-Kamiokande detector. The beam is produced

in pulses, enabling neutrino interactions from the beam in SK to be tagged using timing information. The oscillation signal is a distortion of the neutrino energy spectrum measured at SK compared to the expected spectrum obtained by extrapolating the spectrum measured by the near detectors out to the Kamioka site.

K2K has published an analysis of all beam data taken between June 1999 and November 2004 [72]. K2K measure a deficit of events at SK, observing 112 events compared to an expectation of $158.1^{+9.2}_{-8.6}$ events in the absence of oscillations. A distortion of the energy spectrum is also seen in a sample of 58 single-ring muon-like events with reconstructed energies, as is clear from Figure 2.15. A two flavour oscillation analysis was applied to the data. It was found that the probability that the data were explained by the no oscillation hypothesis was 0.00015% (4σ). Figure 2.16 shows the confidence limits obtained by the K2K analysis and shows they are consistent with the L/E atmospheric neutrino analysis performed by the SK collaboration. The allowed region of Δm_{32}^2 for maximal mixing is $1.9 \times 10^{-3} < \Delta m_{32}^2 < 3.7 \times 10^{-3}$ at the 90% C.L. with a best fit point at $\Delta m_{32}^2 = 2.8 \times 10^{-3} \text{eV}^2$, $\sin^2 2\theta_{23} = 1.0$.

K2K has also searched for ν_e appearance in their beam [73]. In an analysis based on data taken between June 1999 and July 2001 (corresponding to an exposure of 4.8×10^{19} protons on target) one candidate event was observed, compared to an expected background in the absence of neutrino oscillations of 2.4 ± 0.6 events. The limit set by K2K is consistent with, but does not exceed the CHOOZ limit.

2.7.2 MINOS

The MINOS long-baseline accelerator experiment, like K2K, uses a ν_μ neutrino beam to study neutrino oscillations. The NuMI (Neutrinos at the Main Injector) beam is manufactured at Fermilab, Batavia, Illinois. A beam of 120 GeV protons from the Main Injector at Fermilab is fired onto a movable graphite target. This produces an intense flux of secondary pions and kaons. The positively charged secondaries are focused by a pair of magnetic horns and then directed into a 675 m long decay pipe where they decay to produce a ν_μ dominated neutrino beam with a mean energy of 1-3 GeV. The initial unoscillated ν_μ spectrum is measured at a distance of 1 km from the proton target using the 0.98 kt MINOS Near Detector. The beam travels 735 km through the Earth and is then sampled by the 5.42 kt MINOS Far Detector located in the Soudan iron mine, Northern Minnesota. Both MINOS detectors are steel-scintillator tracking-calorimeters with toroidal magnetic fields averaging 1.3 T. The beam is produced in

Figure 2.15: The reconstructed E_ν distribution for the K2K 1-ring μ -like sample. Points with error bars are data. The solid line is the best fit spectrum with neutrino oscillation and the dashed line is the expectation without oscillation. These histograms are normalised by the number of events observed (58) [72].

Figure 2.16: Comparison of K2K results with the SK atmospheric neutrino measurement [67]. Dotted, solid, dashed and dash-dotted lines represent 68%, 90%, 99% C.L. allowed regions of K2K and 90% C.L. allowed region from SK atmospheric neutrino, respectively [72].

Figure 2.17: The reconstructed E_ν spectrum at the MINOS Far Detector. Points with error bars are data. The dark solid line is the best fit spectrum with neutrino oscillation, lighter solid line and the dashed line are the expectations without oscillation for the ‘Beam Matrix’ and ‘ND Fit’ extrapolation methods. The last energy bin contains events with energies between 18 GeV and and 30 GeV [74].

pulses, enabling neutrino interactions from the beam in the Far Detector to be tagged using timing information. The oscillation signal is a distortion of the neutrino energy spectrum measured at the Far Detector compared to the expected spectrum obtained by extrapolating the spectrum measured by the Near Detector out to the Soudan mine site.

MINOS has published an analysis data from an exposure of 1.27×10^{20} protons on the NuMI target [74]. MINOS measure a deficit of events at their Far Detector, observing 215 events below 30 GeV compared to an expectation of 336.0 ± 14.4 events in the absence of oscillations. A distortion of the energy spectrum is also seen, as is clear from Figure 2.17. A two flavour oscillation analysis was applied to the data. Figure 2.18 shows the confidence limits obtained by the MINOS beam analysis and shows they are consistent with analyses performed by the SK and K2K collaborations. The allowed regions of $|\Delta m_{32}^2|$ and $\sin^2 2\theta_{23}$ are $2.31 \times 10^{-3} < |\Delta m_{32}^2| < 3.43 \times 10^{-3}$ and $\sin^2 2\theta_{23} > 0.82$ at the 90% C.L. with a best fit point at $|\Delta m_{32}^2| = 2.74_{-0.26}^{+0.44} \times 10^{-3} \text{ eV}^2$, $\sin^2 2\theta_{23} > 0.87$ at the 68% C.L..

Figure 2.18: Comparison of MINOS beam analysis confidence levels (using the ‘Beam Matrix’ method) with the K2K measurement (light dash-dotted line) [75], the SK atmospheric neutrino L/E analysis (light solid line) [67] and the main SK atmospheric neutrino analysis (light dashed line) [68]. The dark dashed and solid lines and light grey dashed lines represent 68%, 90% C.L. allowed regions of MINOS beam analysis and 90% C.L. allowed region from SK atmospheric neutrino, respectively [74].

2.7.3 Future Experiments

The OPERA (*O*scillation *P*roject with *E*mulsion-*t*Racking *A*pparatus) neutrino detector based in the Gran Sasso Laboratory (LNGS) has been constructed in the path of the CERN to LNGS (CNGS) high-energy neutrino beam. The aim of the experiment is to perform the first detector of neutrino oscillations in appearance mode ($\nu_\mu \rightarrow \nu_\tau$). OPERA consists of a lead/emulsion-film target and additional electronic detectors. The experiment has a base-line of 730 km. The average beam neutrino energy is ~ 17 GeV at the LGNS location. In August 2006 the first sample of neutrino events from the CNGS beam was taken successfully [76].

Two future long-baseline accelerator experiments T2K (Tokai to Kamioka) [77] and NO ν A (NuMI Off-Axis ν_e Appearance Experiment) [78] are based around existing infrastructure. NO ν A, as the name suggests, is an experiment which plans to build a far detector 12 km off-axis from the NuMI beam, near Ash River, Minnesota. This will give NO ν A an 810 km baseline, slightly longer than that of MINOS. Like MINOS, NO ν A will have a near and far detector with as similar structure as possible. The NO ν A far detector will be composed solely of liquid scintillator encased in titanium dioxide-loaded PVC extrusions.

T2K is also an off-axis experiment. It will make use of the 50 GeV proton synchrotron in the Japan proton accelerator research complex (J-PARC) built in Tokai-village, Japan. Construction of this facility is expected to be completed in 2007. The T2K beamline can be adjusted enabling the off-axis angle to be varied from 2.0 to 3.0 degrees. T2K will use the existing Super-Kamiokande water Čerenkov detector as its far detector. The near detector hall will be 280 m downstream of the proton target and will contain two independent detectors. One on-axis detector will be used to measure the beam direction and a second off-axis (in the direction of Kamioka) detector will be used to measure the neutrino energy spectrum and estimate ν_e contamination of the beam.

Both T2K and NO ν A hope to measure both θ_{13} and the atmospheric oscillation parameters more precisely. Positioning the detectors off-axis should give a narrower beam neutrino energy spectrum peaked at a lower energy than the on-axis beam energy spectrum. This will provide a larger fraction neutrino events in the energy range where oscillations take place.

2.8 Short-baseline Accelerator Neutrino Experiments

2.8.1 LSND

The *Liquid Scintillator Neutrino Detector* (LSND) is located at the Los Alamos Neutron Science Center, New Mexico [79]. LSND measured a neutrino beam produced by a beam dump of 800 MeV protons onto a fixed target. The majority of the secondary π^- are absorbed by the target, but the secondary π^+ mostly come to rest and decay by $\pi^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$. Since the majority of decays take place at rest, the neutrino beam has a well defined energy spectrum with a maximum energy of $E \simeq 50$ MeV. The detector was a 160 t liquid scintillator detector located 30 m away from the beam production point. The aim of the experiment was to search for $\bar{\nu}_e$ appearance.

LSND reported an excess of $87.9 \pm 22.4 \pm 6.0$ $\bar{\nu}_e$ events in the range $E \simeq 36 - 60$ MeV, interpreted as being the result of $\bar{\nu}_\mu \leftrightarrow \bar{\nu}_e$ oscillations [80]. However, other short baseline neutrino oscillations experiments have not observed a similar $\bar{\nu}_e$ appearance signal. Both the KARMEN 2 (*Karlsruhe Rutherford Medium Energy Neutrino*) accelerator neutrino experiment [81] and the Bugey [82] reactor neutrino experiment have observed $\bar{\nu}_e$ rates in good agreement with the no oscillation expectation.

If LSND is correct, then it gives a third Δm^2 measurement and one with which it is impossible to define three m^2 values consistent with the three measured Δm^2 scales (solar, atmospheric, LSND). In order to explain all the observations at least one “sterile” neutrino, which does not couple to the weak force, must be introduced [83] or even more exotic solutions such as $\mathcal{CPT}$ violation, where some or all of the oscillation parameters for ν and $\bar{\nu}$ differ, must be invoked [84].

2.8.2 MiniBooNE

MiniBooNE (*Mini Booster Neutrino Experiment*) [86] is intended to be the first phase of the Booster Neutrino Experiment (BooNE) at Fermilab. The main goal of the experiment is to unambiguously confirm or refute the evidence for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations reported by the LSND experiment.

The MiniBooNE neutrino beam is a high-intensity, conventional neutrino beam produced via the decay of mesons and muons in a 50 m long decay region following the target hall, where meson production and focusing occurs. Mesons are produced in the

Figure 2.19: Short-baseline accelerator neutrino experiment results. The top plot shows the MiniBooNE 90% CL limit (thick solid curve) and sensitivity (dashed curve) for events with $475 < E_{\nu}^{QE} < 3000$ MeV within a two neutrino oscillation model. Also shown is the limit from the boosted decision tree analysis (thin solid curve) for events with $300 < E_{\nu}^{QE} < 3000$ MeV. The bottom plot shows the limits from the KARMEN [81] and Bugey [82] experiments. The MiniBooNE and Bugey curves are 1-sided upper limits on $\sin^2 2\theta$ corresponding to $\Delta\chi^2 = 1.64$, while the KARMEN curve is a “unified approach” 2D contour. The shaded areas show the 90% and 99% CL allowed regions from the LSND experiment. Taken from [85].

interactions of 8 GeV protons from the Fermilab Booster accelerator in a thick beryllium target, and then focused by a magnetic focusing horn surrounding the target. The MiniBooNE detector consists of a carbon steel spherical tank, 6.1 m in radius, filled with approximately 800 t of undoped mineral oil. The centre of the detector is 541 m from the neutrino production target, below a 3 m dirt overburden. Neutrino interactions are observed through the detection of Čerenkov and scintillation light produced by neutrino-induced charged tracks travelling through the mineral oil.

As of November 2005, about 6.6×10^{20} protons had been sent to the MiniBooNE target and 7.0×10^5 candidate neutrino interactions have been identified in MiniBooNE [86]. The collaboration has recently released the results of a search for ν_e appearance in this data [85]. Using two largely independent analyses they observe no significant excess of events above background for reconstructed neutrino energies above 475 MeV. The data are consistent with no oscillations within a two neutrino appearance only oscillation model. Figure 2.19 shows the MiniBooNE 90% CL limit and the confidence limits obtained from the KARMEN [81] and Bugey [82] experiments, assuming two flavour oscillations. Much of the LSND region is excluded by the KARMEN and Bugey results, but two regions at $\Delta m^2 \simeq 0.2 - 1.0 \text{ eV}^2$ and $\Delta m^2 \simeq 7 \text{ eV}^2$ remain consistent with the data from all three experiments. If the oscillations of neutrinos and antineutrinos are the same, this result excludes two neutrino appearance-only oscillations as an explanation of the LSND anomaly at 98% CL.

MiniBooNE began running in anti-neutrino mode on the 19th of January 2006 [87]. Running in this mode MiniBooNE can directly confirm or rule out that the LSND $\bar{\nu}_e$ excess [80] is due to $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations. Running in both neutrino and anti-neutrino modes is necessary to test scenarios involving $\mathcal{CP}$ and $\mathcal{CPT}$ violation where oscillations of this type may only occur for anti-neutrinos.

2.9 Summary

The current state of knowledge of neutrino oscillation parameters is summarised in Table 2.1. New experiments coming online in the next 5 years will further improve the knowledge of the atmospheric and solar mixing parameters. T2K, NO ν A and double CHOOZ will either measure or set even tighter limits on θ_{13} . If θ_{13} is measured to be non-zero, then there is also the possibility of measuring the $\mathcal{CP}$ violating phase, δ .

Oscillation Parameter	Central Value	99% CL Range
solar mass splitting	$\Delta m_{21}^2 = (8.0 \pm 0.3) 10^{-5} \text{ eV}^2$	$(7.2 \div 8.9) 10^{-5} \text{ eV}^2$
atmospheric mass splitting	$ \Delta m_{32}^2 = (2.5 \pm 0.2) 10^{-3} \text{ eV}^2$	$(2.1 \div 3.1) 10^{-3} \text{ eV}^2$
solar mixing angle	$\tan^2 \theta_{12} = 0.45 \pm 0.05$	$30^\circ < \theta_{12} < 38^\circ$
atmospheric mixing angle	$\sin^2 2\theta_{23} = 1.02 \pm 0.04$	$36^\circ < \theta_{23} < 54^\circ$
‘CHOOZ’ mixing angle	$\sin^2 2\theta_{13} = 0 \pm 0.05$	$\theta_{13} < 10^\circ$
$\mathcal{CP}$ violating phase	$\delta = ?$	?

Table 2.1: Summary of present information on neutrino masses and mixings from oscillation data. Taken from [39].

Chapter 3

The MINOS Experiment

“Insanity in individuals is something rare – but in groups, parties, nations, and epochs it is the rule.”

— Friedrich Nietzsche, 1844–1900

3.1 Overview

The *Main Injector Neutrino Oscillation Search* (MINOS) is an accelerator beam neutrino experiment conducting a study of the oscillations of muon neutrinos over a long baseline. The main aim of the experiment is to confirm the existence of neutrino oscillations and perform precise measurements of the oscillation parameters.

The analysis presented in this thesis uses only the MINOS Far Detector. This chapter presents an overview of the main MINOS experiment and the physics analyses being carried out by the MINOS collaboration. A complete discussion of the MINOS Far Detector is given in the following chapter.

The MINOS neutrino beam is manufactured by the Main Injector Neutrino (NuMI) beam facility at the Fermi National Accelerator Laboratory (Fermilab) in Illinois, USA. The energy spectrum of the neutrinos is sampled at two points along the beam line using two sampling calorimeters. The 0.98 kt Near Detector samples the beam at a distance of 1 km from the proton target at Fermilab. The 5.42 kt Far Detector, located in the Soudan Underground Laboratory in Minnesota, USA, samples the beam spectrum

Figure 3.1: A cartoon of the NuMI beamline, showing how the beam is produced and monitored before entering the Near Detector Hall.

735 km from the proton target.

3.2 NuMI Beam

The NuMI beam provides neutrinos for the MINOS experiment [88]. Every 1.87 seconds batches of 120 GeV protons are extracted from the Main Injector at Fermilab into the NuMI beamline. This is done rapidly using a series of three kicker magnets. Each batch lasts approximately $8 \mu\text{s}$ and contains approximately 2.5×10^{13} protons (at design intensity). Figure 3.1 shows a cartoon of the NuMI beamline. The protons are directed along a carrier tunnel into the graphite target. The segmented target is 1 m long, 6.4 mm wide and 18 mm high. The shape of the target has been carefully designed to allow most of the protons to interact in the target, but also to allow the secondary particles produced (mainly pions and kaons) to escape through the sides without being reabsorbed.

Two magnetic horns focus the positively charged secondary particles emitted from the target back onto the axis of the primary proton beam. Figure 3.2 shows a photograph of one magnetic horn viewed looked toward the target. Each of these horns consists of a parabolic inner conductor and a cylindrical outer conductor. During each beam spill, current pulses of 185 kA are applied to each conductor to generate a toroidal magnetic field between the conductors. The parabolic shape of the inner conductors of the horns gives them their focusing power by enabling each horn to act as a lens for charged particles.

The focused beam of secondary particles then passes into an evacuated 675 m long

Figure 3.2: The view through one of the two NuMI focusing horns. Pulses of current are applied in time with the beam spills generating a magnetic field which is used to focus secondary particles onto the axis of the beam line.

decay pipe, 2 m in diameter. The pions and kaons decay via $\pi^+ \rightarrow \mu^+ \nu_\mu$ and $K^+ \rightarrow \mu^+ \nu_\mu$. Just in front of the absorber, at the end of the decay pipe, a hadron monitor detects the remaining pions, kaons and protons. At the end of the decay pipe there is a water-cooled absorber with a steel-encased aluminium core. This stops any remaining hadrons. The beam then passes through 240 m of dense Dolomite rock before passing into the Near Detector Hall. This absorbs any remaining muons from π/K decays in the beam, leaving a beam purely composed of neutrinos. Three alcoves have been cut into the rock in the path of the beam at distances between 10 m and 50 m downstream of the absorber. In each one there is a muon detector. The lateral profiles of muons and hadrons observed in these four detectors provide a good measure of the NuMI target integrity and the beam focusing of the horns.

Most neutrinos will be ν_μ (97.8%), with a small (1.8%) $\bar{\nu}_\mu$ contamination from $\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e$ and $\pi^- \rightarrow \mu^- \bar{\nu}_\mu$. There is a small ν_e component (0.4% at the expected appearance maximum) and a very small $\bar{\nu}_e$ component (0.03%) [89]. The design of the beamline is such that the energy spectrum of the neutrino beam can be tuned to make the best possible measurement of the oscillation parameters. The focal length varies with the momentum of the particles. Therefore the momenta of the charged secondaries and hence the neutrino beam spectrum can be selected by varying the positions of the horns relative to both each other and the target and the current flowing through them during the spills. In theory the current in the horns could be reversed to allow experiments with a $\bar{\nu}_\mu$ dominated beam.

3.3 The MINOS Detectors

3.3.1 MINOS Near Detector

The MINOS Near Detector is located 1 km from the NuMI target, in the appropriately named Near Detector Hall, which has an over-burden of 98 m. It is a sampling calorimeter composed of interleaved layers of steel and plastic scintillator. Neutrinos interact in the steel producing a number of secondary particles. The charged particles produced interact with the scintillator strips generating small amounts of light, which can be detected.

By the time it reaches the Near Detector the beam spot has an RMS of 40 cm. The detector is designed so that its fiducial volume is centred on the beam spot. The coil is

Figure 3.3: Four logical sections of the Near Detector (taken from [90]). The location of the veto, target, hadron calorimeter and muon spectrometer detector sections are shown. The “uninstrumented” region only has scintillator in every 5th plane.

offset by 50 cm from the centre to accommodate this. The coil returns along the side of the detector. The size of the detector is driven by the need to have sufficient steel to create a similar magnetic field to that created in the Far Detector (see Section 4.2). The detector is composed of 282 planes of 2.54 cm thick steel with dimensions as shown in Figure 3.4. It has a total mass of 0.98 kt.

The detector is divided into four regions, as shown in Figure 3.3:

- veto region** The first few instrumented planes of the detector are used to veto neutrino interactions which occur upstream of the detector.
- target region** This forty plane section (planes 21 - 60) is the fiducial volume used to select neutrino interactions for use in the analyses.
- shower region** This region (planes 61 - 120) is used to contain any electromagnetic and hadronic showers produced.
- spectrometer region** The last 161 planes of the detector are only instrumented every fifth plane and are used to measure muons produced by interactions in the target region.

Every plane (except plane 0) in the first three regions of the detector is instrumented in the area around the beam-spot (the grey area around the dark circle in Figure 3.4). Every fifth plane is fully instrumented (see Figure 3.3). Each scintillator plane is made up of 4 cm wide strips of scintillator in alternating perpendicular orientations. This allows for 3D reconstruction of events. The scintillator strips are discussed in more

Figure 3.4: An front-on (xy plane) view of the MINOS Near Detector. The shaded grey region shows the extent of the scintillator modules. The area surrounded by the solid black line shows the extent of the steel planes. The dark circle represents the 50 cm diameter beam spot. The diamond shaped hole to the right of the centre of the steel planes is the magnetic coil hole.

detail in Section 4.3.1.

The interaction rate in the Near Detector is expected to be approximately twenty neutrinos per beam spill. These events will overlap closely in space and time. For this reason the Near Detector uses a multi-ranging QIE readout system to digitise detector hits continuously at a rate of 53 MHz, with no dead-time. This gives it a temporal resolution of approximately 19 ns.

3.3.2 MINOS Far Detector

After travelling 735 km from the NuMI target, mostly through the Earth, the NuMI beam spot has expanded to around 1.5 km in diameter. At this point it is sampled by the MINOS Far Detector. The GPS system at the Near Detector is used to generate timestamps of the spill signals which are transmitted to the Far Detector over the Internet, where they are stored and served to the Trigger Processors on request by the SpillServer. In this way, MINOS ensures that all in-spill activity in the Far Detector is recorded. The details of the MINOS Far Detector are discussed in Chapter 4.

3.3.3 MINOS Calibration Detector

The MINOS Calibration Detector (CalDet) was constructed at CERN in Geneva in order to test the designs of the MINOS detectors and their readout systems and to measure their response to a test beam [91]. From 2001 to 2003, CalDet was exposed to beams of p , $\pi^\pm$, $\mu^\pm$, $e^\pm$ of varying energies between 0.2-1.0 GeV at the CERN PS accelerator. The detector comprised sixty 1 m by 1 m square planes of 2.5 cm thick steel (as opposed to the 1 inch / 2.54 cm thick steel used in the Near and Far Detectors), with scintillator planes sandwiched between them. The detector was instrumented using electronics from both the Near and Far Detectors.

3.4 MINOS Physics Analyses

3.4.1 Beam ν_μ disappearance Analyses

The main purpose of the MINOS experiment is to measure neutrino oscillations via the observation of a deficit of neutrinos at the Far Detector compared to the expectations extrapolated from Near Detector observations. This will be observed primarily by analysing ν_μ charged-current interactions. An analysis of the data from the initial 1.27×10^{20} protons on the NuMI target has now been published [74] (see Section 2.7.2).

The hypothesis that the oscillations are of the form $\nu_\mu \rightarrow \nu_\tau$ or $\nu_\mu \rightarrow \nu_e$, as opposed to $\nu_\mu \rightarrow \nu_s$, can be tested by comparing the rates of neutral current interactions at both detectors. The neutral current interaction cross-sections for ν_μ and ν_τ should be the same. Therefore the rates of this process should not be affected by oscillations, if the $\nu_\mu \rightarrow \nu_\tau$ hypothesis is correct. If a deficit is observed at the Far Detector compared to the expectation extrapolated from Near Detector observations this could be an indication that something more complicated is happening. Given that the showers produced in neutral current events are often of low energy, this is a far more difficult analysis in the coarse grained environment of the MINOS detectors.

3.4.2 Contained Vertex Atmospheric Neutrino Analyses

The details of previous contained vertex atmospheric ν_μ analyses using the MINOS Far Detector are discussed in Section 2.6.4. The results of the new higher statistics analysis,

that is the subject of this thesis, will be discussed in future chapters. Far Detector atmospheric ν_e observations will allow normalisation of the atmospheric neutrino flux using MINOS data, this will reduce systematic errors on this parameter in other MINOS atmospheric neutrino analyses.

3.4.3 Neutrino Induced Muon Analyses

An analysis searching for upward-going atmospheric neutrino-induced muons in the MINOS Far Detector has been performed on 854.24 live days of data. During this time, 140 neutrino-induced muons were selected from the data. The experimenters searched for evidence of oscillations in this data set by computing the ratio of the number of low momentum muons to the sum of the number of high momentum and unknown (assumed too high to measure) momentum muons for both data and Monte Carlo expectation in the absence of neutrino oscillations. The ratio of ratios, $\mathcal{R}$, is

$$\mathcal{R} = 0.65^{+0.15}_{-0.12}(\text{stat.})^{+0.08}_{-0.04}(\text{sys.}), \quad (3.1)$$

which is consistent with oscillations. A fit to the data for Δm_{32}^2 and $\sin^2 2\theta_{23}$ excludes the null oscillation hypothesis at the 94% confidence level. The selected muons in data and Monte Carlo were separated into μ^- and μ^+ events and the ratio of μ^- to μ^+ in both samples calculated. In this case, the ratio of ratios, $\mathcal{R}_{CPT}$, is a test of $\mathcal{CPT}$ conservation. The result obtained was

$$\mathcal{R}_{CPT} = 0.72^{+0.24}_{-0.18}(\text{stat.})^{+0.08}_{-0.04}(\text{sys.}), \quad (3.2)$$

which is consistent with $\mathcal{CPT}$ conservation [92].

3.4.4 Other Analyses

In addition to the published analyses there are many other lines of research that are being pursued by the MINOS collaboration. If the mixing angle θ_{13} is non-zero, then $\nu_\mu \rightarrow \nu_e$ oscillations will occur in the NuMI beam. Currently there is only an upper limit on θ_{13} . MINOS hopes to observe ν_e appearance in the Far Detector, which would indicate a non-zero value of θ_{13} . If no ν_e interactions are observed above the expected background, then after 2 years of running, it is expected that MINOS will be able to improve on the CHOOZ limit ($\sin^2 2\theta_{13} < 0.17$ [61]) [93].

The MINOS Far Detector is being used to measure cosmic ray muon fluxes, particularly to examine the μ/μ^+ ratio as a function of muon energy. The high neutrino flux through the MINOS Near Detector will allow neutrino cross sections to be measured accurately, as well as allowing many other interesting measurements and searches to be made.

Chapter 4

The MINOS Far Detector

“The major difference between a thing that might go wrong and a thing that cannot possibly go wrong is that when a thing that cannot possibly go wrong goes wrong it usually turns out to be impossible to get at or repair.

– (Mostly Harmless, 1992)”

— Douglas Adams, 1952–2001

4.1 Overview

At 5.42 kt, the MINOS Far Detector is the largest of the three MINOS detectors. It is situated in the Soudan Underground Laboratory, on the 27th level of the Soudan Underground Mine State Park, Northern Minnesota, USA. This location is 700 m below the Earth’s surface, giving the Far Detector a 2070 mwe overburden. This is large enough to make the detector useful for atmospheric neutrino observations, as well as for its designed purpose, measuring the NuMI beam 735 km from source. Construction was completed in August 2003 (although some data taking had occurred before with a partial readout).

Figure 4.1: Photograph of the MINOS Far Detector (viewed toward Fermilab). The magnetic coil is visible exiting the centre of the detector and dropping down to the floor. The scintillator planes above the detector make up the cosmic ray veto shield. (Photograph courtesy of Jerry Meier.)

4.2 Detector Structure

The majority of the MINOS Far Detector's mass is in its 486 steel planes. The detector is divided up into two supermodules, containing 249 and 237 steel planes respectively, in order to mitigate the effect of thermal expansion. The frame that supports the steel planes is made up of two horizontal rails supported by vertical posts. This frame would become unstable if the rails became too elongated and forced the posts away from the vertical, particularly at each end of the rail. By dividing the rails into two parts the effect of the elongation is reduced, enabling the detector to be stable at a larger range of temperatures, should the cooling system fail.

The steel planes are 8 m wide and 2.54 cm (1 inch) thick, with their centres spaced 5.94 cm apart, except in the gap between supermodules where the gap is approximately 1.2 m. The first plane in each supermodule (0 and 249) is uninstrumented. A 1 cm thick plane of plastic scintillator is attached to the front of all other planes, so the supermodules have 248 and 236 scintillator planes respectively. The detector is often described as a steel-scintillator sandwich. The planes are aligned vertically in order to maximise the sampling of beam neutrino interactions.

Each scintillator plane is divided into 192 parallel ‘strips’ in the xy plane. Each strip is 4.1 cm wide, but their lengths vary depending on their position across the plane. The strips in successive planes are aligned alternately along the orthogonal axes $U = \frac{1}{\sqrt{2}}(y+x)$ and $V = \frac{1}{\sqrt{2}}(y-x)$ (see Figure 4.2). By combining the information from adjoining planes 3D measurements of events can be made. The scintillator strips in each plane are attached to the steel inside 8 modules, containing either 20 or 28 strips covered by a thin aluminium casing. As charged particles travel through a strip, scintillation light is generated and propagates in both directions along the strip. At each end, the light is collected by optical fibres and transported to photo-multiplier tubes (PMTs). The PMTs amplify the signal, which is then digitised by sensitive electronics.

Each supermodule is magnetised using 190 turns of copper wire carrying a current of 80 A. The coil runs through the centre of each supermodule and returns underneath the detector, as can be seen in Figure 4.1. When powered-up the coil produces a toroidal magnetic field in the steel with a mean strength of 1.3 T (in the direction shown in Figure 4.2(a)). The field allows the determination of the charge of muon tracks in the detector and also their momentum, via the curvature of the muon track in the magnetic field. The direction of the field is such as to focus the negatively charged muons produced in beam interactions.

The flux of cosmic muons is strongly peaked toward vertical angles, so the majority of cosmic muons are incident on the top of the detector. Due to the vertical alignment of the detector planes, the coverage of the top surface is below 20%. For this reason, a veto shield has been constructed above the detector to tag cosmic muons entering the detector. The veto shield (visible in Figure 4.1) is made up of spare scintillator modules. The shield modules are aligned to the z -axis and provide over 99% coverage of the top surface of the detector [96].

4.3 Detector Technologies

4.3.1 Plastic Scintillator

The scintillator used in all the MINOS detectors is manufactured by extruding solid polystyrene scintillator into long thin strips that are $1.0 \text{ cm} \times 4.1 \text{ cm}$ in cross-section and up to 8 m in length. The polystyrene contains aromatic molecules that, when excited, release a small amount of energy as optical photons. It is doped with primary

(a) MINOS Far Detector coordinate system (viewed toward Fermilab).

(b) MINOS Far Detector strip orientations (viewed away from Fermilab).

Figure 4.2: (a) from [94] (b) from [95]

(1% PPO) and secondary (0.03% POPOP) fluors which absorb the photons at a fast rate and work together to emit blue photons with a mean wavelength of 460 nm. In order to maximise the light yield, each strip has an outer reflective coating (0.025 cm thick) containing 15% TiO_2 that is mixed with the polystyrene during the extrusion process.

The scintillator has an absorption length of 20 cm, hence a technique for extracting the photons from the interaction point is required. To this end, a groove is cut along the centre of every strip into which a 1.2 mm diameter Kuraray Y-11 wavelength shifting (WLS) fibre is glued, as shown in Figure 4.3(a). The WLS fibres absorb light in the blue part of the spectrum (with peak absorption at 420 nm) and re-emit them isotropically. Photons re-emitted along the central axis of the fibre are trapped and hence transmitted along the length of the fibre, as can be seen in Figure 4.3(b). The re-emitted spectrum has a peak at 520 nm [90]. The decay time for emission of light by the WLS fibres is approximately 8 ns [90]. The attenuation of light along the WLS fibres is characterised by two attenuation lengths. A short attenuation length (λ_{short}), which describes the strong absorption of short wavelength light; a long attenuation length (λ_{long}), which characterises the weaker absorption of the remaining longer wavelength light. The WLS fibre used at the Far Detector has been tested and these attenuation lengths have been measured to be $\lambda_{short} \simeq 1.0$ m and $\lambda_{long} \simeq 7.0$ m [97] [98].

Groups of between 20 and 28 strips scintillator strips are encased within light-tight aluminium (0.05 cm thick). These encased groups of strips are known as modules. At the end of the strips the WLS fibres run together in a manifold before they terminate in a connector. Here the light is channelled into highly transparent clear polystyrene fibres (attenuation length 12 m) and transported to multi-pixel photo-multiplier tubes.

4.3.2 M16 Photo-multiplier Tubes

The scintillator strips are read out by 16-pixel Hamamatsu M16 PMTs [99]. A schematic of the Far Detector optical readout is shown in Figure 4.4. Given the large number of channels in the Far Detector (185856 strip-ends) a multiplexing (optical-summing) technique is used to reduce the amount of instrumentation required. Multiplexing the signals from multiple strips in a plane is possible because the transverse spread of the electromagnetic/hadronic showers from beam neutrinos is limited to a ~ 1 m region. Therefore, up to 8 strips in a plane can be read out by the same PMT pixel with a minimal impact on event reconstruction. The situation is slightly more complicated for atmospheric neutrino interactions, where the events can be much steeper and so

(a) Cross-section of a scintillator strip. The position of the groove that holds the WLS fibre can be seen in the middle at the top. The strip has a thin layer of co-extruded TiO_2 on the outside.

(b) A short length of MINOS scintillator illuminated by a blue LED. The groove for the WLS fibre is clearly visible along the length of the strip.

Figure 4.3

cover more than 1 m in a plane. The number of PMTs required to read out the entire detector in this way is 1452. Groups of three PMTs are housed in custom-built light-tight multiplexing (MUX) boxes. Each group of three PMTs reads out one side of two planes ($3 \times 16 \times 8 = 2 \times 192$). A group of fibres from eight strip-ends is coupled to each PMT pixel, as shown in Figure 4.4(b). The resulting eight-fold ambiguity in the readout scheme can be resolved, as the multiplexing scheme is such that each strip is read out by a unique pair of pixels. Therefore, if the PMTs on opposite sides of a scintillator plane each record signals in only one pixel, a unique strip can be identified and even if the PMTs record multiple hits it is usually possible to clearly define the event region (see Section 7.3 for more details of demultiplexing).

The light from the scintillator is converted into photo-electrons (PEs) at the PMT photo-cathodes. At the peak WLS fibre emission (520 nm), the quantum efficiency for photon conversion is approximately 13%. The PMTs operate at gains of 10^6 and typically produce signals of 1 pC on the anodes. A typical muon will produce ~ 5 PE per strip-end. The results of bench studies show that the integrated charge on the anode increases linearly (to within 5%) with the input signal on the photo-cathode for signals up to 100 PE [99].

A common dynode signal is provided by each PMT. This is used to trigger the readout. The trigger threshold is set at a level of charge corresponding to $\frac{1}{3}$ PE. This is subject to variations in gain between pixels however. Background noise causes dynode triggers at a rate of 5-10 kHz per PMT. This background noise is due to the following effects [100] (in order of importance):

Anomalous WLS Fibre Noise: Rates of ~ 5 kHz per PMT are due to single photo-electron signals caused by photon emission originating in the WLS fibre. The source of these photon emissions is believed to be the long-term relaxation of mechanical stress caused by the gluing of WLS fibres into the scintillator strips during the construction of the detector.

Natural Radioactivity: The detector components and the surrounding rock all contain small concentrations of long lived radioactive isotopes (e.g. ^{228}U , ^{232}Th , ^{40}K), whose complex decay chains produce a constant flux of β and γ radiation on the surface of the detector. One decay product which is able to penetrate deep into the detector is ^{222}Rn gas, the levels of which are monitored constantly. The incident radiation produces energy deposits in the scintillator, resulting in signals of a few photo-electrons. Extensive studies have been made of the natural radiation

(a) Schematic of Far Detector optical readout. Light signals generated in scintillator strips are transported by optical fibres to M16 PMTs housed in custom-built MUX boxes.

(b) photographs showing optical fibres threaded into a MUX box (left) and an M16 PMT (right).

Figure 4.4

at the Far Detector. Based on the results of these studies, the total predicted rate of dynode triggers due to radioactivity is 770 Hz per PMT. This is comprised of 410 Hz per PMT from natural radioactivity in the rock, 210 Hz per PMT from radioactivity in the detector and 150 Hz per PMT from radon gas.

Dark Photons: PMTs register signals even when no input light is present. These signals arise chiefly from thermal emissions of electrons from the photo-cathode which produce single photo-electron signals. The rate of these dark photons has been measured to be approximately 500 Hz per PMT.

4.3.3 Front-End Electronics

For each MUX box, there is a VA Front-end Board (VFB) which provides analogue readout and trigger formation for the three PMTs. An Amplifier/Shaper/Discriminator chip (ASD-lite) triggers readout of the PMTs based on the charge on the last dynode stage. Each VFB also holds three customised Viking VA chips (made by IDE AS of Oslo, Norway), which provide an analogue readout of the PMT anodes [101]. Each VA chip is responsible for reading out one PMT. The VA chips use 22 VA readout channels. One VA channel is allocated to each anode of the PMT, with a further channel used to read out one of the PIN diodes that form part of the detector calibration system (see Section 5.4). The remaining 5 channels are used to correct common mode fluctuations across the channels.

Figure 4.5 illustrates the Far Detector front-end readout scheme. Each VFB is a slave module and is under the direct control of a VA Readout Control card (VARC), housed in a VME (Versa Module Europa) crate some distance away. The VARCs are custom built VME boards. Event Time Controller (ETC) chips, located on the VARCs, direct the readout of the PMTs by their associated VA chips. The VA chips readout each of the PMT anodes, amplify and shape the signals, and then send the signals to the VARC. On the VARC, the signals are digitised by VARC Mezzanine Modules (VMMs). Each ETC-VMM pair deals with two VFBs (*i.e.* six VA chips). Up to six ETC-VMM pairs are housed on each VARC. Each VARC has one Analogue-to-Digital-Converter (ADC) which is used for all the VA chips it handles. There are three VARCs in each VME crate and 8 VME crates on each side of the detector.

When a dynode signal is detected by the ASD-lite chip, a signal is sent to the ETC on the VARC. A coincidence trigger is implemented on the VARC that requires 2 out

(a) VA Front-End Boards (VFBs) **(b) VA Readout Cards (VARCs)**

Figure 4.5: Far Detector front-end readout scheme. The VA electronics is organised onto VA Front-End Boards (VFBs) located on the side of MUX boxes, and VA Readout Control cards (VARCs) housed in VME crates. [94]

of the 36 VA chips connected to the VARC to trigger within 400 ns before the readout is activated. The coincidence trigger reduces the trigger rate from ~ 5 kHz to 300 Hz per PMT [102]. Upon receiving a pair of dynode triggers that satisfy the 2/36 VARC trigger, an ETC timestamps the triggers and directs the VA chips to read out the PMTs that have triggered. It is impossible to set the 30-bit TDC (Time-to-Digital-Converter) value using a straight digitisation of the 40 MHz system clock. Rather, to get the lowest 4 bits, the ETC does a phase-comparison of the PMT trigger-fire to 16 different phases of the 40 MHz clock and chooses the one with the best match. Each trigger is time-stamped using a TDC implemented within the ETC, using its internal delay-locked loops (DLLs). One DLL is used to double the 40 MHz system clock to 80 MHz, while a second generates four clock phases separated by 3.125 ns. Each trigger signal is registered by all four phases, thus time-stamping the trigger in 3.125 ns bins. Two such circuits are used for each trigger, each separated by an external 1.5 ns passive delay. This effectively creates 1.56 ns bins [101].

The signals in each of the VA channels are then digitised in turn by the VMM. During this time any further dynode signals from the PMTs are ignored, meaning that a short dead time is incurred. The VA readout and digitisation operates at 200 kHz per chip, producing a dead time of 5 μ s per VA chip. If multiple VA chips with a common VMM are triggered, they are queued by the ETC to await digitisation. Since each VMM services six VA chips, dead times of up to 30 μ s are possible.

The VMMs digitise the VA signals into a 14-bit ADC, with the dynamic range set to ~ 50 ADC/PE, giving a maximum signal of 300 PE. The digitised hits are passed to a sparsifier chip on the VARC. The sparsifier chip subtracts pedestal thresholds from each signal and sparsifies the data by removing all signals below a pre-programmed channel-dependent threshold (nominally 20 ADC). The sparsified data are then transferred into one of two output buffers on the VARC, which are used alternately to collect data from the VARC and transfer data out of the VARC. The buffers swap roles at a rate of 50 Hz, enabling continuous data taking.

Figure 4.6: The architecture of the MINOS Far Detector DAQ system. For clarity only on PVIC input branch is shown. (In reality there are six BRPs rather than the four shown here.)

