A study of muon neutrino disappearance with the MINOS detectors and the NuMI neutrino beam

John Stuart Marshall
Fitzwilliam 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 an analysis of ν_μ disappearance with the MINOS experiment, which studies the neutrino beam produced by the NuMI facility at Fermi National Accelerator Laboratory. The rates and energy spectra of charged current ν_μ interactions are measured in two similar detectors, located at distances of 1 km and 735 km along the NuMI beamline. The Near Detector provides accurate measurements of the initial beam composition and energy, whilst the Far Detector is sensitive to the effects of neutrino oscillations. The analysis uses data collected between May 2005 and March 2007, corresponding to an exposure of 2.5×10^{20} protons on target.

As part of the analysis, sophisticated software was developed to identify muon tracks in the detectors and to reconstruct muon kinematics. Events with reconstructed tracks were then analysed using a multivariate technique to efficiently isolate a pure sample of charged current ν_μ events. An extrapolation method was also developed, which produces accurate predictions of the Far Detector neutrino energy spectrum, based on data collected at the Near Detector. Finally, several techniques to improve the sensitivity of an oscillation measurement were implemented, and a full study of the systematic uncertainties was performed.

Extrapolating from observations at the Near Detector, 733 ± 29 Far Detector events were expected in the absence of oscillations, but only 563 events were observed. This deficit in event rate corresponds to a significance of 4.3 standard deviations. The deficit is energy dependent and clear distortion of the Far Detector energy spectrum is observed. A maximum likelihood analysis, which fully accounts for systematic uncertainties, is used to determine the allowed regions for the oscillation parameters and identifies the best fit values as $\Delta m^2_{32} = 2.29^{+0.14}_{-0.14} \times 10^{-3}\text{eV}^2$ and $\sin^2 2\theta_{23} > 0.953$ (68% confidence level). The models of neutrino decoherence and decay are disfavoured at the 5.0σ and 3.2σ levels respectively, whilst the no oscillation model is excluded at the 9.4σ level.

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 Marshall

Acknowledgements

I would first and foremost like to thank my family and my girlfriend, Jean Comrie, for their constant help and support, without which this thesis would not have been possible. In the Cambridge HEP group, I would like to thank Mark Thomson for his excellent supervision and guidance, and also Andy Blake for answering so many of my questions and providing expert advice. Special mention must go to John Chapman and Andy Culling, who made our office a fun place to be, and I am also grateful to Pat and David Ward for their comments and advice. The Cambridge HEP group is a fantastic working environment and I would also like to thank the many other members of the group for helping to make my PhD course an enjoyable experience.

The analysis in this thesis was only possible because of the support of the MINOS collaboration, for which I am grateful. In particular, I would like to thank Jim Musser, who embraced my new track reconstruction software and helped introduce it to the collaboration, and Peter Litchfield, who helped to get me started in the charged current analysis. I would also like to acknowledge the many other members of the reconstruction and charged current working groups for their helpful suggestions. Funding for this analysis has been provided by the UK STFC and the Leathersellers' charitable fund and I would like to express my thanks for their generous support.

Finally, I would like to mention my grandmother, Brenda Barratt, who passed away in December 2006. She staunchly supported me throughout her life, and the memory of her selfless nature and positive outlook provides me with a constant source of inspiration. She is greatly missed.

Contents

1 Introduction **1**

2 Neutrino Physics **4**

2.1 Theory of Neutrino Oscillations 5

2.1.1 Alternative Models of Neutrino Propagation 9

Neutrino Decoherence 9

Neutrino Decay 10

2.2 Neutrino Oscillation Experiments 10

2.2.1 Long Range Oscillations 10

Radiochemical Experiments 11

SNO 12

KamLAND 14

2.2.2 Short Range Oscillations 18

Super-Kamiokande 20

K2K 24

2.2.3 Subdominant Short Range Oscillations 24

CHOOZ 26

2.2.4 Summary 26

3 The MINOS Experiment **28**

3.1 The NuMI Beam 30

3.2 The MINOS Detectors 33

3.2.1 Far Detector 36

3.2.2 Near Detector 37

3.3 Detector Calibration 40

3.4 Monte Carlo Simulation 42

3.5 Reconstruction Software 43

4 Track Finding **46**
4.1 How the Algorithm Works . . . 47
4.1.1 The General Aim . . . 47
4.1.2 Hit and Cluster Formation . . . 49
4.1.3 Triplet Formation . . . 49
4.1.4 Simple Associations . . . 50
4.1.5 Preferred Associations . . . 52
4.1.6 Matched Associations . . . 52
4.1.7 Seed Segment Selection . . . 53
4.1.8 2D Track Formation . . . 54
4.1.9 3D Track Formation . . . 55
4.1.10 The Final Track . . . 57
4.1.11 Track Properties . . . 58
4.2 Near Detector Modifications . . . 58
4.3 Results . . . 61
4.3.1 Far Detector Beam MC . . . 62
4.3.2 Near Detector Beam MC . . . 64