4.4 Data Acquisition and Trigger Processing

4.4.1 Far Detector Data Acquisition System Operation

Figure 4.6 shows a schematic overview of the DAQ system. The digitised data in the front-end electronics buffers (time blocks) are read out with single block Direct Memory Access (DMA) transfer by a VME computer in each crate, known as the Read Out Processor (ROP). There are 16 ROPs serving the Far Detector. Read out is synchronised across all front-end crates by the Timing system Central Unit (TCU). The TCU fans out control signals to VME Timing Receiver Cards (TRCs) in each VME crate. The TRCs generate VME interrupts at a programmable rate (typically 50 Hz). Upon receipt of a VME interrupt the ROP swaps front-end buffers and transfers the data into memory.

The ROP assembles the time blocks over a programmable time period, nominally one second, into “time frames”. The time frames are overlapped by one time block to circumvent processing problems at the time frame boundaries. The ROP adds a header containing crate timing information and time frame number and appends some monitoring statistics. Groups of ROPs are interconnected with differential PCI Vertical Interconnection (PVIC) cables to form a chain which is connected by an optical PVIC cable to an off-detector DAQ PC, known as a Branch Readout Processor (BRP). The PVIC buses allow transfer rates of up to 40 MB/s. The Far Detector has six such chains.

All PCs in the MINOS DAQ system are Intel Pentium-based and run the Fermilab release of the GNU/Linux operating system.

Time frames are buffered by the ROPs until requested by their BRP. They are then transferred into the memory of the BRP by DMA across the PVIC. One BRP acts as the master (MBRP), coordinating the transfer of data by issuing transfer instructions to the BRPs. Each BRP is also connected to an output branch. The output branch is a differential PVIC bus which connects the BRPs to a small farm of trigger processors (the TPs). When all the BRPs have completed transferring the next time frame from the memory of the ROPs, the MBRP selects a TP from the trigger farm, based on a selectable least-loaded or round robin algorithm. The MBRP then instructs each BRP to transfer the time frame to the selected TP via DMA over the output PVIC branch. By passing the appropriate target address to each BRP in turn, the time frame for the whole detector can be assembled continuously in the TP memory. Like the ROPs and BRPs, the TPs are able to buffer and queue multiple time frames. Typically, in a physics run, the raw data rate to the trigger farm is 8 MB/s at the Far Detector.

4.4.2 The Trigger Farm

The trigger farm is comprised of five 3.2 GHz Pentium4 PCs. Upon receipt of a time frame a TP orders the data on a crate level. It then performs data integrity checks, to ensure the time frames from all ROPs are correctly structured. At this point any calibration data is processed and the results placed in the output data stream. (See Chapter 5 for more details.) The TP applies flexible software triggering algorithms to the time frame to locate events of physics interest. Since each of the triggers can in principle gather multiple events, the output from the trigger is termed a *snarl*.

The triggers fall into three categories: special triggers for debugging and calibration; bias free triggers based on spill signals or spill times to gather beam events; triggers based on the clustering of hits in the detector to gather out of spill events. In the latter case, a candidate snarl is first identified as a temporal cluster of hits bounded by at least 100 clock ticks (156 ns) of detector inactivity. The non-spill based trigger algorithms are applied to these candidate snarls.

The triggers available at the MINOS Far Detector are described below:

(1) Remote Spill Trigger: At the Far Detector a direct spill signal is not available

so a remote spill trigger is applied. The Near Detector GPS system is used to generate timestamps of the spill signals which are transmitted to the Far Detector over the Internet, where they are stored and served to the TPs on request by the SpillServer. All readout within a configurable time window around each spill is extracted and written out as a spill-event. Since the DAQ has considerable buffering there is ample time to wait for spill information to arrive.

- (2) **Fake Remote Spill Trigger:** Fake spill times are generated randomly between spills to provide random sampling of detector activity.
- (3) **Plane Trigger:** M detector planes in any set of N contiguous planes must contain at least one hit. Nominally M=4, N=5.
- (4) **Energy Trigger:** M contiguous planes of the detector must have a summed raw pulse height greater than E and a total of at least N hits in two or more adjacent planes. Nominally M=4, E=1500 ADC, N=6.
- (5) **Activity Trigger:** There must be activity in any N planes of the detector. Nominally N=20.
- (6) **Special Trigger:** A variety of special runs are available to perform detector and electronics calibration or debugging.

Each event can satisfy more than one trigger. MINOS beam neutrino analyses use the Remote Spill Trigger to select beam neutrino events. The MINOS atmospheric neutrino and cosmic-ray muon analyses use the Plane Trigger to select interesting events. An addition requirement that an event does not also have the Remote Spill Trigger bit set is made for atmospheric neutrino analyses to avoid contamination of the sample by beam neutrino events.

Once processing in the trigger farm is complete, the physics event and monitoring statistics are moved to the output buffer. The 30 μ s of detector activity prior to the physics event is also written out. These data are called the *pre-trigger window*. The pre-trigger window can be used to flag channels that are dead. The integrated trigger rate at the Far Detector is typically 30 Hz and is dominated by cosmic ray and noise events. The output rate to disk from all sources, including calibration data is approximately 10 kB/s. All Data are transmitted via TCP/IP to the Data Collection Process (DCP).

4.4.3 Data Collection Process

The DCP collects processed time frames from the trigger farm. In general, the processing in the trigger farm does not finish in time order. The DCP buffers time frames and time orders them before further processing. The integrity of the data is checked by testing checksum words in the data stream. If data corruption is detected, data-taking is stopped and the offending time frame is written to a debugging file. The overlaps between adjacent time frames, introduced by the ROPs to improve trigger efficiencies (see Section 4.4.1), are removed at this point. The DCP examines adjacent time frames and removes repeated events before archival to file occurs. (The probability of event repetition is 0.16% for cosmic muons at the Far Detector.)

The DCP archives the current run configuration and detector configuration to file with the data. This allows easy reconstruction of exact DAQ conditions in data analysis. Data can be written either in ROOT [103] format or written in the native MINOS DAQ format. The output files are supplied to the MINOS Data Distribution System which distributes data to the quasi-online processes such as online monitoring and event displays. Once a file is closed, it is archived to the Fermilab mass storage facility by Kerberized FTP over the Internet.

4.4.4 Run Control

Overall Structure

The DAQ is co-ordinated and monitored by Run Control. Run Control is based on the client-server model. A server (rcServer) runs on a DAQ PC at each detector location. Clients can connect to the rcServer via the Internet. The rcServer code, along with a number of communications libraries were written in C++ [96]. Clients are purely a user interface and contain no control logic within themselves. The client management software in the rcServer decides which client can issue commands and handles the exchange of control privileges between them. The client in use is a ROOT [103] based Graphical User Interface (rcGui). There is also a passive web interface, from which no commands can be issued, but where the current DAQ state and data transfer rates can be observed.

All connections between DAQ processes and rcServer are made through TCP/IP sockets using a protocol designed specifically for this task. This interface is used to pass all configuration, commands, state reports and monitoring data between the DAQ

processes and rcServer.

Run States

In order to take data the various logical nodes of the DAQ must work in unison. This process is managed by the Run Control Server (rcServer). The DCP, TPs, light injection system (FPC), ROPs, BRPs, MBRP and timing system (TPC) are all critical to data taking. Their states are monitored by rcServer. If any of these nodes are in an **ERROR** state, then data taking cannot take place until the problem has been overcome. The rcServer also monitors some other processes, so called quasi-online processes, which are not strictly part of the DAQ, but are important to watch. Quasi-online processes include the SpillServer (at the Far Detector), the Detector Control System (DCS) and data archiver among others. Data may be taken while these nodes are in error, but a warning will be shown on the Run Control GUI. Each node is responsible for maintaining its own state. This state may in turn depend on the state of one or more of its dependents. For example, the MBRP derives its state from that of the BRPs, which derive their individual states from the states of the ROPs they manage. Each node must conform to the state model which describes the various run states and the allowed transitions between them.

The Run Model used by the MINOS DAQ is described in detail elsewhere [104]. Figure 4.7 shows the finite state machine that each node must follow. There are a set of well defined set of primary states and transitions between them. In subsequent discussions we shall use the term ‘*node*’ to refer to the actual DAQ component itself and *node-object* to refer to the C++ object-instance used to represent a node within the Run Control code. The primary states of the Run Model are defined in Table 4.1. Nodes themselves only know about the states in the top half of the table, **TIMEOUT** and **UNKNOWN** are internal concepts, which apply only to the node-objects within the Run Control code. A node can move through a state transition under the command of rcServer, the exception being the **ERROR** state, which a node enters when it is no longer able to continue with data taking or complete a state transition.

The rcServer maintains a DAQ *Global state*, which is derived from the states of the nodes critical for data taking and conforms to the same state machine described above. The Global state machine was implemented using the State and Command Design Patterns[105]. If necessary an operator can dynamically edit the list of nodes from which

Figure 4.7: The finite state machine for the DAQ run model. Note that the implicit `abort` command, which is valid from any state, is not shown. Neither are the implicit transitions from any state to `ERROR` or `TIMEOUT`.

the Global state derives for debugging purposes. Some nodes' state transitions take a finite length of time and must be synchronised with state changes in other nodes. The Global state contains logic to ensure this is done correctly. Consider the case where, in order to start a run, rcServer issues a single command to the Global state, which then starts nodes in the correct sequence and distributes the required configuration information. If any dependant node is in the `ERROR` state, the Global state will move to `ERROR`. From the `ERROR` state the error recovery logic can be executed by issuing a `reset` command to the Global state. If the dependent nodes are in a mixture of non-`ERROR` states and the Global state is not in a state transition, then the Global state moves to the `UNKNOWN` state. If this occurs, then user intervention is required to bring the DAQ back into a well defined state. In order to facilitate this process, rcServer allows for expert users to issue commands directly to nodes, by-passing the Global state. When a command is issued, nodes are required to acknowledge the command. If a node fails to acknowledge a command within a certain time then rcServer puts the matching node-object into the `TIMEOUT` state. While in this state the node-object ignores state report messages from the node and behaves as if it is in the `ERROR` state until such time as the user issues a `reset` command, where upon the node-object starts to listen to state report messages from the node once more as it attempts the transition to the `IDLE` state.

DAQ State	Description
IDLE	The node is not participating in data taking. This is the default state of all nodes at start-up. A node stays in this state until instructed to prepare for data taking by Run Control.
PREPARED	The node has successfully configured itself for data taking according to the parameters provided by Run Control. Data taking does not start until instructed to do so by Run Control.
RUNNING	The node is currently involved in data taking and will continue to be until instructed to stop by Run Control or an error occurs.
ERROR	The node is no longer able to continue with data taking or a state transition has failed.
TIMEOUT	This state is similar to the ERROR state. Run Control internally records a node's state as TIMEOUT if the node has taken too long to acknowledge a command issued by Run Control. State reports from the node are ignored until the node is reset.
UNKNOWN	Run Control initialises its internal record of a node's state to UNKNOWN when it first connects a socket to the node, before any state reports have been received. Otherwise, if a node's dependent nodes are in a mixture of non- ERROR states and the node is not in a state transition, then the node moves to the UNKNOWN state.

Table 4.1: The primary DAQ states. The top half of the table contains states that are reported by the nodes themselves. The bottom half of the table contains states that can only be set internally in Run Control.

Detector Configuration Management

In the original version of Run Control, parts of the detector could be removed from the readout by changing the ropmask ($= \sum 2^n$, where n is the ROP number and the sum is over all ROPs to be used in the readout). More fine grained control required editing detector configuration files in a text editor, followed by restarting various Run Control processes. During 2004 it was decided that this fine-grained control should be done within Run Control. The resulting Detector Configuration Manager, written by the author, has elements in both rcServer and the rcGui clients. The rcServer manages which detector configuration is currently being used and also keeps a copy of the default

configuration (everything on).

When a client requests to edit the detector configuration the server sends the requesting client the current configuration. The client uses this information to build an interactive graphical display of the front end electronics. The user can then select which parts of the electronics readout they want to have active. The underlying logic of which parts of the detector can be active at a given time is handled by an underlying object-based representation of the detector electronics. For example, if the user elects to deactivate a particular VARC, all the VA chips it reads out will be deactivated also. A component can only be active if the component responsible for it is active also. Equally, if all the VA chips readout by a VARC had been deactivated, then the VARC would be deactivated also. Once the user is satisfied with their changes they can ‘commit’ the revised configuration. This is then sent to the server, where it is stored. All this can occur while a run is taking place, as the information is only checked by the rest of the DAQ processes when a `reset` or `prepare` command is sent.

Consider the case where a piece of the readout electronics is malfunctioning, causing the ROP responsible for it to return an `ERROR` state. A `reset` command cannot repair a broken piece of electronics, so the ROP would remain in an `ERROR` state. The situation can be recovered as follows: the detector configuration is edited to remove the malfunctioning piece of electronics from the readout; a `reset` command is then issued; all nodes parse the new detector configuration; the ROP removes the broken electronics from its readout and returns to the `IDLE` state. Previously recovering from this situation would have involved editing a configuration file by hand and restarting various DAQ processes.

4.5 Detector Control System

The *Detector Control System* (DCS) is used to control and monitor the state of the detector systems[106], in order to ensure safe and efficient detector operations. The DCS is comprised of the following systems: Environmental Controls, Low Voltage Monitoring, High Voltage Monitoring, Magnet Controls and Rack Protection Systems.

Environmental controls cover the monitoring of temperature, humidity and the level of radon gas. Sensors are situated around the Far Detector Hall and on a sub-set of detector planes. The low voltage monitoring system monitors the ± 2.5 V voltage rails situated on each VFB. At the start of each run the voltages are recorded by the VFB and

voltage regulators on each VFB ensure quiet operation. The high voltage monitoring system monitors the eight high voltage mainframes that power the PMTs. The nominal voltage supplied to each PMT is 800 V. Custom designed control software is used to control the high voltage supply. The software reads the voltage on each PMT at regular intervals and attempts to adjust any voltage found to be outside the tolerated range about the nominal. Each magnet coil draws an 80 A current from a 450 V high voltage power supply. The magnet controls use custom control software to ramp the power up and down. Both current and voltage are monitored at the supply. Thermocouples are also placed at several locations on the coils, in order to monitor the coil temperature. Rack Protection Systems (RPSs) are installed in each VME and DAQ rack around the detector. These systems monitor sensors that measure environmental conditions inside the racks (temperature, presence of smoke etc) and the power supply voltages.

Most systems are programmed to sound an alarm or shutdown automatically if they are found to be operating outside nominal conditions. Ethernet connections to the RPS boxes on each allow the power to racks to be cycled by a remote expert user if a problem is found.

4.6 Data Quality

Before data can be used in any physics analysis it must be checked to ensure that it has been acquired under normal conditions. The requirements for ‘good’ data include:

1. All detector systems must be functioning normally.
2. There must be few malfunctioning readout components.
3. The entire detector must be live.

For non-beam analyses we must make the additional requirements that we must know whether or not the NuMI beam (see Section 3.2) is online and if the NuMI beam is online we must know that whether or not we are within a spill (see Section 8.3). Data from both the DCS (for detector systems) and DAQ (raw data rate and quality) are used to perform checks on the data. The DCS and DAQ monitoring information is analysed to identify periods of abnormal running. This is then combined with information obtained from written logs and the daily checks performed at the Far Detector site.

4.6.1 Coil Currents

The DCS records the values of the coil currents every two minutes. During normal operations the coil current in each supermodule remains stable at 80 A. Data from any time periods when the coil current drops below 70 A in one or both supermodules are removed from the physics analysis dataset. This occurs in 2.9% of physics data and is caused by power trips or hardware malfunctions. Figure 4.8 shows the coil currents recorded in each supermodule for the period 1st August 2003 to 28th February 2006. On 21st June 2004 the polarity of both coils was reversed in order to study systematic effects in cosmic muon charge determination. On 1st February 2005 the polarity of both coils was returned to normal.

4.6.2 Dead Electronics

Malfunctions in hardware can lead to localised regions in which the readout is dead. These readout holes are chiefly caused by malfunctioning readout electronics and mainly affect single PMTs or VA chips dispersed through the detector. However, much larger dead regions are created by trips in the eight mainframe HV supplies that power the PMTs. The rate of hits on each VA chip for each time frame is stored in the DAQ monitoring blocks during data taking (see Section 4.4). A VA chip is flagged as dead if the rate drops below 50 Hz. Figure 4.9 shows the number of dead chips in each Supermodule plotted versus time between 1st August 2003 and 28th February 2006. During normal running the number of dead chips is below 10. Occasionally spikes in the number of dead chips occur due to HV trips. A HV trip is tagged if the number of dead chips in a supermodule rises above 20.

Figure 4.10(a) shows the distribution of dead VA chips in data where the HV was found to be functioning normally. The mean number of dead VA chips per time frame over the whole data set is calculated to be 0.66 chips (0.05% of VA chips in the detector readout electronics). The mean number of dead VA chips each day is shown versus time in Figure 4.10(b), this shows a clear drop in the number of dead VA chips as a function of time. The situation has improved to such an extent that lately the most common situation is that there are no dead VA chips in the readout electronics.

(a) Supermodule 1 coil currents

(b) Supermodule 2 coil currents

Figure 4.8: Coil currents versus days since 1st August 2003 (day 0). During the period from 21st June 2004 (day 325) until 1st February 2005 (day 550) the polarity of both coils was reversed, this period is indicated by the blue points. Periods of normal coil polarisation are indicated by red points.

(a) Dead chips in Supermodule 1

(b) Dead chips In Supermodule 2

Figure 4.9: The number of dead chips versus days since 1st August 2003 (day 0). The blue points represent times when no high voltage trip has been flagged, while the red points represent times when a high voltage trip has been flagged.

(a) Distribution of dead VA chips, when no HV trips have been flagged

(b) Mean number of dead VA chips versus days since 1st August 2003 (day 0)

Figure 4.10: Distributions of the number of dead VA chips in the MINOS Far Detector.

4.6.3 Busy Electronics

As discussed in Section 4.3.3, each time a VA chip reads out a PMT, at least 5 μ s of dead time is incurred. The readout is activated when two dynode triggers are recorded that satisfy the 2/36 VARC trigger conditions. Due to the nature of the readout, if the triggers occur in VA chips serviced by different VMMs (85% of the time), both VA chips will go dead for 5 μ s. However if the VA chips triggered are serviced by the same VMM (15% of the time), one VA chip will be queued while the other is directed to read out, producing dead times of 5 μ s and 10 μ s respectively. This gives an average dead time of 5.375 μ s each time a chip is read out. The mean rate of hits on each VA chip is 300 Hz [102]. Combining these two pieces of information we find that each chip is busy on average 0.16% of the time. This corresponds to around 2.3 busy chips in the Far Detector at any given moment during a run.

A 30 μ s pre-trigger window of detector hits is written out with each snarl. This can be analysed offline to determine which VA chips were missing from the readout during the triggered event. Figure 4.11(a) shows the distribution of busy VA chips (not including the VA chips which read out the veto shield strips) obtained using the pre-trigger windows of events. The peak at two arises from the VARC trigger conditions, which require at least two coincident dynode triggers. Events with one busy chip occur when one of the dynode triggers is sparsified away. The vast majority of events contain less than ten busy chips, with the mean number of busy chips calculated to be 2.3 (corresponding to 0.16% of VA chips in the detector readout electronics). This value is in excellent agreement with the predicted number of busy chips. Due to the fact that all shield activity is recorded in the detector, independent of any VARC trigger, VA chips which readout shield strips are not included in these plots (see Section 4.7). Figure 4.11(b) shows the mean number of busy VA chips plotted against time. This slowly decreases with time. This decrease reflects the reduction in detector noise from photons emitted due to long-term relaxation of mechanical stress in the WLS fibres over time, as was discussed in Section 4.3.2. The drops at around 200 days and at around 400 days in follow retuning of PMT dynode trigger thresholds.

4.7 The Veto Shield

A veto shield was constructed above the detector to enable the selection of atmospheric neutrino events from the data [107]. The veto shield is used to tag cosmic muons entering

(a) Distribution of busy VA chips

(b) Mean number of busy VA chips versus days since 1st August 2003 (day 0)

Figure 4.11: Distributions of the number of busy VA chips in the MINOS Far Detector, made using hits in the pre-trigger window.

the detector from above or the sides. If the tagging efficiency is high enough then the cosmic muon background to the atmospheric neutrino event selection can be reduced to a manageable level.

The same kinds of scintillator modules as those used in the main detector were used to build the veto shield. The modules were placed along the top and sides of the detector and were aligned along the z direction in a layout illustrated in Figure 4.2(a). The veto shield is divided into four overlapping sections in the z direction. Sections 1 and 2 cover the first supermodule, while sections 3 and 4 cover the second supermodule. A double layer of modules directly above the detector comprises the centre of each shield section. The rest of each section is made up of single layers diagonally above and beside the detector. The highest flux of cosmic muons is incident on the top of the detector. The double layer of modules ensures the highest tagging efficiency in this region.

The shield module readout is integrated with the main detector readout. As with the main detector, fibres run from each scintillator strip end into M16 PMTs housed in MUX boxes. The multiplexing differs from that for the main detector however. The multiplexing is the same on both sides of the detector, with groups of four or eight adjacent scintillator strips readout by a single PMT pixel. These groups of strips are known as *shield planks*. Each plank has a resolution in the xy plane of 10 cm. The PMTs are read out by VA chips and digitised by VMMs and share common VARC with the main detector readout. The trigger conditions differ from those in the main detector in order to optimise the muon tagging efficiency. All shield activity in a time frame is recorded, this is achieved by not including the shield in the VARC trigger.

There are a number of reasons why cosmic muons may not be tagged by the shield as they enter the detector:

1. A shield channel may be dead, due to electronics or damage to the scintillator strips during installation.
2. A shield channel may be busy due to noise or radioactivity.
3. A cosmic muon may pass through the gap between two shield modules.
4. The activity produced in the shield by a cosmic muon may be below the 1-2 PE threshold.

The efficiency of individual planks can be determined using the *CandShield* package [108]. This package is able to project a reconstructed track back through the shield and determine which plank(s) it should have passed through and hence where in the shield

energy deposits should be visible. By comparing the number of expected shield hits for each plank with the number of in-time hits that were actually observed, an estimate of the tagging efficiency for each plank can be obtained. If there are any planks with low tagging efficiencies, then this will lead to an increase in the number of cosmic muon background events in any selected atmospheric neutrino sample, so it is important that any problems with the shield are known. Tagging inefficiencies can arise from single PE noise. In order to reduce this the dynode threshold is set to a higher level than in the main detector and is equivalent to 1-2 PE.

Figure 4.12 shows the overall average tagging efficiencies for each shield plank from August 2003 to February 2006. Most planks have tagging efficiencies of over 80%, as can be more clearly seen from Figure 4.13. The efficiencies of some planks is lower, but in the majority of cases these are in regions where the shield is two layers thick, so this helps to compensate. Low efficiency regions along the sides of the detector are less of a problem as most of the background cosmic-muon flux is steeply downward going. The effects of the veto shield efficiency on background rejection in an atmospheric neutrino analysis will be discussed in Section 8.6.5.

Figure 4.12: The tagging efficiencies for each shield plank. The position of each point represents the position of the centre of the shield plank in the xy plane, the colour-scale indicates the tagging efficiency of the particular plank.

Figure 4.13: Distributions of tagging efficiencies for each shield plank, divided by shield section.

Chapter 5

Calibrating The MINOS Far Detector

“A common mistake that people make when trying to design something completely fool-proof is to underestimate the ingenuity of complete fools.”

— Douglas Adams, 1952–2001

5.1 Overview

This chapter describes the steps taken to ensure that timing and pulse height information in hits read out by the MINOS detectors are as accurate as possible. Section 5.2 describes the calibration of hit times in the MINOS Far Detector. This is extremely important to the atmospheric neutrino analysis in this thesis as the dominant background to upward-going atmospheric neutrino events is downward-going cosmic muon tracks mistakenly identified as upward-going tracks.¹ The requirement of the timing calibration is that the error in the timing calibration should be small compared to the overall timing resolution of the detector. Section 5.3 briefly describes the calibration of beam neutrino spill times used for the main MINOS beam neutrino analyses [74], which is also important for efficient vetoing of beam neutrino ‘background’ events from the atmospheric neutrino

¹Currently no timing calibration is performed for the MINOS Near Detector as the front end electronics mean that the timing resolution is at best 19 ns and cosmic muon backgrounds are not an issue for beam neutrino studies.

analysis. Finally, the calibration of hit pulse heights is described in Section 5.4.

5.2 Timing Calibration

The overall correction, $\Delta T(s, t, Q)$, made to the hit times from each strip-end is made up of three components:

$$\Delta T(s, t, Q) = \Delta T_{timewalk}(Q) + \Delta T_{jump}(c(s), t > t_{jump}) + \Delta T_0(s). \quad (5.1)$$

In the above equation:

- $\Delta T_{timewalk}(Q)$ is the time walk correction for a hit with ADC count Q , which compensates for pulse height dependent variations in signal arrival times and is discussed further in Section 5.2.2.
- $\Delta T_{jump}(c(s), t > t_{jump})$ is the shift in time offset correction for a hit registered at time t in a strip-end s due to a hardware swap at time t_{jump} in electronics channel c , which reads out s . This correction also includes a correction for any long term drift in the timing of the channel. Corrections for hardware swaps are discussed in Section 5.2.3, while corrections for long term timing drift are discussed in Section 5.2.4.
- $\Delta T_0(s)$ is the combined time offset from scintillator strip and readout channel for strip-end (s) and is discussed further in Section 5.2.5.

5.2.1 Calibration Software

The timing calibration software was written by the author and based on earlier work[94]. The calibration procedure is closely coupled to the detector readout. The calibration uses straight through-going muon tracks with little or no shower activity, where the longitudinal position of the track in each strip it passes through can be accurately determined. Only runs where sufficient statistics have been accumulated for each strip are used.

In order to ensure that the muon tracks used in the calibration are well measured, they are required to span at least twenty scintillator planes and satisfy straight line fits in both U and V views with RMS deviations of < 1 cm. This means that the 3D position of each hit can be reconstructed by combining the straight line fits. These criteria select around

750,000 muons/month from the data, this corresponds to 400-500 hits/strip-end/month.

Before discussing the timing calibration it is useful to define a number of quantities. (NB the calibration is performed on hit times multiplied by the speed of light.)

East - West Times: Each scintillator strip which is hit in an event can be read out on both the east and west sides of the detector. In the case of perfect calibration, the two hit times should be identical after accounting for propagation time. In the raw data this is not the case. A variety of factors, such as the position of the hit along the strip, the lengths of readout fibres and response times of the electronics, mean that there can be large differences between the two times. The east - west times used in the calibration are calculated as follows. First the hit times are corrected for the position of the hit along the strip and corrections for the differing fibre lengths are applied. Next a charge weighted average east-west time multiplied by the speed of light for each electronics channel, ΔcT_{EW} , is calculated (note this time is an average over all the strips which are readout by a PMT during a run).

VA Chip Times: These times are calculated using information from just one side of the detector. In the event that a track crosses all the planes readout by a given VARC, the times for each VA chip on the VARC can be fitted against the distance along the track, S ($S = cT + S_0$). The residual for each VA chip is then averaged over a run. This is the VA chip time for the channel for that run, ΔcT_{VAchip} .

VARC Times: If a track crosses all the VA chips read out by two adjacent VARCs, the VA chip times can be fitted as before for each VARC. In this case though the intercepts of the straight line fits to the hit on each VARC can be compared. This gives the VARC time for the channels for the run. Most VARCs have two neighbours and hence each channel can have two VARC times, ΔcT_{VARC0} and ΔcT_{VARC1} .

In order to speed up the later stages of the calibration these values are calculated once for each channel, for each run and then saved to a file.

5.2.2 Time Walk Corrections

The decay time for emission of light by the WLS fibres is approximately 8 ns [90]. The higher the number of excited molecules, the more completely the decay curve will be sampled and the more likely there will be a prompt photon. This means that in general the first photons from hit strips with higher total pulse heights will arrive earlier than

than those with lower total pulse heights. As the dynode trigger on each PMT is set to approximately $\frac{1}{3}$ PE (compared to ~ 5 PE per strip-end for a typical muon), then the trigger times for such hits will be earlier also. This is the source of optical time walk in the MINOS Far Detector.

The time walk corrections used in this thesis were calculated in an iterative process [94] where electronics swaps were calculated, followed by strip-to-strip timing offsets, followed by time walk corrections, which were then used to calculate a more accurate set of electronics swaps, followed by a more accurate set of strip-to-strip timing offsets and so on, until the corrections were stable. The time walk correction was calculated as a function of the pulse height in ADCs ($65 \text{ ADC} = 1 \text{ PE}$). Timing fits were applied to muon tracks using only the hits from track strips with pulse heights greater than 5000 ADC, where the time walk effects are considered to be small. By calculating the mean time displacements from these fits for track hits with pulse heights below 5000 ADC as a function of the hit ADC, the time walk correction was calculated. The parameterisation obtained for the mean time walk correction for a strip-end reading out pulse height Q is as follows:

$$\frac{\Delta T_{timewalk}}{\text{ns}} = 20.66 - 2.0178 \left(\ln \frac{Q}{\text{ADC}} \right) - 0.2337 \left(\ln \frac{Q}{\text{ADC}} \right)^2 + 0.0218 \left(\ln \frac{Q}{\text{ADC}} \right)^3 \quad (5.2)$$

5.2.3 Electronics Swaps

Each individual electronics component will have very slightly different response times. Therefore when an electronics component is swapped with another, the timing of the channels it read out will be altered. These jumps in timing must be removed before channels can be calibrated relative to each other. The electronics swap calibration is done using a sliding two month sample of data (*i.e.* data for January and February are checked for jumps followed by February and March and so on). This ensures that jumps near the beginning and end of months are not missed.

Each run/electronics channel combination has an associated object called a *TimePoint*. Each *TimePoint* contains: ΔcT_{EW} , $\Delta cT_{\text{VAchip}}$, ΔcT_{VARC0} , ΔcT_{VARC1} , their associated errors and the number of hits used to calculate each value. For each channel the *TimePoints* for ten consecutive runs are grouped together in another object called a *TimePeriod*. As shown for an example channel ([15 : 1 : 4 : 0 : 1]) in Figure 5.1, here a number of jumps are visible. *TimePeriods* contain additional information on the mean

Figure 5.1: Uncalibrated ΔcT_{EW} values for channel $[15 : 1 : 4 : 0 : 1]$ for runs in August and September 2003. Colours of points indicate initial *TimePeriod* divisions. Gaps are due to pulse height calibration and test runs.

and standard deviation of the ΔcT_{EW} (μ_{EW}, σ_{EW}), ΔcT_{VAchip} (μ_{VA}, σ_{VA}), ΔcT_{VARC0} ($\mu_{VARC0}, \sigma_{VARC0}$), ΔcT_{VARC1} ($\mu_{VARC1}, \sigma_{VARC1}$) values they contain.

Once all the *TimePoints* have been placed in *TimePeriods*, the software loops over the electronics channels. For each channel it attempts to reorganise the *TimePoints* between *TimePeriods*, so that any jumps in ΔcT_{EW} occur at the boundaries between two *TimePeriods*. It does this by exchanging *TimePoints* between two adjacent *TimePeriods* and attempting to jointly minimise the σ values of the two *TimePeriods*. This is done once for all neighbouring pairs of *TimePeriods* for a channel. Next the software calculates the minimum size jump in μ_{EW} it could resolve between each neighbouring pair of *TimePeriods*, $\text{MinJump}_{EW} = 5.0 \times \sqrt{\sigma_{EW1}/N_{EW1} + \sigma_{EW2}/N_{EW2}}$, where N_{EW1} and N_{EW2} are the numbers of hits used in the calculation of σ_{EW1} and σ_{EW2} respectively. This calculation controls how many jumps the software will find. Decreasing the size of the constant will increase the likelihood of a jump being found. The value has been adjusted to strike a balance between finding all important jumps and not overloading the database with too many small jump corrections, which would slow down reconstruction. If the actual size of the jump in μ_{EW} is less than MinJump_{EW} the two *TimePeriods* are merged into one. This is then repeated for all adjacent pairs of *TimePeriods*. The exchanging and merging steps are then repeated up to twice more, less if only one *TimePeriod* remains at any point. The result of this process is shown for the example channel is shown in Figure 5.2

Figure 5.2: Uncalibrated ΔcT_{EW} values for channel $[15 : 1 : 4 : 0 : 1]$ for runs in August and September 2003. The colours of the points indicate *TimePeriod* divisions after sorting and merging of *TimePoints*. Gaps are due to pulse height calibration and test runs.

The software now has a list of one or more *TimePeriods* for each channel. The software understands that each channel has an opposite channel on the other side of the detector which will have the same ΔcT_{EW} values for each run and hence the same jumps. In order to resolve the ambiguity as to which side of the detector the jump took place (cases where electronics involved in the readout of both channels were swapped at the same time will be discussed later) the software looks for a jump in μ_{VA} for the channel and its opposite. This is done by calculating $\text{MinJump}_{VA} = 1.0 \times \sqrt{\sigma_{VA1}/N_{VA1} + \sigma_{VA2}/N_{VA2}}$ (where the definitions are as above except with East-West times replaced with VA chip times) and requiring the size of the jump in $\Delta cT_{VA\text{chip}}$ to be greater than MinJump_{VA} .

If only one side has a jump in μ_{VA} , this indicates the jump has most likely been caused by the changing of one VA chip on that side. The channel number, size and which run it was first observed are then recorded in a file. If both sides register a jump then the ratio $(\text{Jump Size})/(\text{MinJump}_{VA})$ is calculated for each. If one is more than 50 times larger than the other then the jump is assumed to have occurred on that side. If not then it is assumed that electronics involved in the readout of both channels were swapped at the same time. (There are procedures in place at the Far Detector to try to ensure that at least one run is taken between swapping electronics on each side of a strip should that be necessary.) In this case jumps are recorded for both the channel and its opposite channel using the size of the jumps in μ_{VA} rather than that in μ_{EW} .

Figure 5.3: Uncalibrated ΔcT_{VAchip} values for channel $[15 : 1 : 4 : 0 : 1]$ for runs in August and September 2003. The colours of the points indicate *TimePeriod* divisions after sorting and merging of *TimePoints*. Gaps are due to pulse height calibration and test runs.

In some cases a jump is observed in μ_{EW} , but neither of the channels show a jump in μ_{VA} . This could be caused by the swapping of a VARC on one side. In this case μ_{VARC0} and μ_{VARC1} values are examined for jumps. (VARC swaps have been rare so far in the experiment and no simultaneous VARC swaps on both sides of a strip have ever been done according to the engineers at the Far Detector.) If one or both VARC times on one side of the detector show a jump then the jump is recorded on that side of the detector. If both or neither sides record a jump in VARC times then a warning is given and a detective has to try and unravel the mystery of what actually happened.

In the case of the example channel, the ΔcT_{VAchip} values in Figure 5.3 show clear jumps at the runs predicted from the ΔcT_{EW} values, while the ΔcT_{VAchip} values for the opposite channel (channel $[14 : 1 : 4 : 0 : 1]$), shown in Figure 5.4, show no jumps. Therefore the software will decide that the jumps all occur in channel $[15 : 1 : 4 : 0 : 1]$. This is shown in Figure 5.5.

Once this has been done for all ‘*channel/opposite channel*’ pairs, the ΔcT_{EW} , ΔcT_{VAchip} , ΔcT_{VARC0} , ΔcT_{VARC1} values are recalculated using the corrections for the jumps already found. The whole jump finding process is now repeated. New smaller jumps which were previously masked may now be visible and these additional corrections inserted into the list of those that were found previously. Figure 5.6 shows the ΔcT_{EW} values for the

Figure 5.4: Uncalibrated ΔcT_{VAChip} values for the channel on the opposite side of the detector to channel [15 : 1 : 4 : 0 : 1] (channel [14 : 1 : 4 : 0 : 1]) for runs in August and September 2003. The colours of the points indicate *TimePeriod* divisions after sorting and merging of *TimePoints*. Gaps are due to pulse height calibration and test runs.

Figure 5.5: Uncalibrated ΔcT_{EW} values for channel 15 : 1 : 4 : 0 : 1 for runs in August and September 2003. The colours of the points indicate *TimePeriod* divisions after sorting and merging of *TimePoints*. The vertical arrows between the horizontal lines indicate the magnitudes and directions of the jumps identified. Gaps are due to pulse height calibration and test runs.

Figure 5.6: Partially calibrated ΔcT_{EW} values for channel $[15 : 1 : 4 : 0 : 1]$ for runs in August and September 2003. The colours of the points indicate *TimePeriod* divisions after sorting and merging of *TimePoints*. Gaps are due to pulse height calibration and test runs.

example channel corrected for the jumps found in the first iteration. The colours of the points still indicate the *TimePeriods* found in the first iteration. As is clear from the plot, subsequent iterations merge all *TimePoints* for this channel (in this time range) into a single *TimePeriod*. The values are still offset from zero as the strip-strip timing calibrations have not yet been applied.

5.2.4 Timing Drift

For each channel, all the jumps in timing should now have been removed. There is still the possibility that the timing may still drift slowly over time. This drift is calibrated out in a similar way to the jumps caused by electronics swaps. *TimePoints* for one month's worth of runs are grouped together in a single *TimePeriod*. Adjacent month long *TimePeriods* are compared. In order to decide whether not a correction should be applied, the size of the shift in mean ΔcT_{EW} is tested against the same criteria as for electronics jumps. Any shifts which pass this check are checked in the same way as the electronics jumps in order to decide which side of the detector they occurred on. The found shifts in timing are added as extra jumps to the calibration constants.

5.2.5 Strip-to-Strip Calibration

The $\Delta T_0(s)$ correction removing the timing offsets between scintillator strip-ends, is calculated as a number of separate contributions.

$$\Delta T_0(s) = \Delta T_{\text{PMT}} + \Delta T_{\text{VARC}} + \Delta T_{\text{strip-end}} + \Delta T_{\text{East/West}} , \quad (5.3)$$

where ΔT_{PMT} is a correction for the timing offsets between PMTs on the same VARC; ΔT_{VARC} is a correction for timing offsets between VARCs on the same side of the detector; $\Delta T_{\text{strip-end}}$ is a correction for remaining timing offsets between strip-ends on the same side of the detector and $\Delta T_{\text{East/West}}$ is a correction for timing offsets between the east and west sides of the detector.

In order to calculate ΔT_{PMT} , muon tracks which leave hits on all the planes read out by a given VARC are selected. Timing fits are applied to each track, using only the hits read out by the current VARC. The value of ΔT_{PMT} for each PMT on the VARC is given by the mean residual of the fits for the PMT. The time offsets between PMTs on a VARC are due to variations in the rise time of dynode signals in each of the VFBS and the propagation time of signals sent between the VFBS and VARCs.

The calculation for ΔT_{VARC} uses muons which pass through all the planes read out by two adjacent VARCs. Separate timing fits are calculated for the track hits read out by each of the two VARCs, using the ΔT_{PMT} corrections previously calculated. The separate fits are then extrapolated back to the boundary between the two VARCs and the time difference calculated. Using many tracks the mean time difference between each pair of VARCs is calculated. The ΔT_{VARC} corrections are calculated using the mean time differences, together with the constraint that the mean ΔT_{VARC} correction for each side of the detector is zero. The offsets between VARCs are due to the propagation time of the signals sent from the TCU to the TRCs and the signals sent from the TRCs to the VARCs.

Timing fits to muon tracks are calculated, using hits from one side of the detector and with times corrected using the ΔT_{VARC} and ΔT_{PMT} corrections previously calculated. The $\Delta T_{\text{strip-end}}$ corrections for that side of the detector are calculated from the mean residuals for each strip-end. The process is then repeated for the other side of the detector. The remaining offsets between strip-ends are due to variations in the rise time of dynode signals between pixels on each PMT and small residuals in the measured lengths of the readout fibres.

The final stage is to use double-ended strips on muon tracks to project the corrected times for the east and west strip-ends back to the track. The mean time difference between east and west times is then calculated and a constant applied to correct this difference back to zero. The overall $\Delta T_0(s)$ correction for each strip-end can now be calculated as the sum of all the individual terms.

5.2.6 Validating The Calibration

The timing calibration is validated by recreating the calibration files described in Section 5.2.1, but this time using the constants calculated in the steps described above to correct the times for each strip-end. Figure 5.7 and Figure 5.8 show the results of calibration on data from 1st August 2003 to 31st October 2004 and from 1st November 2004 to 28th February 2006 respectively. It is clear from these plots that calibration is stable over time. In the un-calibrated data plots for channels in each VME crate on the east side of the detector the time differences appear to slope upward. This is followed by a downward slope in time differences for the channels on the west side of the detector which read out the same set of planes. This is because the signals generated by the TRCs in each VME crate propagate in opposite directions on opposite sides of the detector. Therefore the smallest time offsets on the east side of the detector correspond to the largest time offsets on the west side of the detector.

As defined in Section 5.2.1, ΔT_{EW} is the charge-weighted average of East - West times. Figure 5.9 shows the distribution of ΔT_{EW} before calibration. The uncalibrated timing differences vary between -20 ns and +12 ns. Figure 5.10 shows the calibrated distribution of ΔT_{EW} values. The majority of timing differences lie in the region between -4 ns and +4 ns and are very strongly peaked around zero. Outlying points are caused by malfunctioning components, which produce a small degree of instability in the measured times. Figure 5.11 shows the result of a Gaussian fit to the calibrated ΔT_{EW} distribution. An *RMS* of 0.433 ns is obtained. The mean statistical error in each ΔT_{EW} value is calculated to be 0.09 ns. This is small compared to the overall *RMS* of the ΔT_{EW} distribution, implying that the *RMS* is still dominated by errors in the calibration. The *RMS* of the ΔT_{EW} distribution is a measure of the quadrature sum of the mean calibration errors on the east and west sides of the detector. The mean calibration error on a single strip-end can therefore be estimated as $\sigma_{\Delta T} = 0.433/\sqrt{2} = 0.31$ ns. This indicates that this method has successfully calibrated the timing system at the sub-nanosecond level.

Figure 5.7: ΔT_{EW} distributions plotted versus electronics channel (channel = $108 \times crate + 36 \times varc + 6 \times vmm + 3 \times vaadc + vachip$) for the period 1st August 2003 to 31st October 2004. The red points represent uncalibrated data, while the blue points represent calibrated data.

Figure 5.8: ΔT_{EW} distributions plotted versus electronics channel (channel = $108 \times crate + 36 \times varc + 6 \times vmm + 3 \times vaadc + vachip$) for the period 1st November 2004 to 28th February 2006. The red points represent uncalibrated data, while the blue points represent calibrated data.

Figure 5.9: Distribution of ΔT_{EW} before calibration.

Figure 5.10: Distribution of ΔT_{EW} after calibration.

Figure 5.11: Gaussian fit to calibrated ΔT_{EW} distribution.

The overall resolution of the timing system is a convolution of the intrinsic resolution of the timing system (1.56 ns) with fluctuations in arrival times of photons at the photocathode.² For each track in a sample of stopping muons, the calibrated hit times are fitted to $s = cT + cT_0$, where s is the distance along the track. The *RMS* value of the deviations for each stopping muon track is then calculated. Figure 5.12 shows the distribution of *RMS* deviations for stopping muons which span at least twenty scintillator planes and satisfy straight line fits in both U and V views with *RMS* deviations of < 1 cm. The mean of this distribution is a measurement of the overall timing resolution. A Gaussian fit to this distribution gives a resolution of 2.47 ns. Thus, as $0.31 \text{ ns} \ll 2.47 \text{ ns}$ the goal of the calibration has been achieved.

²The time walk correction in the calibration corrects for the variation of first photon arrival time with total pulse height, but fluctuations about this mean correction will still occur for individual hits and so this smearing effect will still be present in any single hit timing measurements.

Figure 5.12: Distribution of RMS deviations for muon timing fits. The mean of the fitted Gaussian is quoted in ns.

5.3 Spill Time Calibration

A GPS receiver is located at both the Near and Far Detector. At each site the GPS receiver is used as a clock for the front-end electronics and DAQ computers. The GPS receivers are located underground and connected by optical fibre to surface antennas. The worst-case resolution of each GPS is approximately 200 ns during normal operation. The GPS units are used to capture spills in two ways. The Near Detector GPS receiver is used to tag the time at which Main Injector protons arrive at the target. The GPS receiver at the Far Detector records the time of each DAQ trigger. These times can be compared offline, taking into account antenna delays, hardware offsets and the 2.449 ms time of flight of beam neutrinos, in order to identify Far Detector events which are consistent with beam neutrinos. Due to uncertainties in hardware delays, there is a $< 1 \mu\text{s}$ uncertainty on the absolute time offset between the two detectors.

The second method used by MINOS to capture in-spill events relies on the buffering capacity of the DAQ. The spill time from the Near Detector GPS is sent via the Internet to Soudan. This information is used by the Trigger Processors at the Far Detector to put a window around the time of the spill, in which all data are recorded (see Section 4.4.2). This allows for events with very low energy deposits ($\sim 300 \text{ MeV}$) to be recorded. The dominant causes of inefficiency in this method are temporary network outages and transient network latencies between Fermilab and Soudan. Despite this the two methods reliably capture 99% of all spills [109].

5.4 Pulse Height Calibration

The MINOS experiment intends to measure Δm_{32}^2 to an accuracy of better than 10%. For this reason energy-scale calibration goals were set at 2% relative systematic uncertainty and 5% absolute uncertainty. The calibration uses both the optical light-injection system, which measures the behaviour of the readout instrumentation, and cosmic-ray muon tracks, to measure the response of the scintillator. The calibration is performed independently for the Near Detector and Far Detector. Consider the raw pulse height $Q_{raw}(s, x, t)$ measured in strip s at time t at a distance x along the strip. The calibration is a multi-stage process which converts $Q_{raw}(s, x, t)$ into a fully corrected signal Q_{corr} . Each stage applies a multiplicative calibration constant:

$$Q_{corr} = Q_{raw} \times D(t) \times L(s, Q_{raw}) \times S(s, t) \times A(s, x) \times M(d) \quad (5.4)$$

In the above equation:

- $D(t)$ is the drift correction, which compensates for PMT, electronics and scintillator response changing with temperature and age and is discussed further in Section 5.4.1.
- $L(s, Q_{raw})$ is the function which linearises the response of each channel with pulse height. The linearity correction is discussed further in Section 5.4.2.
- $S(s, t)$ is the strip-to-strip correction which removes differences in response, strip-to-strip and channel-to-channel. This is discussed further in Section 5.4.4.
- $A(s, x)$ is the attenuation correction. This describes the attenuation of light depending on event position along each strip and is discussed further in Section 5.4.6.
- $M(d)$ is an overall scale factor that converts corrected pulse height into the same absolute energy unit for the Near Detector, Far Detector and Calibration detector. This is discussed further in Section 5.4.7

5.4.1 Gain Calibration

During normal data taking, the light injection system periodically pulses the fibre from every strip-end to monitor the stability and gain of each channel (there are vetoes in place to stop this occurring during a beam spill). Each strip-end is pulsed at a rate of 300 pulses per hour. The intensity of the pulses is tuned to so that the PMT pixel in question receives approximately 50 PE per pulse. The pulses are monitored by PIN

diodes to allow corrections to be made for time instabilities.

These data are used to provide corrections for variations in PMT gains, electronics and other transient instabilities. Most short-term variations can be eliminated through good environmental controls (minimising temperature variations etc). Long term drifts in the response are still observed, possibly due to seasonal environmental changes and the effects of ageing detector components. These drifts equate to changes in gain of $\lesssim 4\%$ per year.

In order to track these changes, the data from each month are collated and used to compute the average response per photo-electron per channel. This is done using photon statistics [91] [110]. By comparing the RMS widths of many pulses to the mean, the number of ADC counts per photo-electron is found for each channel. The offline software uses these “gains” for Monte Carlo simulations, cross-talk removal and other reconstruction purposes.

5.4.2 Linearity Calibration

The PMT pixel responses become non-linear for light levels $O(100 \text{ PE})$. In addition, the Far Detector electronics have a non-linear response on a similar scale. Both components are linearised with a single correction. The LI system is used to convert the ADC scale into a scale which is linear in light. The fibres are pulsed at a series of light levels chosen to cover the entire dynamic range of the ADC readout. The PMT response is then mapped out using the PIN diodes. These “gain curves” are run once a month, interspersed with normal data taking.

5.4.3 Cosmic-ray Muon Drift Calibration

The light injection system measures the time variations of the PMTs and the electronics, but it cannot monitor changes in the scintillator and WLS fibres. Changes in scintillator and WLS fibre responses can be caused by temperature variations and ageing. These changes can be monitored using cosmic-ray muons events, which occur at a rate of $\sim 0.5 \text{ Hz}$ at the Far Detector. This rate is sufficient to track the response of the detector on a daily basis. This “drift” calibration is performed by measuring the total pulse height per plane deposited by through-going cosmic-ray muons. The integrated response of this measure is expected to be constant with time. The drift is calculated by computing the

daily median pulse height per plane and comparing this to an initial value:

$$D(t) = \frac{\text{Median Response}(t_0)}{\text{Median Response}(t)} \quad (5.5)$$

This correction includes changes due to the scintillator, WLS fibres, PMTs and the electronics. The values obtained show that the Far Detector response drifts at a level of $\sim 2\%$ per year. Long term drifts due to aging of scintillator, as well as short-term variations, correlated with environmental changes, are observed.

5.4.4 Strip-to-Strip Calibration

The strip-to-strip (channel-by-channel) time-dependent response of the detector, $S(s, t)$, is measured using through-going cosmic-ray muon data. This calibration relates the mean response of each strip-end to the detector average.

$$S(s, t) = \frac{\text{Mean Response of the Detector}(t)}{\text{Mean Response of the Strip} - \text{End}(s, t)} \quad (5.6)$$

There are a number of causes of variations in response across the detector:

- Differing scintillator light output.
- Varying WLS fibre light collection efficiency.
- Varying readout fibre attenuation.
- Scale factors not removed by linearity calibration (e.g. PIN gain and PMT quantum efficiency).

Cosmic-ray muon tracks are used to measure the mean light level, post-linearisation, at each strip-end. To remove known spatial and angular dependencies, attenuation and path length corrections are applied to each hit. In this way, the mean response of a muon travelling perpendicularly through the centre of the strip can be calculated. This response is then used to calculate the calibration constant for the strip-end as above. A statistical approach is used to calculate the frequency with which a track just clips the corner of a strip [111]. One slight subtlety is that due to the low light level at the PMT face ($\sim 2 - 10$ PE) the mean light level calculation must account for the Poisson probability of producing 0 PE. As the spatial resolution of the detectors is insufficient to predict missed strips from event topology, an iterative method is used to estimate the most probable light yield of a strip. This is then used to calculate the zero contribution

probability [112].

5.4.5 Attenuation Calibration

Cosmic-ray muon data could be used to correct the variation in light caused by attenuation along the WLS fibre in a scintillator strip, but there is a more accurate source of information. Namely, the module mapper measurements taken at each module production facility and at Soudan. The module mapper uses a collimated 5 mCi source ($E_\gamma = 662$ keV) to illuminate scintillator strips in modules. The $4\text{ cm} \times 4\text{ cm}$ collimated beam was moved across the face of the module being tested using a computer controlled $x - y$ drive. The resulting measurements provided detailed maps of the response of all scintillator modules to ionising radiation every 8 cm along the length of each strip.

The mapper data for each strip are fitted to a double exponential function:

$$A(x) = A_1 \exp -x/L_1 + A_2 \exp -x/L_2, \quad (5.7)$$

where x is the length along the strip and L_1 and L_2 stand for the two attenuation lengths (see Section 4.3.1). The parameters from these fits are used to correct observed data.

5.4.6 MIP Calibration

The calibration steps considered so far, have produced temporally and spatially uniform detector responses. A relative calibration is then necessary to normalise the energy scales at the Near, Far and Calibration detectors to be the same to within the 2% goal. Stopping cosmic-ray muons are used for this task because they are abundant enough at all detectors and their energy depositions in each plane can be accurately determined from range measurements.

Muon energy loss by ionisation is described by the Bethe-Bloch equation, which determines the stopping power for muons in a particular material as a function of muon momentum, as shown in Figure 5.13. Using this along with momentum from range measurements it is possible to measure the relative responses of the three detectors. The accuracy of this calibration was improved using a track-window technique [113]. Among other things, this technique only uses the region of the Bethe-Bloch curve for muon momenta between 0.5 GeV and 1.1 GeV, avoiding using the region near the end of

Figure 5.13: The green stars show the Bethe-Bloch calculation of the stopping power for muons in polystyrene scintillator. The black dots and red triangles show the response of stopping muons in the Far Detector data and GEANT3 Monte Carlo simulations respectively. (Both real and simulated data points have been calibrated to give the expected stopping power at the minimum ionising point.) Taken from [109].

a track where ionisation increases rapidly (see Figure 5.13). The calorimetric responses of the three detectors, $1/M(d)$, are used to normalise the detectors' energy scales to within in the 2% target.

5.4.7 Energy Calibration

In order to determine the energy scale, the Calibration Detector has measured the response of the MINOS detectors to different particle types with known energies[114][115]. The results have been used to tune the models used to simulate particle showers. For this analysis, the overall energy scale is set by measuring the mean response of the Far

Detector to cosmic-ray muons. The detector simulation has then been tuned to reproduce this response. The reconstruction of neutrino shower energies is then tuned to the detector simulation. This is discussed in the following chapter.

Chapter 6

Monte Carlo Simulation

“The degree of one’s emotion varies inversely with one’s knowledge of the facts – the less you know the hotter you get.”

— Bertrand Russell, 1872–1970

6.1 Overview

The Monte Carlo simulation of atmospheric neutrino and cosmic-ray muon events in the MINOS Far Detector is used to develop the reconstruction and selection of atmospheric neutrino events and rejection of backgrounds. The Monte Carlo is also used in the oscillation analysis to calculate the expectation in the case of particular oscillation parameters, which can then be compared with the data. This chapter discusses how the cosmic-ray and atmospheric neutrino fluxes are simulated, then how the response of the MINOS Far Detector to events is simulated and finally evaluates the results of simulations by comparing simulated distributions with those from data.

6.2 Cosmic-Ray Muon Simulation

6.2.1 Cosmic-Ray Muon Flux

As mentioned briefly in Chapter 2, the Earth is constantly bombarded by a stream of cosmic-rays which strike the top of the Earth's atmosphere at a rate of $\sim 1000 \text{ m}^{-2}\text{s}^{-1}$. Protons dominate the primary cosmic-ray flux (95%), with helium nuclei comprising (4.5%) at energies above 2 GeV/nucleon. C, N and O, plus heavier elements, make up the remaining 0.5% of the flux [43] [116]. Cosmic-rays have been observed with energies from 10^7 eV up to at least 10^{20} eV . For cosmic-rays with energies below 10^{12} eV , direct measurements have been made using a series of air-borne [117] [118] [119] and space-borne experiments [120]. Figure 6.1 shows a collation of such results for the major components of the primary cosmic-ray flux [121]. Arrays of ground based detectors [122] [123] [124] [125] are used to study higher energy cosmic-rays. These cosmic-rays are studied indirectly by observing the large air showers produced by their interactions with molecules in the atmosphere. The energy of the primary cosmic-ray is then inferred from the multiplicity of particles in the air shower or the amount of radiation emitted in the atmosphere.

Figure 6.2 shows the “all-particle” spectra measured by these air shower experiments [43]. For energies below 10^{10} eV , the cosmic-ray flux is modulated by the *solar wind*. The solar wind is a steady stream of charged particles emitted from the surface of the sun, whose magnetic field shields the inner solar system from the flux of low energy cosmic-rays arriving from outside the solar system. For energies above 10^{10} eV , the observed energy spectrum appears to follow a power law distribution $N(E) \sim E^{-\gamma}$ over a wide energy range, with the exponent, γ measured to be around 2.7. The spectrum steepens to $\gamma \sim 3.0$ above $5 \times 10^{15} \text{ eV}$. This feature is commonly known as the *knee* of the spectrum. The bulk of the cosmic-rays below the *knee* are believed to originate within the Milky-Way galaxy, with the power law shape resulting from an acceleration process (e.g. supernova explosions) that approaches its maximum at the *knee*. The energy spectrum flattens out again above $3 \times 10^{18} \text{ eV}$, this feature is known as the *ankle* of the spectrum. At these energies, cosmic-rays are no-longer confined by the galaxy's magnetic field. The *ankle* is therefore believed to represent the point when the spectrum becomes dominated by cosmic-rays of extra-galactic origin [127].

Cosmic-rays interact hadronically with air molecules in the upper atmosphere to produce a cascade of secondary pions and kaons. The pions and kaons in turn decay to produce a high flux of muons. These cosmic muons are typically produced at alti-

Figure 6.1: Major components of the primary cosmic-ray spectrum, taken from [121]

Figure 6.2: The “all-particle” spectrum of Cosmic-ray interactions above 10^{12} eV. The energy spectrum has been multiplied by $E^{2.5}$ to demonstrate the features of the spectrum. Taken From [43], for references see [126].

tudes of ~ 20 km. The muon energy spectrum and angular distribution reflects that of the primary cosmic-ray spectrum and the underlying physics of hadron production and decay. A number of experiments have measured the cosmic muon flux [128] [129] [130]. All experiments agree on the energy spectrum, but their measurements of the absolute flux vary at the 20% level. The muon spectrum is observed to fall steeply (*c.f.* shape of primary spectrum). At the Earth's surface, the mean cosmic muon energy is measured to be 4 GeV, but muon energies of over 100 GeV have been observed.

Muons of a given energy cannot penetrate the Earth further than a certain depth. As the depth below ground increases progressively fewer muons have sufficient energy to penetrate to this depth. This means that moving deeper underground, the cosmic muon flux falls off rapidly with depth. For example, at the MINOS Far Detector, which is at a depth of 2070 mwe, the flux is reduced by a factor of $\sim 10^5$ relative to the surface, meaning cosmic muons are incident on the detector at a rate of ~ 0.5 Hz.

As well as forming the primary background to atmospheric neutrino analyses in the Far Detector, cosmic muons are also useful for validating both the detector calibration and simulation. It is therefore important that the cosmic muon flux at the Far Detector is reasonably well simulated.

6.2.2 Simulating the Underground Cosmic Muon Flux

The cosmic muon flux at the surface can be parameterised as follows [131]:

$$N(E, \cos \theta) = \frac{0.14 \times E^{-2.7}}{\text{cm}^2 \text{sr GeV s}} \left(\frac{1}{1 + \frac{1.1 \times E \cos \theta}{115 \text{ GeV}}} + \frac{0.054}{1 + \frac{1.1 \times E \cos \theta}{850 \text{ GeV}}} \right), \quad (6.1)$$

where E is the muon energy and θ is the zenith angle of the incident muon's path. This parameterisation is valid for $E > 100$ GeV, where muon decays can be neglected, and for $\theta < 60^\circ$, where the curvature of the Earth can be neglected. The two terms in equation (6.1) represent the contributions to the muon flux from pion and kaon decays respectively. The fluxes of μ^+ and μ^- are generated in the ratio $\mu^+/\mu^- = 1.25$ reflecting the excess of positive mesons produced by the mainly positively charged primary cosmic-ray flux. Although this charge ratio is in accordance with surface measurements of low energy muons ($E < 100$ GeV), it is likely that it underestimates the ratio of muons at the Far Detector, since these muons correspond to surface energies of 500 GeV and above, where the muon flux has a larger contribution from kaon decays [131].

A knowledge of the density and composition of the local rock, along with an understanding of muon energy loss through the rock, is required in order to extrapolate the muon flux from the surface to the Far Detector. The rock above the Far Detector cavern is primarily a local type known as *Lake Vermillion Greenstone*, with a measured density of 2.8 g cm^{-3} [132]. This rock is interspersed with pockets of iron ore, so the average rock density varies with direction. The Far Detector cosmic muon simulation uses a rock density map calculated by the Soudan 2 experiment using cosmic muon data. This rock map is combined with a digitised map of the surface topography to find the overburden for a given direction [132]. The average muon energy loss through the rock is parameterised as follows [131]:

$$\frac{dE}{dx} = A + \frac{E}{L}, \quad (6.2)$$

where A ($\simeq 1.9 \text{ MeV/g cm}^{-2}$) represents energy losses due to ionisation and the second term represents losses due to radiative processes, with $L \simeq 2.5 \times 10^4 \text{ g cm}^{-2}$. Although in practice A and L both vary slowly with energy and dE/dx undergoes stochastic fluctuations, all three terms can be approximately treated as constant. In this approximation, integrating equation (6.2) gives:

$$E_X = (E_0 + AL)e^{-X/L} - AL, \quad (6.3)$$

where E_X is the muon energy at a depth X and E_0 is the surface energy of the muon. The incident muon flux at the Far Detector is calculated as a function of E_X and zenith angle, θ , by propagating the surface flux as follows:

$$N(E_X, X, \cos \theta) dE_X = N(E_0, 0, \cos \theta) dE_0 = N(E_0, 0, \cos \theta) e^{X/L} dE_X \quad (6.4)$$

For each value of θ , the digitised map of the rock above the Far Detector cavern at the Soudan mine is used to determine the overburden X [133]. Equation (6.3) is then used to convert the muon energy E_X into the equivalent surface energy E_0 . Finally, equation (6.4) is used to determine the muon flux at energy E_X and zenith angle θ from the surface flux at energy E_0 and zenith angle θ .

A sample of simulated cosmic muons for use in analyses is generated by following the steps below. For each snarl in the Monte Carlo sample the initial energy and direction of the muon is selected at random from the probability distribution given by equation (6.4). Then the initial position of the muon is given by a point chosen at random on the surface

of a box ($9.4\text{ m} \times 9.4\text{ m} \times 34.0\text{ m}$) placed around the detector. If the muon will not pass through the detector the event is discarded at this stage (as such an event would not trigger the real detector). Finally, the muon is propagated through the detector using *gminos* and the detector response is simulated using *PhotonTransport* and *DetSim* (this process is described in Section 6.4).

6.3 Atmospheric Neutrino Simulation

The atmospheric neutrino flux is simulated by combining a parameterisation of the primary cosmic-ray flux with models of neutrino yields from interactions between cosmic-rays and molecules in the atmosphere. A number of different flux models have been developed [63], but for the oscillation analysis that is the subject of this thesis the flux model of *Barr et al.* [64] was used. It is this model that will be described in detail, although other models are developed in a broadly similar way.

Barr et al. simulate the atmospheric neutrino flux at the MINOS Far Detector using a full 3-dimensional treatment of cosmic-ray interactions with molecules in the atmosphere. In the first part of the simulation used to develop the model, cosmic-rays are injected at the top of the atmosphere all over the globe. The cosmic-rays are allowed to interact and the secondary particles from these interactions are tracked through the atmosphere. The flux of neutrinos is then measured using a ‘detector’ centred on the Far Detector site. The model will now be described in more detail:

Primary Flux The primary cosmic-ray flux is parameterised as a function of energy by applying power law fits to measurements of each type of nucleus incident on the atmosphere. Recent measurements of the dominant primary proton flux by AMS[134] and BESS[135] are found to agree to 5% for energies below 100 GeV and measurements of the sub-dominant primary flux of α -particles by AMS and BESS are found to agree to 15% [63]. It should be noted that only some 25% of nucleons in the primary-flux are carried by helium and heavier nuclei, so such a difference corresponds to an uncertainty of $\sim 3\%$ in the flux of nucleons. The cosmic-ray spectrum is modulated by solar activity. During periods of high solar activity the solar wind suppresses the low energy ($E \lesssim 10\text{ GeV}$) portion of the cosmic-ray spectrum and therefore reduces the flux of low energy atmospheric neutrinos. The strength of solar activity varies cyclically, with a period of roughly 11 years. The effect of the solar cycle is discussed in Section 6.3.1.

Geomagnetic Fields The Earth's magnetic field also has a filtering effect on the cosmic-ray flux incident on the atmosphere. The ability of a cosmic-ray to penetrate this magnetic field depends on its *rigidity* ($rigidity \equiv momentum/charge$). Cosmic-rays with sufficient rigidity are able to penetrate the Earth's magnetic field and interact in the atmosphere, the rest are deflected back out into space. Due to the shape of the magnetic field, the cut-off rigidity depends on the magnetic latitude and on the angle at which the cosmic-rays approach the Earth. At high magnetic latitudes, cosmic-rays are incident parallel to the lines of magnetic flux and are deflected weakly. At low magnetic latitudes, cosmic-rays are incident perpendicular to the lines of magnetic flux and are deflected strongly. The Earth's magnetic field also causes an east-west asymmetry in the cosmic-ray flux. This occurs since the vast majority of cosmic-rays are positively charged. Cosmic-rays arriving from the west are focused toward the Earth, while cosmic-rays arriving from the east are defocused away from the Earth. This also has the effect of making the rigidity cut off lower for cosmic-rays arriving from the east than those arriving from the west. At higher magnetic latitudes this effect is reduced, as the focusing/defocusing begins to be directed parallel to the Earth's surface. The MINOS Far detector is at a magnetic latitude of 56.3° , therefore there is a large up-down asymmetry in the rigidity cut-off (10 GeV cosmic-rays incident from below compared to 1 GeV for those incident from above), but only a small east-west asymmetry [63]. In order to remove cosmic-rays simulated below the rigidity cut-off, those which produce at least one neutrino in the Far Detector are traced back through the atmosphere. A trajectory is considered valid if a particle can be propagated to a distance of at least $30 R_E$ with a total path length of less than $300 R_E$, where R_E is the radius of the Earth.

Hadron Production The *Target* simulation package [136] is used to generate the cascades of secondary hadrons produced when cosmic-rays interact with molecules in the atmosphere. For unstable particles, interactions and decays are competing processes. Each particle is assigned a decay and interaction length chosen at random from probability distributions. The two lengths are compared to determine whether the particle interacts or decays. The models of hadron production have been tuned on data gathered from fixed target experiments in which protons are accelerated onto light nuclei (ZGS [137], CERN Rome group [138], E802 [139] and HARP [140]). The phase space explored by these experiments provides only limited coverage of the phase space of hadron production in the atmosphere. The simulations must therefore use interpolations or extrapolations into unmeasured

regions of phase space. This leads to large uncertainties in the hadron production models, which dominate the overall uncertainty in the atmospheric neutrino flux at all energies. Several new experiments: HARP at CERN (3-15 GeV primary energy) [140], E910 at Brookhaven (6-18 GeV) [141], MIPP at Fermilab (5-120 GeV) [142] and NA49 at CERN (100 and 158 GeV) [143] have been carried out and results are being produced. These results should improve the situation for future atmospheric neutrino (and long baseline) analyses.

Figure 6.3 shows the angular distribution of the atmospheric ν_μ flux predicted by *Barr et al.* for a range of energies [144]. The plot clearly shows that the ν_μ flux falls rapidly with increasing neutrino energy. This reflects the shape of the primary cosmic-ray flux. The large asymmetries caused by geomagnetic effects are clearly visible in the zenith angle distributions for energies below 1 GeV. The downward-going flux shows the effect of the local geomagnetic field, while the upward-going flux is an average of global geomagnetic effects. The peak in the zenith angle distributions around the horizon for all energies is caused by the fact that the 3-dimensional calculations include the possibility for a vertical shower to descend into the atmosphere close to the detector and produce a neutrino sideways. Another reason for this peak is that hadrons propagate for longer through the less dense higher regions of the atmosphere at large zenith angles. This has the effect of enhancing hadronic decays relative to hadronic interactions. As neutrino energy increases the flux becomes more up-down symmetrical. At high energies, the effect of muons hitting the Earth is visible in the more gentle peak in the zenith angle distribution near the horizon than for lower energies. This is reproduced by the 1-dimensional simulation also. The effect of hitting the Earth is to remove the chance of emitting a second neutrino from muon decay ($\pi^- \rightarrow \mu^- \bar{\nu}_\mu \rightarrow e^- \bar{\nu}_e \nu_\mu \bar{\nu}_\mu$). This effect also heavily suppresses electron neutrinos at high energies.

The uncertainty in the atmospheric neutrino flux over the energy range 1-10 GeV is estimated to be 5% from errors in the primary flux [144] and 10% from errors in hadron production [145]. The uncertainty in the overall flux can be estimated by comparing the flux of *Barr et al.* with those of other models. These are found to differ at the level of 15%. It should be noted that the error in the ratio of $\nu_\mu/\bar{\nu}_\mu$ and $(\nu_\mu + \bar{\nu}_\mu)/(\nu_e + \bar{\nu}_e)$ fluxes have significantly lower errors, since single processes contribute the numerator and denominator in both cases. This leads to a large cancellation of the uncertainties in both hadronic models and primary fluxes [146]. This cancellation gets worse at higher energies where the chance of the muon produced in the pion decay itself decaying before hitting the ground is much reduced.

Figure 6.3: Zenith angle distribution of simulated atmospheric ν_μ flux for a range of energies. The solid lines indicate the results from using a 3-dimensional model. The dashed lines shows the results when a 1-dimensional model is used. Taken from [144].

6.3.1 Normalising Atmospheric Neutrino Monte Carlo over the Solar Cycle

The cosmic-ray spectrum is modulated by solar activity. During periods of high solar activity the solar wind suppresses the low energy ($E \lesssim 10$ GeV) portion of the cosmic-ray spectrum and therefore reduces the flux of low energy atmospheric neutrinos. The strength of solar activity varies cyclically, with a period of roughly 11 years. The MINOS Far Detector started taking data in 2003, just after the last maximum in the solar cycle. For previous analyses of MINOS atmospheric neutrino data [29] [94] [96] it was therefore a valid approximation to use atmospheric neutrino Monte Carlo generated using the neutrino flux at Solar Maximum. The dataset for this analysis spans 31 months, a period over which solar activity will have dropped significantly according to the model. The approximation of using Solar Maximum Monte Carlo is not valid in this case.

One way of monitoring the level of solar activity over time is to measure the flux of atmospheric neutrons produced by cosmic-ray interactions with molecules in the atmosphere. Since neutrons and neutrinos are produced with similar energies relative to the primary cosmic-ray energies, they should possess a similar dependence on the level of solar activity. There are a number of experiments which measure the atmospheric neutron rate. Information from the CLIMAX neutron detector [147] was used by the Soudan 2 experiment to normalise their Monte Carlo [66]. The CLIMAX detector is sited in Colorado and has been recording atmospheric neutron rates since 1951. CLIMAX data up to November 2006 is readily available online. Figure 6.4 shows the variation in atmospheric neutron rates between 1st of January 1980 and 30th November 2006. The peaks in the neutron rates represent periods of Solar Minimum and the troughs represent periods of Solar Maximum.

The best measurements of cosmic-ray fluxes at Solar Minimum were made during the minimum which occurred in the 1990s. It is these data which are used in the Monte Carlo atmospheric neutrino simulations. The neutron counts which correspond to the Solar Maximum and Solar Minimum were determined by fitting from just before the last Solar Minimum (January 1995) until November 2006 to a sine function, $Rate = A + B \times \sin(C \times (month + D))$, then taking the maximum (Solar Minimum) and minimum (Solar Maximum) values of this function. Figure 6.5 shows the fitted mean hourly neutron rates for each month. The best fit rate was $Rate = 3970.0 + 316.8 \times \sin(0.047 \times (month - 166.8))$ giving a neutron rate of 4286.9 Counts/Hour at Solar Minimum and a neutron rate of 3653.3 Counts/Hour at Solar Maximum.

Figure 6.4: Long term variation in mean daily atmospheric neutron rates from CLIMAX. The data shown is from 1st of January 1980 to 30th November 2006 and is available from [147].

Figure 6.5: Fit to mean hourly neutron rates from CLIMAX for each month. The data shown is from January 1995 to November 2006 and is available from [147].

The average neutron rate over the MINOS dataset from August 2003 to February 2006 was calculated by taking mean of the neutron rates for each month weighted by the MINOS Far Detector live time for that month. The average neutron rate over this dataset was found to be 3933.74 Counts/Hour. Making the same assumption as Soudan 2, that the neutron and neutrino fluxes vary linearly in the same way between Solar Minimum and Solar Maximum. The above results imply that this dataset should be fitted to a combination of the Solar Maximum and Solar Minimum Monte Carlo fluxes in the ratio 55.7% to 44.3%. These percentages were used to reweight the existing Solar Maximum atmospheric neutrino flux Monte Carlo to the average MINOS dataset flux.

6.3.2 Simulation of Atmospheric Neutrino Cross-Sections

For an incident neutrino to be observed in the MINOS Far Detector it must first interact. Neutrino interactions are simulated using the *Neugen3* package [148]. This provides models for interactions between neutrinos and nucleons bound inside nuclei and quarks inside nucleons. The main interactions simulated are: Quasi-Elastic Scattering (QE), which dominates for energies below 1 GeV; Resonances (RES), which dominate for energies between 1 GeV and 3 GeV and Deep Inelastic Scattering (DIS), which dominates for energies above 3 GeV. The cross-section for each interaction type is parameterised by *Neugen3* and tuned to fit the experimental data. For each neutrino, *Neugen3* compares the total cross-section with a random number in order to determine whether an interaction will take place. If an interaction will occur, then *Neugen3* selects the type of interaction and simulates the hadronic final state.

For QE and RES interactions, the interaction is simulated as the incident neutrino interacting with a nucleon. Nucleons bound within a nucleus are modelled as a Fermi gas. Each nucleon is modelled as having a Fermi momentum of ~ 230 MeV and a binding energy of ~ 30 MeV. Since nucleons are degenerate within a Fermi gas, they can only be scattered into unoccupied states above the Fermi surface. Therefore neutrino interactions which produce nucleons below the Fermi momentum are suppressed (this is known as *Pauli Blocking*). In QE events the neutrino scatters elastically off the nucleon. In RES events, the neutrino and nucleon interact to form a resonance, which then decays into a final state usually containing one or more pions. A total of 17 resonances are simulated by *Neugen3*, dominated by the production of $\Delta(1223)$.

For DIS interactions, the incident neutrino scatters off a quark within a nucleon, rather than a nucleon. The interaction is expressed in terms of structure functions

Figure 6.6: Cross-section for ν_μ CC interactions on iron as a function of neutrino energy, generated using the *Neugen3* model of neutrino interactions. The red, green and blue lines show the exclusive QE, RES and DIS cross-sections respectively. The total cross-section is shown by the black line.

calculated from parton distribution functions. The multiplicity of the hadronic final state is then parameterised as a function of its invariant mass. Once an interaction has taken place, the resulting hadrons may be absorbed or re-scattered as they exit the nucleus. *Neugen3* uses a collection of routines known as *Intranuke* [149] to simulate these intra-nuclear effects.

For the energies relevant to atmospheric neutrino analyses in the MINOS Far Detector, all three processes make a significant contribution to the overall cross-section. The uncertainty in the atmospheric neutrino event rate due to errors in the interaction model is estimated to be around 10% [150]. Figure 6.6 and Figure 6.7 show the ν_μ and $\bar{\nu}_\mu$ charged current interaction cross-sections on ^{56}Fe nuclei, as simulated by *Neugen3*, as a function of neutrino energy. The total cross-section is plotted along with the exclusive QE, RES and DIS cross-sections. The majority of atmospheric neutrinos observed in the Far Detector have multi-GeV energies. At these energies all three interaction types make a significant contribution to the overall cross-section.

Figure 6.7: Cross-section for $\bar{\nu}_\mu$ CC interactions on iron as a function of neutrino energy, generated using the *Neugen3* model of neutrino interactions. The red, green and blue lines show the exclusive QE, RES and DIS cross-sections respectively. The total cross-section is shown by the black line.

6.4 Detector Simulation

The process of simulating the detector response described below should be compared to the description of the scintillator and readout given in Section 4.3.

6.4.1 Active Detector Simulation

The principal simulation software used by MINOS is the *gminos* package [151]. The *gminos* software provides a model of the detector composition and geometry, and interfaces to GEANT3 [152] for the physics of particle transport through materials. *gminos* reads in a list of particle four-momenta, inserts these particles into the simulated detector, and outputs a record of the energy deposits in the active scintillator simulated by GEANT3.

The next stage of the simulation is to model the response of the active scintillator to these energy deposits. This is done by the *PhotonTransport* package [153]. This package simulates the creation of photons in the scintillator, from the simulated energy deposits, and propagates them to the PMT photo-cathodes, whereupon their conversion into photo-electrons is simulated. For each recorded energy deposit in the scintillator,

PhotonTransport generates a number of blue photons (N) according to Birks's Law [154] as follows:

$$N = C \frac{L_0}{C_0} \frac{dE}{1 + k_B \frac{dE}{dx}}, \quad (6.5)$$

where: dE is the energy deposited in the scintillator, L_0 is the light output of the scintillator, C_0 is the calibration constant applied to correct for the strip response, C is an overall normalisation used to tune the light output and k_B is Birks's constant (taken to be 0.133 m GeV^{-1}).

PhotonTransport next simulates the capture of blue photons from the scintillator by the WLS fibres and the conversion of some of them into green photons. (The reality of this process is described in the Section 4.3.1.) The green photons are distributed along the WLS fibres using a set of PDFs relating the positions and times of green photons to those of the blue photons. The PDFs were constructed from the results of detailed simulations of photon emission, propagation and absorption in the scintillator strips and WLS fibres.

The software then propagates each green photon down the WLS fibres and through the clear fibres to the PMT photo-cathodes, accounting for attenuation along the way. At the PMT photo-cathodes the green photons are converted into photo-electrons with a flat probability of 13% (this is equal to the measured quantum-efficiency of the M16 PMTs used in the Far Detector). Background noise is also simulated in the form of additional photo-electrons generated according to the measured rates. The simulated photo-electrons are then passed to the *DetSim* package [153] for the next stage of the simulation.

6.4.2 Detector Electronics Simulation

The *DetSim* package is responsible for simulating the amplification of the photo-electrons through the M16 PMTs, the digitisation of the resulting charge by the VA electronics and the implementation of DAQ triggers.

The first step is to calculate the total number of photo-electrons on a photo-cathode by summing up the individual PE contributions on that photo-cathode from each photon generated by *PhotonTransport*. At this stage optical cross-talk between pixels on the same PMT is taken into account by redistributing the photo-electrons on each PMT

photo-cathode according to a probability matrix calculated from the results of M16 bench studies [155].

The amplification of each signal by a PMT is simulated in two stages. Stage one is to simulate the amplification of the signal by the first pair of dynodes. This process includes smearing the resulting charge to account for fluctuations due to secondary dynode emissions. The second stage is to calculate the final charge on the photo-anode using the measured PMT response curve. The dynode trigger threshold is set to the equivalent of 0.3 photo-electrons and the time of the dynode trigger is set as the time when the first photon arrived on *any* of the photo-cathodes on the PMT.

DetSim then simulates the response of the front end electronics and DAQ. The modelling of VA readout electronics' response to the charge on the photo-anodes returns digitised charges and times for each hit. As described in Section 4.3.3, the process of digitising the dynode trigger times involves phase comparisons between two system clocks. To take into account the possibility of aliasing¹ during the phase matching, the raw hit times are smeared by a configurable amount before being converted into TDCs. The amount of smearing is tuned by comparison with data (see Section 6.4.3). Digitisation is followed by the application of the 2/36 VARC trigger is applied to the signals. Any signals that satisfy the trigger are amplified using the measured VA response curves and smeared using the measured pedestal widths. The signals are digitised and those with less than 20 ADC are sparsified away. Next, the digitised charges and times are used along with the electronics channel ID to create raw detector hits. The set of raw hits is passed through a simulation of a trigger processor and the DCP. The 4/5 plane trigger algorithm, used to select interesting physics events, is applied to the hits. The resulting snarls are then written out in the same format as the real data.

6.4.3 Electronics Timing Smearing

It is very important that timing based variables are well simulated in the Monte Carlo as a large part of the analysis in this thesis relies heavily on such variables to discriminate between upward-going neutrino events and downward-going stopping cosmic muon events (the selection is discussed in Chapter 8). As can be seen from Figure 6.8, the nominal Monte Carlo RMS Down distribution peaks at a lower value than that of data. This is due to the fact that there is an unsimulated source of errors in hit times. The

¹one position in the sine wave being mis-identified as another.

Figure 6.8: Comparison of timing RMS_{Down} for data and unsmeared Monte Carlo. Cosmic muon Monte Carlo is shown by the red points. Data is shown by the black points.

30-bit TDC values written out in the Raw Digits cannot be set using a straight digitisation of the 40 MHz system clock. Section 4.3.3 describes the mechanism by which the TDC value is set for each hit. There exists a possibility of aliasing occurring during the phase matching which is part of this process. To account for this possibility the raw hit times in Monte Carlo can be smeared by a configurable amount before being converted into TDCs. For the current Monte Carlo simulation this value has been tuned to be 0.93 ns. The RMS_{Down} distribution smeared Monte Carlo is shown with the data in Figure 6.9. The peaks of the distributions are now in agreement, although the data distribution is still slightly wider than the Monte Carlo distribution. The remaining differences are probably due to errors in the timing calibrations in data. The systematic error on the track direction identification caused by this level of difference is below the 1% level and small compared to the 3% systematic error assigned to the relative efficiency for selecting upward versus downward-going events in the final oscillation fit (see Section 9.5.2).