5 Track Fitting **67**
5.1 How the Algorithm Works . . . 68
5.1.1 The Kalman Filter . . . 68
5.1.2 The Propagator Matrix . . . 70
5.1.3 The Noise Covariance Matrix . . . 71
5.1.4 Interface to the Kalman Filter . . . 73
5.1.5 Track Strip Identification in Vertex Showers . . . 75
5.1.6 Calculation of Track Properties . . . 76
5.2 Far Detector Results . . . 78
5.2.1 Far Detector Beam MC . . . 79
5.2.2 Comparison with Previous Reconstruction . . . 83
5.2.3 Far Detector Cosmic Muons . . . 85
5.3 Near Detector Spectrometer Demultiplexing . . . 86
5.4 Near Detector Results . . . 90

6 Event Selection **96**
6.1 Data Quality . . . 97
6.1.1 Detector Status . . . 97
6.1.2 Beam Monitoring . . . 98

6.2 The Analysis Dataset . . . 99
6.3 Selection of Charged Current ν_μ Interactions . . . 101
6.3.1 Triggers . . . 101
6.3.2 Muon Track and Detector Fiducial Volume . . . 102
6.3.3 Track Properties . . . 105
6.4 Removal of Neutral Current Background . . . 107
6.5 Examination of Near Detector Data . . . 113
6.6 Neutral Current Background Systematic . . . 117
6.6.1 Muon Removal Technique . . . 118
6.6.2 Muon Removal Results . . . 121
6.6.3 Final Results . . . 128

7 Extrapolation and Fitting Far Detector Data 130
7.1 The Need for an Extrapolation . . . 131
7.2 How the Extrapolation Works . . . 133
7.2.1 Reweighting the Near Detector MC . . . 133
7.2.2 Gaussian Matching of Events . . . 135
7.2.3 Flattening the MC . . . 137
7.2.4 Transport to the Far Detector . . . 138
7.2.5 Event Selection and Efficiencies . . . 139
7.2.6 Final Far Detector Prediction . . . 140
7.2.7 Basic Oscillation Fit . . . 144
7.3 Performance Tests . . . 145
7.3.1 Simple MC Tests . . . 145
7.3.2 Mock Data Challenges . . . 149
MDC II . . . 149
MDC III . . . 150
7.4 Systematic Uncertainties . . . 150
7.4.1 Approach and Results . . . 153
7.4.2 Cross-section Distortions . . . 157
7.4.3 Enhancement of NC . . . 160
7.4.4 Extrapolation without PID cut . . . 161
7.5 Full Fit to Oscillation and Nuisance Parameters . . . 162
7.6 Data Predictions . . . 165

8 Data/MC Comparisons **168**
8.1 Near Detector . . . 169
8.1.1 PID related quantities . . . 169
8.1.2 Track related quantities . . . 171
8.1.3 Shower related quantities . . . 176
8.1.4 Neutrino energy . . . 178
8.2 Far Detector . . . 178
8.2.1 PID related quantities . . . 179
8.2.2 Track related quantities . . . 179
8.2.3 Shower related quantities . . . 181
8.3 Conclusions . . . 183

9 Sensitivity Improvements **186**
9.1 Binning in the Oscillation Fit . . . 187
9.2 Deweighted Shower Energy . . . 188
9.2.1 Detector Response to Hadronic Showers . . . 188
9.2.2 Linear Calibration . . . 189
9.2.3 Deweighted Calibration . . . 190
9.2.4 Results . . . 191
9.3 Use of Event Energy Resolution . . . 195
9.4 Removal of Events in Coil Hole . . . 198
9.5 Pulse Height Omitted by Shower Finder . . . 202
9.6 Collective Improvement . . . 204
9.7 Associated Systematics . . . 204
9.8 Implementation . . . 209

10 Results **212**
10.1 Selected Events . . . 213
10.2 Results without Sensitivity Improvements . . . 214
10.3 Results with Sensitivity Improvements . . . 215
10.4 Comparison with Other Results . . . 224
10.5 Decoherence and Decay Models . . . 224

11 Summary and Outlook **227**

Bibliography **1**

List of Figures **9**

List of Tables **15**

"There are two possible outcomes:
if the result confirms the hypothesis, then you've made a measurement.
If the result is contrary to the hypothesis, then you've made a discovery."

— Enrico Fermi (1901-1954)

Chapter 1

Introduction

In December 1930, Wolfgang Pauli wrote an open letter to a meeting of nuclear physicists in Tübingen, Germany[1]. In this letter, he proposed a 'desperate remedy' to save the law of conservation of energy. The letter addressed Chadwick's 1914 observation of the continuous energy spectrum of electrons emitted in β-decay[2]. If, as believed, β-decay was a two body process, then energy conservation unavoidably predicted a monochromatic electron spectrum. Pauli's bold solution was to postulate the existence of a new electrically neutral particle, which he called a 'neutron'. This neutron must be emitted alongside the electron in β-decay, avoiding experimental detection whilst carrying energy 'such that the sum of the energies of neutron and electron is constant'. In addition to being electrically neutral, experimental constraints meant that the neutron needed to be spin-$\frac{1}{2}$, light ('not larger than 0.01 proton mass') and very penetrating ('perhaps a ten times larger ability to get through material than a gamma ray').

A little over a year later, Chadwick announced the possible existence of a neutron[3]. However, scattering data showed that this clearly had too great a mass to be the same particle predicted by Pauli; Chadwick had discovered what is now considered as the modern neutron. Pauli's ideas were not neglected and were taken up by Enrico Fermi, who renamed Pauli's new particle the 'neutrino' (or 'little neutral one') and incorporated it into his groundbreaking 1934 quantitative theory of nuclear β-decay [4].

Fermi's theory accounted for all the observed properties of β-decay and so provided strong evidence for the existence of neutrinos. However, the tiny neutrino interaction cross-sections meant that it was not until the experiments of Fred Reines and Clyde

Cowan in the 1950s that direct detection of the neutrino became a possibility. Reines and Cowan attempted to observe the flux of neutrinos from nuclear reactors, but their first attempts suffered from large cosmic ray backgrounds and provided tantalising but insufficient evidence of neutrino interactions[5]. After redesigning their experiment and constructing an underground detector near the Savannah River reactor in South Carolina, they finally demonstrated the existence of the neutrino in 1956[6].

After the first direct observation of the neutrino, a number of important experimental results quickly followed. In 1958, Goldhaber *et al.* established that the neutrino has negative helicity[7], whilst in 1959 Davis *et al.* suggested that the neutrino and anti-neutrino differ in their reactions with nucleons[8]. Three years later, Lederman *et al.* demonstrated that the neutrinos produced in pion decays were different to those involved in β-decay[9]. This was the discovery of a second 'flavour' of neutrino; the muon neutrino, as opposed to the electron neutrino postulated by Pauli.

The 1974 discovery of the J/Ψ[10, 11] confirmed the existence of the charm quark and so completed the discovery of two 'generations' of quarks and leptons. Each generation contained particle and anti-particle versions of two quarks, a charged lepton and an associated neutrino. The generations differed only by the masses of their constituents; all other quantum numbers were unchanged. The identification of a third generation of charged lepton, the tau in 1975[12], therefore implied the existence of an associated tau neutrino. This was discovered by the DONUT collaboration[13] in 2000. Precise measurements of the width of the Z resonance by the LEP experiments at CERN provide strong evidence that there are only three light active neutrino flavours[14].

In the Standard Model of particle physics, the three flavours of neutrino are assumed to be massless. All attempts to directly measure the neutrino masses, by examining the kinematics of decays producing neutrinos, have produced only upper limits[15–18]. However, there is now compelling evidence that neutrinos undergo periodic transitions between flavours as they propagate. This phenomenon is known as 'neutrino oscillations', and was first suggested by Pontecorvo in 1957[19]. Importantly, if neutrinos oscillate, it must mean that neutrino flavours are mixtures of the neutrino mass eigenstates and that at least one of the neutrinos has a non-zero mass.

Evidence for oscillations has been observed in neutrinos produced in the Sun[20], in the Earth's atmosphere[21, 22], in nuclear reactors[23] and by particle accelerator beams[24]. The MINOS experiment is currently searching for evidence of oscillations in an intense beam of muon neutrinos manufactured by the NuMI facility at the Fermi National Accelerator Laboratory in Illinois, USA. The experiment observes beam neutrino interactions in two functionally identical detectors, located at distances of 1 km

and 735 km from the neutrino production target. The Near Detector measures the initial composition and energy spectrum of the neutrino beam, whilst observations at the Far Detector are sensitive to the effects of oscillations. MINOS aims to confirm the existence of neutrino oscillations and to make high precision measurements of the underlying parameters. This thesis presents an analysis of recent MINOS beam neutrino data.