Figure 6.9: Comparison of timing RMS_{Down} for data and smeared Monte Carlo. Cosmic muon Monte Carlo is shown in the red histogram. Data is shown by the black points.

Figure 6.10: Total pulse-height from double-ended track strips. Cosmic muon Monte Carlo is shown in the red histogram. Data is shown by the black points.

6.5 Evaluating The Simulation

6.5.1 Evaluating The Detector Simulation

In order to evaluate the simulation of the detector it is instructive to compare low level quantities. Each detector hit contains a measurement of timing and pulse height. Problems with the simulation of either of these quantities would introduce systematic biases to the atmospheric neutrino analysis, so it is important that they are well simulated. Samples of cosmic-ray muons which cross at least 20 scintillator planes and satisfy straight line fits in each view (UZ , VZ) with deviations of less than 1 cm are used to compare the data to Monte Carlo, as such tracks are usually well reconstructed and enable the 3D position of hits to be easily determined.

Figure 6.10 shows a comparison of the summed pulse height from double-ended track strips between data and Monte Carlo. There is reasonable agreement between the two distributions, although pulse heights in Monte Carlo are systematically higher by around 1 PE.

In order to check that the attenuation of signals along the WLS fibres is correctly

Figure 6.11: Mean PH per strip end vs WLS fibre length. Cosmic muon Monte Carlo is shown by the red points. Data is shown by the blue points.

simulated, the mean pulse-height at each strip end as a function of the distance of the muon from the strip-end is plotted for data and cosmic muon Monte Carlo. This is shown in Figure 6.11. The data and Monte Carlo agree to within 10% over the entire length of the WLS fibre. The small differences occur because the Monte Carlo applies constant attenuation lengths for each strip, whereas the data represents the average over many attenuation curves. The fact that the curve in Monte Carlo is flatter than in data, may suggest that the shorter (or both) of the two the attenuation lengths (see Section 4.3.1) used in the Monte Carlo may in fact be set too low. A drop in signal occurs 30 cm from the strip ends as it is possible for light to escape from the ends of the scintillator before it is captured by the embedded WLS fibre.

6.5.2 Data-Monte Carlo Comparison of Cosmic Muon Flux

Having assured ourselves that the detector response is well simulated, we can now look at reconstructed quantities in order to evaluate the cosmic muon flux simulation. The Far Detector cosmic muon simulation is validated by comparing the simulated cosmic muon flux with that of the data. It is first necessary to estimate the absolute flux in cosmic muon Monte Carlo by normalising it to the data. A sample of 3729162 cosmic muons which span at least 20 planes was selected from a sample of 9875030 generated Monte

Figure 6.12: Comparison of the number of planes in tracks between data and cosmic muon Monte Carlo. Cosmic muon Monte Carlo is shown in the red histogram. Data is shown by the black points.

Carlo cosmic muon events. Three months of data events (total live-time 5961161 s) yielded 1883040 cosmic muons which span at least 20 planes. Assuming the selection efficiencies in data and Monte Carlo are equal, this gives an absolute muon rate of 0.84 Hz (100000 Monte Carlo snarls $\equiv$ 33.2 hours). This is in agreement with previous measurements of the absolute muon flux [96] [94] [133].

Figure 6.12 shows the total number of planes crossed by the selected muon tracks. There is good agreement between data and Monte Carlo. Figure 6.13 shows the measured flux as a function of reconstructed zenith angle for data and Monte Carlo. The flux falls off sharply with increasing zenith angle due to the rapid increase in rock overburden, which filters out all but the highest energy cosmic muons from the surface flux. The drop off around $\cos \theta_{zenith} = 1$ is due to detector acceptance effects. The upward-going events in cosmic muon Monte Carlo are due to reconstruction errors, the excess of upward-going events in data is due to neutrino-induced muon events.

The measured flux is shown as a function of the reconstructed azimuthal angle (defined so $\phi_{Azimuth} = 0$ in the positive z -direction) in Figure 6.14 for data and Monte Carlo. The shape of the distribution is principally determined by the acceptance of the

Figure 6.13: Comparison of the reconstructed zenith angle distribution for data and cosmic muon Monte Carlo. Cosmic muon Monte Carlo is shown in the red histogram. Data is shown by the black points.

Far Detector. As trajectories approach the xy plane, progressively fewer events cross enough planes to pass the 20 plane cut. This suppresses the measured flux for azimuthal angles around $\pm 90^\circ$. The small differences between the data and Monte Carlo observed in both the zenith and azimuthal angle distributions suggest that systematic errors are incurred in transferring the rock density map from Soudan 2 to MINOS.

6.6 Monte Carlo Samples

Table 6.1 gives a full list of the Monte Carlo samples used in the atmospheric neutrino analysis in this thesis. The atmospheric neutrino samples use the flux model of *Barr et al.* and the *Neugen3* interaction model. The propagation of hadrons is simulated using the *GCALOR* package [156] [157] within *GEANT3*, which has been found to give a good description of low energy hadronic interactions in the MINOS detectors [115]. A large sample of almost twenty million cosmic muons is used to study the principle backgrounds in the atmospheric neutrino selection, with additional Monte Carlo samples of neutrons emitted by muon spallation in the rock [158], neutrino-induced upward-going muons and

Figure 6.14: Comparison of the reconstructed azimuthal angle distribution for data and cosmic muon Monte Carlo. Azimuthal angle is defined such that $\phi_{Azimuth} = 0$ in the positive z -direction. Cosmic muon Monte Carlo is shown in the red histogram. Data is shown by the black points.

ν_τ produced by $\nu_\mu \rightarrow \nu_\tau$ oscillations used to investigate other possible backgrounds.

Monte Carlo Sample	Physics Models Used	Events	Exposure / ktyr
atmospheric neutrinos			
atmospheric ν_e/ν_μ	(Barr <i>et al.</i> , <i>Neugen3</i> , <i>GICALOR</i>)	1.0×10^6	3605.0
atmospheric ν_τ	(Barr <i>et al.</i> , <i>Neugen3</i> , <i>GICALOR</i>)	1.0×10^5	360.5
cosmic μ	(full spectrum)	1.99×10^7	4.1
ν -induced muons	[133]	4.75×10^4	1080.0
μ -induced neutrons	[158]	8.38×10^4	67.7

Table 6.1: Monte Carlo samples used for the atmospheric neutrino analysis. The physics models used to generate each sample are listed in the second column. The third and fourth columns list the number of events generated and the equivalent Far Detector exposure.

Chapter 7

Reconstructing Events

“No one ever went broke underestimating the intelligence of mankind.”

— Phineas Taylor Barnum, 5 July 1810 – 7 April 1891

7.1 Reconstruction Overview

The MINOS offline software framework uses a ROOT-based [103] application called ‘*loon*’¹. The reconstruction of events is done by a chain of algorithms each of which build a set of candidate reconstructed objects (digits, strips, tracks, showers, events), which are then passed to the next algorithm in the chain and so on, until whole events have been reconstructed. As only Far Detector data is used in this thesis, the description of event reconstruction in this chapter will focus on the reconstruction of Far Detector events, but a similar reconstruction path is followed at the Near Detector.

7.2 From Raw Data to Digits

The raw data from the detector readout is stored as ROOT files (see Section 4.4). For each snarl the raw data is converted into a list of *CandDigit* objects. Each *CandDigit* contains a digitised measurement of the pulse height in a VA channel and a digitised

¹Named after the state bird of Minnesota.

measurement of the time which this occurred within the time frame. Due to the nature of the readout each *CandDigit* also contains a list of the eight strip-ends which are read out by the PMT pixel connected to the VA channel (See Section 4.3.3). As discussed in Chapter 6, the Monte Carlo simulation writes out fake raw data in the same format as the real detector readout, so the reconstruction process is the same for both the data and Monte Carlo.

At this stage there are many unwanted digits in the snarl which are due to single-photon noise in the PMTs or natural radioactivity in the detector. The list of *CandDigit* objects is passed through a filter which attempts to remove these noise digits, to prevent obfuscation of the triggered event in later reconstruction steps. The *NoiseFilter* works by first identifying the main body of the event, the part which satisfies the 4/5 plane trigger. Digits outside of this region are rejected if they are displaced from this body of digits by more than ten planes. Digits are also rejected if they occur more than 50 ns before or 500 ns after the earliest digit in the region satisfying the 4/5 plane trigger. The remaining list of *CandDigits* is then passed to the demultiplexer algorithm in a *CandDigitList* object.

7.3 Demultiplexing

Demultiplexing refers to the process of deciding which of the eight strip-ends connected to each PMT pixel the light detected on the pixel came from. For this analysis the *AltDeMux* [159] algorithm was used to demultiplex the event. This algorithm was developed for the atmospheric neutrino analysis, but has since been adopted as the demuxer used in the default reconstruction chain for all MINOS analyses.

The first stage of the *AltDeMux* algorithm is the identification of cross-talk. This is done by making maps of the pulse heights collected in the pixels of each PMT. Cross-talk can then be tagged by analysing the pulse heights and relative positions of the digits on each PMT. Thresholds are applied to the absolute pulse height of all digits and the relative pulse height of digits on adjacent pixels on a PMT. Once the cross-talk digits in an event have been tagged the remaining digits can be used to determine which strips have been hit in the event.

AltDeMux first uses the digits on each side of the detector to produce a list of all the possible solutions in each plane. The measured times of digits on each side of the

Figure 7.1: Cartoon showing demux alternatives in the U-View for a single straight through-going muon. The filled circles represent ‘Golden Hits’, while the open circles represent “demux” alternatives. Taken from [159].

detector are used to constrain the event region in each plane and reduce the number of possible solutions. To demultiplex the event, *AltDeMux* forms structures of strips that are reconstructed unambiguously within a plane. These ‘*Golden Hits*’ are formed from configurations with a unique solution, or a unique solution which places strips close together in a plane.

The Golden Hits are used to classify obvious event types, such as multiple muons or straight through-going muons and to define suitable hit positions on planes without demultiplexed hits, so called ‘*target*’ regions. This classification influences the exact strategy used to demultiplex the rest of the digits. For example, in the case of an event identified as a single straight through-going muon, the target regions would be linear extrapolations based on both the hits and a straight line fit to the Golden Hits. This does not force the event to be a straight line. The event is just given the chance to grow into a straight line. Figure 7.1 shows possible “demux” alternatives in this case.

The next stage is to look for further east-west combinations of digits that lie near the target regions. This stage demultiplexes most of the remaining hits. The target regions

are then recalculated using all of the demultiplexed digits and any undemultiplexed hits are placed at the position nearest to a target.

In the penultimate stage of the algorithm the solution found is validated. Each plane is checked in turn and, using the information known about the whole event, target regions for the hits in the current plane are identified. The algorithm then attempts to demultiplex the current plane around these targets. If the new solution is significantly better than the existing solution, then the existing solution is replaced by the new one.

Finally, the solution is checked for any obvious mistakes and the cross-talk is re-tagged using the pixel spot information (which is now available due to the completed demultiplexing). The list of demultiplexed *CandDigits* is then written out in a *CandDeMuxDigitList* object.

7.4 Strip making

The *CandDeMuxDigitList* object is read in by the *FarDetStrip* algorithm. This algorithm constructs a list of *CandStrip* objects from the demultiplexed *CandDigit* objects. The first stage of the algorithm is to match cross-talk digits with the non-cross-talk digits which caused them. *CandStrip* objects (hereafter strips) are constructed from the non-cross-talk digits. The pulse height from the associated cross-talk digits is added to the pulse height recorded for each end of the strip. A *CandStripList* object containing all the reconstructed strips is then written out.

7.5 Track Finding

The aim of the track finder is to find a *seed track*. It is expected to accurately find all the ‘clean’ sections (away from any showers) of the track, but not necessarily to find exactly the right strips within a large vertex shower. For this analysis the algorithms within the *CandTrackCam* package[160] were used to produce a list of track candidates. The input to this process is the *CandStripList* object and if the algorithm is successful a *CandTrackList* object, which contains one or more *CandTrack* objects is written out. The process of track finding will now be described in more detail.

The first step is to form *Clusters* from hits on adjacent strips on the same plane.

Figure 7.2: The three possible methods of forming *simple associations* between *Triplets* in the Track Finder. Taken from [160].

Initially, U and V view planes are considered separately. In each view, groups of three *Clusters* on different planes which overlap transversely are grouped together to form short *TrackSegment* objects called *Triplets*. *TrackSegments* can be associated with each other or combined to form longer *TrackSegments*. The algorithm attempts to associate each *Triplet* centred on a given plane with all the *Triplets* centred on nearby planes in the same view. These *simple associations* are shown in Figure 7.2. *Simple associations* are formed between *Triplets* where the end of one *Triplet* overlaps with the beginning of another by one or two clusters (1 and 2 in the diagram), or if two *Triplets* do not overlap, an association may be made if the *Triplets* have compatible² beginning/end directions, as shown in the third example in Figure 7.2.

Preferred associations are made from those *simple associations* which give the most track-like results. The next step in the algorithm is to find long chains of *Triplets* with *preferred associations*. If the *Triplets* in these chains have one beginning *preferred association* and one end *preferred association*, then the *Triplets* are joined together to make a longer *TrackSegment*. Otherwise, *matched associations* are made between the *TrackSegments*.

The algorithm compares all possible paths along chains of *TrackSegments* joined by *matched associations*. Each possible path is given a score. Straighter paths with more *Clusters* are given the highest scores. The path with the highest score is made into a

²Compatible means that when extrapolating a straight line in the end (beginning) direction from the end (beginning) of one *Triplet* to the beginning (end) plane of the other *Triplet*, the line passes suitably close to the beginning (end) of the other *Triplet*. Here, the level of proximity required depends upon the difference in z position.

single *TrackSegment* called a 2D-Track. The unused *TrackSegments* are then reexamined to see if another 2D-Track can be formed. This is repeated until no further 2D-Tracks can be formed. Possible joins between 2D-Tracks are now examined. This process allows for joins to be made across the super-module gap. The 2D-Tracks are now compared between views. The best matches between views are determined using the degree of overlap.

Clusters in shower-like sections of each 2D-Track are now refined by performing straight-line fits between track-like *Clusters* on either side of the shower-like region³. Replacement *Clusters* are made for each plane in a shower-like region containing those strips within a certain tolerance of the linear fit position. Local linear fits over track *Clusters* in small plane ranges are used to determine which strips from the *Clusters* should be included in new *Track* objects. The algorithm then tidies up the *Track* objects for each view by filling in any obvious gaps in each *Track*, using linear fits over the gap section. It then makes any obvious extensions to the beginning/end of the *Track*. This is useful for picking up any strips which could not form a *Triplet* and for picking up strips that were missed because they were on the other side of the super-module gap. Each $U - V$ pair of *Tracks* is used to produce a *CandTrack* object.

Figure 7.3 shows the efficiency of track reconstruction versus the true number of planes crossed by the muon for atmospheric neutrino Monte Carlo events. It is clear that the performance of the track finder is very good, track reconstruction efficiency saturates at almost 100% above 12 planes. The efficiency drops off for low numbers of planes crossed as the track finder requires at least three planes to be hit in a single view in order to make *Triplets*, without which the rest of the track finding cannot be accomplished. Figure 7.4 shows the track reconstruction efficiency plotted as a function of true muon momentum for atmospheric neutrino Monte Carlo events. In this plot the reconstruction efficiency saturates at 92%. This is due to the fact that some muons are produced at very steep angles and so despite having sufficiently high energy do not cross enough planes for a track to be reconstructed.

³If track-like *Clusters* are only found on one side of the shower-like region, then a linear fit through the track-like *Clusters* is extrapolated through the shower. If the 2D-Track in one view is entirely within a shower-like region, a linear fit through shower-like track *Clusters* is attempted.

Figure 7.3: Track reconstruction efficiency versus the number of planes crossed by the muon for atmospheric neutrino Monte Carlo events.

Figure 7.4: Track reconstruction efficiency versus true muon momentum for atmospheric neutrino Monte Carlo events.

7.6 Track Fitting

The *CandTrackList* object containing the found tracks (*CandTrack* objects) is read in by the fitting algorithm. The *CandFitTrackCam* package [161] is an implementation of a Kalman Filter [162]. The purpose of the track fitting package is to take the hit strips identified by the track finder and reconstruct the momentum and charge sign of the particle which produced the track, using its curvature in the magnetic field of the detector. If a track is partially contained, this is the only effective way to estimate the momentum.

The Kalman filter is a set of simple equations that estimate the state of a process in an efficient iterative way, that minimises the mean of the squared error. The filter fully accounts for correlated errors, such as multiple scattering and energy loss. After repeated iterations back and forward along the track, the filter converges to give the best and most complete description of the particle (muon), detailing its position, direction and momentum at every point. This allows for the possibility of improving the track strip identification, especially within large vertex showers, where the track finder may have problems, as shown in Figure 7.5. Outside of the vertex shower finder and fitter strips both agree well with the truth information. Inside the vertex shower the found track becomes inaccurate. However, the strips in the fitted track are still in good agreement with the true muon strips.

For each track plane, the transverse positional information for the plane is taken as the charged-weighted mean transverse position of the group of adjacent strips containing the track strip(s) on that plane and the error in this estimate is taken to be the width of the group, divided by $\sqrt{12}$. Along with the z position, this information now represents a measurement of the track position in each plane.

The knowledge of the Kalman filter about the properties of the track in each plane is represented by the state vector, $\mathbf{x} = (u, v, du/dz, dv/dz, q/p)$. This is initialised for the vertex plane using information from the found track. The algorithm uses the Kalman filter equations to predict the state vector for plane k from the state vector for plane $k - 1$. This is then compared to any measured values for plane k and a ‘filtered’ state vector for plane k is produced. This is then used to predict a state vector for plane $k + 1$ and so on, until the end of the track is reached.

The algorithm now has a rough estimate of the gradient of the track on each plane.

Figure 7.5: Improved track strip identification within large vertex showers. Green points represent strips with recorded hits. The plots on the right show the location of the true muon track in a Monte Carlo ν CC event. The plots on the left show the U and V views of the reconstructed track. Strips ringed in blue are part of the found track. Strips ringed in red are part of the fitted track. Taken from [161].

It uses these gradients to search for a vertex shower. If no shower plane⁴ is found within nine planes of the track vertex, then the algorithm assumes there is no vertex shower. This means the track finder is likely to have done well and all found track strip information can be used as normal. The algorithm uses the Kalman filter equations to move from the end of the track back to the vertex. If a vertex shower is found, provided there are enough planes outside the shower (tuned to be six) for the Kalman filter to build up information about the track, then all the track finder information for the planes within the shower is discarded. The normal Kalman filter propagation then stops at the edge of the vertex shower and the *Swimmer* package [95] is then used to find the most likely track strips on the next plane within the shower.

For a given particle mass, charge, initial position and momentum, the *Swimmer* traces the particle's trajectory through the detector until a user-provided condition (z -position or momentum) is met. The measurement data for any strips found by the *Swimmer* are used by the Kalman filter to generate a new state vector for the next plane within the shower. The *Swimmer* is then used again, and in this way the algorithm moves back toward the track vertex. This propagation continues until a plane window (tuned to four planes) is reached where the *Swimmer* finds no hits, or the edge of the detector is reached.

The filtered state vector at each plane is used to pick a new set of track strips and the measurement information is stored in the same way as for the first iteration. A second iteration is performed, using this new measurement information. The algorithm uses the Kalman filter equations to move from the track vertex to the end and back again.

At this point the algorithm tries to fill in any gaps in the track. If there is a plane between the vertex and end of the track where hit strips are recorded, but there are no track strips, then a state vector for the plane is estimated by using the *Swimmer* to interpolate between the two nearby planes. A final set of track strips is then selected using the state vectors from the second iteration plus those state vectors interpolated from others. Finally, the new track strips are used to build a *CandFitTrack* object, which has a estimate of charge and momentum from curvature in addition to updated values of the *CandTrack* object's properties. The *CandFitTrack* objects from the snarl are written out in a *CandFitTrackList* object.

Figure 7.6 shows the momentum from range resolution obtained by the track fitter

⁴If there are more than a certain number strips containing hits with pulse height above 2 PE in a plane it is considered to be a shower plane. This limit has been tuned to `int(4 + |0.25 × gradient|)` and allows the cut to be less harsh for steeper tracks, but is capped at seven such strips.

Figure 7.6: Muon momentum resolution versus true muon momentum for fully contained atmospheric neutrino events.

Figure 7.7: Muon momentum resolution versus true muon momentum for partially contained atmospheric neutrino events.

for fully contained atmospheric neutrino events. The momentum from range resolution can be parameterised as $\sigma_p/p = 0.06 \oplus 0.07/\sqrt{p(\text{GeV})}$. The momentum resolution is worse at low energies due to the increased effect of sampling in the detector, fluctuations in muon energy loss and multiple scattering. The resolution of momentum from range measurements for partially contained events is shown in Figure 7.7. The momentum resolution is found to be approximately constant at $\sigma_p/p = 0.27$.

7.7 Shower Finding

Previous atmospheric neutrino analyses in the MINOS Far Detector [29] [94] [96] used the combined track and shower reconstruction of the *AtNuReco* package [94]. Due to the adoption of the *CandTrackCam* [160] and *CandFitTrackCam* [161] packages for track finding and fitting respectively, a new separate shower finding algorithm was written to replicate the shower finding of the main *AtNuReco* algorithms, without repeating the track finding.

Some degree of repetition is unavoidable however, as with track finding the first thing the shower finding algorithm, *AlgAtmosShowerList*, does is to make *Clusters* from the hits on each plane. The *Clusters* are then checked to determine how track-like or shower-like they are, using the same criteria as the *CandTrackCam* package. The *CandFitTrackList* is then used to build two *TrackSegments* (one in each view) from the *Clusters* for each of the *CandFitTrack* objects in the list.

AlgAtmosShowerList now loops over all *Clusters* and attempts to form *ShowerSegments*. When the algorithm finds a *Cluster* has not been used in an existing *ShowerSegment*, it uses the *Cluster* as the seed for a new *ShowerSegment*. The algorithm then loops over other unused *Clusters* on nearby planes in the same view as the seed *Cluster*. Any *Clusters* which have a shower-like association with the seed *Cluster* are added to the new *ShowerSegment*. The algorithm attempts to add further *Clusters* to the new *ShowerSegment* by checking nearby unused *Clusters* for shower-like associations with those *Clusters* added in the previous iteration. In this way the algorithm continues to build up the *ShowerSegment* by looping over the *Clusters* in an increasing span of planes, until such time as an iteration finishes in which no new *Clusters* have been added. The *ShowerSegment* is then stored and the process starts again with the next unused *Cluster*, until all *Clusters* belong to a *ShowerSegment*.

The algorithm then assesses the shower-like nature of each *ShowerSegment* formed. Any *ShowerSegments* which cross no shower planes (defined as in the track fitter) are abandoned at this point. The *ShowerSegment* containing the most strips is used as the seed for a 2D shower. The algorithm then attempts to add further *ShowerSegments* to the 2D shower, by checking for shower-like associations. This is repeated until no more *ShowerSegments* can be added. A new seed *ShowerSegment* is then sought, ignoring any *ShowerSegments* already included in 2D showers. If one is found, then the process is repeated, and so on until, there are no more seed *ShowerSegments* are found in a view.

Figure 7.8: Shower reconstruction efficiency versus true E_{shw} .

Each U/V pair of *TrackSegments* made earlier is now checked to see if any *Clusters* each contains are also in a 2D shower. These 2D showers are compared between views and those with the largest plane overlaps are used to form a *ShowerCam* object. Then any 2D showers not already matched are checked for overlaps between views the pair with the largest plane overlaps are used to form another *ShowerCam* object and so on until no more *ShowerCam* objects can be made. Any *ShowerCam* objects which begin within nine planes of a track vertex are flagged as vertex showers. The *ShowerCam* objects are then passed to a second algorithm which uses them to make *CandShower* objects, which are then written out in a *CandShowerList* object. Figure 7.8 shows that shower reconstruction efficiency reaches 100% by around 4.0 GeV.

7.8 Veto Shield Reconstruction

The reconstruction of shield planks proceeds as follows. First, the raw digits from the PMT pixels which read out shield strips are paired up, one from the north end of the strips and one from the South end of the shield strips, by *AlgFarDetShieldPlankList*.

(The orientation of modules in the veto shield is discussed in Section 4.7.) This information is passed to *AlgFarDetShieldPlank*, which uses the *ShieldGeom* class from the *CandShield* package [108] to reconstruct the raw digits and find out the location of the shield planks in the xy plane and their extents in the z direction. This information is used to calibrate the hit times for the resulting *FarDetShieldPlank* objects, which are written out as a *FarDetShieldPlankList*.

7.9 Event Reconstruction

The final stage of the reconstruction is the building of events from the tracks and showers found in a snarl. In theory there can be more than one event in a snarl, but with the exception of cosmic-ray multiple muon events, this situation does not occur at the MINOS Far Detector outside of beam spills. For an atmospheric neutrino analysis using the MINOS Far Detector the process of event reconstruction is therefore relatively simple.

The *CandFitTrackList* and *CandShowerList* are read in. Vertex showers are associated with the matching tracks. The assumption is made at this point that the longest track in the snarl is a muon (occasionally in Monte Carlo it is seen to be a pion and in such cases the muon track is too short to find). In order to avoid tagging Bremsstrahlung showers from the muon as hadronic showers, the vertex shower is required to be within 2 m of the track vertex. The neutrino energy is estimated from the sum of the muon energy and the energy of the vertex shower. The estimate of the shower energy is based on the pulse height deposited in the strips which make up the shower.

In the case of a vertex shower, it is often the case that some shower strips will also be on the track and so the total shower pulse height will contain contributions from the energy deposited by the muon. In order to correct for this extra pulse height, an estimate of the energy deposited by the muon in the strip using the Bethe-Bloch equation and an estimate of the energy of the muon at that point on the track. The calculated muon pulse height is subtracted from the strip pulse height and the result is added to the total shower pulse height. If the corrected strip pulse height is less than zero then nothing is added to the shower total.

The shower energy in GeV is then measured using a parameterisation of the pulse height in MIPs developed using charged current ν_μ and $\bar{\nu}_\mu$ Monte Carlo events. The

CandFitTrack identified as the muon, along with the *CandShower* identified as the vertex shower in the event are used to build a *CandEvent* object which is written out in a *CandEventList*.

7.10 Shower Energy Calibration

A large sample of atmospheric ν_μ and $\bar{\nu}_\mu$ charged current Monte Carlo events with vertex showers were reconstructed. For each event the total vertex shower pulse height was calculated as described above. The events were divided up in bins of true hadronic energy. For each true hadronic energy bin, the mean vertex shower pulse height was then calculated. To remove the effects of outliers a truncated mean was then calculated by removing contributions from events with vertex shower pulse heights more than two standard deviations from the original mean. The remaining events were then used to calculate a mean true hadronic energy for the bin. Mean true hadronic energy was then fitted against mean vertex shower pulse height using a quadratic function for low shower pulse heights and a linear function for high shower pulse heights. The functions were constrained to be continuous at the change over point, which was a free parameter in the fit. The resulting parameterisation was then used to calculate reconstructed shower energies given vertex shower pulse height.

Figure 7.9 shows the reconstructed shower energy as a function of true shower energy for atmospheric ν_μ and $\bar{\nu}_\mu$ charged current Monte Carlo events. It is clear from Figure 7.10 that the shower energy is calibrated to a reasonable degree. A shower energy resolution of $\sigma/E = 0.34 \oplus 0.52/\sqrt{E(\text{GeV})}$ has been achieved. The resolution is lowered by the wide variation in detector acceptance with shower propagation angles and shower composition.

Figure 7.9: Reconstructed vertex shower energy versus true E_{shw} for atmospheric ν_μ and $\bar{\nu}_\mu$ charged current Monte Carlo events. The line $E_{reco} = E_{true}$ is shown for comparison.

Figure 7.10: $\sigma_{E_{shw}}/E_{true}$ versus true E_{shw} for atmospheric ν_μ and $\bar{\nu}_\mu$ charged current Monte Carlo events.

Chapter 8

Atmospheric Neutrino Event Selection

“A casual stroll through the lunatic asylum shows that faith does not prove anything.”
— Friedrich Nietzsche, 1844–1900

8.1 Selection Overview

The event selection used in this thesis differs from those used by previous atmospheric neutrino analyses at MINOS [29] [94] [96] in that there are two paths via which each event can pass the selection. In essence events can pass the selection either by being upward-going or by the highest end of the track being contained within the fiducial volume. In reality though more stringent cuts must be applied in order to reduce cosmic muon contamination of the samples, but these criteria define the two main paths. The selected events are divided up into four categories: Fully-Contained (FC), Partially-Contained Downward-going (PCDN), Partially-Contained Upward-going (PCUP) and Upward Muon (UPMU) events. The event selection is summarised in Figure 8.1.

Selection plots in this chapter are shown with Monte Carlo normalised to the data exposure 12.23 ktyr and oscillated using the best fit oscillation parameters from the MINOS measurement of the NuMI beam ($|\Delta m_{32}^2| = 2.74 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$) [74] in order to show better the level of agreement between data and Monte Carlo distribu-

Figure 8.1: An overview of the event selection used in this analysis

tions of selection variables. Contributions from ν_τ and $\bar{\nu}_\tau$ events are shown in the plots also.

8.2 Data Quality

Most of the data quality checks are done before snarls are considered for reconstruction. Snarls recorded in time periods where the Far Detector magnetic field was off (Section 4.6.1) or parts of the readout were offline (Section 4.6.2) are removed at this stage.

The Light Injection system used to calibrate the readout (see Chapter 5) contributes two further sources of background. If a trigger occurs too soon after a nearby region of the detector was ‘flashed’ by the LI system, then the electronics dead time incurred while the LI induced hits are read out (up to 30 μ s) may cause a through-going event to be identified as a contained vertex event. For example, if a muon enters the detector in a dead region, then the first hits it leaves in the detector will be when it leaves the dead region, which may be some way inside the fiducial volume.

The second source of background from the LI system was caused by a fault in the power supplies for the system. This resulted in lower than expected light levels from the

system. These low light levels were sometimes insufficient to activate the trigger-PMT for the electronics crate being flashed. This meant the Trigger Processor analysing the hits had no way of knowing that the pulse was caused by the LI system. The hits would then be written out as snarls in the normal data stream, rather than being used in the calibration as intended. Fortunately, the low power meant that the pulse height in the scintillator strips was usually too low to be read out on the opposite side of the detector to that being flashed. Thus, these LI induced snarls are characterised by a large fraction of pulse height from single-ended strips and is easily vetoed. The problem has now been fixed.

In earlier versions of the trigger algorithms events which occurred too near the end of a time frame boundary were likely to be truncated. For this reason events starting less than 100 μs before a time frame boundary are vetoed.

8.3 Removal of Beam Events

From January 2005 the NuMI beam started commissioning and from March 2005 was delivering around 10^{13} protons on target per spill. This adds a source of background neutrino events to the atmospheric analysis. The flagging of beam snarls in MINOS is quite good. Most spills (99% [109]) are recorded in a database table and an offline software package has been developed to allow comparison of trigger times with the nearest spill time at the Far Detector.

The plot of (spill arrival time - trigger time) in Figure 8.2 shows two peaks. There is an obvious main peak around 40 μs and a much smaller peak around zero. The large peak is due to spill-triggered snarls. In this case the trigger processor opens up a window which starts 40 μs before the spill time and ends 40 μs afterwards (in case of any errors in the spill time predictions). Thus the trigger time is peaked 40 μs before the spill. This means that while the NuMI beam is on, any event occurring within 40 μs of a spill arrival time is contained within the spill-triggered snarl. This is the reason for the depletion of 4/5 plane triggered snarls with trigger times within 40 μs of a spill arrival time. Any spill-triggered snarls are vetoed from the atmospheric neutrino analysis at this point.

The smaller peak is due to beam events occurring when the SpillServer at the Far Detector was offline. In this situation beam events can still be picked up by the 4/5

Figure 8.2: Distribution of time from trigger time to the nearest spill. Snarls with spill triggers are plotted in red. Snarls with 4/5 plane triggers are plotted in green.

plane trigger when the spill arrives, so (spill arrival time - trigger time) will be peaked around zero. These ‘missed’ beam snarls can be easily removed by vetoing any snarl with a trigger time within 15 μs of a spill arrival time. At some points the error in the predicted spill times were large (greater than 100 ns nominal uncertainty) in these cases all snarls are vetoed to prevent contamination.

Finally, at some times (particularly in January 2005 when the NuMI beamline was being commissioned) the Near Detector process which fills the spill time database table and communicates with the SpillServer at the Far Detector died while the NuMI beam was active, meaning that spill events could not be vetoed. Again in order to prevent beam contamination all snarls are vetoed. This results in the loss of around 4 days (0.06 ktyr) worth of data.

8.4 Pre-selection

The first part of the pre-selection focuses on the quality of the reconstruction for each snarl. For an event to be pre-selected it must satisfy the following criteria:

- There must be a reconstructed track.
- The track must span at least 8 planes.
- The track must leave hits on at least 7 scintillator planes.
- There must be track hits on at least 3 planes in each view.
- The difference in the number of planes hit in each view must be less than 10.
- At least 5 of the planes in the track must be considered *track-like*¹.
- The track must have a *UVasymmetry* (defined below) of less than 0.5.

In addition a cut on the total pulse height in any single plane in the snarl is applied. To pass, no plane in the event must have a total pulse height of more than 500 PE. This vetoes steep cosmic muon events in which the muon has travelled down a single scintillator plane for a large distance.

The *UVasymmetry* of a track is defined as:

$$UVasymmetry = 2 \times \left| \frac{N_{Uplanes} - N_{Vplanes}}{N_{Uplanes} + N_{Vplanes}} \right|, \quad (8.1)$$

where $N_{Uplanes}$ is the number of *U*-view planes in the track and $N_{Vplanes}$ is the number of *V*-view planes in the track. Events where the track *UVasymmetry* is large are more likely to be reconstruction failures. The effect of the cut on this variable is shown in Figure 8.3, as well as ensuring remaining events are more likely to be well reconstructed, this cut also removes a significant amount of cosmic muon Monte Carlo events which could have otherwise appeared to be signal events.

The second stage of the pre-selection makes a preliminary classification of the event to see whether it is either upward-going from timing and/or the highest end of the track is contained.

8.4.1 Upward-going Pre-selection

The upward-going branch of the selection relies on timing information to identify upward-going events. The majority of reconstructed events are cosmic muons. The main background to the upward-going neutrino event sample are stopping cosmic muon events where the direction of muon propagation has been mis-reconstructed as upward. This

¹For a track propagating purely in the z direction a track-like plane has less than 80 PE deposited on the plane and 80% of the pulse height on the plane has to be within two strip widths of the track position. This window is widened for steeper tracks, as is the maximum pulse height condition.

Figure 8.3: Track $UVasymmetry$, events to the right of the red line are vetoed from the analysis.

can happen in one of two ways. If the event has poor timing information then, although the correct hits in the detector have been added to the track, the reconstructed direction the muon travelled along the track may be wrong. Secondly, the initial direction cosines of the track may be incorrect. This can happen due to the wrong strips being added to the track at the vertex and/or hits around the vertex being demultiplexed into the wrong strips. Backgrounds due to mistakes in the direction cosines are most common for very horizontal events. The upward-going pre-selection concentrates on removing those events where these reconstruction ambiguities are possible and is summarised as follows:

- Track vertex y direction cosine > 0.0 ,
- Quality of track vertex direction reconstruction,
- Quality of timing information.

The reason for the first cut is clear and is the loosest possible definition of the sample. The quality of the track vertex direction reconstruction is checked in two ways. As part of the reconstruction a straight line is fitted to the track strips in each view. The RMS deviation of the track strips from this line is calculated in each view (RMS_{Linear}). The RMS deviation of the track strips to another straight line passing through the centre of the event with gradient taken from the vertex direction cosines is calculated for each

Figure 8.4: Difference in RMS track strip position deviation from the best straight line fit to a track and a line with gradient taken from the track vertex direction cosines and passing through the centre of the event, events to the right of the red line are vetoed from the upward-going branch of the analysis.

view (RMS_{Vertex}). If the track is relatively straight (defined as $RMS_{Linear} < 1.0$ m), then RMS_{Vertex} is checked. If $RMS_{Vertex} > 1.0$ m or $RMS_{Vertex} - RMS_{Linear} > 0.5$ m, this means the vertex direction is inconsistent with the rest of the track, despite the fact that the track is close to being straight. Such discrepancies are usually due to reconstruction failures, so these events are vetoed. Figure 8.4 shows the difference in the two RMS values for events where $RMS_{Linear} < 1.0$ m. For Monte Carlo, the vetoed region is dominated by cosmic muon events. This cut is useful for removing cosmic muons which enter the detector with a large shower, as shown in Figure 8.5. This can sometimes confuse the track finder/fitter into extending the track through the shower along the side of the detector by mistake.

The second method for checking the vertex direction cosines is to remove the first strip in each view from the calculation of the track direction cosines. If the old and new y direction cosines differ in sign, then the direction reconstruction is insufficient to classify the event as upward-going and the event is vetoed from the upward-going selection.

Figure 8.5: An example of a data event where a cosmic muon entering the detector within a large shower. The VZ view is shown.

In the reconstruction the hits on the track are tested against two hypotheses $s_i = cT_o + cT_i$ and $s_i = cT_o - cT_i$, where s_i is the distance of the i th hit along the track, T_i is the recorded hit time, T_o is an offset time and c is the speed of light, as shown in Figure 8.6. The RMS of the track hits to each hypothesis is calculated. In the track finder the hypothesis with the lowest RMS is used to give the direction of motion of the particle along the track. A loose assessment of the timing resolution of the events is made using the difference in RMS between the two hypotheses. If the best fit direction has changed between the track finder and the track fitter (a good indication that the resolution is poor) then the difference in the RMS of the ‘correct’ and ‘incorrect’ hypotheses (as determined by the track finder) will be positive. The difference in RMS values is shown in Figure 8.7. Remaining cosmic muon Monte Carlo events are equally likely to have negative as positive differences, while neutrino events are concentrated on the positive side of the plot. Any events on the positive side of the plot are vetoed from the upward-going selection.

8.4.2 Contained-Highest-End Pre-selection

The atmospheric neutrino signal events are hidden in a huge background of downward-going cosmic muon events. This branch of the analysis assumes that all tracks are due to cosmic muons. In general cosmic muon events are downward-going, so irrespective of timing information the end of the track with the highest y -coordinate is treated as the vertex and all initial cuts are applied to this end of the track. In this way cosmic muon events which have mis-reconstructed as upward-going can be vetoed along with the rest of the cosmic muon background.

The contained-highest-end pre-selection can be summarised as follows:

- Require the highest end of the track to be contained within the fiducial volume (defined below).
- Trace cut on the highest end of the track of 0.15 m (defined below, see Figure 8.10).
- No activity before the highest end of the track within 5 planes of the end of a supermodule.

The fiducial volume chosen reflects the octagonal prism shape of the Far Detector. In the UV plane the fiducial area is an octagon 7.6 m across (compared to 8.0 m for a steel plane) centred on the coil hole, with a circle 40 cm in radius removed from the centre to veto events where vertex shower energy could be lost into the coil hole. Along the z -axis,

Figure 8.6: Cartoon showing the hits on an example track compared with the lines $s = cT_o - cT$ (dashed line) and $s = cT_o + cT$ (solid line). Red points represent hits from strip-ends on the East side of the detector. Blue points represent hits from strip-ends on the West side of the detector. Taken from [94].

Figure 8.7: Difference in RMS between fits of track hits to $s = cT_o + cT$ and $s = cT_o - cT$, events to the right of the red line are vetoed from the upward-going branch of the analysis.

Figure 8.8: Definition of the ‘Distance To Edge’ variable, as opposed to the radial co-ordinate.

Figure 8.9: ‘Distance To Edge’ for highest end of track, events to the left of the red line are vetoed from the Contained-Highest-End branch of the analysis.

Figure 8.10: The muon track is projected back (forward) from the vertex (end) to the edge of the detector to determine the muon entry (exit) point in the detector. The trace cut is made on the z component of the distance between the track vertex (end) and muon entry (exit) point. Taken from [94].

a veto region at each end of the supermodules is used to stop cosmic muons sneaking in through the ends of the detector. The fiducial region is defined as $5 < \text{plane} < 244$ and $254 < \text{plane} < 481$. The cut of 0.2 m on the distance of the highest end of the track to the edge of the detector (as defined in Figure 8.8), which makes up part of this containment cut is shown in Figure 8.9. The negative values in the plot correspond to points reconstructed outside the detector volume. This is often due to a demultiplexing error where the demultiplexer returns an unphysical combination of strips, in other cases it can be caused by the fact that the track vertex/end is extrapolated by half a plane width in the z direction further than the first/last track strip. For steep tracks this can move the vertex/end of the track outside the detector volume.

The trace of a track, as defined in Figure 8.10, is an extrapolation back along a straight line in the vertex (end) direction from the vertex (end) of the track to the edge of the detector. The z -component of this line is used to give a measure of how many planes a cosmic muon track would have had to pass through without leaving a measurable energy deposit in a strip. A cut is made on this variable to remove tracks with a z -component of less than 0.15 m, the effect of this cut is illustrated in Figure 8.11.

For the third pre-selection cut, the planes either side of the track are checked for activity. The extremal plane in each direction is called the *DigitTrackVertexPlane* or

Figure 8.11: Trace of highest end of track, events to the left of the red line are vetoed from the Contained-Highest-End branch of the analysis.

DigitTrackEndPlane depending on whether it is closer to the vertex or end of the track. If no new planes with activity are found then the variables are set to the reconstructed track vertex or end plane. For an event to pass the pre-selection when the vertex (end) of the track is highest then the *DigitTrackVertexPlane* (*DigitTrackEndPlane*) is also required to be in the fiducial region (*i.e.* $5 < \text{plane} < 244$ and $254 < \text{plane} < 481$).

8.5 Upward-Going Events

The selection of upward-going events can be summarised as follows:

- Reconstruction quality cuts
- Topology cuts
- Timing cuts
- Vertex containment cuts

Events are then divided according to the containment of the end of the track into fully contained upward-going (FCUP) events and partially contained upward-going events.

Figure 8.12: Fraction of non-fiducial track pulse height, events to the right of the red line are vetoed from the upward-going branch of the analysis.

8.5.1 Reconstruction Quality Cuts

For events propagating through the detector near to the ends of the scintillator strips in one view, determining the longitudinal position of hits along a strip precisely is sometimes difficult as hits in the opposite view may be missing. This means the calibrated hit times may not be as accurate as for other tracks propagating more centrally through the detector. For this reason shorter tracks (less than 50 planes) with less than 50% of track pulse height within the fiducial volume are removed from the sample. The effect of this cut is shown in Figure 8.12. The remaining cosmic muon sample peaks above 50%, while the atmospheric neutrino distribution is strongly suppressed here.

The selection of upward-going atmospheric neutrino events relies on good quality timing information. This requires that tracks are well reconstructed and are greater than a minimum length, as shorter tracks and tracks with certain topologies may not have sufficient timing resolution to determine the direction that the particle being tracked was moving. For this reason tracks are required to span at least 10 planes and have a range of at least 1.0 m. In order to ensure the timing information available is reasonably accurate tracks are also required to have at least 10 *track-like* planes. The effect of this

Figure 8.13: Number of *track-like* planes in tracks, events to the left of the red line are vetoed from the upward-going branch of the analysis.

cut on the data and Monte Carlo distributions is shown in Figure 8.13.

8.5.2 Topology Cuts

For cosmic muons that enter the detector parallel to the scintillator strips or bend sharply in the magnetic field, errors in the reconstruction of events can generate scatter in the timing at the beginning of the track. This can cause events to be mis-reconstructed as up-going, such topologies are illustrated in Figure 8.14. These events typically contain large energy deposits and/or several hits above the track vertex. As Figure 8.14(a) shows, events may also change direction sufficiently to deposit energy in two different positions within a plane, making correct demultiplexing of the observed digits very difficult. In order to remove such events two variables can be defined. The largest 3D displacement of strips from the muon vertex within ± 4 planes from the vertex, R_{max} , is calculated by joining together the U and V strips in adjacent planes. Also, the total pulse height per plane is calculated for each plane in this region and the maximum pulse height is called the Q_{max} value of the vertex.

(a) Cosmic muon background events often contain several early hits above the reconstructed track.

(b) Cosmic muon background events often contain large energy deposits in the early planes.

Figure 8.14: Typical background topologies. Figures taken from [94].

An initial topology cut that $R_{max} < 3.0$ m is applied to all events in the sample. For the remaining events stricter requirements that $R_{max} < 1.25$ m and $Q_{max} < 300$ PE are applied. However, in order not to remove many higher energy events, where larger Q_{max} and R_{max} values may be due to high energy showers, these stricter cuts are only applied to events with tracks spanning 25 planes or less. The distribution of Q_{max} values for tracks spanning less than 25 planes is shown in Figure 8.15. Higher values of Q_{max} are more common in data than in Monte Carlo. This is because, as mentioned previously, real cosmic muons often enter the detector accompanied by showers of other particles, these showers are not simulated in the current MINOS cosmic muon Monte Carlo. Cutting on these variables helps improve the agreement between the two samples.

8.5.3 Timing Cuts

The the next stage of the selection is to apply the main timing cuts. When the *Cand-FitTrack* objects are formed, a fit to:

$$cT_i = \frac{s_i}{\beta} + s_0, \quad (8.2)$$

Figure 8.15: Q_{max} distribution, events to the right of the red line are vetoed from the upward-going branch of the analysis.

where s_i is the distance of the i th hit along the track, T_i is the recorded hit time, s_0 is an offset and the velocity, $v = \beta c$; is used to determine the velocity of the particles tracked. The sign of $\frac{1}{\beta}$ is determined by the z component of the particle's velocity. Using the y and z direction cosines of the track this can be converted so that the sign is determined by the y component of the particle's velocity, this new variable is termed $\frac{1}{\beta_{up}}$. To increase the probability that the selected events actually contain upward-going muon tracks the selection requires that $0.5 < \frac{1}{\beta_{up}} < 2.0$, *i.e.* that the value corresponds to an upward velocity of around the speed of light. As Figure 8.16 shows, the effect of this cut is to remove most of the remaining cosmic muon MC events.

In the following RMS_{Up} refers to the RMS deviation of hits on the track from the line $s = cT_o + cT$ (where s increases along the track from vertex to end) and RMS_{Down} refers to the RMS deviation of hits on the track from the line $s = cT_o - cT$ (see Section 8.4.1), as the muon tracks in all remaining events appear upward-going from their vertex directions. In order to be selected as part of the upward-going sample events must pass the following cuts $RMS_{Up}/(\text{track range}) < 0.66$, $RMS_{Up} < 1.33$ m (see Figure 8.17), $RMS_{Down} > 1.0$ m (see Figure 8.18) and $RMS_{Up} - RMS_{Down} < -0.5$ m (see Figure 8.19).

Figure 8.16: Fitted $\frac{1}{\beta_{up}}$ values for all remaining events, events to the left of the red line are vetoed from the upward-going branch of the analysis.

Figure 8.17: Distribution of RMS_{Up} values for events remaining in the upward-going sample, events to the right of the red line are vetoed from the upward-going branch of the analysis.

Figure 8.18: Distribution of RMS_{Down} values for events remaining in the upward-going sample, events to the left of the red line are vetoed from the upward-going branch of the analysis.

Figure 8.19: Distribution of $RMS_{Up} - RMS_{Down}$ values for events remaining in the upward-going sample, events to the right of the red line are vetoed from the upward-going branch of the analysis.

8.5.4 Vertex Containment Cuts

The remaining events at this stage should be dominated by upward-going neutrino induced muons. The next set of cuts divide the events into two groups; one where the neutrino interactions occur inside the fiducial volume of the Far Detector and another where the interactions occur elsewhere. The cut used to define these two samples is that the track vertex lies within the fiducial volume. In the UV plane the fiducial region is an octagon 7.6 m across (compared to 8.0 m for a steel plane) centred on the coil hole, with a circle 40 cm in radius removed from the centre to veto events where vertex shower energy would be lost into the coil hole. Along the z -axis, a region at each end of the supermodules is used to tag events sneaking in through the ends of the detector. The fiducial region is defined as $5 < \text{plane} < 244$ and $254 < \text{plane} < 481$. Any events with track vertices outside the fiducial region are placed in the uncontained vertex sample. Contained events are also required to have a vertex trace z -component greater than 0.15 m. After an investigation of events entering the fiducial volume via the coil hole additional regions consisting of cylinders 1.0 m in radius, centred on the coil hole in the UV plane and extending 15 steel planes into the detector at each end and 9 planes into each supermodule at the supermodule gap, was used tag these events. Any events with track vertices within this region were also placed in the uncontained vertex sample. Currently nothing further is done with the uncontained vertex events. This sample is dominated by upward-going neutrino induced muons and could be used in a future oscillation analysis.

8.5.5 Fully Contained/Partially Contained Separation

The final stage of this branch of the selection is to divide the contained vertex events into Fully-Contained Upward-Going (FCUP) events (those where the entire event is contained within the fiducial volume) and Partially-Contained Upward-Going (PCUP) events (those where the track leaves the fiducial volume). No events are thrown away in this stage of the selection. The containment of the end of the track in each event is established in the same way as above. Checks are made that both the co-ordinates of the end of the track and the Digit-based extremal variables described above lie within the fiducial volume. In addition, in order to be classified FCUP, the z -component of the trace from this lowest end of the track to the edge of the detector is required to be greater than 0.15 m.

8.5.6 Summary of Upward-Going Selection

A summary of the upward-going selection is given in Table 8.1.

8.6 Contained-Highest-End Events

The selection of events via the contained highest end of track branch can be summarised as follows:

- Highest End Containment Cuts
- Topology Cuts

Events are then divided according to the containment of the lowest end of the track into fully contained (FC) events and partially contained downward-going candidate events. The timing of partially contained downward-going candidate events is then checked to remove events where the timing information does not agree with this hypothesis. This removes a background from upward-going neutrino-induced muon events with poor timing. The remaining events are then classified as partially contained downward-going (PCDN).

8.6.1 Highest End Containment Cuts

As in the pre-selection, the planes either side of the track are checked for activity. In addition to finding the extremal plane hit on either side of the track, the extremal strips in the UV plane both before and after the track are used to calculate the *DigitTrackVertexR* and *DigitTrackVertexDistToEdge* (similarly for the end of the track), where these variables are defined as in Figure 8.8. A check is made on the containment of the highest end of the track (as described previously) using these digit based variables also. A harsher cut on the trace of the highest end of each track is also applied, as shown in Figure 8.20. Tracks are required to have a trace with a z -component greater than 0.50 m.

A second kind of trace variable is also used in this part of the selection. The best fit directions from a linear fit to the whole track are used to calculate the path of this second trace. Any strips within a three strip window of this line are examined and the *DigitTrace* at either end is defined from the extremal strips, rather than from the vertex

Selection Cut	Data	Monte Carlo Expectation					
		Total MC (unoscillated)	$\nu_\mu/\bar{\nu}_\mu$ CC	$\nu_e/\bar{\nu}_e$ CC, NC	Cosmic μ^-/μ^+	Upward μ^-/μ^+	Neutrons
Pre-Selection	7729	7543.0 ± 723.8	206.6 ± 31.0	12.9 ± 1.9	7073.8 ± 722.2	245.4 ± 36.8	4.3 ± 1.0
Reco Quality	1474	1421.0 ± 124.6	160.1 ± 24.0	3.3 ± 0.5	1036.6 ± 117.7	220.0 ± 33.0	1.0 ± 0.4
Topology Cuts	1416	1392.3 ± 122.6	156.3 ± 23.4	2.8 ± 0.4	1021.7 ± 116.2	210.5 ± 31.5	1.0 ± 0.4
Timing Cuts	194	284.3 ± 31.7	94.1 ± 14.1	0.5 ± 0.1	—	189.5 ± 28.4	0.2 ± 0.2
Contained Vertex Cuts	56	85.9 ± 12.5	83.6 ± 12.5	0.4 ± 0.1	—	1.9 ± 0.3	—
FCUP	29	40.7 ± 6.0	39.9 ± 6.0	0.3 ± 0.0	—	0.5 ± 0.1	—
PCUP	27	45.2 ± 6.6	43.7 ± 6.6	0.1 ± 0.0	—	1.4 ± 0.2	—

Table 8.1: Summary of the Upward-Going Selection. The quoted uncertainties in the atmospheric neutrino and upward muon Monte Carlo expectations are due to the 15% normalisation uncertainty. The quoted uncertainties in the cosmic muon Monte Carlo expectations are due to the statistical error combined in quadrature by a 10% normalisation uncertainty. A dash indicates no Monte Carlo events passed the cuts at this stage. In addition, the expected number of $\nu_\tau/\bar{\nu}_\tau$ CC events in the samples if ν_μ and $\bar{\nu}_\mu$ oscillate with the MINOS beam analysis best fit parameters ($|\Delta m_{32}^2| = 2.74 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$ [74]) is found to be 0.8 ± 0.1 PCUP events and 0.5 ± 0.1 FCUP events, based on a Monte Carlo sample equivalent to 360.5 ktyr.

Figure 8.20: z component of High End Trace, events to the left of the red line are vetoed from the contained highest end branch of the analysis.

and end of the track itself. The value of the *DigitTraces* are very similar to the normal traces for most events. The value of the cut is in catching events where the tracking has failed partway along the track for some reason, but which still pass the containment and digit containment cuts. The z -component of the *DigitTrace* of the highest end of the track is required to be greater than 0.50 m for an event to remain in the sample. Figure 8.21 clearly shows events which passed the previous trace cut ‘leaking’ into the vetoed region.

8.6.2 Topology Cuts

As mentioned in Section 8.5, events with certain topologies, particularly steep or highly curved tracks (see Figure 8.14), are more likely to have reconstruction errors or to be due to cosmic muons which have entered the fiducial volume down a steel plane, then scattered out, leaving an observable track. These kinds of cosmic muon events are the dominant remaining background in the sample at this point. In addition to the R_{max} and Q_{max} variables defined in Section 8.5, two extra variables are used to remove background in this sample. They are the charge-weighted mean ($\langle \Delta_T \rangle$) and RMS ($\langle \Delta_T^2 \rangle^{\frac{1}{2}}$) transverse

Figure 8.21: z component of High End Trace using digits, events to the left of the red line are vetoed from the contained highest end branch of the analysis.

displacements of strips from the transverse co-ordinate of muon track vertex (end) and are calculated using strips in the region ± 4 planes from the vertex (end). Events are required to satisfy $\langle \Delta_T \rangle < 0.25$ m about the highest end of the track, as shown in Figure 8.22, and $\langle \Delta_T^2 \rangle^{\frac{1}{2}} < 0.50$ m about the highest end of the track, as shown in Figure 8.23.

Events with shorter tracks (20 planes or less) are required to satisfy $R_{max} < 0.8$ m (Figure 8.24) and $Q_{max} < 300$ PE, as shown in Figure 8.25. These events are also required to satisfy $|\frac{1}{\beta}| < 3.0$. The $\frac{1}{\beta}$ cut is less harsh for this sample than in the upward-going selection, as it contains many shorter fully contained neutrino events, for which timing resolution may not be as good.

The remaining background is reduced by varying the Q_{max} cut with the direction of the muon track. The cosmic muon background is characterised by large energy deposits at the vertex and directions peaked toward vertical angles and small angles to the planes. In addition the vast majority of the remaining cosmic muons are found to span less than 25 planes. In contrast, the atmospheric neutrino signal is characterised by smaller energy deposits at the vertex and a more uniform angular distribution. In order to separate the signal and background, the events spanning less than 25 planes with $\cos \Theta_y > 0.7$ or

Figure 8.22: Mean transverse displacement of strips from the highest end of the track ($\langle\Delta_T\rangle$), events to the right of the red line are vetoed from the contained highest end branch of the analysis.

Figure 8.23: RMS transverse displacement of strips from the highest end of the track ($\langle\Delta_T^2\rangle^{1/2}$), events to the right of the red line are vetoed from the contained highest end branch of the analysis.

Figure 8.24: Distribution of R_{max} about the highest end of the track, events to the right of the red line are vetoed from the contained highest end branch of the analysis.

Figure 8.25: Distribution of Q_{max} about the highest end of the track, events to the right of the red line are vetoed from the contained highest end branch of the analysis.

Figure 8.26: Distribution of vertex pulse height Q_{max} as a function of $\cos \Theta_y$, events in the crossed-out regions are vetoed from the contained highest end branch of the analysis.

$|\cos \Theta_z| < 0.5$ are required to satisfy an additional cut of $Q_{max} < 70$ PE. The effects of these cuts are shown in Figure 8.26 and Figure 8.27 respectively.

8.6.3 Fully Contained/Partially Contained Separation

The next stage of this branch of the selection is to divide the remaining events into Fully-Contained (FC) events (those where the entire event is contained within the fiducial volume) and Partially-Contained (PC) events (those where the lowest end of the track leaves the fiducial volume). The PC events must then be checked to ensure that the timing information is good enough to support the hypothesis that it is the true track vertex which lies within the fiducial volume and the end of the track which leaves the

Figure 8.27: Distribution of vertex pulse height Q_{max} as a function of $\cos \Theta_z$, events in the crossed-out regions are vetoed from the contained highest end branch of the analysis.

Figure 8.28: ‘Distance To Edge’ for lowest end of track, events to the left of the red line are moved to the partially contained downward-going branch of the analysis.

fiducial volume. If this is the case then the events may be classified as Downward-Going Partially-Contained (PCDN).

The containment of the lowest end of the track in each event is established in the same way as the highest end above. Checks are made that both the co-ordinates of the lowest end of the track and the Digit-based extremal variables described above lie within the fiducial volume. Figure 8.28 shows the effect of a cut on the distance of the lowest end of the track from the edge of the detector. A large fraction of the events in the vetoed region are cosmic muon and upward muon events. In addition, in order to be classified FC, the z -component of the trace from this lowest end of the track to the edge of the detector is required to be greater than 0.15 m. The trace cut on the lower end does not have to be as tough as for the higher end of the track, as the cut here is simply a means of dividing the sample into two rather than removing remaining background. Figure 8.29 shows the distribution of this variable.

Figure 8.29: The z -component of the trace for lowest end of track, events to the left of the red line are moved to the partially contained downward-going branch of the analysis.

8.6.4 Downward-Going Partially Contained Event Selection

Only those events with sufficiently good timing information (agreeing with the PCDN hypothesis) can make it into the PCDN sample. For that reason the selection criteria require that PC events span at least 10 planes, contain at least 6 track-like planes and have a track range of at least 1.0 m. As described in the upward-going pre-selection, if the best fit direction from timing has changed between the track finder and the track fitter (a good indication that the resolution is poor) then the difference in the RMS of the ‘correct’ and ‘incorrect’ hypotheses (as determined by the track finder) will be negative. Any such events are vetoed from the PCDN sample. Figure 8.30 shows the distribution of $\frac{1}{\beta_{up}}$ values for all remaining events. The selection next requires that $-2.0 < \frac{1}{\beta_{up}} < -0.5$. The timing information for the remaining events is now correct to a high level of confidence. Now all that remains is to check that the vertex direction of the track in each event is downward-going and to check that the vertex does indeed lie within the fiducial volume.

Figure 8.30: Fitted $\frac{1}{\beta_{up}}$ values for all remaining events, events in the crossed-out region are vetoed from the partially contained downward-going branch of the analysis.

8.6.5 Veto Shield Cuts

The veto shield is used to reduce the number of background events in the contained highest end selection branch. In order to determine if an event should be vetoed by the shield the highest end of the track is compared with each hit shield plank in the snarl in turn. If the highest end occurs under the section of the shield containing the hit then it is possible for the shield hit to veto the event. The shield hit time corrected for travel time along the scintillator, T_{shield} , is compared with the time of the highest end of the track corrected for the travel time for a muon moving at c between the highest end of the track and the shield plank, T_{det} (effectively a prediction of the shield hit time using the detector timing information). If these two times coincide to within ± 50 ns then the event is vetoed by the shield. Figure 8.31 shows a plot of the distribution of smallest $T_{\text{shield}} - T_{\text{det}}$ values, in shield sections above the highest end of the track for each event in a stopping muon dominated data sample created by relaxing the trace, topology and timing cuts. Figure 8.32 shows the same plot for data events passing the FC and PCDN selections up to this point.

The efficiency of the veto shield for tagging cosmic muons is measured using selected

Figure 8.31: Distribution of the smallest $T_{\text{shield}} - T_{\text{det}}$ values, in shield sections above the highest end of the track for a stopping muon dominated sample. The events in the region between the two red lines would be vetoed.

Figure 8.32: Distribution of the smallest $T_{\text{shield}} - T_{\text{det}}$ values, in shield sections above the highest end of the track for events in the FC and PCDN samples. The events in the region between the two red lines are vetoed.

FC and PCDN events. The trace, topology and timing cuts are backed off leaving a sample of 41027 data events. After the veto shield cuts 1572 of these events remain. Using the MINOS beam analysis best fit parameters ($|\Delta m_{32}^2| = 2.74 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$ [74]) the atmospheric neutrino expectation from Monte Carlo is 307.9 events. Taking this into account, the overall efficiency is calculated to be $96.9 \pm 0.2\%$. A larger sample of selected events was obtained by using all events which passed the contained highest end pre-selection cuts. This sample contained 95214 data events, which were reduced down to 3037 events by the veto shield cut. Using the MINOS beam analysis best fit parameters the atmospheric neutrino expectation from Monte Carlo is 333.1 events. Taking this into account, the overall efficiency is calculated to be $97.2 \pm 0.1\%$.

Accidental coincidence of shield activity generated by natural radioactivity or noise will result in a small fraction of signal events being vetoed. This loss of signal was measured by combining a sample of MC atmospheric $\nu_\mu/\bar{\nu}_\mu$ CC events with shield hits extracted from the pre-trigger window of events in the data. The overall loss of signal was calculated to be $2.2 \pm 0.1\%$ [94].

8.6.6 Summary of Contained Highest-End Selection

A summary of the contained highest end selection is given in Table 8.2.

8.7 Combining the Two Branches of The Selection

The selections described above are performed for all MC and data events. After the selection it is possible to see whether an event has passed either branch, or both or indeed neither. Due to the nature of the classifications the only overlap which occurs is between the FC and FCUP samples. The results are summarised in Table 8.3 below. The expected number of $\nu_\tau/\bar{\nu}_\tau$ CC events in the samples if ν_μ and $\bar{\nu}_\mu$ oscillate with the MINOS beam analysis best fit parameters ($|\Delta m_{32}^2| = 2.74 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$ [74]) is found to be 0.8 ± 0.1 FC (only) events, 0.1 ± 0.0 FCUP (only) events, 0.5 ± 0.1 events which pass both the FC and FCUP cuts, 0.8 ± 0.1 PCUP events and 0.03 ± 0.0 PCDN events. Giving a total background of 2.2 ± 0.3 events.

Atmospheric neutrino interactions are expected to occur uniformly in time and uni-

Selection Cut	Data	Monte Carlo Expectation					
		Total MC (unoscillated)	$\nu_\mu/\bar{\nu}_\mu$ CC	$\nu_e/\bar{\nu}_e$ CC, NC	Cosmic μ^-/μ^+	Upward μ^-/μ^+	Neutrons
Pre-Selection	95214	99065.1 ± 9869.3	386.5 ± 58.0	40.8 ± 6.1	98541.2 ± 9869.1	81.5 ± 12.2	15.1 ± 2.2
High Containment	41027	43560.3 ± 4324.0	358.6 ± 53.8	36.9 ± 5.5	43086.6 ± 4323.6	73.7 ± 11.1	4.5 ± 1.0
Trace Cuts	5031	5839.0 ± 558.7	324.7 ± 48.7	36.1 ± 5.4	5416.9 ± 556.5	58.7 ± 8.8	2.6 ± 0.7
Topology Cuts	708	828.3 ± 73.3	295.6 ± 44.3	29.8 ± 4.5	446.4 ± 57.7	54.8 ± 8.2	1.7 ± 0.6
Low-end contained (FC)	443	557.7 ± 54.7	224.1 ± 33.6	26.6 ± 4.0	305.6 ± 43.0	1.1 ± 0.2	0.3 ± 0.2
Low-end uncontained	265	270.6 ± 28.3	71.5 ± 10.7	3.1 ± 0.5	140.8 ± 24.9	53.8 ± 8.1	1.4 ± 0.5
PCDN	69	113.0 ± 17.2	43.6 ± 6.5	0.3 ± 0.04	68.9 ± 15.9	0.2 ± 0.03	—
Veto Shield Cuts							
FC Shield Cut	194	255.2 ± 33.2	219.6 ± 32.9	26.1 ± 3.9	9.2 ± 1.3	0.03 ± 0.0	0.3 ± 0.2
PCDN Shield Cut	47	45.1 ± 6.4	42.7 ± 6.4	0.3 ± 0.03	2.1 ± 0.5	0.004 ± 0.0	—

Table 8.2: Summary of the Contained-Highest-End Selection. The quoted uncertainties in the atmospheric neutrino and upward muon Monte Carlo expectations are due to the 15% normalisation uncertainty. The quoted uncertainties in the cosmic muon Monte Carlo expectations are due to the statistical error combined in quadrature with a 10% normalisation uncertainty. A dash indicates no Monte Carlo events passed the cuts at this stage. Monte Carlo values in the final two rows are estimated from the veto shield efficiencies calculated in Section 8.6.5. In addition, the expected number of $\nu_\tau/\bar{\nu}_\tau$ CC events in the samples if ν_μ and $\bar{\nu}_\mu$ oscillate with the MINOS beam analysis best fit parameters ($|\Delta m_{32}^2| = 2.74 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$ [74]) is found to be 0.03 ± 0.0 PCDN events and 1.2 ± 0.2 FC events.

Classification	Data	Monte Carlo Expectation (no oscillations)				
		Cosmic μ^-/μ^+	$\nu_\mu/\bar{\nu}_\mu$ CC	$\nu_e/\bar{\nu}_e$ CC, NC	Upward μ^-/μ^+	Neutrons
FC (only)	174	9.2 ± 1.3	188.6 ± 28.3	25.9 ± 3.9	0.0	0.3 ± 0.2
FCUP (only)	9	—	8.8 ± 1.3	0.1 ± 0.0	0.6 ± 0.1	—
FC && FCUP	20	—	31.0 ± 4.7	0.2 ± 0.0	0.0	—
PCUP	27	—	43.7 ± 6.6	0.1 ± 0.0	1.4 ± 0.2	—
PCDN	47	2.1 ± 0.5	42.7 ± 6.4	0.3 ± 0.0	0.0	—
Total	277	11.3 ± 1.8	314.8 ± 47.2	26.3 ± 3.4	2.0 ± 0.3	0.3 ± 0.2

Table 8.3: Summary of Event Selection. MC values normalised to unoscillated expectations for data exposure 12.23 ktyr. The quoted uncertainties in the atmospheric neutrino and upward muon Monte Carlo expectations are due to the 15% normalisation uncertainty. The quoted uncertainties in the cosmic muon Monte Carlo expectations are due to the statistical error combined in quadrature with a 10% normalisation uncertainty. A dash indicates no Monte Carlo events passed the cuts at this stage.

formly throughout the detector. The distribution of reconstructed neutrino interaction points in the xy plane for selected data events is shown in Figure 8.33. This shows no concentrations of events in particular regions of the detector. Similarly the distribution of interaction points along the z -axis shown in Figure 8.34 shows no unexpected features. The dip around 15 m is due to the super module gap. Figure 8.35 shows the number of data events observed in successive twelve week periods. The solid line reflects the exposure time of the detector during that time and is normalised to the Monte Carlo expectation.

Figure 8.33: Distribution of interaction points in the xy plane for selected data events. The octagon represents the edges of the detector. The other lines show the locations of the veto shield planks.

Figure 8.34: Distribution of interaction points along the z -axis for selected data events. The points indicate data events observed in each 2 m region. The solid line histogram shows the unoscillated distribution of selected Monte Carlo events. The dashed line shows the distribution of selected Monte Carlo events oscillated using $\Delta m_{32}^2 = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$ (the best fit parameters from the MINOS observations of the NuMI beam [74]).

Figure 8.35: Distribution of neutrino event rate over time. Each point represents the number of data events observed in an 84 day period. The solid histogram varies according to the exposure in each period and is normalised to the total Monte Carlo expectation in the absence of oscillations. The dashed histogram shows the exposure normalised in the Monte Carlo expectation for $\Delta m_{32}^2 = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$ (the best fit parameters from the MINOS observations of the NuMI beam [74]).

Figure 8.36: Selection Efficiency as a function of E_ν for $\nu/\bar{\nu}$ CC events.

8.7.1 Selection Efficiencies

The efficiency of the selection is determined using true contained vertex atmospheric $\nu/\bar{\nu}$ CC events, which produce muons crossing at least 8 planes. The efficiency is the fraction of these events which pass the selection cuts. Figure 8.36 shows the selection efficiency as a function of true E_ν . Figure 8.37 shows the selection efficiency as a function of muon zenith angle. For the central part of the distribution the selection efficiency hovers around 70%, the efficiency is lower for steeply upward and downward-going tracks, partly due to the difficulty of reconstructing such events. The overall efficiency for upward-going events is $70.1 \pm 0.2\%$, while the efficiency for downward-going events is only $65.8 \pm 0.2\%$. This is due to the harsher cuts which must be applied in order to remove the large cosmic muon background to this sample. In particular the efficiency for events with $\cos \Theta_{zenith} > 0.7$ is much lower than for events with $\cos \Theta_{zenith} < -0.7$. This is due to the harsh topology cuts discussed in Section 8.6.2.

8.8 Identifying A High Resolution Sample

A sub-sample of high resolution events (those with well reconstructed directions) is selected from the main sample of selected events. This high resolution sample is used to measure properties of the atmospheric neutrino flux. It is therefore important that

Figure 8.37: Selection Efficiency as a function of muon zenith angle for $\nu/\bar{\nu}$ CC events.

the direction of muon propagation along the reconstructed track has been correctly identified. A sample satisfying this condition is selected by applying further requirements on the properties of the muon track. Tracks are required to cross at least ten planes, travel at least 1.0 m through the detector, have at least ten track-like planes and have $\text{RMS}_{\text{Vtx}} - \text{RMS}_{\text{End}} < -0.1$ m. As mentioned previously, events with more hits in the track, or which span a larger distance will have better timing resolution. In the following chapter a Bayesian technique will be used to determine the L/E resolution of these high resolution events, it is therefore necessary to require that partially contained events in the high resolution sample have a successfully fitted track. Figure 8.38 shows the distribution of track length in planes for selected events. Figure 8.39 shows the distribution of track-like track planes for selected events. Figure 8.40 shows the distribution of track ranges for selected events. Figure 8.41 shows the distribution of $\text{RMS}_{\text{Vtx}} - \text{RMS}_{\text{End}}$ for selected events. These cuts were tuned on the atmospheric neutrino Monte Carlo sample. The cuts select 70.3% of $\nu_\mu/\bar{\nu}_\mu$ CC events in the main selected sample and 99.5% of those $\nu_\mu/\bar{\nu}_\mu$ CC events selected as high resolution have a correctly reconstructed track direction.

The high resolution sample is divided up into up-going and down-going events based on the muon zenith angle. The resulting separation of events is shown in Table 8.4. There is a clear deficit of events in the upward-going data sample compared to the unoscillated Monte Carlo expectation, 77 observed and 118.2 ± 17.2 expected. The deficit of data events in the downward-going sample is much less pronounced, 105 observed compared

Figure 8.38: Distribution of track planes for selected events. The solid blue line represents the total Monte Carlo expectation in the absence of oscillations.

to an unoscillated expectation of 116.3 ± 16.3 , and is within the normalisation error on the Monte Carlo expectation.

Classification	Data	Monte Carlo Expectation (no oscillations)				
		Cosmic μ^-/μ^+	$\nu_\mu/\bar{\nu}_\mu$ CC	$\nu_e/\bar{\nu}_e$ CC, NC	Upward μ^-/μ^+	Neutrons
UP	77	—	114.6 ± 17.2	1.4 ± 0.2	2.2 ± 0.3	—
DOWN	105	7.1 ± 1.1	107.8 ± 16.2	1.2 ± 0.2	0.2 ± 0.0	—
UNCERTAIN	95	4.2 ± 0.7	92.2 ± 13.8	24.0 ± 3.6	0.4 ± 0.1	0.3 ± 0.2

Table 8.4: Separation of selected events into high and low resolution samples. MC values normalised to unoscillated expectations for data exposure 12.23 ktyr. The quoted uncertainties in the atmospheric neutrino and upward muon Monte Carlo expectations are due to the 15% normalisation uncertainty. The quoted uncertainties in the cosmic muon Monte Carlo expectations are due to the statistical error combined in quadrature with a 10% normalisation uncertainty. A dash indicates no Monte Carlo events passed the cuts at this stage.

Figure 8.39: Distribution of track-like track planes for selected events. The solid blue line represents the total Monte Carlo expectation in the absence of oscillations.

Figure 8.40: Distribution of track ranges for selected events. The solid blue line represents the total Monte Carlo expectation in the absence of oscillations.

8.9 Charge Identification

The charge of the muon in each event is determined by the track fitter during the reconstruction (see Section 7.6). Events with a well measured charge can be selected from the high resolution sample by requiring that $|(Q/p)/\sigma_{(Q/p)}| > 2.0$. Events with $(Q/p)/\sigma_{(Q/p)} > 2.0$ are classified as $\bar{\nu}_\mu$ events, while those with $(Q/p)/\sigma_{(Q/p)} < -2.0$ are classified as ν_μ events. Figure 8.42 shows the distribution of $(Q/p)/\sigma_{(Q/p)}$ for high resolution events. The $(Q/p)/\sigma_{(Q/p)} < -2.0$ requirement selects 88.1% of high resolution ν_μ events from Monte Carlo as ν_μ , with a 1.4% $\bar{\nu}_\mu$ contamination. The $(Q/p)/\sigma_{(Q/p)} > 2.0$ requirement selects 89.1% of high resolution $\bar{\nu}_\mu$ events from Monte Carlo as $\bar{\nu}_\mu$, with a 6.8% ν_μ contamination. The higher level of contamination of the $\bar{\nu}_\mu$ sample reflects the fact that the ν_μ cross-section is higher than the $\bar{\nu}_\mu$ cross-section.

Table 8.5 shows the division of selected high resolution events by reconstructed charge, while Table 8.6 shows the selected events divided by reconstructed direction as well as reconstructed charge.

Figure 8.41: Distribution of $\text{RMS}_{V_{tx}} - \text{RMS}_{\text{End}}$ for selected events. The solid blue line represents the total Monte Carlo expectation in the absence of oscillations.

Figure 8.42: Distribution of $(Q/p)/\sigma_{(Q/p)}$ for high resolution selected events.

Classification	Data	Monte Carlo Expectation (no oscillations)				
		Cosmic μ^-/μ^+	$\nu_\mu/\bar{\nu}_\mu$ CC	$\nu_e/\bar{\nu}_e$ CC, NC	Upward μ^-/μ^+	Neutrons
ν_μ	112	3.3 ± 0.7	134.7 ± 20.2	0.8 ± 0.1	1.3 ± 0.2	—
$\bar{\nu}_\mu$	55	3.7 ± 0.7	68.5 ± 10.3	1.2 ± 0.2	0.7 ± 0.1	—
?	15	1.0 ± 1.0	19.2 ± 2.3	0.5 ± 0.1	0.4 ± 0.1	—

Table 8.5: Separation of selected events into ν_μ and $\bar{\nu}_\mu$ samples. MC values normalised to unoscillated expectations for data exposure 12.23 ktyr. The quoted uncertainties in the atmospheric neutrino and upward muon Monte Carlo expectations are due to the 15% normalisation uncertainty. The quoted uncertainties in the cosmic muon Monte Carlo expectations are due to the statistical error combined in quadrature with a 10% normalisation uncertainty. A dash indicates no Monte Carlo events passed the cuts at this stage.

Classification	Data	Monte Carlo Expectation (no oscillations)				
		Cosmic μ^-/μ^+	$\nu_\mu/\bar{\nu}_\mu$ CC	$\nu_e/\bar{\nu}_e$ CC, NC	Upward μ^-/μ^+	Neutrons
UP ν_μ	52	—	69.9 ± 10.5	0.4 ± 0.1	1.2 ± 0.2	—
UP $\bar{\nu}_\mu$	22	—	35.7 ± 5.4	0.7 ± 0.1	0.7 ± 0.1	—
UP ?	3	—	9.0 ± 1.4	0.3 ± 0.0	0.3 ± 0.0	—
DOWN ν_μ	60	3.3 ± 0.7	64.7 ± 9.7	0.3 ± 0.0	0.1 ± 0.0	—
DOWN $\bar{\nu}_\mu$	33	3.7 ± 0.7	32.8 ± 4.9	0.6 ± 0.1	0.05 ± 0.0	—
DOWN ?	12	1.0 ± 1.0	10.2 ± 1.5	0.3 ± 0.0	0.05 ± 0.0	—
UNCERTAIN	95	4.2 ± 0.7	92.2 ± 13.8	23.6 ± 3.5	0.4 ± 0.1	0.3 ± 0.2

Table 8.6: Separation of selected events by identified charge and direction. MC values normalised to unoscillated expectations for data exposure 12.23 ktyr. The quoted uncertainties in the atmospheric neutrino and upward muon Monte Carlo expectations are due to the 15% normalisation uncertainty. The quoted uncertainties in the cosmic muon Monte Carlo expectations are due to the statistical error combined in quadrature with a 10% normalisation uncertainty. A dash indicates no Monte Carlo events passed the cuts at this stage.

8.9.1 Splitting Tracks

In order to quantify the systematic error on muon charge identification, fitted tracks with at least 20 planes were divided into two parts, called split tracks. Each split track

was then refitted and the results stored along the original track information. The length of each part was varied to allow original tracks to contribute split tracks of different lengths to the sample. Original tracks which passed the high resolution and charge sign identification cuts were compared with any of their split tracks which also passed these cuts. This was done for two samples: through going muons and stopping muons using data and cosmic muon Monte Carlo. Figure 8.43 shows the fraction of split stopping muon tracks where the reconstructed charge of the split track agrees with that of the original track. Figure 8.44 shows the same quantity for a through going muon sample. For both samples the level of agreement improves as the number of planes in the track increases. Figure 8.45 and Figure 8.46 show the difference in the fraction of split tracks where the reconstructed charge of the split track agrees with that of the original track between the data and Monte Carlo for the stopping muon and through going muon samples respectively. Even for the shorter split tracks the difference in the level of agreement is always better than 5%. The stopping muon sample most closely reflects the energies of the muon tracks in contained vertex atmospheric neutrino events. Using this sample a 1% systematic uncertainty on charge identification can be assigned for tracks crossing at least 20 planes and a 2% systematic uncertainty can be assigned for tracks crossing fewer than 20 planes.

8.10 Summary

A total of 277 candidate contained vertex $\nu/\bar{\nu}_\mu$ CC data events are selected from a 12.23 ktyr exposure, compared to an expectation of 354.4 ± 47.4 events from Monte Carlo in the absence of neutrino oscillations. Table 8.3 summarises how many events pass each branch of the selection. Applying further cuts to give a high resolution sub-sample yields a total of 182 data events which have clearly identified directions. In data 77 events are identified as upward going and 105 data events are identified as downward going, compared to expectations of 118.2 ± 17.7 events and 116.3 ± 16.2 events respectively from Monte Carlo in the absence of neutrino oscillations. Table 8.4 summarises the separation of selected events into high and low resolution samples. A total of 167 data events have clearly identified charge, 112 are identified as ν_μ events, 55 are identified as $\bar{\nu}_\mu$ events, compared to expectations of 139.8 ± 20.2 events and 74.1 ± 10.3 events respectively from Monte Carlo in the absence of neutrino oscillations. Table 8.5 summarises the separation of selected events into ν_μ and $\bar{\nu}_\mu$ samples.

Figure 8.43: Fraction of split stopping muon tracks where the reconstructed charge agrees with that of the whole track versus number of planes cross by the split track. Data is shown by the blue points. Cosmic muon Monte Carlo is shown by the red points.

Figure 8.44: Fraction of split through going muon tracks where the reconstructed charge agrees with that of the whole track versus number of planes cross by the split track. Data is shown by the blue points. Cosmic muon Monte Carlo is shown by the red points.