In Chapter 2, the theory governing neutrino oscillations is described and current experimental evidence is reviewed, setting the context for the MINOS experiment. Chapter 3 provides details of the NuMI neutrino beam and the MINOS detectors. It also describes the detector calibration, Monte Carlo (MC) simulation and reconstruction software. Chapters 4 and 5 provide a comprehensive description of the software developed to reconstruct muon tracks in the MINOS detectors. Chapter 4 describes the first stage of track reconstruction, which uses pattern recognition to identify the basic features of the muon trajectory (a 'seed track'). Chapter 5 then describes the track fitting stage, in which information from the seed track is used alongside knowledge of muon propagation and energy loss in order to fully reconstruct the muon kinematics.

Chapter 6 begins the description of the oscillation analysis, detailing the event selection procedures required to efficiently identify a high quality sample of charged current muon neutrino interactions in the MINOS detectors. The analysis dataset is specified and details of a data-driven technique to improve MC modelling of the neutral current background are provided. Chapter 7 contains a detailed description of the extrapolation method that is used to obtain accurate predictions of the Far Detector neutrino energy spectrum, based on the observations of beam neutrino interactions in the Near Detector. The most important systematic uncertainties in the analysis are identified and an oscillation fit that fully accounts for these uncertainties is described.

Chapter 8 presents a detailed comparison of events in the MINOS data and MC simulation, making use of the MC event weights calculated by the extrapolation method. Chapter 9 provides details of techniques to improve the sensitivity of the analysis to the oscillation parameters. These techniques help to obtain the most accurate measurement of the parameters from the analysis dataset. Chapter 10 brings together the ideas and techniques developed throughout the thesis in order to extract an optimised measurement of the oscillation parameters from the MINOS data. Results are also presented for alternative 'exotic' models of neutrino propagation. The final Chapter gives a summary of the thesis and presents its conclusions, discussing future results from MINOS and related experiments.

Chapter 2

Neutrino Physics

Neutrino physics is currently best described by the electroweak interaction, which is incorporated in the Standard Model of particle physics. In this model there are three generations (or flavours) of leptons, each represented by a pair of particles called a weak doublet. The doublets each contain a massless neutrino and an associated charged lepton: the electron neutrino (ν_e) and electron are paired together, as are the muon neutrino (ν_μ) and muon and, finally, the tau neutrino (ν_τ) and tau lepton. The interactions of the leptons are confined within these doublets, so that the number of leptons of each flavour is conserved.

In the Standard Model, neutrinos couple only to the weak force, so their interactions are mediated by the massive $W^\pm$ and Z bosons. In charged current (CC) interactions, the neutrino is converted into its partner charged lepton via exchange of a $W^\pm$ boson with a quark or lepton. Observation of the partner charged lepton allows the flavour of the incoming neutrino to be identified. In neutral current (NC) interactions, the neutrino is scattered via exchange of a Z boson and the neutrino flavour cannot be determined. The cross-sections for neutrino interactions with matter are extremely small, due to the high masses of these bosons.

Despite the small cross-sections, increasingly sophisticated techniques have been developed to observe neutrino interactions and to reconstruct their kinematic properties. Such observations have challenged the Standard Model assumption that neutrinos are massless. In particular, there is compelling evidence that neutrinos undergo periodic transitions between flavours as they propagate through space; a phenomenon known as

'neutrino oscillations'. If neutrinos oscillate, then neutrino flavours must be mixtures of the neutrino mass eigenstates and at least one of the neutrinos must have a non-zero mass.

In this Chapter, the theory governing neutrino oscillations is described, and details of alternative models of neutrino propagation are provided. The current experimental evidence for neutrino oscillations is then reviewed, setting the context for the MINOS experiment, which aims to confirm the existence of oscillations and to make precise measurements of the underlying parameters.

2.1 Theory of Neutrino Oscillations

If neutrinos have masses, there will exist a set of mass eigenstates $|\nu_i\rangle\,(i = 1, 2, 3)$. The mass eigenstates diagonalise the free Hamiltonian and so describe the evolution of neutrinos in time and space. Interactions with matter, however, are described by the flavour eigenstates, $|\nu_\alpha\rangle\,(\alpha = e, \mu, \tau)$, which couple to the weak force. The possibility of neutrino mixing means that the flavour and mass eigenstates need not be identical. In general, the flavour eigenstates can be written as a linear superposition of the mass eigenstates:

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

where U is the unitary lepton mixing matrix, named the PMNS matrix in recognition of contributions to the physics of mixing and oscillations made by Pontecorvo[19] and Maki, Nakagawa and Sakata[25].

In a vacuum, the mass eigenstates each propagate as a free particle, so a state $|\nu_i\rangle$ with position four-vector $\mathbf{x}$ and four-momentum $\mathbf{p}_i$ evolves as:

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

A neutrino produced in a flavour eigenstate $|\nu_\alpha\rangle$ will therefore evolve as:

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

Using the unitarity of the PMNS matrix to invert Equation 2.1 then allows the evolution

of the flavour eigenstate to be expressed as:

$$|\nu_\alpha(\mathbf{x})\rangle = \sum_{\beta=e,\mu,\tau} \left[\sum_{i=1,2,3} e^{-i\mathbf{p}_i \cdot \mathbf{x}} U_{\alpha i} U^*_{\beta i} \right] |\nu_\beta\rangle \tag{2.4}$$

Assuming that the mass eigenstate components all have a common energy E, and that this energy is much greater than the neutrino masses, the momentum of the $|\nu_i\rangle$ component can be approximated as:

$$p_i = \sqrt{E^2 - m_i^2} \simeq E - m_i^2/2E \tag{2.5}$$

Working in natural units and assuming highly relativistic neutrinos (for which $t \simeq L$), the phase at a distance L from the point of production can then be written as:

$$\mathbf{p}_i \cdot \mathbf{x} = Et - p_i L \simeq \frac{m_i^2 L}{2E} \tag{2.6}$$

A more rigorous approach, which avoids the assumption of a common energy for the mass eigenstates, has been demonstrated to produce this same result[26]. Substituting for the phase in Equation 2.4 then gives:

$$|\nu_\alpha(\mathbf{x})\rangle = \sum_{\beta=e,\mu,\tau} \left[\sum_{i=1,2,3} e^{-im_i^2 L/2E} U_{\alpha i} U^*_{\beta i} \right] |\nu_\beta\rangle \tag{2.7}$$

Equation 2.7 demonstrates that, if the neutrino masses are different, the phases of the mass eigenstates will evolve at different rates. If the PMNS matrix has non-zero off-diagonal terms, a neutrino produced in one flavour eigenstate will therefore develop components of other flavour eigenstates and so may be detected as a different flavour.

The probability that a neutrino produced at the origin in flavour eigenstate $|\nu_\alpha\rangle$ is observed at $\mathbf{x}$ in flavour eigenstate $|\nu_\beta\rangle$ is given by:

$$P(\nu_\alpha \to \nu_\beta) = |\langle \nu_\beta | \nu_\alpha(\mathbf{x}) \rangle|^2 \tag{2.8}$$

Using Equation 2.7, this can be evaluated as:

$$\begin{aligned} P(\nu_\alpha \to \nu_\beta) = \delta_{\alpha\beta} &- 4 \sum_{i>j} \Re(U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j}) \sin^2 \left(\frac{\Delta m_{ij}^2 L}{4E} \right) \\ &+ 2 \sum_{i>j} \Im(U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j}) \sin \left(\frac{\Delta m_{ij}^2 L}{2E} \right) \end{aligned} \tag{2.9}$$

where $\delta_{\alpha\beta}$ is the Kronecker delta and $\Delta m^2_{ij} = m^2_i - m^2_j$ is the splitting between the squared masses of the i^{th} and j^{th} mass eigenstates. The transition probability in Equation 2.9 is periodic with distance from the point of neutrino production. If the mass eigenstates have different masses, and if the flavour and mass eigenstates are not aligned, then neutrinos will oscillate between flavours as they propagate.

The corresponding transition probabilities for anti-neutrinos are calculated in [27]. The results indicate that a complex PMNS matrix would produce different probabilities for neutrino and anti-neutrino transitions between the same flavour eigenstates. This would represent a source of $\mathcal{CP}$ violation.

Neutrino oscillation experiments attempt to measure changes in the flavour composition of neutrino sources over long distances. If oscillations occur, the signature is a flavour composition that differs from the Standard Model prediction and which displays periodic variations with L/E. The period of the oscillations specifies the mass splittings, whilst the amplitude of the oscillations specifies the mixing angles.

To interpret the data from neutrino oscillation experiments, it is useful to use a specific representation of the PMNS matrix. The standard representation is obtained by expressing U as a product of three rotation matrices based on the mixing angles between the mass eigenstates (θ_{12}, θ_{23} and θ_{13}) and a complex phase factor ($e^{i\delta}$):

$$\begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{2.10}$$

where $c_{ij} \equiv \cos\theta_{ij}$, $s_{ij} \equiv \sin\theta_{ij}$ and θ_{ij} is the mixing angle between the i^{th} and j^{th} mass eigenstates. The phase δ is known as the 'Dirac phase', and its size determines the extent of $\mathcal{CP}$ violation in the neutrino sector. Two further phases are required if neutrinos are Majorana particles (identical to their anti-particles), but these phases affect only the diagonal of the matrix and so are unobservable in neutrino oscillations. The Majorana phases are therefore neglected here.