Figure 8.45: Difference between fraction of stopping muon split tracks with matching charges in data and Monte Carlo versus number of planes cross by the split track.

Figure 8.46: Difference between fraction of through going muon split tracks with matching charges in data and Monte Carlo versus number of planes cross by the split track.

Chapter 9

Searching For Atmospheric Neutrino Oscillations

“One of the penalties for refusing to participate in politics is that you end up being governed by your inferiors.”

— Plato, 427–347 BC

9.1 Overview

Chapter 8 described the selection of CC atmospheric ν_μ and $\bar{\nu}_\mu$ interactions in the MINOS Far Detector. From this sample of events a sub-sample of events with well determined directions and a sub-sample of events with well determined charges were selected. This chapter details how these selected events can be used in a fit to the hypothesis of $\nu_\mu \rightarrow \nu_\tau$ oscillations.

Each event has a measurement of the muon momentum (including the initial direction), $\mathbf{p}_\mu$, and a measurement of the energy of the associated shower, E_{shw} . A measurement of the neutrino energy is given by $E_{\nu_{reco}} = E_{\mu_{reco}} + E_{shw}$. The distance the neutrino travelled before interacting, L_{reco} , can be estimated from the direction the neutrino was travelling when it interacted with the detector. The best measurement of the neutrino direction is the direction of the muon at the event vertex. From this $\left(\frac{L}{E}\right)_{reco}$ can be calculated. It should be noted that for low energy and/or high y ($= 1 - E_\mu/E_\nu$) events

the opening angle between the observed muon direction and the true neutrino direction will be large. This means that particularly when the muon direction is close to the horizon, $(\frac{L}{E})_{reco}$ may differ by a large amount from the true value, $(\frac{L}{E})_{true}$.

As the sample of selected events, although larger than in previous analyses, is still small, it is important to extract as much information as possible from each event. The Bayesian methods described in this chapter attempt to estimate the accuracy of $(\frac{L}{E})_{reco}$ as an estimator of $(\frac{L}{E})_{true}$ for each event by calculating a probability distribution for $\log(\frac{L}{E})_{true}$ given the observed $\mathbf{p}_\mu$ and E_{shw} values. The RMS of this PDF about its mean, $\sigma_{\log(L/E)}$, can be used as an estimate of the accuracy of the (L/E) reconstruction. Events can then be binned according to $\sigma_{\log(L/E)}$ in the oscillation fit. Another version of this Bayesian method has been used before with a smaller dataset [29]. The larger dataset used in this thesis makes it possible to apply this technique to ν_μ and $\bar{\nu}_\mu$ events separately for the first time.

9.2 Double Ratios

9.2.1 Determination of Up/Down Ratio

Selected events with well measured directions, as described in Section 8.8 and summarised in Table 8.4, are used to calculate the atmospheric neutrino up/down ratio. This ratio provides a measure of the consistency of the data with neutrino oscillations. The ratio determined from data is $R_{u/d}^{\text{data}} = 0.73_{-0.11}^{+0.13}(\text{stat.})$. The quoted statistical errors on the ratio are calculated using Poisson statistics using the methods described in [163]. All systematic errors are quoted as part of the Monte Carlo ratio. The Monte Carlo expectation in the absence of oscillations is calculated to be $R_{u/d}^{\text{MC}} = 1.02 \pm 0.05(\text{sys.})$. The 5% systematic error on this ratio is dominated by the uncertainties in the atmospheric neutrino flux [146]. As mentioned previously the Monte Carlo sample is sufficiently large that its contribution to the statistical error may be safely neglected. The expected value differs from 1.0 due to the differing selection efficiencies as a function of $\cos \Theta_{\text{zenith}}$ and the presence of background events. The ratio between the observed and expected up/down ratios is calculated to be:

$$\frac{R_{u/d}^{\text{data}}}{R_{u/d}^{\text{MC}}} = 0.72_{-0.11}^{+0.13}(\text{stat.}) \pm 0.04(\text{sys.}). \quad (9.1)$$

Figure 9.1: The up/down ratio as a function of Δm_{32}^2 . The red line assumes $\sin^2 2\theta_{23} = 1.0$, while the black line assumes $\sin^2 2\theta_{23} = 0.8$. The central dashed line indicates the observed ratio in data. The shaded region shows indicates the $\pm 1\sigma$ statistical uncertainty calculated using Poisson statistics using the methods described in [163].

This is 2.1σ away from the expectation for no oscillations.

Figure 9.1 shows a comparison of $R_{u/d}^{\text{MC}}$ with the measured $R_{u/d}^{\text{data}}$ value with 1σ error bars for a range of oscillation parameter values. The red line shows the variation of $R_{u/d}^{\text{MC}}$ with Δm_{32}^2 when $\sin^2 2\theta_{23} = 1.0$, while the black line shows the variation with Δm_{32}^2 when $\sin^2 2\theta_{23} = 0.8$. $R_{u/d}^{\text{MC}}$ initially falls with increasing Δm_{32}^2 , as the oscillation probability for upward-going events increases, while the oscillation probability for downward-going events remains small. As Δm_{32}^2 increases further, the oscillation probability for downward-going events increases, while the oscillation probability for upward-going events varies rapidly, averaging out to 0.5. This leads to the observed increase of $R_{u/d}^{\text{MC}}$ with larger Δm_{32}^2 values. The observed ratio is consistent with $R_{u/d}^{\text{MC}}$ for a wide range of oscillation parameters.

Figure 9.2 shows the E_ν distribution for high resolution selected data events, together with the expected distributions from Monte Carlo in the absence of oscillations and ν_μ and $\bar{\nu}_\mu$ oscillations with $\Delta m_{32}^2 = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$. This is the best fit point from the MINOS measurement of the NuMI beam [74]. The nominal Monte Carlo normalisation is calculated using the exposures in Table 6.1 and the solar cycle reweighting discussed in Section 6.3.1. Figure 9.3 shows the $\cos \Theta_{\text{zenith}}$ distribution for high resolution selected data events, together with the expected distributions from Monte Carlo in the absence of oscillations and ν_μ and $\bar{\nu}_\mu$ oscillations with $\Delta m_{32}^2 = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$. A χ^2 fit of each Monte Carlo distribution to the data $\cos \Theta_{\text{zenith}}$

Figure 9.2: Distribution of E_ν for high resolution events. The solid blue line shows the total Monte Carlo expectation in the absence of oscillations with nominal normalisation. The dashed blue line indicates the total Monte Carlo expectation for ν_μ and $\bar{\nu}_\mu$ oscillations with $\Delta m_{32}^2 = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$ and nominal normalisation. The black points indicate the numbers of data events observed.

Figure 9.3: Distribution of $\cos \Theta_{zenith}$ for high resolution events. The solid blue line shows the total Monte Carlo expectation in the absence of oscillations with nominal normalisation. The dashed blue line indicates the total Monte Carlo expectation for ν_μ and $\bar{\nu}_\mu$ oscillations with $\Delta m_{32}^2 = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$ and nominal normalisation. The black points indicate the numbers of data events observed.

Figure 9.4: The $\bar{\nu}_\mu/\nu_\mu$ ratio as a function of $\overline{\Delta m^2_{32}}$. The red line assumes $\Delta m^2_{32} = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = \sin^2 2\overline{\theta}_{23} = 1.0$, while the black line assumes $\Delta m^2_{32} = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = \sin^2 2\overline{\theta}_{23} = 0.8$. The central dashed line indicates the observed ratio in data. The shaded region shows indicates the quadrature sum of the $\pm 1\sigma$ statistical and systematic uncertainties.

distribution is performed with the normalisation as a nuisance parameter with systematic uncertainty 0.15. In the absence of oscillations $\chi^2 = 9.93$ for six degrees of freedom, with a normalisation of 79.3% of nominal, with oscillations as described above $\chi^2 = 6.64$ for six degrees of freedom, with a normalisation of 105.4% of nominal. The data favour oscillations.

9.2.2 Determination of $\nu_\mu/\bar{\nu}_\mu$ Ratio

The ratio of $\bar{\nu}_\mu$ to ν_μ events provides a means of testing the consistency between neutrino and anti-neutrino oscillations. This ratio is determined using the selected events with well measured charges as defined previously in Section 8.9 and summarised in Table 8.5. The ratio determined from data is $R_{\bar{\nu}/\nu}^{\text{data}} = 0.49_{-0.08}^{+0.10}(\text{stat.}) \pm 0.01(\text{sys.})$. The quoted statistical errors on the ratio are calculated using Poisson statistics using the methods described in [163]. The systematic uncertainty is the uncertainty associated with charge sign identification and was calculated from the ‘split track’ analysis described in Section 8.9.1. The expected ratio from Monte Carlo is $R_{\bar{\nu}/\nu}^{\text{MC}} = 0.53 \pm 0.07(\text{sys.})$. The systematic uncertainty on this ratio is dominated by the relative uncertainty in the ν_μ to $\bar{\nu}_\mu$ cross-sections of 13.5% [29], but contains an additional contribution from the uncertainty

Figure 9.5: Distribution of $\cos \Theta_{zenith}$ for high resolution events identified as ν_μ events. The solid blue line shows the total Monte Carlo expectation in the absence of oscillations with nominal normalisation. The dashed blue line indicates the total Monte Carlo expectation for ν_μ and $\bar{\nu}_\mu$ oscillations with $\Delta m_{32}^2 = \overline{\Delta m_{32}^2} = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = \sin^2 2\overline{\theta}_{23} = 1.0$ and nominal normalisation. The black points indicate the numbers of data events observed.

in the relative $\bar{\nu}_\mu$ to ν_μ fluxes (4%). The ratio of ratios is calculated to be:

$$\frac{R_{\bar{\nu}/\nu}^{\text{data}}}{R_{\bar{\nu}/\nu}^{\text{MC}}} = 0.93_{-0.15}^{+0.19}(\text{stat.}) \pm 0.12(\text{sys.}). \quad (9.2)$$

This is consistent with neutrinos and anti-neutrinos having the same oscillation parameters. This ratio is almost independent of the values of the oscillation parameters provided they are the same for neutrinos and anti-neutrinos.

To allow for the possibility of independent ν_μ and $\bar{\nu}_\mu$ oscillations from here onward $(\Delta m_{32}^2, \sin^2 2\theta_{23})$ are used to refer to the oscillation parameters controlling two flavour $\nu_\mu \rightarrow \nu_\tau$ oscillations and $(\overline{\Delta m_{32}^2}, \sin^2 2\overline{\theta}_{23})$ are used to refer to the oscillation parameters controlling two flavour $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ oscillations. Figure 9.4 shows a comparison of $R_{\bar{\nu}/\nu}^{\text{MC}}$ with the measured $R_{\bar{\nu}/\nu}^{\text{data}}$ value with 1σ error bars for a range of oscillation parameter values. The red line shows the variation of $R_{\bar{\nu}/\nu}^{\text{MC}}$ with $\overline{\Delta m_{32}^2}$ assuming $\Delta m_{32}^2 = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = \sin^2 2\overline{\theta}_{23} = 1.0$, while the black line shows the variation of $R_{\bar{\nu}/\nu}^{\text{MC}}$ with $\overline{\Delta m_{32}^2}$ assuming $\Delta m_{32}^2 = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = \sin^2 2\overline{\theta}_{23} = 0.8$. In both cases, the observed ratio is consistent with the expectations over a wide range of $\overline{\Delta m_{32}^2}$ values.

Figure 9.6: Distribution of $\cos \Theta_{zenith}$ for high resolution events identified as $\bar{\nu}_\mu$ events. The solid blue line shows the total Monte Carlo expectation in the absence of oscillations with nominal normalisation. The dashed blue line indicates the total Monte Carlo expectation for ν_μ and $\bar{\nu}_\mu$ oscillations with $\Delta m_{32}^2 = \overline{\Delta m_{32}^2} = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = \sin^2 2\overline{\theta_{23}} = 1.0$ and nominal normalisation. The black points indicate the numbers of data events observed.

Figure 9.5 shows the $\cos \Theta_{zenith}$ distribution for high resolution selected data events identified as ν_μ events, together with the expected distributions from Monte Carlo in the absence of oscillations and ν_μ and $\bar{\nu}_\mu$ oscillations with $\Delta m_{32}^2 = \overline{\Delta m_{32}^2} = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = \sin^2 2\overline{\theta_{23}} = 1.0$. A χ^2 fit of each Monte Carlo distribution to the data $\cos \Theta_{zenith}$ distribution is performed with the normalisation as a nuisance parameter with systematic uncertainty 0.15. In the absence of oscillations $\chi^2 = 3.86$ for six degrees of freedom, with a normalisation of 83.7% of nominal, with oscillations as described above $\chi^2 = 4.09$ for six degrees of freedom, with a normalisation of 107.7% of nominal. The fits to both hypothesis are similar, although the normalisation has had to be shifted by over 1σ to reach the unoscillated best fit.

Figure 9.6 shows the $\cos \Theta_{zenith}$ distribution for high resolution selected data events identified as $\bar{\nu}_\mu$ events, together with the expected distributions from Monte Carlo in the absence of oscillations and ν_μ and $\bar{\nu}_\mu$ oscillations with $\Delta m_{32}^2 = \overline{\Delta m_{32}^2} = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = \sin^2 2\overline{\theta_{23}} = 1.0$. A χ^2 fit of each Monte Carlo distribution to the data $\cos \Theta_{zenith}$ distribution is performed with the normalisation as a nuisance parameter with systematic uncertainty 0.15. In the absence of oscillations $\chi^2 = 11.21$ for six degrees of freedom, with a normalisation of 91.1% of nominal, with oscillations as described above

$\chi^2 = 8.56$ for six degrees of freedom, with a normalisation of 100.3% of nominal. Here the data favour oscillations.

9.3 L/E Distribution

In the approximation of two-flavour $\nu_\mu \rightarrow \nu_\tau$ oscillations the survival probability is as follows:

$$P(\nu_\mu \rightarrow \nu_\mu) = 1.0 - \sin^2 2\theta_{23} \sin^2 \left(\frac{1.27 \Delta m_{32}^2 L}{E} \right), \quad (9.3)$$

where E = neutrino energy in GeV and L = neutrino propagation distance in km. The neutrino propagation distance L is calculated using the following equation:

$$L = -R \cos \Theta_{zenith} + \sqrt{R^2 \cos^2 \Theta_{zenith} + h \times (2R + h)}, \quad (9.4)$$

where R is the radius of the Earth (here assumed to be a sphere of radius 6370 km), h is the production height of neutrinos in the atmosphere (here assumed to be a constant 15 km) and Θ_{zenith} is the zenith angle. For plotting reconstructed events the muon zenith angle is used. For calculating the oscillation probability and in the Bayesian algorithm below the ‘true’ neutrino zenith angle is used.

The reconstructed $\log(L/E)$ distribution for high resolution events is shown in Figure 9.7. Rather than fitting this distribution a more sophisticated technique is used.

9.4 Bayesian Analysis

9.4.1 Bayes’s Theorem

Before his death in 1761, the Reverend Thomas Bayes studied how to compute a distribution for the parameter of a binomial distribution. This work was never published in his lifetime. A friend of Reverend Bayes, Richard Price, found the work after Bayes’s death. Price edited and presented the work to the Royal Society of London in 1763. It was published in the same year in the *Philosophical Transactions of the Royal Society of London* under the title “*Essay Towards Solving a Problem in the Doctrine of*

Figure 9.7: Reconstructed $\log(L/E)$ distribution for high resolution events. The solid blue line shows the total Monte Carlo expectation in the absence of oscillations. The dashed blue line indicates the total Monte Carlo expectation for ν_μ and $\bar{\nu}_\mu$ oscillations with $\Delta m_{32}^2 = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$. The black points indicate the numbers of data events observed.

Chances” [164]. Proposition 5 of this essay gives a simple description of conditional probability, and shows that it can be expressed independently of the order in which things occur:

“If there be two subsequent events, the probability of the second b/N and the probability of both together P/N , and it being first discovered that the second event has also happened, from hence I guess that the first event has also happened, the probability I am right is P/b .”

It does not appear that Bayes himself focused on this result, but his preliminary results imply the result now called Bayes’s Theorem. These results were independently replicated and extended by Pierre-Simon Laplace in an essay published in 1774 [165].

Bayes’s theorem relates the posterior and prior probabilities of stochastic events A and B .

$$P(A|B) = \frac{P(B|A) P(A)}{P(B)} \propto L(A|B) P(A) \quad (9.5)$$

where $P(A)$ is the prior or marginal probability that A occurs, $P(B|A)$ is the posterior or conditional probability of B given that A has occurred, $L(A|B)$ is the likelihood of A

given that B has occurred.

9.4.2 Applying Bayes's Theorem To Atmospheric Neutrinos

Consider an atmospheric ν_μ CC interaction recorded in the MINOS Far Detector with reconstructed muon momentum $\mathbf{p}_\mu$ and reconstructed shower energy E_{shw} , but which was in truth a neutrino with energy E_ν which had travelled a distance L_ν between its creation in the atmosphere and interaction in the detector. Using Bayes's Theorem the following statement can be made:

$$P(L_\nu; E_\nu | \mathbf{p}_\mu; E_{shw}) P(\mathbf{p}_\mu; E_{shw}) = P(\mathbf{p}_\mu; E_{shw} | L_\nu; E_\nu) P(L_\nu; E_\nu) \quad (9.6)$$

The prior probability $P(\mathbf{p}_\mu; E_{shw})$ is a non-zero constant as the event has been observed. The prior probability $P(L_\nu; E_\nu)$ can be determined from a prior probability density function (prior PDF) $f(L_\nu; E_\nu)$ determined using atmospheric neutrino models such as [64].

The true values of E_μ , E_{shw} and the opening angle between the neutrino and muon, θ , depend on the kinematic variables y and W^2 as well as E_ν , where y is the fraction of visible energy in the hadron and W^2 is the invariant mass of the hadronic system. For a given combination of E_ν , y and W^2 the kinematics of the muon are fully determined. Equation (9.7) shows how the posterior PDF for reconstructing different $\mathbf{p}_\mu$ and E_{shw} values for a given L_ν , E_ν combination, $f(\mathbf{p}_\mu; E_{shw} | L_\nu; E_\nu)$, can be calculated.

$$f(\mathbf{p}_\mu; E_{shw} | L_\nu; E_\nu) = \int \int \int f(\mathbf{p}_\mu; E_{shw} | L_\nu; E_\nu; y; W^2; \Omega_\nu) f(y; W^2; \Omega_\nu | L_\nu; E_\nu) dy dW^2 d\Omega_\nu \quad (9.7)$$

Ω_ν is the incident neutrino direction with respect to that of the muon. Here the posterior PDF $f(\mathbf{p}_\mu; E_{shw} | L_\nu; E_\nu; y; W^2; \Omega_\nu)$ represents the resolution with which quantities can be measured in the Far Detector. The second posterior PDF in equation (9.7), $f(y; W^2; \Omega_\nu | L_\nu; E_\nu)$, can be constructed by generating large numbers of Monte Carlo events.

$$f(L_\nu; E_\nu | \mathbf{p}_\mu; E_{shw}) = \frac{f(\mathbf{p}_\mu; E_{shw} | L_\nu; E_\nu) f(L_\nu; E_\nu)}{f(\mathbf{p}_\mu; E_{shw})} \quad (9.8)$$

The PDF calculated in equation (9.8) can be used to estimate the accuracy of the measurement of $\left(\frac{L}{E}\right)_{true}$ obtained using $\left(\frac{L}{E}\right)_{reco}$. It should be noted that the differ-

ent ν -nucleus interaction types; quasi-elastic (QE), resonance (RES) and deep-inelastic scattering (DIS); give different kinematic distributions and detector responses. For this reason separate contributions to the $(\frac{L}{E})$ PDF are calculated for each interaction type. The dependence of cross-section on neutrino interaction type is discussed in Section 6.3.2.

9.4.3 Implementation of The Bayesian Method

The first step is to construct the posterior PDFs. These PDFs can be divided into two groups: those based on purely MC generator quantities and those based on reconstructed quantities. For the former group, *Neugen3* [148] was used to generate two billion ν_μ interactions and two billion $\bar{\nu}_\mu$ interactions. Each sample was generated in two parts with flat energy spectra between 0.2 GeV and 10 GeV and between 10 GeV and 100 GeV. Due to the huge sample size only the minimum information required was recorded for each interaction.

In making the PDFs events were divided up by $\nu_\mu/\bar{\nu}_\mu$ and interaction type. PDFs were binned in E_ν in 0.1 GeV bins between 0.2 GeV and 10 GeV and 0.5 GeV bins between 10 GeV and 100 GeV. For QE interactions fixing E_ν and y also fixes W^2 . Therefore only a 1 dimensional PDF in y , $f(y|E_\nu; QE)$, is required for each E_ν bin. For RES and DIS interactions 2 dimensional PDFs of W^2 versus y , $f(y; W^2|E_\nu; RES)$ and $f(y; W^2|E_\nu; DIS)$, are created for each E_ν bin. Figure 9.8 shows example PDFs for the $1.5 \text{ GeV} < E_\nu < 1.6 \text{ GeV}$ bin for $\bar{\nu}_\mu$ events. In addition for each E_ν bin a PDF, $f(interaction\ type|E_\nu)$, giving the relative posterior probabilities of each interaction type is created. In this way information on the relative cross-section of each interaction type is included.

The PDFs using reconstructed quantities, which represent the resolution of the detector, were constructed using true CC ν_μ and $\bar{\nu}_\mu$ events from the atmospheric neutrino MC sample which passed the contained vertex analysis detailed in Chapter 8. PDFs were created for the following quantities μ momentum from range ($p_{\mu_{range}}$), μ momentum from curvature ($p_{\mu_{curve}}$) and shower energy (E_{shw}). The PDFs for shower energy and $p_{\mu_{range}}$ were constructed in the following way. PDFs were binned according to the true muon or shower energy. For each bin of true energy a PDF of $E_{reco} - E_{true}$ was created from reconstructed events. Momentum from curvature was treated slightly differently. As discussed in Chapter 7, the track fitter actually fits q/p . This can lead to larger discrepancies in reconstructed momentum values than happens with the range measurement. This meant that insufficient Monte Carlo statistics were available to prop-

(a) y distribution for $\bar{\nu}_\mu$ quasi-elastic events with $1.5 \text{ GeV} < E_\nu < 1.6 \text{ GeV}$

(b) y vs w^2 distribution for $\bar{\nu}_\mu$ resonance events with $1.5 \text{ GeV} < E_\nu < 1.6 \text{ GeV}$

Figure 9.8

erly cover the range of variations with well populated PDFs. For this reason, instead of using PDFs constructed directly from reconstructed events, the events were used to create a parametrisation of the error in the reconstructed momentum from curvature values. The parametrisation chosen was $\sigma_{p_\mu} = 1.3226 \times p_{\mu_{\text{curve}}}^2 \times \sigma_{(q/p)}$, where $p_{\mu_{\text{curve}}}$ is the reconstructed momentum from curvature and $\sigma_{(q/p)}$ is the estimated error on $\left(\frac{q}{p}\right)$ at the vertex of the track, as returned by the track fitter (see Section 7.6). This was then used as the standard deviation for a Gaussian PDF of reconstructed - true momentum. The momentum from curvature PDFs are only used for partially contained events. This parameterisation gives a good approximation to a Gaussian error as can be seen from the pull distribution shown in Figure 9.9.

For each data event the following algorithm was used to generate the $\log\left(\frac{L}{E}\right)$ PDF. The algorithm loops over possible neutrino-muon opening angles, θ , in steps of 1° . For a given value of θ , the first stage is to calculate an E_ν PDF. To do this the algorithm loops over E_ν values in 0.1 GeV steps.

For each E_ν value, the algorithm checks a range of E_μ values. For fully contained events possible true E_μ values within $\pm 2.5 \text{ GeV}$ of the reconstructed E_μ value in steps of 25 MeV are checked. In order to reflect the larger uncertainties in momentum from curvature, for partially contained events a range of values within $\pm 2.5\sigma_{p_\mu}$ are checked in steps of $0.0025\sigma_{p_\mu}$. This ensures the possible values are well covered in both cases. Each possible ‘true’ θ , E_ν , E_μ combination fixes the values of the kinematic variables,

Figure 9.9: $(p_{reco} - p_{true})/\sigma_{p_\mu}$ for momentum from curvature of partially contained ν_μ and $\bar{\nu}_\mu$ CC atmospheric neutrino MC events.

as follows:

$$y = 1 - \frac{E_\mu}{E_\nu},$$

$$W^2 = 2E_\nu \times (y \times m_p + E_\mu \cos \theta) - (1 - y^2) \times E_\nu^2 - E_\mu^2 + m_p^2,$$

where m_p is the mass of the proton. The contribution to the E_ν PDF from this set of parameters is now calculated, leading to the following contributions from each simulated interaction type. Each DIS possibility contributes:

$$\begin{aligned} f_i^{\text{DIS}}(E_\nu|\theta) &= P(\text{DIS}|E_\nu) \times P(p_{\mu(\text{true-reco})}|E_{\mu\text{true}}) \times P(y; W^2|\text{DIS}; E_\nu) \\ &\quad \times P(E_{\text{had}(\text{true-reco})}|\text{DIS}; E_{\text{hadtrue}}) \end{aligned}$$

Each resonance possibility contributes:

$$\begin{aligned} f_i^{\text{RES}}(E_\nu|\theta) &= P(\text{RES}|E_\nu) \times P(p_{\mu(\text{true-reco})}|E_{\mu\text{true}}) \times P(y; W^2|\text{RES}; E_\nu) \\ &\quad \times P(E_{\text{had}(\text{true-reco})}|\text{RES}; E_{\text{hadtrue}}) \end{aligned}$$

In the case of quasi-elastic events kinematic constraints on W^2 mean that most E_ν/E_μ combinations do not contribute for a given θ . For those that closely match the con-

straints, the following contribution is added to the PDF:

$$f_i^{\text{QE}}(E_\nu|\theta) = P(\text{QE}|E_\nu) \times P(p_{\mu(\text{true-reco})}|E_{\mu\text{true}}) \times P(y|\text{QE}; E_\nu) \\ \times P(E_{\text{had}(\text{true-reco})}|\text{QE}; E_{\text{hadtrue}})$$

The complete PDF, $f(E_\nu|\theta)$, is given by the sum of these contributions:

$$f(E_\nu|\theta) = \sum_i f_i^{\text{DIS}}(E_\nu|\theta) + \sum_i f_i^{\text{RES}}(E_\nu|\theta) + \sum_i f_i^{\text{QE}}(E_\nu|\theta)$$

After the loop over possible true E_ν values, an E_ν PDF for the current θ value is used in the next step of the process. The algorithm now loops over possible neutrino-muon azimuthal angles, ϕ , for the current θ value. From each θ - ϕ pairing a possible true neutrino zenith angle and hence a value of L can be calculated using the calculation described in equation (9.4) above. From this L value and the current E_ν PDF, contributions to the $\log(\frac{L}{E})$ PDF can be calculated. The algorithm now moves to the next θ value and calculates a new E_ν PDF and so on until the loop over θ values is complete. The final step is to calculate $\sigma_{\log(L/E)}$ from the $\log(\frac{L}{E})$ PDF. This is the posterior probability density function for true $\log(\frac{L}{E})$ values given reconstructed $\mathbf{p}_\mu$ and E_{shw} values.

Figure 9.10 shows an example PDF for a downward-going muon track. The true value (blue line) is close to the reconstructed value (red line) and the $\sigma_{\log(L/E)}$ value of the PDF is correspondingly small (green line). The same is also true for the case of a steeply upward-going muon track, as shown in Figure 9.11. However, for muon tracks near the horizon, as shown in Figure 9.12, the PDF reflects the large range of possible neutrino directions which could produce a horizontal muon. The large $\sigma_{\log(L/E)}$ value indicates the lower $\log(\frac{L}{E})$ resolution of such events.

Figure 9.13 shows the distribution of $\sigma_{\log(L/E)}$ for high resolution data and Monte Carlo events. The high resolution events are divided up into five bins $\sigma_{\log(L/E)}$ as follows: $\sigma_{\log(L/E)} > 0.60$, $0.60 \geq \sigma_{\log(L/E)} > 0.40$, $0.40 \geq \sigma_{\log(L/E)} > 0.30$, $0.30 \geq \sigma_{\log(L/E)} > 0.25$, $\sigma_{\log(L/E)} \leq 0.25$. Figure 9.14 shows the $\log(\frac{L}{E})_{\text{reco}}$ distributions for high resolution Monte Carlo events binned in $\sigma_{\log(L/E)}$ for no oscillations and for oscillations with $\Delta m_{32}^2 = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = \sin^2 2\overline{\theta}_{23} = 1.0$. As the value of $\sigma_{\log(L/E)}$ decreases oscillations become more concentrated on the right of the plots. Figure 9.15 shows the ratio of the oscillated to unoscillated distributions. The resolution of the ‘oscillation dip’ improves as $\sigma_{(L/E)}$ decreases. Dividing up the sample in this way will give improved sensitivity to oscillations.

Figure 9.10: Example $\log\left(\frac{L}{E}\right)$ PDF: Downward-going muon. The PDF is shown by the black line. The blue line indicates the true value, the red line indicates the reconstructed value, the green line shows $\pm 1\sigma_{\log(L/E)}$.

Figure 9.11: Example $\log\left(\frac{L}{E}\right)$ PDF: Upward-going muon. The PDF is shown by the black line. The blue line indicates the true value, the red line indicates the reconstructed value, the green line shows $\pm 1\sigma_{\log(L/E)}$.

Figure 9.12: Example $\log\left(\frac{L}{E}\right)$ PDF: muon near horizon. The PDF is shown by the black line. The blue line indicates the true value, the red line indicates the reconstructed value, the green line shows $\pm 1\sigma_{\log(L/E)}$.

Figure 9.13: Distribution of $\sigma_{\log(L/E)}$ for high resolution data and Monte Carlo events. The solid blue line indicates the Monte Carlo expectation in the absence of oscillations. The data are shown by the black points.

Figure 9.14: $\log\left(\frac{L}{E}\right)_{reco}$ distributions for high resolution Monte Carlo events binned in $\sigma_{\log(L/E)}$. The solid blue line indicates the Monte Carlo expectation in the absence of oscillations. The dashed red line indicates the Monte Carlo expectation for ν_μ and $\bar{\nu}_\mu$ oscillations with $\Delta m_{32}^2 = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$.

Figure 9.15: Oscillated/Unoscillated Monte Carlo events vs $\log \left(\frac{L}{E} \right)_{reco}$ binned in $\sigma_{\log(L/E)}$. The solid blue line indicates the ratio of the Monte Carlo expectation for ν_μ and $\bar{\nu}_\mu$ oscillations, with $\Delta m_{32}^2 = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$, to the Monte Carlo expectation in the absence of oscillations.

9.5 Fits to Neutrino Oscillations

9.5.1 Likelihood Function

A maximum likelihood fit to the data is performed using the following negative log-likelihood function:

$$\begin{aligned}
 -\ln \mathcal{L} = & N_{MC} - N \ln N_{MC} \\
 & - \sum_{k=1}^5 (N_k^{up} \ln p_k^{up} + N_k^{dn} \ln p_k^{dn}) \\
 & - \sum_{data} \ln (f(L/E)_{ki}^{up}) \\
 & - \sum_{data} \ln (f(L/E)_{ki}^{dn}) \\
 & + \sum_j \frac{\alpha_j^2}{2\sigma_{\alpha_j}^2}
 \end{aligned} \tag{9.9}$$

In the above function the subscript k refers to the k th bin in $\sigma_{\log(L/E)}$. This function can be divided up into five terms:

Normalisation Term The first two terms in the likelihood function ($N_{MC} - N \ln N_{MC}$) represent the Poisson probability of observing a total of N events in data given an expectation of N_{MC} from Monte Carlo. Both high and low resolution events contribute to this term.

“Up-Down” Term The next term represents the relative probabilities of observing upward and downward-going events. Only high resolution events where the direction can be determined contribute to this term. N_k^{up} is the number of upward-going data events in the k th resolution bin and N_k^{dn} is the number of downward-going data events in the k th resolution bin, p_k^{up} and p_k^{dn} are the probabilities that a Monte Carlo event in the k th resolution bin is upward-going or downward-going respectively ($p_k^{up} + p_k^{dn} = 1$).

Shape Terms The next two terms represent shape fits to the $\log(L/E)$ distributions for upward and downward-going events respectively. Only high resolution events where the direction can be determined contribute to these terms. $f(L/E)_{ki}^{up}$ is the value of the i th bin of the L/E distribution for upward-going Monte Carlo events in the k th resolution bin, similarly $f(L/E)_{ki}^{dn}$ is the value of the i th bin of the

L/E distribution for downward-going Monte Carlo events in the k th resolution bin. The sum is over the data events in the k th resolution bin. The distributions are normalised such that $\sum_i f(L/E)_{ki}^{up} = 1$ and $\sum_i f(L/E)_{ki}^{dn} = 1$.

Nuisance Term The values of the Monte Carlo derived terms above vary according to the oscillation parameters (Δm_{32}^2 and $\sin^2 2\theta_{23}$) and the values of various systematic parameters. Deviations of the values of these systematic parameters from the nominal are taken into account in the final “nuisance” term of the likelihood function, where α_j are the particular values of the systematic parameters and σ_{α_j} are the estimated systematic uncertainties.

The form of the likelihood function in equation (9.9) is the same as that used in [29] and was chosen to simplify the inclusion of the systematic errors in the fit. The MINUIT [166] function minimisation package was used to find the ‘best’ values of the nuisance parameters (*i.e.* those which gave the lowest negative log-likelihood value) for each particular (Δm_{32}^2 , $\sin^2 2\theta_{23}$) point.

9.5.2 Systematic Uncertainties

Six systematic effects are included as nuisance parameters in the oscillation fit. The overall uncertainty on the expected neutrino event rate is taken to be 15% [167]. This systematic uncertainty affects the normalisation term only, the atmospheric neutrino Monte Carlo is scaled according to $(1.0 + \alpha_{\text{norm}})$ in the fit, where α_{norm} is the normalisation nuisance parameter. The second nuisance parameter, $\alpha_{u/d}$ is the relative uncertainty on the efficiency for selecting upward versus downward-going events. The associated systematic uncertainty on this parameter is taken to be 3%. This nuisance parameter only affects the “Up-Down” term in the likelihood function, the number of upward-going Monte Carlo events in bin k is scaled by $(1.0 + \alpha_{u/d})$ when calculating p_k^{up} and p_k^{dn} .

A 3% systematic uncertainty is assigned to the muon energy scale and a 5% systematic uncertainty is assigned to the hadronic energy scale. The associated nuisance parameters, α_μ and α_{had} , only affect the shape of the $\log(L/E)$ distributions and so their effects are confined to Shape Likelihoods. When calculating which $\log(L/E)$ bin a Monte Carlo event should contribute to the reconstructed neutrino energy is altered as follows $E_\nu = (1.0 + \alpha_\mu)E_\mu + (1.0 + \alpha_{\text{had}})E_{\text{shw}}$. This new neutrino energy is then used to calculate $\log(L/E)$.

In order to accommodate the uncertainties in the shape of the neutrino energy spectrum, the spectrum is allowed to scale according to $1.0 + 0.1\alpha_{\text{spec}}(E_\nu - 2.0)$ for $E_\nu < 2.0$ and $1.0 + 0.1\alpha_{\text{spec}}(\ln E_\nu - \ln 2.0)$ for $E_\nu \geq 2.0$, where the nuisance parameter α_{spec} is normally distributed (*i.e.* $\sigma_{\alpha_{\text{spec}}} = 1.0$). Studies performed for previous analyses [29] show that these variations cover the differences in the neutrino spectra obtained from different flux models [64] [168] [169]. Finally, to allow for uncertainties in the modelling of neutrino cross sections, an uncertainty of 20% in the relative cross section for quasi-elastic interactions is included. This associated nuisance parameter $\alpha_Q E$ is used to scaled QE events by $(1.0 + \alpha_Q E)$ when calculating p_k^{up} , p_k^{dn} , $f(L/E)_k^{up}$ and $f(L/E)_k^{dn}$.

9.5.3 Fit Mechanics

Using the minimisation package MINUIT has many advantages and greatly improves the run time of the oscillation fits, but it can have problems with variables which do not alter the likelihood in a continuous fashion. For example in this fit, the likelihood value would remain constant as α_μ or α_{had} is varied until such time as either was large enough to cause an event to move from one bin to an adjacent bin. In the limit of infinite Monte Carlo statistics, the changes would be continuous, but despite the high statistics available, as MINUIT probes the minima of the likelihood surface this problem can occur, causing MINUIT to fail to converge. In order to mitigate this effect as an event approaches the boundary between two bins the contribution the event makes to the PDF is allowed to ‘leak’ into the next bin, this helps to smooth out the steps in the likelihood. The contribution to each bin, w_n is determined as follows:

$$\begin{aligned} w_n(x) &= 1.0 - \frac{1}{1 + \exp\left(\frac{x-l_n}{\delta}\right)} \quad (x < c_n) \\ &= \frac{1}{1 + \exp\left(\frac{x-h_n}{\delta}\right)} \quad (x \geq c_n), \end{aligned} \tag{9.10}$$

where $x = \log(L/E)/(\text{bin width})$, $c_n = (\text{bin centre value})/(\text{bin width})$, $l_n = c_n - 0.5$, $h_n = c_n + 0.5$ and δ is a numerical constant. An event at the boundary between two bins will give half of its contribution to each bin. The effect of this is shown in Figure 9.16, here $\delta = 0.05$ is used in order to illustrate the shape of the distributions. In the oscillation fit the value of δ was tuned so that the fluctuations in likelihood for small variations of α_μ or α_{had} were smoothed out, while the variation of the likelihood appeared unchanged over larger ranges of α_μ or α_{had} . The value chosen was $\delta = 0.0075$.

Figure 9.16: Smoothing $\log(L/E)$ distributions using bin weights as in equation (9.10) with $\delta = 0.05$. The x -axis in this plot is scaled to $\log(L/E)/(\text{bin width})$. Integer values represent bin centres.

9.5.4 Fit To The Combined $\nu_\mu/\bar{\nu}_\mu$ Sample

The entire sample of selected events was fitted to the hypothesis of $\nu_\mu \rightarrow \nu_\tau$ oscillations as described above, assuming that ν_μ and $\bar{\nu}_\mu$ oscillate in the same way. The negative log-likelihood was evaluated for 201 values of $\log(\Delta m_{32}^2/\text{eV}^2)$ between -5.0 and 0.0 and 41 values of $\sin^2 2\theta_{23}$ between 0.0 and 1.0. The best fit point was ($\Delta m_{32}^2 = 1.0 \times 10^{-3} \text{eV}^2$, $\sin^2 2\theta_{23} = 0.8$). Table 9.1 shows the values of the nuisance parameters at the best fit point. Each is within $\pm 1\sigma$ of the nominal value. Over the whole oscillation parameter space considered, all nuisance parameters except the normalisation systematic remain well within this range. This indicates that there is nothing unexpected in the data pulling these parameters away from their nominal values.

The confidence limits on the oscillation parameters are determined by calculating the difference in the negative log-likelihood $-\Delta \ln \mathcal{L}$ at each point relative to the best fit point. In the limit of Gaussian errors, the (68%, 90%, 99%) confidence limits are given by $-\Delta \ln \mathcal{L}$ values of (1.2, 2.3, 4.6). The resulting confidence level contours are shown in Figure 9.17. The contour from the published MINOS atmospheric neutrino analysis [29] is shown by the dashed lines in Figure 9.18, as before the solid lines show the confidence levels obtained above. This analysis used an ν_e result from Soudan 2 [170] to normalise the atmospheric neutrino flux at solar maximum and so only had a 10% systematic uncertainty on the normalisation. This cross-check was unavailable for the

Figure 9.17: Confidence Level contours for a combined $\nu_\mu/\bar{\nu}_\mu$ fit. The best fit point is indicated by a star. The 99%, 90% and 68% confidence regions are shown in red, green and blue lines respectively.

analysis in this thesis as the dataset extends over a longer period and the assumption of a solar maximum flux is no longer valid (see Section 6.3.1). The data are consistent with the previous MINOS atmospheric neutrino result as well as with the results from the MINOS measurement of the NuMI beam [74], Super Kamiokande [67] [68] [69] [70] and K2K [72]. By examining the $-\Delta \ln \mathcal{L}$ values at $\sin^2 2\theta_{23} = 0.0$ it was determined that the null oscillations hypothesis is rejected at the 98% confidence level. As the statistics are quite low is possible that the Gaussian approximation provides an incomplete coverage of the confidence limits. Previously it was found that a Feldman and Cousins approach and a likelihood approach gave similar results [171].