Current experimental evidence (detailed in Section 2.2) suggests that the squared mass splitting $|\Delta m^2_{32}|$ is much greater than the squared mass splitting $|\Delta m^2_{21}|$. Furthermore, the evidence suggests that the mixing angles θ_{12} and θ_{23} are large, whilst θ_{13} is small. In this regime, the three-flavour oscillations effectively decouple into three sets

of two-flavour oscillations, as described in [28]. Each of these two-flavour oscillations is associated with one of the rotation matrices in the decomposed PMNS matrix:

- **Long Range Oscillations:** Over large distances, such that $\Delta m^2_{21} L/E = \mathcal{O}(1)$, the dominant mode of oscillations is $\nu_e \leftrightarrow \nu_\mu/\nu_\tau$. This mode of oscillations provides an explanation for observations of ν_e disappearance in solar neutrino and long baseline reactor experiments. The oscillations are controlled by the parameters Δm^2_{21} and θ_{12}, and the oscillation probability is:

$$P(\nu_e \to \nu_\mu/\nu_\tau) \simeq \sin^2(2\theta_{12}) \sin^2\left(\frac{\Delta m^2_{21} L}{4E}\right) \tag{2.11}$$

- **Short Range Oscillations:** Over short distances, such that $\Delta m^2_{32} L/E = \mathcal{O}(1)$, the dominant mode of oscillations is $\nu_\mu \leftrightarrow \nu_\tau$. This mode of oscillations provides an explanation for observations of ν_μ disappearance in atmospheric neutrino and long baseline accelerator experiments. The oscillations are controlled by the parameters Δm^2_{32} and θ_{23}, and the oscillation probability is:

$$P(\nu_\mu \to \nu_\tau) \simeq \sin^2(2\theta_{23}) \sin^2\left(\frac{\Delta m^2_{32} L}{4E}\right) \tag{2.12}$$

- **Subdominant Short Range Oscillations:** Over short distances, a subdominant mode of oscillations $\nu_e \leftrightarrow \nu_\mu/\nu_\tau$ is possible. This mode of oscillations is associated with searches for ν_e disappearance in short range reactor experiments and is controlled by the squared mass splitting Δm^2_{32} and the small mixing angle θ_{13}. The oscillation probability is:

$$P(\nu_e \to \nu_\mu/\nu_\tau) \simeq \sin^2(2\theta_{13}) \sin^2\left(\frac{\Delta m^2_{32} L}{4E}\right) \tag{2.13}$$

To perform a study of neutrino oscillations, intense neutrino sources and large detectors are required in order to obtain an adequate event rate. The initial flux of neutrinos from the source must also be accurately measured or simulated, so that Standard Model predictions can be made for the flux and flavour composition. At the detector, observation of CC interactions provides information about the flavour composition, whilst observation of NC interactions yields information about the overall flux.

The optimal sensitivity to a particular mass splitting is achieved when the characteristic L/E value for the neutrino flux at the detector satisfies $L/E \simeq \pi/\Delta m^2$. With significantly smaller L/E values, the transition probabilities are too small for oscillations

to be detected. With significantly larger L/E values, the transition probabilities vary rapidly with either L or E. In this situation, the transition probability simply averages to $\frac{1}{2}\sin^2 2\theta$ and all information about the squared mass splitting is lost.

2.1.1 Alternative Models of Neutrino Propagation

Assuming that neutrinos oscillate, the probability that a neutrino produced in a $|\nu_\mu\rangle$ flavour eigenstate is subsequently observed as a $|\nu_\mu\rangle$ after propagating a distance L is given by the 'survival probability':

$$P(\nu_\mu \to \nu_\mu) = 1 - \sin^2(2\theta_{23}) \sin^2\left(\frac{\Delta m^2_{32} L}{4E}\right) \tag{2.14}$$

However, different models of neutrino propagation exist, which provide alternative explanations for the apparent changes in neutrino flavour composition. These models produce different expressions for the survival probability, which can be used to predict the composition of a $|\nu_\mu\rangle$ source. By comparing the predictions for different models with experimental measurements, the model that provides the best description of the data can be identified.

Neutrino Decoherence

Neutrino decoherence[29,30] explains neutrino flavour changes as due to the loss of coherence of the neutrinos' quantum mechanical phases. This can occur in the Standard Model, due to mass eigenstates travelling at slightly different velocities and separating out over very large distances (affecting neutrinos arriving at Earth from supernovae, for instance). To reduce the range over which decoherence effects are observable requires the introduction of new physics, such as quantum gravity.

Neutrino decoherence effects could occur alongside neutrino oscillations. Incorporating the effects of decoherence into the $|\nu_\mu\rangle$ survival probability requires the introduction of one new parameter (μ^2) alongside the standard oscillation parameters:

$$P(\nu_\mu \to \nu_\mu) = 1 - \frac{1}{2}\sin^2 2\theta_{23}\left[1 - e^{-\frac{\mu^2 L}{2E}} \cos\left(\frac{\Delta m^2_{32} L}{2E}\right)\right] \tag{2.15}$$

In the absence of first-principle calculations, μ^2 should be considered as a purely phenomenological parameter. If $\mu^2 = 0$, the survival probability returns to that expected with oscillations. If $\Delta m^2_{32} = 0$, then a pure decoherence model is obtained and the $|\nu_\mu\rangle$

survival probability is as shown below:

$$P(\nu_\mu \to \nu_\mu) = 1 - \frac{1}{2}\sin^2 2\theta_{23}\left(1 - e^{-\frac{\mu^2 L}{2E}}\right) \tag{2.16}$$

Neutrino Decay

Neutrino decay[31,32] explains differences between observed and predicted neutrino fluxes by allowing at least one neutrino component to decay into a sterile neutrino. Neutrino decay will typically occur alongside oscillations, but the squared mass splitting between the mass eigenstates could be sufficiently small to allow the effects of decay to dominate.

In the presence of a decaying mass eigenstate $|\nu_2\rangle$, with lifetime τ_2, the survival probability for a $|\nu_\mu\rangle$ flavour eigenstate is:

$$P(\nu_\mu \to \nu_\mu) = \sin^4\theta_{23} + \cos^4\theta_{23}e^{-\alpha L/E} + 2\sin^2\theta_{23}\cos^2\theta_{23}e^{-\alpha L/2E}\cos\left(\frac{\Delta m^2_{32}L}{2E}\right) \tag{2.17}$$

where $\alpha = m_2/\tau_2$. If the decay of $|\nu_2\rangle$ is into a state with which it does not mix, and if the squared mass splitting Δm^2_{32} is small enough to make the effects of oscillations negligible, the survival probability becomes:

$$P(\nu_\mu \to \nu_\mu) = \left(\sin^2\theta_{23} + \cos^2\theta_{23}e^{-\frac{\alpha L}{2E}}\right)^2 \tag{2.18}$$

2.2 Neutrino Oscillation Experiments

2.2.1 Long Range Oscillations

The first evidence for neutrino oscillations was provided by observations of solar neutrinos, which are produced by nuclear fusion reactions in the Sun. The best current physical model of the Sun is the Standard Solar Model (SSM)[33]. The SSM can be used to determine the rate of the different fusion reactions and to calculate the resulting flux of solar neutrinos.

The vast majority of solar neutrinos are produced via the '*pp* chain', illustrated in Figure 2.1, which combines four hydrogen nuclei (protons) into a helium nucleus and two electron neutrinos. Most of the neutrinos are produced in the first step of the *pp* chain, but these have a maximum energy of only 0.425 MeV, making them difficult to detect. However, different branches of the *pp* chain produce neutrinos of higher energies.

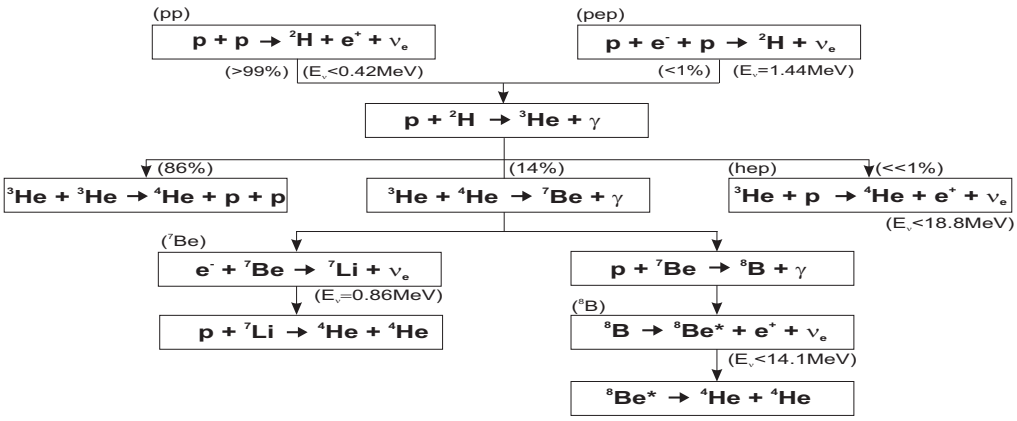

Figure 2.1: The *pp* chain of solar fusion reactions.