Figure 9.19 shows a slice through the likelihood surface at $\sin^2 2\theta_{23} = 1.0$. The large peak around $\log(\Delta m_{32}^2/\text{eV}^2) = -1.2$ corresponds to the point where the oscillation minimum coincides with the left-hand peak in the $\log(L/E)$ distribution. The slightly smaller peak around $\log(\Delta m_{32}^2/\text{eV}^2) = -3.7$ corresponds to the point where the expected oscillation maximum coincides with the best fit oscillation minimum. The exact features in the data cause the structure between these two peaks. Figure 9.20 shows the $\log\left(\frac{L}{E}\right)_{\text{reco}}$ distributions for high resolution Monte Carlo and data events binned in $\sigma_{\log(L/E)}$. The red dashed line indicates the Monte Carlo expectation for the best fit values of the oscillation parameters.

Figure 9.18: Confidence Level contours for a combined $\nu_\mu/\bar{\nu}_\mu$ fit. The best fit point is indicated by a star. The 99%, 90% and 68% confidence regions are shown in red, green and blue lines respectively. Solid lines indicate the confidence levels for this thesis, dashed lines indicate the confidence levels from [29].

Figure 9.19: Slice through likelihood surface at $\sin^2 2\theta_{23} = 1.0$ for a combined $\nu_\mu/\bar{\nu}_\mu$ fit. The 99%, 90% and 68% confidence levels are shown by the red, green and blue lines respectively. The black line shows the height of the likelihood surface.

Figure 9.20: $\left(\frac{L}{E}\right)_{reco}$ distributions for high resolution Monte Carlo and data events binned in $\sigma(L/E)$. The solid blue line indicates the Monte Carlo expectation in the absence of oscillations. The dashed red line indicates the Monte Carlo expectation for ν_μ and $\bar{\nu}_\mu$ oscillations with the best fit oscillation parameters, $\Delta m_{32}^2 = 0.001 \text{ eV}^2$ and $\sin^2 2\theta_{23} = \sin^2 2\bar{\theta}_{23} = 0.8$. The data are shown by the black points.

Nuisance Parameter	σ_j	α_j	α_j/σ_j
Normalisation	0.15	-0.037	-0.244
Muon Momentum Error	0.03	0.002	0.051
Hadronic Energy Scale Error	0.05	0.010	0.197
Up/Down Relative Efficiency	0.03	0.002	0.073
Spectrum Scaling	1.00	-0.098	-0.098
Relative QE Cross-Section	0.20	-0.022	-0.110

Table 9.1: Values of nuisance parameters at the best fit point for the combined $\nu_\mu/\bar{\nu}_\mu$ fit. The second column shows the estimated systematic uncertainty in each parameter. The third column shows the absolute values of the best fit nuisance parameters, while the fourth column shows each best fit nuisance parameter as a fraction of the systematic uncertainty in that parameter.

The expected sensitivity was determined by averaging the $-\Delta \ln \mathcal{L}$ distributions from 100 simulated experiments generated using the best fit oscillation parameters $\Delta m_{32}^2 = 1.0 \times 10^{-3} \text{ eV}^2$, $\sin^2 2\theta_{23} = 0.8$ as the input parameters. Figure 9.21 shows the expected sensitivity contours. The black points show the location of the best fit points for the 100 Monte Carlo tests; 68 best fit points lie within the 68% confidence region, 92 points lie within the 90% confidence region and all points lie within the 99% confidence region. Figure 9.22 shows a slice through expected sensitivity surface at $\sin^2 2\theta_{23} = 1.0$. The large peak around $\log(\Delta m_{32}^2/\text{eV}^2) = -1.2$ corresponds to the point where the oscillation minimum coincides with the left-hand peak in the $\log(L/E)$ distribution. The peak around $\log(\Delta m_{32}^2/\text{eV}^2) = -3.7$ observed in the data fit is absent here. This part of the contour depends strongly on the exact $\log(L/E)$ values events of observed upward-going events and so averages out over many tests.

9.5.5 Fits To ν_μ Dominated And $\bar{\nu}_\mu$ Dominated Samples

In order to gauge the sensitivity of a separated fit to ν_μ and $\bar{\nu}_\mu$, it is instructive to fit a ν_μ dominated sample and an $\bar{\nu}_\mu$ dominated sample separately. This was done by fitting the sample of events which passed the ν_μ charge identification cut to the hypothesis of $\nu_\mu \rightarrow \nu_\tau$ oscillations as described above. The same fit was then applied to the sample of events which passed the $\bar{\nu}_\mu$ charge identification cut. If the oscillation parameters for $\nu_\mu \rightarrow \nu_\tau$ oscillations and $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ oscillations differ, then a shift in the best fit point between the ν_μ dominated and $\bar{\nu}_\mu$ dominated fits would be expected. A smaller shift

Figure 9.21: Expected Sensitivity of combined $\nu_\mu/\bar{\nu}_\mu$ fit for 12.23 ktyr exposure. The expected 99%, 90% and 68% confidence levels are shown by the red, green and blue lines respectively. The black points show the location of the best fit points for the 100 Monte Carlo tests used to generated the sensitivity contour.

Figure 9.22: Slice through expected sensitivity surface at $\sin^2 2\theta_{23} = 1.0$ of combined $\nu_\mu/\bar{\nu}_\mu$ fit for 12.23 ktyr exposure. The expected 99%, 90% and 68% confidence levels are shown by the red, green and blue lines respectively. The black shows the height of the expected sensitivity surface.

Figure 9.23: Confidence Level contours for a ν_μ dominated fit. The best fit point is indicated by a star. The 99%, 90% and 68% confidence regions are shown in red, green and blue respectively.

relative to the best fit point of the combined fit would also be expected.

The resulting confidence level contours for the fit to the ν_μ dominated sample are shown in Figure 9.23. The best fit point was $(\Delta m_{32}^2 = 9.44 \times 10^{-4} \text{ eV}^2, \sin^2 2\theta_{23} = 0.8)$. Table 9.2 shows the values of the nuisance parameters at the best fit point. Each is within $\pm 1\sigma$ of the nominal value. The result is consistent with the combined $\nu_\mu/\bar{\nu}_\mu$ fit, but the reduced statistics allow a much larger region of parameter space. Figure 9.24 shows a slice through the likelihood surface at $\sin^2 2\theta_{23} = 1.0$. By examining the $-\Delta \ln \mathcal{L}$ values at $\sin^2 2\theta_{23} = 0.0$ it was determined that the null oscillations hypothesis is rejected at the 93% confidence level.

The resulting confidence level contours for the fit to the $\bar{\nu}_\mu$ dominated sample are shown in Figure 9.25. The best fit point was $(\Delta m_{32}^2 = 8.91 \times 10^{-4} \text{ eV}^2, \sin^2 2\theta_{23} = 0.7)$. Table 9.3 shows the values of the nuisance parameters at the best fit point. Each is within $\pm 1\sigma$ of the nominal value. Again the result is consistent with the combined $\nu_\mu/\bar{\nu}_\mu$ fit, but the $\bar{\nu}$ statistics are too low to rule out most of the parameter space. Figure 9.26 shows a slice through the likelihood surface at $\sin^2 2\theta_{23} = 1.0$. By examining the $-\Delta \ln \mathcal{L}$ values at $\sin^2 2\theta_{23} = 0.0$ it was determined that the null oscillations hypothesis is rejected at the 72% confidence level. The fits to ν_μ dominated and $\bar{\nu}_\mu$ dominated samples are

Figure 9.24: Slice through likelihood surface at $\sin^2 2\theta_{23} = 1.0$ for a ν_μ dominated fit. The 99%, 90% and 68% confidence levels are shown by the red, green and blue lines respectively. The black line shows the height of the likelihood surface.

Nuisance Parameter	σ_j	α_j	α_j/σ_j
Normalisation	0.15	0.007	0.048
Muon Momentum Error	0.03	0.013	0.440
Hadronic Energy Scale Error	0.05	0.001	0.011
Up/Down Relative Efficiency	0.03	0.004	0.126
Spectrum Scaling	1.00	-0.141	-0.141
Relative QE Cross-Section	0.20	-0.007	-0.035

Table 9.2: Values of nuisance parameters at the best fit point for the ν_μ dominated fit. The second column shows the estimated systematic uncertainty in each parameter. The third column shows the absolute values of the best fit nuisance parameters, while the fourth column shows each best fit nuisance parameter as a fraction of the systematic uncertainty in that parameter.

consistent with each other, this suggests that the data will not be able to resolve different oscillation parameters for $\bar{\nu}_\mu$ and ν_μ .

9.5.6 Separated $\nu_\mu/\bar{\nu}_\mu$ Fit

In this fit any event which failed the charge identification cut is classified as low resolution and only used in the normalisation. The negative log-likelihood is altered to reflect the

Figure 9.25: Confidence Level contours for an $\bar{\nu}_\mu$ dominated fit. The best fit point is indicated by a star. The 99%, 90% and 68% confidence regions are shown in red, green and blue respectively.

Figure 9.26: Slice through likelihood surface at $\sin^2 2\theta_{23} = 1.0$ for an $\bar{\nu}_\mu$ dominated fit. The 99%, 90% and 68% confidence levels are shown by the red, green and blue lines respectively. The black line shows the height of the likelihood surface.

Nuisance Parameter	σ_j	α_j	α_j/σ_j
Normalisation	0.15	-0.051	-0.339
Muon Momentum Error	0.03	-0.007	-0.224
Hadronic Energy Scale Error	0.05	-0.006	-0.114
Up-Down Relative Efficiency	0.03	-0.001	-0.026
Spectrum Scaling	1.00	-0.417	-0.417
Relative QE Cross-Section	0.20	0.015	0.076

Table 9.3: Values of nuisance parameters at the best fit point for the $\bar{\nu}_\mu$ dominated fit. The second column shows the estimated systematic uncertainty in each parameter. The third column shows the absolute values of the best fit nuisance parameters, while the fourth column shows each best fit nuisance parameter as a fraction of the systematic uncertainty in that parameter.

separation of the sample into ν_μ and $\bar{\nu}_\mu$ events:

$$\begin{aligned}
-\ln \mathcal{L} = & N_{MC} - N \ln N_{MC} \\
& - \sum_{k=1}^5 (N_k^+ \ln p_k^+ + N_k^- \ln p_k^-) \\
& - \sum_q \sum_{k=1}^5 (N_{kq}^{up} \ln p_{kq}^{up} + N_{kq}^{dn} \ln p_{kq}^{dn}) \\
& - \sum_{data} \ln (f(L/E)_{kqi}^{up}) \\
& - \sum_{data} \ln (f(L/E)_{kqi}^{dn}) \\
& + \sum_j \frac{\alpha_j^2}{2\sigma_{\alpha_j}^2}
\end{aligned} \tag{9.11}$$

This function can be divided up into five terms:

Normalisation Term Defined as in Section 9.5.1.

Charge Term The next term represents the relative probabilities of observing ν_μ and $\bar{\nu}_\mu$ events. Only high resolution events where the muon charge can be determined contribute to this term. N_k^+ is the number of data events identified as $\bar{\nu}_\mu$ events in the k th resolution bin and N_k^- is the number of data events identified as ν_μ events in the k th resolution bin, p_k^+ and p_k^- are the probabilities that a Monte Carlo

event in the k th resolution bin is identified as a $\bar{\nu}_\mu$ event or a ν_μ event respectively ($p_k^+ + p_k^- = 1$).

“Up-Down” Term The next term represents the relative probabilities for each charge bin $(-1, +1)$ of observing upward and downward-going events. Only high resolution events where the muon charge can be determined contribute to this term. N_{kq}^{up} is the number of upward-going data events identified as having charge q in the k th resolution bin and N_{kq}^{dn} is the number of downward-going data events identified as having charge q in the k th resolution bin. p_{kq}^{up} and p_{kq}^{dn} are the probabilities that a Monte Carlo event in the k th resolution bin identified as having charge q is upward-going or downward-going respectively ($p_{kq}^{up} + p_{kq}^{dn} = 1$).

Shape Terms The next two terms represent shape fits to the $\ln(L/E)$ distributions for upward and downward-going events respectively. Only high resolution events where the muon charge can be determined contribute to these terms. $f(L/E)_{kqi}^{up}$ is the value of the i th bin of the L/E distribution for upward-going Monte Carlo events identified as having charge q in the k th resolution bin, similarly $f(L/E)_{kqi}^{dn}$ is the value of the i th bin of the L/E distribution for downward-going Monte Carlo events identified as having charge q in the k th resolution bin. The sum is over the data events in the k th resolution bin identified as having charge q . The distributions are normalised such that $\sum_i f(L/E)_{kqi}^{up} = 1$ and $\sum_i f(L/E)_{kqi}^{dn} = 1$.

Nuisance Term Defined as in Section 9.5.1, except for the following changes. As the variations observed in the previous fits were small $\alpha_Q E$, α_μ and α_{had} were fixed at their nominal values in order to speed up the fitting process. In addition to the other nuisance parameters used in the previous fits, an additional nuisance parameter, α_q , associated with the systematic uncertainty in the charge identification, σ_q , is used in the fit. From the split track studies in Section 8.9.1 a value of 2% is assigned to σ_q . This nuisance parameter only affects the charge term in the likelihood function, the number of ν_μ Monte Carlo events in bin k is scaled by $(1.0 + \alpha_q)$ when calculating p_k^+ and p_k^- .

The entire sample of selected events was fitted to the hypothesis of independent $\nu_\mu \rightarrow \nu_\tau$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ oscillations. The negative log-likelihood was evaluated for 201 values of $\log(\Delta m_{32}^2/\text{eV}^2)$ between -5.0 and 0.0 and 201 values of $\log(\overline{\Delta m_{32}^2}/\text{eV}^2)$ between -5.0 and 0.0. Most atmospheric neutrino and long-baseline analyses favour the maximal mixing hypothesis, so $\sin^2 2\theta_{23}$ and $\sin^2 2\overline{\theta}_{23}$ were fixed at 1.0 in order to reduce the number of parameter combinations.

Figure 9.27: Confidence Level contours for separated $\nu/\bar{\nu}$ fit. The best fit point is indicated by a star. The 90% and 68% confidence regions are shown in red and green respectively. All values of $\log(\overline{\Delta m_{32}^2}/\text{eV}^2)$ tested are allowed at the 99% confidence level.

The resulting confidence level contours are shown in Figure 9.27. The best fit point was $(\Delta m_{32}^2 = 9.44 \times 10^{-4} \text{ eV}^2, \overline{\Delta m_{32}^2} = 1.58 \times 10^{-2} \text{ eV}^2)$. Here the best fit normalisation is 106.5% of nominal. All values of $\log(\overline{\Delta m_{32}^2}/\text{eV}^2)$ tested are allowed at the 99% confidence level. The constraints on Δm_{32}^2 are stronger than those on $\overline{\Delta m_{32}^2}$, this reflects the relative numbers of events in the sample. The data do not rule out large differences in the values of the two parameters, although more statistics would be required to make a stronger statement.

9.6 Projected Sensitivities

The MINOS experiment is due to run until 2009. In which time a Far Detector exposure of 25 ktyr will be accumulated. The projected sensitivity of the combined $\nu_\mu/\bar{\nu}_\mu$ oscillation fit is calculated for a detector exposure of 25 ktyr by averaging the $-\Delta \ln \mathcal{L}$ distributions from 50 simulated experiments generated using the best fit point from the MINOS measurement of the NuMI beam ($\Delta m_{32}^2 = 0.00274 \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$) [74] as the input parameters. Figure 9.28 shows the projected sensitivity contours. The black points show the location of the best fit points for the 50 Monte Carlo tests; 39 best fit

Figure 9.28: Projected Sensitivity of combined $\nu_\mu/\bar{\nu}_\mu$ fit for 25 kyr exposure. The projected 99%, 90% and 68% confidence levels are shown by the red, green and blue lines respectively. The black points show the location of the best fit points for the 50 Monte Carlo tests used to generated the sensitivity contour.

points lie within the 68% confidence region, 46 points lie within the 90% confidence region and all points lie within the 99% confidence region. Figure 9.29 shows a slice through projected sensitivity surface at $\sin^2 2\theta_{23} = 1.0$. The oscillation parameters are found to be in the range $1.05 \times 10^{-3} \text{ eV}^2 < \Delta m_{32}^2 < 8.32 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta_{23} > 0.7$ at the 90% confidence level.

Figure 9.29: Slice through projected sensitivity surface at $\sin^2 2\theta_{23} = 1.0$ of combined $\nu_\mu/\bar{\nu}_\mu$ fit for 25 ktyr exposure. The projected 99%, 90% and 68% confidence levels are shown by the red, green and blue lines respectively. The black shows the height of the projected sensitivity surface.

Chapter 10

Conclusions

“If you have an apple and I have an apple and we exchange these apples then you and I will still each have one apple. But if you have an idea and I have an idea and we exchange these ideas, then each of us will have two ideas.”

— George Bernard Shaw, 1856 - 1950

The completed MINOS Far Detector has been taking data since 1st August 2003 and as of 1st March 2006 had accumulated 12.23 ktyr of exposure to the flux of neutrinos generated by the interaction of cosmic rays with the Earth’s atmosphere. The Far Detector and its systems are now well understood. Detailed monitoring information is available. This information is used in a series of well-defined data quality checks to select the dataset for the atmospheric neutrino analysis. The start of beam operations in January 2005 created an additional source of ν_μ events in the Far Detector. A GPS timing signal is used to synchronise the clocks at the Near and Far Detectors and so enable the definition of a time window when each ‘spill’ of beam neutrinos is expected to arrive at the Far Detector. Records of these times enable 99% of beam neutrino events to be vetoed. To ensure purity of the atmospheric neutrino sample, periods when the beam was active, but spill arrival times were not recorded are vetoed from the atmospheric neutrino sample.

DAQ run control software has been developed which enables far greater configuration of the detector readout and easier diagnosis of any problems with detector electronics and online software processes.

The data has been calibrated to correct for variations in response across the detector. The process of calibrating the Far Detector timing system at the sub-nanosecond level has been automated and the constants produced validated. Accurate timing information is particularly important for the atmospheric neutrino analysis as this information is used to measure the direction of atmospheric neutrino events.

Two algorithms have been developed to identify $\nu/\bar{\nu}_\mu$ CC interactions in the Far Detector. The signature of these events is a muon track with an interaction vertex contained within the detector fiducial volume. The $\nu/\bar{\nu}_\mu$ CC signal events must be selected from a background of cosmic muon events.

The first algorithm is based on the fact that the background cosmic muon events are predominantly composed of downward-going tracks. Therefore selecting only upward-going events vastly reduces the background. The remaining background events are cosmic muon events which stop inside the detector, but have poor timing information and have been mis-reconstructed as upward-going and neutrino-induced muons from neutrino interactions outside the detector. The first source of background is removed by cutting on the quality of the timing information for the event. The second source of background is removed by requiring the muon track vertex to be contained within the fiducial volume.

There are many shorter low energy neutrino events which are fully contained within the fiducial volume, but are lost in this first selection algorithm due to poor timing information. The second algorithm relies on the containment of the highest end of the track. This method initially ignores timing information and assumes that all events are cosmic muons (*i.e.* downward-going). Here the main background is cosmic muons which have entered the fiducial volume between planes and so appear contained. This is removed by a series of containment cuts on the highest end of the track. The remaining events are then divided into events with tracks fully contained within the fiducial volume and events with partially contained tracks. This is done by looking at the containment of the lowest end of the track. The partially contained sample must then be checked to ensure that the track vertex is contained within the fiducial volume. This is done using timing information and removes a background of neutrino-induced muon events. The veto shield is then used to tag the majority of the remaining cosmic muon background in both samples. Events which pass either or both of the selection algorithms are added to the final sample of selected events. A total of 277 candidate contained vertex $\nu/\bar{\nu}_\mu$ CC data events were selected. This compares to an expectation of 354.4 ± 47.4 events in the absence of neutrino oscillations.

By applying further cuts on the quality of timing information of selected events a high resolution sub-sample where the direction of the muon could be clearly identified was selected. This yielded 77 upward-going data events, 105 downward-going data events and 95 data events where the direction was ambiguous. The ratio between the measured and expected up/down ratio is:

$$\frac{R_{u/d}^{\text{data}}}{R_{u/d}^{\text{MC}}} = 0.72_{-0.11}^{+0.13}(\text{stat.}) \pm 0.04(\text{sys.}). \quad (10.1)$$

This is 2.1σ away from the expectation for no oscillations.

Further cuts were applied to the high resolution events to select those where the charge of the muon could be clearly identified. A total of 112 high resolution data events were identified as ν_μ CC interactions, 55 high resolution data events were identified as $\bar{\nu}_\mu$ CC interactions, with the remaining 15 high resolution data events being ambiguous. From the upward-going sample 52 data events were identified as ν_μ and 22 data events were identified as $\bar{\nu}_\mu$. From the downward-going sample 60 data events were identified as ν_μ and 33 data events were identified as $\bar{\nu}_\mu$. The ratio between the measured and expected $\bar{\nu}_\mu/\nu_\mu$ ratio is:

$$\frac{R_{\bar{\nu}/\nu}^{\text{data}}}{R_{\bar{\nu}/\nu}^{\text{MC}}} = 0.93_{-0.15}^{+0.19}(\text{stat.}) \pm 0.12(\text{sys.}). \quad (10.2)$$

This is consistent with neutrinos and anti-neutrinos having the same oscillation parameters to within the errors quoted.

The best measurement of the neutrino direction is the direction of the muon at the event vertex. From this $(\frac{L}{E})_{\text{reco}}$ can be calculated. It should be noted that for low energy and/or high y events the opening angle between the observed muon direction and the true neutrino direction will be large. This means that particularly when the muon direction is close to the horizon $(\frac{L}{E})_{\text{reco}}$ may differ by a large amount from the true value, $(\frac{L}{E})_{\text{true}}$. As the sample of selected events is small, it is important to extract as much information as possible from each event. A Bayesian method is used to estimate the accuracy of $(\frac{L}{E})_{\text{reco}}$ for each event by calculating a probability distribution for $\log(\frac{L}{E})_{\text{true}}$ given the observed $\mathbf{p}_\mu$ and E_{shw} values. The RMS of this PDF, $\sigma_{\log(L/E)}$, can be used as an estimate of the accuracy of the (L/E) reconstruction. Events can then be binned according to $\sigma_{\log(L/E)}$ in the fit to the hypothesis of $\nu_\mu \rightarrow \nu_\tau$ oscillations.

A maximum likelihood analysis was used to determine the oscillation parameters

$(\Delta m_{32}^2, \sin^2 2\theta_{23})$ assuming two-flavour $\nu_\mu \rightarrow \nu_\tau$ oscillations and that ν_μ and $\bar{\nu}_\mu$ oscillate in the same way. The best fit point was $(\Delta m_{32}^2 = 1.0 \times 10^{-3} \text{ eV}^2, \sin^2 2\theta_{23} = 0.8)$. The null oscillations hypothesis was rejected at the 98% confidence level.

In order to gauge the sensitivity to independent ν_μ and $\bar{\nu}_\mu$ oscillations, the maximum likelihood analysis was applied to a ν_μ dominated sample and an $\bar{\nu}_\mu$ dominated sample separately. For the ν_μ dominated sample, the best fit point was $(\Delta m_{32}^2 = 9.44 \times 10^{-4} \text{ eV}^2, \sin^2 2\theta_{23} = 0.8)$, with the null oscillations hypothesis being rejected at the 93% confidence level. For the $\bar{\nu}_\mu$ dominated sample, the best fit point was $(\Delta m_{32}^2 = 8.91 \times 10^{-4} \text{ eV}^2, \sin^2 2\theta_{23} = 0.7)$, with the null oscillations hypothesis being rejected at the 72% confidence level. The results of the two fits were consistent with $\Delta m_{32}^2 = \overline{\Delta m_{32}^2}$ and $\sin^2 2\theta_{23} = \overline{\sin^2 2\theta_{23}}$.

The data were separated into three sub-samples, events with unknown charge, events identified as ν_μ interactions and events identified as $\bar{\nu}_\mu$ interactions. A maximum likelihood analysis was applied the data assuming independent two-flavour $\nu_\mu \rightarrow \nu_\tau$ oscillations and $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ oscillations and maximal mixing. The best fit point was $(\Delta m_{32}^2 = 9.44 \times 10^{-4} \text{ eV}^2, \overline{\Delta m_{32}^2} = 1.58 \times 10^{-2} \text{ eV}^2)$.

The MINOS experiment is due to run until 2009, in which time a Far Detector exposure of 25 ktyr will be accumulated. After this exposure it is projected that the oscillation parameters can be determined to be $1.05 \times 10^{-3} \text{ eV}^2 < \Delta m_{32}^2 < 8.32 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta_{23} > 0.7$ at the 90% confidence level.

There a number of improvements that could be made to increase the sensitivity of the analysis. For example, systematic errors on the neutrino flux shape and normalisation could be reduced by including a measurement of the atmospheric $\nu_e/\bar{\nu}_e$ flux in the analysis. Currently only events which interact inside the fiducial volume are considered. Combining these events with upward-going muons due to neutrino interactions in the rock would improve the overall sensitivity further.

Appendix A

Selected Events

The analysis in this thesis is based on data collected at MINOS Far Detector between 1st August 2003 and 28th February 2006. The 277 candidate contained vertex $\nu_\mu/\bar{\nu}_\mu$ CC events are listed in Table A.1. The values of E_μ and E_{shw} are quoted in GeV. Class refers to the branch(es) of the selection in Chapter 8 which the event passed. RESN refers to whether the event passed the high resolution cuts in Chapter 8. For those events which passed this selection the U/D values gives the direction of the track; either upward-going (1) or downward-going (-1). Q indicates the μ charge sign for those events which passed the charge identification cuts in Chapter 8. $\log(L/E)$ values are quoted where L is in km and E is in GeV, $\sigma_{(L/E)}$ is the uncertainty in $\log(L/E)$, as calculated in Chapter 9.

Table A.1: A complete list of selected data events

Run	Snarl	E_μ	E_{shw}	Class	RESN	U/D	Q	$\log(L/E)$	$\sigma_{(L/E)}$
18581	63807	0.67	0	FC	1	-1	-1	1.64	0.91
18866	62010	1.65	0	FC	1	-1	1	1.33	0.39
18902	36351	8.73	0.48	PCUP	1	1	1	2.79	0.25
19125	93841	1.71	0	PCDN	1	-1	-1	1.62	0.62
19147	68821	4.13	0	PCDN	0	0	0	0.67	-
19251	92948	1.47	28.73	FC+FCUP	0	0	0	1.93	-
Continued on next page									

Table A.1 – continued from previous page

Run	Snarl	E_μ	E_{shw}	Class	RESN	U/D	Q	$\log(L/E)$	$\sigma_{(L/E)}$
19459	14542	0.91	4.60	FC	0	0	0	3.19	-
19794	48889	0.50	0	FC	0	0	0	1.83	-
19953	1269	1.20	0.10	FC	1	1	1	3.49	0.71
20012	88409	0.62	0.142	FC	0	0	0	3.41	-
20131	33435	0.58	0	FC	1	-1	-1	1.72	0.96
20147	26192	0.28	0	FC	0	0	0	3.76	-
20629	31711	2.36	0	FC+FCUP	1	1	-1	3.37	0.33
20747	17532	2.13	0	PCUP	1	1	1	3.63	0.23
20818	65032	1.22	0	FCUP	1	1	-1	3.73	0.60
20884	27361	20.64	0.25	PCDN	1	-1	0	0.27	0.33
20998	6938	3.49	3.73	FC	1	-1	-1	1.21	0.67
21162	2737	0.68	0.68	FC	0	0	0	3.90	-
21178	11871	0.83	0.27	PCDN	1	-1	-1	1.40	0.48
21464	2273	1.17	0.04	FC+FCUP	1	1	-1	3.74	0.62
21498	29870	2.37	0	FC	1	-1	-1	1.91	0.70
21561	13512	0.72	0	FC	1	1	-1	4.14	0.62
21565	15664	1.29	1.00	FC	1	-1	-1	1.07	0.53
21565	28374	1.22	1.44	FC	1	-1	-1	1.29	0.86
21649	50029	1.26	0	FC+FCUP	1	1	-1	3.69	0.61
21686	56630	1.33	0	FC+FCUP	1	1	-1	3.77	0.43
21737	63416	1.18	3.25	FC	1	1	-1	2.45	0.98
21772	49880	4.49	1.39	PCUP	1	1	-1	2.71	0.47
21906	116585	0.53	0	FC	0	0	0	1.61	-
21924	63327	9.38	2.67	PCDN	1	-1	0	0.96	0.52
22050	51358	23.34	2.46	PCDN	1	-1	-1	-0.04	0.27
22101	121177	1.47	0.06	PCDN	1	-1	-1	1.45	0.49

Continued on next page

Table A.1 – continued from previous page

Run	Snarl	E_μ	E_{shw}	Class	RESN	U/D	Q	$\log(L/E)$	$\sigma_{(L/E)}$
22201	10289	0.85	0	FC	0	0	0	2.17	-
22213	45399	2.62	0.07	PCDN	1	-1	1	1.43	0.41
22312	21640	0.31	1.93	FC	0	0	0	1.16	-
22315	65804	0.67	1.73	FC	0	0	0	3.15	-
22424	136567	1.61	0	PCDN	1	-1	1	1.00	0.23
22424	149574	0.40	0	FC	0	0	0	1.94	-
22446	82423	3.88	0.72	FC	1	-1	-1	1.04	0.33
22509	24855	0.88	0	FC	1	1	-1	3.80	0.87
22713	10258	0.49	0.70	FC	0	0	0	2.80	-
22746	143907	3.78	0	PCDN	1	-1	0	1.15	0.56
22850	92736	0.61	0	FC	0	0	0	1.58	-
22868	10353	0.71	0	FC	1	-1	-1	1.85	1.01
22874	14147	0.58	1.15	FC	0	0	0	1.39	-
22919	7988	1.62	0.14	FC	1	-1	1	2.31	0.77
22919	67340	5.45	1.54	PCDN	1	-1	1	0.98	0.27
22936	84388	0.94	0.94	FC	1	-1	-1	1.23	0.83
22992	106945	3.78	2.19	PCDN	1	-1	-1	1.59	0.61
22995	67760	0.36	0	FC	0	0	0	4.17	-
23285	33554	172.19	1.09	PCDN	1	-1	0	-0.76	0.22
23285	59052	0.56	0	FC	0	0	0	3.05	-
23297	43710	1.43	1.24	FC	1	1	-1	3.26	0.71
23584	107525	1.24	0.18	PCDN	1	-1	-1	1.18	0.25
23604	95164	0.65	1.95	FC	1	-1	-1	1.47	1.07
23943	85316	0.58	0.06	FC	0	0	0	1.89	-
23952	24080	0.82	0	FC	1	1	-1	3.57	1.01
23998	9213	1.10	0.46	FC	1	1	0	3.62	0.70

Continued on next page

Table A.1 – continued from previous page

Run	Snarl	E_μ	E_{shw}	Class	RESN	U/D	Q	$\log(L/E)$	$\sigma_{(L/E)}$
24070	4	7.05	0.13	PCDN	1	-1	0	0.50	0.27
24088	8719	0.70	0	FC	1	1	-1	3.80	1.00
24141	4689	1.99	3.76	FCUP	1	1	-1	3.23	0.26
24816	796	0.41	0	FC	0	0	0	3.84	-
24829	36847	102.01	0	PCDN	1	-1	0	-0.45	0.37
24829	87042	0.87	0	FC	1	1	-1	3.53	0.99
24844	97277	1.46	0	FC+FCUP	1	1	1	3.78	0.26
24871	120533	1.08	0	FC	1	1	1	3.91	0.37
24874	37215	4.17	0	PCUP	1	1	1	3.21	0.23
24881	44348	0.85	0.63	FC	0	0	0	3.66	-
24941	72237	1.00	3.57	FC	1	-1	-1	0.68	0.59
24982	23731	2.66	6.17	FC	1	1	-1	2.34	0.77
24988	8655	1.36	0	FC	1	-1	1	1.50	0.58
25048	25582	1.15	0.37	FC	1	-1	-1	1.08	0.37
25066	70672	4.45	0	PCDN	1	-1	1	1.67	0.54
25142	38534	0.67	0.99	FC	0	0	0	3.10	-
25154	96745	0.55	1.98	FC	0	0	0	3.32	-
25189	77584	3.68	0.61	PCUP	1	1	1	3.09	0.27
25225	74481	0.92	0	FC	0	0	0	2.07	-
25225	79537	0.73	0	FC	0	0	0	1.98	-
25240	23184	0.82	0.31	FC	0	0	0	1.43	-
25283	18347	6.26	0.61	FC	1	-1	-1	1.18	0.39
25304	10577	2.03	0.36	PCUP	1	1	1	3.29	0.37
25313	5784	0.68	0.06	FC	0	0	0	3.62	-
25355	119017	0.54	5.63	FC	0	0	0	0.59	-
25393	105572	0.45	0	FC	0	0	0	4.14	-

Continued on next page

Table A.1 – continued from previous page

Run	Snarl	E_μ	E_{shw}	Class	RESN	U/D	Q	$\log(L/E)$	$\sigma_{(L/E)}$
25411	35592	1.76	0	FC+FCUP	1	1	-1	3.65	0.29
25538	73415	1.16	0.95	FC	1	1	1	3.55	0.45
25581	34328	0.48	0.29	FC	0	0	0	1.60	-
25589	98333	0.56	0	FC	0	0	0	3.61	-
25610	47904	0.45	2.28	FC	0	0	0	3.52	-
25627	47998	1.37	0.25	PCDN	1	-1	-1	1.96	0.82
25773	114723	238.57	6.10	PCDN	1	-1	0	-0.46	0.39
25825	52614	0.83	0	FC	0	0	0	1.90	-
25831	21795	0.43	0	FC	0	0	0	1.77	-
25840	57979	1.07	0	PCUP	1	1	-1	3.16	0.79
25988	89674	4.97	0.06	PCUP	1	1	1	2.68	0.33
25991	67835	0.85	0	FC	1	-1	1	2.15	0.95
25991	68957	3.72	0	FC	1	-1	-1	0.96	0.17
25994	47474	1.50	0	FC	0	0	0	2.34	-
26000	60501	0.61	0	FC	0	0	0	1.87	-
26003	33461	0.74	0	FC	0	0	0	1.56	-
26277	28086	11.34	5.50	PCDN	1	-1	-1	0.44	0.24
26451	19477	1.47	0	FC	1	-1	1	1.13	0.23
26471	25432	0.80	0	FC	1	-1	1	2.62	0.98
26537	46199	0.66	0	FC	0	0	0	1.45	-
26974	83043	0.32	0	FC	0	0	0	2.54	-
27033	50421	0.62	4.32	FC	0	0	0	3.14	-
27184	116702	0.75	0	FC	1	-1	-1	2.16	1.05
27310	44407	0.58	3.74	FC	0	0	0	2.19	-
27325	123806	0.78	2.70	FC	0	0	0	3.40	-
27459	33093	0.37	0.05	FC	0	0	0	2.23	-

Continued on next page

Table A.1 – continued from previous page

Run	Snarl	E_μ	E_{shw}	Class	RESN	U/D	Q	$\log(L/E)$	$\sigma_{(L/E)}$
27566	28709	0.44	0	FC	0	0	0	3.86	-
27581	61928	0.72	0.13	FC	1	-1	-1	1.31	0.56
27684	21584	3.56	0.38	FC	1	-1	-1	0.61	0.10
27694	103569	1.62	0	FC	1	-1	-1	1.18	0.30
27721	13243	1.13	0.05	FC	1	-1	1	1.19	0.30
27724	85333	7.85	0.06	PCUP	1	1	1	2.33	0.35
27740	63058	0.82	0	FC	0	0	0	3.85	-
27781	54144	0.49	0	FC	0	0	0	3.40	-
27796	6882	0.56	2.12	FC	0	0	0	3.19	-
27796	43312	0.34	0	FC	0	0	0	3.99	-
27834	56265	1.45	0	PCDN	1	-1	1	1.14	0.27
27858	51881	0.36	0	FC	0	0	0	2.37	-
27879	79396	1.57	1.29	FC	1	1	-1	2.57	0.88
27894	80994	1.22	2.80	FC	1	-1	1	1.00	0.68
27950	94297	1.72	0.67	FC	1	-1	-1	1.93	0.84
28022	88565	0.41	3.86	FC	0	0	0	3.17	-
28031	84961	1.34	0	FC	0	0	0	1.40	-
28040	30360	2.11	2.41	PCDN	1	-1	-1	0.57	0.20
28113	15701	1.46	0	FC+FCUP	1	1	-1	3.62	0.52
28220	5608	0.43	2.34	FC	0	0	0	2.28	-
28226	69989	1.75	0	FCUP	1	1	-1	3.61	0.33
28273	47784	2.76	0	PCDN	1	-1	-1	1.83	0.63
28330	76058	0.37	4.74	FC	0	0	0	3.00	-
28401	6778	0.74	0.07	FC	1	-1	-1	1.60	0.92
28410	22364	0.56	0	FC	0	0	0	1.93	-
28451	45021	0.40	0.35	FC	0	0	0	3.14	-

Continued on next page

Table A.1 – continued from previous page

Run	Snarl	E_μ	E_{shw}	Class	RESN	U/D	Q	$\log(L/E)$	$\sigma_{(L/E)}$
28535	65406	12.18	7.07	PCDN	1	-1	-1	0.92	0.30
28630	108505	1.31	0	FC+FCUP	1	1	1	3.59	0.55
28937	69178	0.67	0	FC	1	1	-1	3.69	1.05
28955	8240	0.80	0	FC	1	-1	-1	1.59	0.86
28955	15294	3.20	7.19	FC	1	-1	-1	0.69	0.53
29052	63525	1.83	0	FC+FCUP	1	1	1	3.79	0.17
29103	35113	41.49	0	PCDN	1	-1	0	-0.28	0.29
29103	62927	1.72	0	FC+FCUP	1	1	-1	2.59	0.81
29104	25295	2.83	6.52	PCDN	1	-1	1	0.41	0.24
29168	27987	0.61	0.36	FC	0	0	0	1.21	-
29174	57022	1.78	11.90	FC+FCUP	1	1	-1	2.83	0.35
29178	32507	2.73	0.09	FC	1	-1	-1	1.05	0.23
29185	441	3.01	0.45	FC	1	-1	-1	2.06	0.69
29209	57325	0.42	0	FC	0	0	0	2.01	-
29223	70659	5.22	8.54	PCUP	1	1	-1	2.64	0.21
29241	24666	0.44	0	FC	0	0	0	4.15	-
29297	65373	0.48	0.18	FC	0	0	0	1.94	-
29326	70747	3.10	0.11	FC	1	-1	1	1.75	0.59
29359	18505	1.08	0	FC+FCUP	1	1	-1	2.91	0.97
29360	75148	0.71	0	FC	1	1	-1	4.04	0.81
29400	9266	1.03	0.37	FC	0	0	0	1.42	-
29401	73086	1.40	0	PCUP	1	1	-1	3.80	0.27
29405	5448	4.02	0.36	PCUP	1	1	1	3.27	0.23
29442	2686	2.12	0	PCUP	1	1	-1	2.69	0.69
29442	95637	0.44	4.52	FC	0	0	0	2.75	-
29466	33223	2.23	1.19	PCDN	1	-1	0	1.27	0.69

Continued on next page

Table A.1 – continued from previous page

Run	Snarl	E_μ	E_{shw}	Class	RESN	U/D	Q	$\log(L/E)$	$\sigma_{(L/E)}$
29475	49153	0.45	0	FC	0	0	0	3.28	-
30372	30679	1.97	2.70	FC+FCUP	1	1	1	3.06	0.44
30612	5693	0.41	1.04	FC	0	0	0	1.89	-
30613	4708	0.56	0	FC	0	0	0	2.80	-
30637	67770	4.93	0	FC	1	-1	1	0.86	0.13
30642	97399	2.75	0	PCUP	1	1	-1	3.45	0.22
30671	103627	2.73	0	PCDN	1	-1	-1	0.87	0.21
30677	55672	0.43	0	FC	0	0	0	1.99	-
30680	103859	0.40	1.86	FC	0	0	0	1.61	-
30681	26686	0.36	0.99	FC	0	0	0	3.72	-
31216	48743	2.08	0	FC+FCUP	1	1	-1	3.06	0.68
31286	110416	0.49	0	FC	0	0	0	1.74	-
31375	52901	4.92	0	PCDN	1	-1	1	0.62	0.20
31399	72409	2.83	2.92	PCUP	1	1	-1	2.76	0.45
31403	106303	1.29	0	FC	1	-1	-1	1.49	0.69
31423	46039	0.73	0	FC	1	-1	-1	2.72	1.06
31423	71388	2.93	0.25	FC	1	-1	1	1.10	0.27
31428	106871	0.99	2.22	FC	1	1	-1	2.54	1.02
31453	51127	0.43	1.23	FC	0	0	0	2.36	-
31745	76681	7.69	1.44	PCUP	1	1	-1	1.80	0.50
31833	32965	1.04	0	FC	1	-1	-1	1.95	0.95
31833	61980	2.42	0	PCDN	1	-1	1	1.13	0.24
31836	124767	0.50	1.20	FC	0	0	0	1.23	-
31913	99456	0.45	0	FC	0	0	0	1.79	-
31929	44485	0.71	6.65	FC	0	0	0	2.86	-
31944	51547	0.33	0	FC	0	0	0	2.91	-

Continued on next page

Table A.1 – continued from previous page

Run	Snarl	E_μ	E_{shw}	Class	RESN	U/D	Q	$\log(L/E)$	$\sigma_{(L/E)}$
31954	45894	1.35	5.56	PCDN	1	-1	-1	0.90	0.69
31975	129943	2.26	0	PCDN	1	-1	-1	0.88	0.21
32018	120421	0.73	0.13	FC	1	1	-1	2.98	1.07
32158	34707	1.45	0.74	FC	1	-1	1	0.89	0.21
32167	102367	0.32	0	FC	0	0	0	4.22	-
32183	25584	5.27	5.95	PCDN	1	-1	1	0.15	0.19
32194	42630	1.22	5.63	FC+FCUP	1	1	-1	3.00	0.72
32227	30744	0.97	0.72	FC	1	1	-1	3.71	0.60
32231	100485	0.83	0	FC	1	-1	-1	1.56	0.83
32454	63617	1.80	2.18	FC+FCUP	1	1	-1	3.36	0.27
32460	70494	1.13	0	FC	1	-1	-1	1.16	0.31
32470	65088	2.76	0.05	PCDN	1	-1	0	0.89	0.33
32471	107926	4.48	0	PCDN	1	-1	-1	0.69	0.19
32476	91416	0.39	0	FC	0	0	0	2.53	-
32481	111317	0.70	0.57	FC	0	0	0	3.55	-
32486	62343	3.01	0	FC+FCUP	1	1	-1	3.42	0.16
32503	353	0.84	0	FC	1	-1	1	1.28	0.35
32527	76966	6.56	0.50	PCDN	1	-1	-1	0.59	0.21
32586	80936	0.41	3.37	FC	0	0	0	1.47	-
32592	87273	2.04	1.40	FC	1	-1	-1	1.20	0.65
32608	98071	1.24	2.05	FC	1	-1	-1	1.06	0.77
32608	146846	2.73	0.23	FC	1	-1	-1	1.05	0.26
32611	155812	5.37	6.25	PCUP	1	1	1	2.28	0.38
32623	18475	1.09	0	FC	1	1	-1	3.03	0.96
32635	117559	1.76	2.08	PCUP	1	1	1	3.30	0.29
32635	139530	0.54	0	FC	0	0	0	1.81	-

Continued on next page

Table A.1 – continued from previous page

Run	Snarl	E_μ	E_{shw}	Class	RESN	U/D	Q	$\log(L/E)$	$\sigma_{(L/E)}$
32644	83222	0.92	3.69	PCUP	1	1	-1	3.12	0.58
32654	66555	1.42	1.73	FC+FCUP	1	1	-1	3.38	0.47
32654	172531	1.89	0	FC	1	-1	1	1.67	0.65
32666	53808	0.38	3.32	FC	0	0	0	1.81	-
32666	262583	31.04	10.91	PCUP	1	1	0	2.09	0.29
32672	333044	0.79	0.23	FC	1	-1	1	1.40	0.66
32678	284600	2.88	7.04	FC	1	-1	1	0.99	0.63
32693	246451	24.68	12.09	PCUP	1	1	-1	2.20	0.23
32713	151220	0.82	2.08	FC	1	-1	1	0.99	0.70
32725	88260	0.65	0	FC	1	-1	-1	1.54	0.79
32731	296285	0.61	0.40	FC	0	0	0	1.51	-
32737	182097	1.32	0.71	FCUP	1	1	1	3.76	0.23
32737	242131	0.50	5.16	FC	0	0	0	2.76	-
32746	100729	0.74	10.22	FC	0	0	0	2.82	-
32746	137373	0.38	2.26	FC	0	0	0	1.24	-
32785	208696	0.31	2.54	FC	0	0	0	3.01	-
32820	48880	1.85	0.27	PCDN	1	-1	1	1.01	0.26
32823	211279	1.67	0.44	FCUP	1	1	1	3.23	0.59
32836	192163	6.68	4.79	PCDN	1	-1	-1	0.33	0.24
32862	107208	1.49	0.11	FCUP	1	1	-1	3.40	0.70
32868	150167	4.97	0.03	PCUP	1	1	1	3.16	0.25
32871	181373	0.93	1.15	FC	1	1	-1	3.11	1.00
32907	41892	0.76	0.24	FC	1	-1	-1	1.29	0.63
32907	239050	2.68	1.18	FC	1	-1	1	0.64	0.14
32910	42483	9.28	9.70	PCDN	1	-1	-1	0.45	0.37
32951	120166	0.81	0	FC	1	-1	-1	1.35	0.56

Continued on next page

Table A.1 – continued from previous page

Run	Snarl	E_μ	E_{shw}	Class	RESN	U/D	Q	$\log(L/E)$	$\sigma_{(L/E)}$
33003	111323	1.51	1.49	FC	1	-1	1	1.68	0.82
33006	134190	2.06	0.09	FC	1	-1	-1	1.44	0.60
33035	68870	2.67	0	PCUP	1	1	-1	3.52	0.21
33086	90299	33.08	0.08	PCUP	1	1	0	1.55	0.37
33119	26399	1.93	0	PCDN	1	-1	-1	0.98	0.21
33145	7553	0.48	0	FC	0	0	0	1.79	-
33156	11280	0.99	0	FC	1	-1	-1	1.45	0.71
33165	78424	0.56	0	FC	0	0	0	2.51	-
33178	120546	0.36	0.52	FC	0	0	0	3.87	-
33181	238023	0.97	0.39	FC	1	-1	-1	1.70	0.97
33268	261361	1.71	0.08	FCUP	1	1	-1	3.84	0.19
33280	93186	0.83	0	PCDN	1	-1	1	1.79	0.62
33280	131606	0.38	0	FC	0	0	0	1.94	-
33297	75059	5.54	0	PCUP	1	1	-1	2.64	0.42
33300	4606	1.05	0.46	FC	1	1	-1	3.37	0.92
33300	112735	0.42	0.08	FC	0	0	0	1.91	-
33300	133292	0.62	0	FC	0	0	0	1.43	-
33300	246175	2.65	0	PCDN	1	-1	1	0.79	0.25
33369	62126	1.06	0.91	FC	1	1	-1	2.47	0.99
33405	254468	0.49	0.93	FC	0	0	0	1.22	-
33461	170413	6.55	0.54	PCDN	1	-1	1	0.56	0.24
33492	218513	3.07	0	FCUP	1	1	1	3.35	0.17
33508	76220	0.52	0.11	FC	0	0	0	3.37	-
33511	272104	1.66	0.28	FC	1	-1	-1	1.00	0.23
33627	10855	1.94	0.08	FC+FCUP	1	1	-1	2.95	0.75
33668	199751	0.62	0	FC	1	-1	0	1.92	1.05

Continued on next page

Table A.1 – continued from previous page

Run	Snarl	E_μ	E_{shw}	Class	RESN	U/D	Q	$\log(L/E)$	$\sigma_{(L/E)}$
33668	204867	11.92	8.89	PCDN	1	-1	-1	1.02	0.47
33692	42091	5.77	5.46	PCUP	1	1	-1	2.16	0.49
33710	104030	2.03	0.73	PCDN	1	-1	1	1.58	0.57
33710	121286	0.40	0	FC	0	0	0	2.19	-
33710	200482	0.52	0	FC	0	0	0	1.78	-
33716	205905	17.65	0.68	PCDN	1	-1	0	0.57	0.37
33719	39401	1.59	0.54	PCUP	1	1	1	3.69	0.21
33729	108895	7.96	0.07	PCDN	1	-1	-1	0.40	0.21
33732	165188	0.49	0.34	FC	0	0	0	4.06	-
33756	27720	3.06	0.38	FCUP	1	1	-1	3.48	0.13
33782	143914	0.55	0	FC	0	0	0	4.21	-

Bibliography

- [1] J. Chadwick, Ver. Dtsch. Physik. Ges. **16**, 383 (1914).
- [2] W. Pauli, a letter to the Turbingen Conference on radioactivity, December, 1930.
- [3] J. Chadwick, Nature **129**, 312 (1932).
- [4] E. Fermi, Z. Phys. **88**, 161 (1934).
- [5] H. Yukawa, Proc. Phys. Math. Soc. Jap. **17**, 48 (1935).
- [6] S. Sakata and T. Inoue, Prog. Theo. Phys. **1**, 143 (1946).
- [7] F. Reines and C. L. Cowan, Phys. Rev. **92**, 830 (1953).
- [8] F. Reines and C. L. Cowan, Science **124**, 103 (1956).
- [9] B. Pontecorvo, Sov. Phys. JETP **7**, 172 (1958).
- [10] M. Goldhaber, L. Grodzins, and A. W. Sunyar, Phys. Rev. **109**, 1015 (1958).
- [11] R. P. Feynman and M. Gell-Mann, Phys. Rev. **109**, 193 (1958).
- [12] C. N. Lee, T. D. amd Yang, Phys. Rev. **105**, 1671 (1957).
- [13] R. Davis and D. S. Harmer, Bull. Am. Phys. Soc. **4**, 217 (1959).
- [14] M. Schwartz, Phys. Rev. Lett. **4**, 306 (1960).
- [15] G. Danby *et al.*, Phys. Rev. Lett. **9**, 36 (1962).
- [16] Z. Maki, M. Nakagawa, and S. Sakata, Prog. Theo. Phys. **28**, 870 (1962).
- [17] C. V. Achar *et al.*, Phys. Lett. **18**, 196 (1965).
- [18] F. Reines *et al.*, Phys. Rev. Lett. **15**, 429 (1965).

- [19] R. Davis, Phys. Rev. Lett. **12**, 303 (1964).
- [20] R. Davis, D. S. Harmer, and K. C. Hoffman, Phys. Rev. Lett. **20**, 1205 (1968).
- [21] B. T. Cleveland *et al.*, Astrophys. J. **496**, 505 (1998).
- [22] B. Pontecorvo, Sov. Phys. JETP **26**, 984 (1968).
- [23] B. Pontecorvo and V. Gribov, Phys. Lett. **B26**, 493 (1969).
- [24] Gargamelle Neutrino, F. J. Hasert *et al.*, Phys. Lett. **B46**, 138 (1973).
- [25] T. J. Haines *et al.*, Phys. Rev. Lett. **67**, 1986 (1986).
- [26] KAMIOKANDE, K. S. Hirata *et al.*, Phys. Lett. **B205**, 416 (1988).
- [27] KAMIOKANDE-II, K. S. Hirata *et al.*, Phys. Rev. Lett. **63**, 16 (1989).
- [28] LEP COLLABORATIONS, **CERN-PH-EP/2005-051** (2005), hep-ex/0511027.
- [29] MINOS, P. Adamson *et al.*, Phys. Rev. **D73**, 072002 (2006), hep-ex/0512036.
- [30] M. Beuthe, Physics Reports **375**, 105 (2003).
- [31] F. Boehm and P. Vogel, *Physics of Massive Neutrinos, 2nd Edition* (Cambridge University Press, 1992).
- [32] The LEP Collaborations and the LEP Electroweak Working Group, J. Dress, 2001, as reported at the XX International Symposium on Lepton and Photon Interactions at High Energy, Rome, Italy.
- [33] M. Kobayashi and T. Maskawa, Prog. Theor. Phys. **49**, 652 (1973).
- [34] L. Wolfenstein, Phys. Rev. **D17**, 2369 (1978).
- [35] S. P. Mikheev and S. A. Yu., Sov. J. Nucl. Phys. **42**, 913 (1985).
- [36] E. Lisi, A. Marrone, and D. Montanino, Phys. Rev. Lett. **85**, 1166 (2000).
- [37] V. Barger, J. G. Learned, S. Pakvasa, and T. J. Weiler, Phys. Rev. Lett. **82**, 2640 (1999).
- [38] N. Arkani-Hamed, S. Dimopoulos, G. R. Dvali, and J. March-Russell, Phys. Rev. **D65**, 024032 (2002), hep-ph/9811448.

- [39] A. Strumia and F. Vissani, (2006), hep-ph/0606054.
- [40] R. N. Mohapatra and A. Y. Smirnov, Ann. Rev. Nucl. Part. Sci. **56**, 569 (2006), hep-ph/0603118.
- [41] V. Lobashev *et al.*, Phys. Lett. **B460**, 227 (1999).
- [42] C. Kraus *et al.*, Nuclear Physics **A721**, C533 (2003).
- [43] W.-M. Yao *et al.*, Journal of Physics G **33**, 1+ (2006).
- [44] KATRIN, G. Drexlin, Nucl. Phys. Proc. Suppl. **145**, 263 (2005).
- [45] K. Assamagan *et al.*, Phys. Rev. **D53**, 6065 (1996).
- [46] ALEPH, R. Barate *et al.*, Eur. Phys. J. **C2**, 395 (1998).
- [47] WMAP, D. N. Spergel *et al.*, Astrophys. J. Suppl. **170**, 377 (2007), astro-ph/0603449.
- [48] J. Bahcall, A. Serenelli, and S. Basu, Astrophys. J. **621**, L85 (2005).
- [49] J. Bahcall and A. Serenelli, Astrophys. J. **626**, 530 (2005).
- [50] S. Degl’Innocenti *et al.*, Astro. Phys. **7**, 77 (1997).
- [51] B. J. and M. H. Pinsonneault, Phys. Rev. Lett. **92**, 121301 (2004).
- [52] GALLEX, W. Hampel *et al.*, Phys. Lett. **B447**, 127.
- [53] SAGE, J. N. Abdurashitov *et al.*, J. Exp. Theor. Phys. **95**, 181 (2002), astro-ph/0204245.
- [54] GNO, M. Altmann *et al.*, Phys. Lett. **B616**, 174 (2005).
- [55] KAMIOKANDE, Y. Fukuda *et al.*, Phys. Rev. Lett. **77**, 1683 (1996).
- [56] Super-Kamiokande, J. Hosaka *et al.*, Phys. Rev. **D73**, 112001 (2006).
- [57] SNO Collaboration, B. Aharmim *et al.*, Phys. Rev. **C72**, 055502 (2005).
- [58] Super-Kamiokande, S. Fukuda *et al.*, Physics Letters **B539**, 179 (2002).
- [59] KamLAND, J. Shirai, Nucl. Phys. Proc. Suppl. **118**, 15 (2003).
- [60] KamLAND, T. Araki *et al.*, Phys. Rev. Lett. **94**, 081801 (2005), hep-ex/0406035.

- [61] M. Apollonio *et al.*, Eur. Phys. J. **C27**, 331 (2003), hep-ex/0301017.
- [62] T. Kajita and Y. Totsuka, Rev. Mod. Phys. **73**, 85 (2001).
- [63] T. K. Gaisser and M. Honda, Ann. Rev. Nucl. Part. Sci. **52**, 153 (2002), hep-ph/0203272.
- [64] G. D. Barr, T. K. Gaisser, P. Lipari, S. Robbins, and T. Stanev, Phys. Rev. **D70**, 023006 (2004), astro-ph/0403630.
- [65] Soudan 2, J. Thron *et al.*, Nucl. Instr. Meth. **A283**, 642 (1989).
- [66] Soudan 2, M. C. Sanchez *et al.*, Phys. Rev. **D68**, 113004 (2003), hep-ex/0307069.
- [67] Super-Kamiokande, Y. Ashie *et al.*, Phys. Rev. Lett. **93**, 101801 (2004), hep-ex/0404034.
- [68] Super-Kamiokande, Y. Ashie *et al.*, Phys. Rev. **D71**, 112005 (2005), hep-ex/0501064.
- [69] Super-Kamiokande, J. Hosaka *et al.*, Phys. Rev. **D74**, 032002 (2006), hep-ex/0604011.
- [70] Super-Kamiokande, K. Abe *et al.*, Phys. Rev. Lett. **97**, 171801 (2006), hep-ex/0607059.
- [71] K2K, S. H. Ahn *et al.*, Phys. Lett. **B511**, 178 (2001), hep-ex/0103001.
- [72] K2K, M. H. Ahn *et al.*, Phys. Rev. **D74**, 072003 (2006), hep-ex/0606032.
- [73] K2K, M. H. Ahn *et al.*, Phys. Rev. Lett. **93**, 051801 (2004), hep-ex/0402017.
- [74] MINOS, D. G. Micheal *et al.*, Phys. Rev. Lett. **97**, 191801 (2006), hep-ex/0607088.
- [75] K2K, E. Aliu *et al.*, Phys. Rev. Lett. **94**, 081802 (2005), hep-ex/0411038.
- [76] OPERA, R. Acquafredda *et al.*, (2006), hep-ex/0611023.
- [77] T2K, Y. Hayato, Nucl. Phys. Proc. Suppl. **143**, 269 (2005).
- [78] NOvA, D. S. Ayres *et al.*, (2004), hep-ex/0503053.
- [79] LSND, C. Athanassopoulos *et al.*, Nucl. Instr. Meth. **A388**, 149 (1997).
- [80] LSND, A. Aguilar *et al.*, Phys. Rev. D **64**, 112007 (2001).

- [81] KARMEN, B. Armbruster *et al.*, Phys. Rev. D **65**, 112001 (2002).
- [82] BUGEY, B. Achkar *et al.*, Nuclear Physics **B434**, 503 (1995).
- [83] A. Strumia, Physics Letters **B539**, 91 (2002).
- [84] L. B. G. Barenboim and J. Lykken, Physics Letters **B534**, 106 (2002).
- [85] MiniBooNE, A. A. Aguilar-Arevalo *et al.*, Phys. Rev. Lett. **98**, 231801 (2007), arXiv:0704.1500 [hep-ex].
- [86] M. Sorel, J. Phys. Conf. Ser. **39**, 320 (2006), hep-ex/0602018.
- [87] MiniBooNE, M. O. Wascko, Nucl. Phys. Proc. Suppl. **159**, 79 (2006), hep-ex/0602051.
- [88] NuMI, J. Hylen *et al.*, Internal NuMI Report (2003).
- [89] NuMI, J. Hylen *et al.*, Internal NuMI Report **NuMI-B-285** (1997).
- [90] MINOS, Internal NuMI Report **NuMI-L-337** (1998).
- [91] P. Adamson *et al.*, Nucl. Instrum. Meth. **A556**, 119 (2006).
- [92] MINOS, P. Adamson *et al.*, Phys. Rev. **D75**, 092003 (2007), hep-ex/0701045.
- [93] A. B. Pereira e Sousa, FERMILAB-THESIS-2005-67.
- [94] A. Blake, FERMILAB-THESIS-2005-77.
- [95] MINOS Collaboration, *The MINOS Off-line Software User's Manual: Release R1.25*, 2007.
- [96] C. L. F. Howcroft, FERMILAB-THESIS-2004-71.
- [97] C. Andreopoulos *et al.*, MINOS technical note **NuMI-ANA-994** (1999).
- [98] R. Hatcher, MINOS technical note **NuMI-SIM-478** (1999).
- [99] K. Lang *et al.*, Nucl. Instrum. Meth. **A545**, 852 (2005).
- [100] S. Avvakumov *et al.*, Nucl. Instrum. Meth. **A545**, 145 (2005).
- [101] J. Oliver *et al.*, IEEE Trans. Nucl. Sci. **51**, 2193 (2003).
- [102] A. S. T. Blake, MINOS technical note **MINOS-doc-1567** (2006).

- [103] R. Brun and F. Rademakers, Nucl. Inst. Meth **A389**, 81 (1997).
- [104] T. Nicholls, MINOS DAQ note (2001).
- [105] E. Gamma, R. Helm, R. Johnson, and J. Vlissides, *Design Patterns: Elements of Reusable Object Oriented Software* (Addison-Wesley, 1995).
- [106] E. Beall, MINOS technical note **NuMI-GEN-1089** (2005).
- [107] MINOS Collaboration, MINOS technical note **NuMI-GEN-1037** (2004).
- [108] J. P. Ochoa, MINOS technical note **MINOS-doc-1156** (2005).
- [109] D. G. Michael *et al.*, MINOS technical note **MINOS-doc-1967** (2006), To be submitted to Nucl. Instr. Meth. A.
- [110] N. Tagg *et al.*, Nucl. Instrum. Meth. **A539**, 668678 (2005).
- [111] C. B. Smith, FERMILAB-THESIS-2002-58.
- [112] P. A. Symes, FERMILAB-THESIS-2005-76.
- [113] J. J. Hartnell, FERMILAB-THESIS-2005-51.
- [114] P. L. Vahle, FERMILAB-THESIS-2004-35.
- [115] M. A. Kordosky, FERMILAB-THESIS-2004-34.
- [116] The BESS collaboration, J. Wang *et al.*, Astrophys. J. **564**, 244 (2002).
- [117] The CAPRICE collaboration, M. Boezio *et al.*, Astrophys. J. **518**, 457 (1999).
- [118] The IMAX collaboration, W. Menn *et al.*, Astrophys. J. **533**, 281 (2000).
- [119] The BESS collaboration, T. Sanuki *et al.*, Astrophys. J. **545**, 1135 (2000).
- [120] The AMS collaboration, A. J. et al, Phys. Lett. **B490**, 27 (2000).
- [121] J. A. Simpson, Annual Review of Nuclear and Particle Science **33**, 323 (1983).
- [122] M. Takeda *et al.*, Phys. Rev. Lett. **81**, 1163 (1998), astro-ph/9807193.
- [123] A. V. Glushkov *et al.*, Energy spectrum of primary cosmic rays in the energy region of 10^{17} to 10^{20} eV by Yakutsk Array data, in *Proc. 28th ICRC (international cosmic ray conference)*, pp. 389–392, 2003.

- [124] HIRES, D. J. Bird *et al.*, *Astrophys. J.* **424**, 491 (1994).
- [125] High Resolution Fly's Eye, T. Abu-Zayyad *et al.*, *Astropart. Phys.* **23**, 157 (2005), astro-ph/0208301.
- [126] J. R. Hoerandel, *Astropart. Phys.* **19**, 193 (2003), astro-ph/0210453.
- [127] G. Sigl, *Annals of Physics* **303**, 117 (2003).
- [128] J. Kremer *et al.*, *Phys. Rev. Lett.* **83**, 4241 (1999).
- [129] M. Motoki *et al.*, *Astropart. Phys.* **19**, 113 (2003), astro-ph/0205344.
- [130] L3, P. Achard *et al.*, *Phys. Lett.* **B598**, 15 (2004), hep-ex/0408114.
- [131] T. Gaisser, *Cosmic rays and particle physics* (Cambridge University Press, 1990).
- [132] S. M. Kasahara, UMI-97-21625.
- [133] B. J. Rebel, FERMILAB-THESIS-2004-33.
- [134] AMS, J. Alcaraz *et al.*, *Phys. Lett.* **B490**, 27 (2000).
- [135] T. Sanuki *et al.*, *Astrophys. J.* **545**, 1135 (2000), astro-ph/0002481.
- [136] R. Engel, T. K. Gaisser, T. Stanev, and P. Lipari, Hadronic interactions for atmospheric cascades, in *Proceedings of the 27th International Cosmic Ray Conference (ICRC 2001)*, Hamburg, Germany, 7-15 Aug 2001.
- [137] Y. Cho *et al.*, *Phys. Rev.* **D4**, 1967 (1971).
- [138] T. Eichten *et al.*, *Nucl. Phys.* **B44**, 333 (1972).
- [139] E-802, T. Abbott *et al.*, *Phys. Rev.* **D45**, 3906 (1992).
- [140] HARP, G. Prior, *Nucl. Phys.* **A752**, 24 (2005).
- [141] I. Chemakin *et al.*, *Phys. Rev.* **C65**, 024904 (2002), nucl-ex/0108007.
- [142] MIPP, H. Meyer, *Nucl. Phys. Proc. Suppl.* **142**, 453 (2005).
- [143] NA49, G. Barr, Hadroproduction in proton carbon collisions at the na49 experiment, in *Proceedings of the 28th International Cosmic Ray Conferences (ICRC 2003)*, Tsukuba, Japan, 31 Jul - 7 Aug 2003.

- [144] G. Barr, Nucl. Phys. Proc. Suppl. **143**, 89 (2005).
- [145] T. K. Gaisser, Phys. Scripta **T121**, 51 (2005), astro-ph/0502380.
- [146] G. D. Barr, T. K. Gaisser, S. Robbins, and T. Stanev, Phys. Rev. **D74**, 094009 (2006), astro-ph/0611266.
- [147] Space physics data system: Neutron monitor datasets.
- [148] H. Gallagher, Nucl. Phys. Proc. Suppl. **112**, 188 (2002).
- [149] R. Merenyi *et al.*, Phys. Rev. **D45**, 743 (1992).
- [150] G. P. Zeller, Low energy neutrino cross sections: Comparison of various Monte Carlo predictions to experimental data, in *Proceedings of the 2nd International Workshop on Neutrino - Nucleus Interactions in the Few GeV Region (NUINT 02)*, 2003, hep-ex/0312061, Irvine, California, 12-15 Dec 2002.
- [151] R. Hatcher, MINOS technical note **NuMI-SIM-479** (1999).
- [152] S. Giani *et al.*, *GEANT detector description and simulation tool*, 1994, CERN Program Library Long Writeup.
- [153] N. Tagg, *DetSim and PhotonTransport: MINOS detector and front-end simulation*, 2007.
- [154] J. Birks, *Theory and practice of scintillation counting* (Pergamon Press, 1964).
- [155] M. Kordosky *et al.*, MINOS technical note **NuMI-SIM-1044** (2004).
- [156] C. Zeitnitz and T. Gabriel, Nucl. Instr. Meth. **A349**, 106 (1994).
- [157] C. Zeitnitz, GCALOR simulation package.
- [158] C. P. Ward, MINOS technical note **NuMI-SIM,ATM-NU-1085** (2005).
- [159] M. A. Thomson, MINOS technical note **NuMI-2535** (2005).
- [160] J. S. Marshall, MINOS technical note **MINOS-doc-2104** (2006).
- [161] J. S. Marshall, MINOS technical note **MINOS-doc-1248** (2005).
- [162] G. Welch and G. Bishop, *An Introduction to the Kalman Filter* (Technical Report TR 95-041, University of North Carolina, Department of Computer Science, 1995).

- [163] M. D. Messier, MINOS technical note **MINOS-doc-2299** (2006).
- [164] Bayes, T. and Price, R., Philosophical Transactions (1683-1775) **53**, 370 (1763).
- [165] P.-S. Laplace, Statistical Science **1**, 364 (1986), originally published in 1774, Savants Étranges 6:621-656.
- [166] F. James, *MINUIT: Function Minimization and Error Analysis*, 1994, CERN Program Library Long Writeup.
- [167] B. Speakman, MINOS technical note **MINOS-doc-2668** (2007).
- [168] V. Agrawal, T. K. Gaisser, P. Lipari, and T. Stanev, Phys. Rev. D **53**, 1314 (1996).
- [169] G. Battistoni, A. Ferrari, T. Montaruli, and P. R. Sala, Astroparticle Physics **19**, 269 (2003).
- [170] Soudan-2, W. W. M. Allison *et al.*, Phys. Rev. **D72**, 052005 (2005), hep-ex/0507068.
- [171] M. Thomson, MINOS technical note **NuMI-TRANS-GEN-1095** (2005).

List of Figures

2.1	Neutrino propagation	7
2.2	Normal and inverted neutrino spectra	15
2.3	The p-p fusion chain	20
2.4	SNO Phase II flux results	26
2.5	SNO neutrino oscillation confidence limits	27
2.6	SNO global solar neutrino oscillation confidence limits	28
2.7	The KamLAND experiment	29
2.8	Prompt energy spectrum of e^+ (from $\bar{\nu}_e$) in KamLAND	30
2.9	KamLAND neutrino oscillation confidence limits.	31
2.10	KamLAND L_0/E Plot	31
2.11	CHOOZ results	33
2.12	Super-Kamiokande L/E analysis	36
2.13	Super-Kamiokande L/E analysis CL allowed regions	37
2.14	Comparison of SK combined and L/E analyses	38
2.15	K2K reconstructed E_ν distribution	41
2.16	K2K allowed regions	42
2.17	MINOS beam analysis reconstructed Far Detector E_ν distribution	43
2.18	MINOS beam analysis allowed regions	44

2.19	Short-baseline Accelerator Neutrino Experiment Results	47
3.1	NuMI beamline	51
3.2	A NuMI horn	52
3.3	Near Detector yz plan	54
3.4	Near Detector xy plan	55
4.1	Photograph of the MINOS Far Detector	60
4.2	MINOS Far Detector XY view	62
4.3	MINOS Far Detector scintillator strips	64
4.4	MINOS Far Detector optical readout	66
4.5	Far Detector Front-End readout scheme.	68
4.6	Far Detector DAQ system architecture	70
4.7	MINOS DAQ state machine	75
4.8	Coil current in each Supermodule	80
4.9	Dead chips in each Supermodule	81
4.10	Dead chip distributions	82
4.11	Busy chip distributions	84
4.12	Tagging efficiencies for individual shield planks	87
4.13	Distribution of tagging efficiencies for shield planks	88
5.1	Electronics Swaps 1: Initial <i>TimePeriods</i>	93
5.2	Electronics Swaps 2: Sorted <i>TimePeriods</i>	94
5.3	Electronics Swaps 3: ΔcT_{VAchip} values in example channel	95
5.4	Electronics Swaps 3: ΔcT_{VAchip} values in opposite channel	96
5.5	Electronics Swaps 4: Timing Jumps Identified	96

5.6	Electronics Swaps 5: Timing Jumps Removed	97
5.7	ΔT_{EW} distributions, before and after calibration versus channel	100
5.8	ΔT_{EW} distributions, before and after calibration versus channel	101
5.9	Distribution of ΔT_{EW} before calibration	102
5.10	Distribution of ΔT_{EW} after calibration	102
5.11	Gaussian fit to calibrated ΔT_{EW} distribution	103
5.12	Distribution of RMS deviations for muon timing fits.	104
5.13	Bethe-Bloch equation	109
6.1	Primary cosmic-ray spectrum below 10^{12} eV	113
6.2	The “all-particle” spectrum of cosmic-ray interactions above 10^{12} eV . . .	114
6.3	Simulated atmospheric μ fluxes	120
6.4	CLIMAX data: Long term variation in mean daily atmospheric neutron rates	122
6.5	CLIMAX data: Fit to mean monthly atmospheric neutron rates	122
6.6	Cross-section for ν_μ CC interactions on iron as a function of neutrino energy	124
6.7	Cross-section for $\bar{\nu}_\mu$ CC interactions on iron as a function of neutrino energy	125
6.8	Comparison of timing RMS_{Down} for data and unsmeared Monte Carlo . .	128
6.9	Comparison of timing RMS_{Down} for data and smeared Monte Carlo . . .	129
6.10	Total pulse-height from double-ended track strips	130
6.11	Mean PH per strip end vs WLS fibre plot	131
6.12	Comparison of the number of planes in tracks between data and cosmic muon Monte Carlo	132
6.13	Comparison of the reconstructed zenith angle distribution for data and cosmic muon Monte Carlo	133

6.14	Comparison of the reconstructed azimuthal angle distribution for data and cosmic muon Monte Carlo	134
7.1	Demultiplexing hits in the MINOS Far Detector	138
7.2	<i>Simple associations</i> between <i>Triples</i> in the Track Finder	140
7.3	Track reconstruction efficiency versus planes cross by muon	142
7.4	Track reconstruction efficiency versus true muon momentum	142
7.5	Track Fitter: Improved track strip identification	144
7.6	Muon momentum resolution versus true muon momentum for fully contained atmospheric neutrino events	146
7.7	Muon momentum resolution versus true muon momentum for partially contained atmospheric neutrino events	146
7.8	Shower reconstruction efficiency versus true E_{shw}	148
7.9	Reconstructed vertex shower energy versus true E_{shw}	151
7.10	$\sigma_{E_{shw}}/E_{true}$ versus true E_{shw}	152
8.1	Overview of Selection	154
8.2	Pre-Selection: Distribution of time from trigger time to the nearest spill .	156
8.3	Pre-Selection: track <i>UVasymmetry</i>	158
8.4	Pre-Selection: Δ (Linear Fit RMS)	159
8.5	Cosmic muon entering the detector within a large shower	160
8.6	Cartoon showing timing fits to direction hypotheses	162
8.7	Pre-Selection: Delta RMS	163
8.8	Pre-Selection: ‘Distance To Edge’ definition	164
8.9	Pre-Selection: ‘Distance To Edge’ cut	164
8.10	Definition of the Trace of a track	165
8.11	Pre-Selection: Trace of highest end of track	166

8.12 Upward-Going Selection: Fraction of non-fiducial track pulse height . . .	167
8.13 Upward-Going Selection: Number of <i>track-like</i> planes in tracks	168
8.14 Typical background topologies	169
8.15 Upward-Going Selection: Q_{max} distribution	170
8.16 Upward-Going Selection: $\frac{1}{\beta_{up}}$	171
8.17 Upward-Going Selection: RMS_{Up} distribution	171
8.18 Upward-Going Selection: RMS_{Down} distribution	172
8.19 Upward-Going Selection: $RMS_{Up} - RMS_{Down}$ distribution	172
8.20 Contained-Highest-End Selection: z component of High End Trace	176
8.21 Contained-Highest-End Selection: z component of High End Trace using digits	177
8.22 Contained-Highest-End Selection: $\langle \Delta_T \rangle$	178
8.23 Contained-Highest-End Selection: $\langle \Delta_T^2 \rangle^{\frac{1}{2}}$	178
8.24 Contained-Highest-End Selection: Distribution of R_{max}	179
8.25 Contained-Highest-End Selection: Distribution of Q_{max}	179
8.26 Contained-Highest-End Selection: Q_{max} vs Track y direction cosine . . .	180
8.27 Contained-Highest-End Selection: Q_{max} vs Track z direction cosine . . .	181
8.28 Contained-Highest-End Selection: ‘Distance To Edge’ cut on lowest end of track	182
8.29 Contained-Highest-End Selection: Trace Cut on lowest end of track . . .	183
8.30 Partially Contained Downward-Going Selection: $\frac{1}{\beta_{up}}$	184
8.31 Distribution of the smallest $T_{shield} - T_{det}$ values for stopping muon events	185
8.32 Distribution of the smallest $T_{shield} - T_{det}$ values for FC and PCDN events	186
8.33 Selected Data Events: Interaction points in the xy plane	190
8.34 Selected Events: Interaction point distribution along the z -axis	191

8.35	Selected Events: Distribution of neutrino event rate over time	191
8.36	Selection efficiency vs E_ν for $\nu/\bar{\nu}$ CC events	192
8.37	Selection efficiency vs muon zenith angle for $\nu/\bar{\nu}$ CC events	193
8.38	Distribution of track planes for selected events	194
8.39	Distribution of track-like track planes for selected events	195
8.40	Distribution of track ranges for selected events	196
8.41	Distribution of $\text{RMS}_{\text{Vtx}} - \text{RMS}_{\text{End}}$ for selected events	197
8.42	Distribution of $(Q/p)/\sigma_{(Q/p)}$ for high resolution selected events	197
8.43	Fraction of split stopping muon tracks where the reconstructed charge agrees with that of the whole track	200
8.44	Fraction of split through going muon tracks where the reconstructed charge agrees with that of the whole track	201
8.45	Difference between fraction of stopping muon split tracks with matching charges in data and Monte Carlo	202
8.46	Difference between fraction of through going muon split tracks with match- ing charges in data and Monte Carlo	203
9.1	The up/down ratio as a function of Δm_{32}^2	206
9.2	Distribution of E_ν for high resolution events	207
9.3	Distribution of $\cos \Theta_{\text{zenith}}$ for high resolution events	207
9.4	The $\bar{\nu}_\mu/\nu_\mu$ ratio as a function of $\overline{\Delta m_{32}^2}$	208
9.5	Distribution of $\cos \Theta_{\text{zenith}}$ for high resolution events identified as ν_μ events	209
9.6	Distribution of $\cos \Theta_{\text{zenith}}$ for high resolution events identified as $\bar{\nu}_\mu$ events	210
9.7	Reconstructed $\log(L/E)$ distribution for high resolution events	212
9.8	Example y and W^2 PDFs	215
9.9	Pull distribution for p_μ from curvature	216

9.10	Example $\log\left(\frac{L}{E}\right)$ PDF: Downward-going muon	218
9.11	Example $\log\left(\frac{L}{E}\right)$ PDF: Upward-going muon	218
9.12	Example $\log\left(\frac{L}{E}\right)$ PDF: muon near horizon	219
9.13	Distribution of $\sigma_{\log(L/E)}$ for high resolution data and Monte Carlo events	219
9.14	$\log\left(\frac{L}{E}\right)_{reco}$ distributions for high resolution Monte Carlo events binned in $\sigma_{\log(L/E)}$	220
9.15	Oscillated/Unoscillated Monte Carlo events vs $\log\left(\frac{L}{E}\right)_{reco}$ binned in $\sigma_{\log(L/E)}$	221
9.16	Smoothing $\log(L/E)$ distributions	225
9.17	Confidence Level contours for combined $\nu_\mu/\bar{\nu}_\mu$ fit	226
9.18	Confidence Level contours for combined $\nu_\mu/\bar{\nu}_\mu$ fit	227
9.19	Slice through likelihood surface at $\sin^2 2\theta_{23} = 1.0$ for combined $\nu_\mu/\bar{\nu}_\mu$ fit .	227
9.20	$\left(\frac{L}{E}\right)_{reco}$ distributions for high resolution Monte Carlo and data events binned in $\sigma_{(L/E)}$	228
9.21	Expected Sensitivity of combined $\nu_\mu/\bar{\nu}_\mu$ fit for 12.23 ktyr exposure	230
9.22	Slice through expected sensitivity surface at $\sin^2 2\theta_{23} = 1.0$ of combined $\nu_\mu/\bar{\nu}_\mu$ fit for 12.23 ktyr exposure	230
9.23	Confidence Level contours for a ν_μ dominated fit	231
9.24	Slice through likelihood surface at $\sin^2 2\theta_{23} = 1.0$ for a ν_μ dominated fit .	232
9.25	Confidence Level contours for an $\bar{\nu}_\mu$ dominated fit	233
9.26	Slice through likelihood surface at $\sin^2 2\theta_{23} = 1.0$ for an $\bar{\nu}_\mu$ dominated fit	233
9.27	Confidence Level contours for separated $\nu/\bar{\nu}$ fit	236
9.28	Projected Sensitivity of combined $\nu_\mu/\bar{\nu}_\mu$ fit for 25 ktyr exposure	237
9.29	Slice through projected sensitivity surface at $\sin^2 2\theta_{23} = 1.0$ of combined $\nu_\mu/\bar{\nu}_\mu$ fit for 25 ktyr exposure	238

List of Tables

2.1	Experimental values of oscillation parameters	49
4.1	Primary MINOS DAQ States	76
6.1	Monte Carlo samples used in this thesis	135
8.1	Summary of the Upward-Going Selection	175
8.2	Summary of the Contained-Highest-End Selection	188
8.3	Summary of Event Selection	189
8.4	Separation of selected events into high and low resolution samples	194
8.5	Separation of selected events into ν_μ and $\bar{\nu}_\mu$ samples	198
8.6	Separation of selected events by identified charge and direction	198
9.1	Values of nuisance parameters at the best fit point for the combined $\nu_\mu/\bar{\nu}_\mu$ fit	229
9.2	Values of nuisance parameters at the best fit point for the ν_μ dominated fit	232
9.3	Values of nuisance parameters at the best fit point for the $\bar{\nu}_\mu$ dominated fit	234
A.1	Complete List of Selected Data Events	1