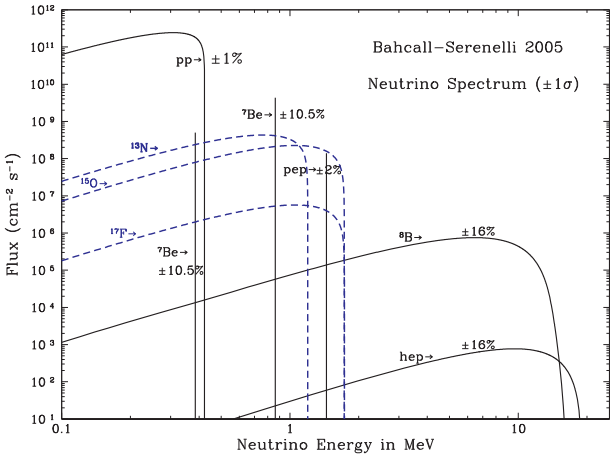

Figure 2.2: The solar neutrino energy spectrum, as predicted by the Standard Solar Model. Taken from [27].

The solar neutrino energy spectrum, as predicted by the SSM, is shown in Figure 2.2. Despite representing only 0.01% of the neutrino flux, electron neutrinos produced in the decay $^8\mathrm{B}\rightarrow{}^8\mathrm{Be}^* + e^+ + \nu_e$ have energies up to 14.1 MeV. These neutrinos are therefore ideal candidates for experimental observation.

Radiochemical Experiments

The first measurement of the solar neutrino flux was performed by the Homestake experiment[34]. This experiment detected neutrinos by observing the conversion of chlorine nuclei to argon in a large tank of C_2Cl_4, via the reaction $^{37}\mathrm{Cl}+\nu_e \rightarrow{}^{37}\mathrm{Ar}+e^-$, which has an energy threshold of 0.841 MeV. The argon atoms were extracted in a helium purg-

ing and filtration process and then counted by monitoring their decays in a gas-filled proportional counter.

The final measurement of the solar ν_e flux was $2.56 \pm 0.16\,(stat.) \pm 0.16\,(syst.)$ SNU, where one SNU (Solar Neutrino Unit) is equivalent to one interaction per 10^{36} atoms per second. This observation was significantly smaller than the SSM prediction for Homestake, which at the time was 8.5 ± 1.8 SNU[35]. The deficit between the measurement and the SSM prediction was named the 'solar neutrino anomaly'.

Initial attempts to explain the anomaly suggested that the SSM was incorrect and that the interior pressures and temperatures of the Sun must be poorly modelled. However, these claims were proven incorrect by advances in helioseismology (the study of wave propagation in the Sun), which allowed the temperature inside the Sun to be inferred; these measurements were found to be in agreement with the SSM.

More recently, other radiochemical experiments have been constructed: SAGE[36], GALLEX[37] and GNO[38] (the successor to GALLEX). These experiments detect solar neutrinos through the reaction $^{71}\mathrm{Ga} + \nu_e \rightarrow {}^{71}\mathrm{Ge} + e^-$, which has a reduced threshold energy of 0.233 MeV. This provides sensitivity to the large flux of low energy neutrinos produced in the first step of the *pp* chain.

SAGE used 60 tonnes of liquid Gallium, whilst GALLEX and GNO used 101 tonnes of Gallium Chloride solution. The SSM predicted flux for these experiments was 131^{+12}_{-10} SNU, but all three experiments observed a large deficit with respect to this prediction. SAGE measured the flux as $70.8^{+5.3}_{-5.2}\,(stat.)^{+3.7}_{-3.2}\,(syst.)$ SNU, whilst the combined GALLEX-GNO result was $69.3 \pm 4.1\,(stat.) \pm 3.6\,(syst.)$ SNU.

Neutrino oscillations offer a solution to the problem, as $\nu_e \leftrightarrow \nu_\mu/\nu_\tau$ transitions could produce the observed deficit. Oscillations could not be confirmed by the radiochemical experiments, as they are only sensitive to ν_e disappearance and cannot measure the associated ν_μ or ν_τ appearance. However, measurements at the Sudbury Neutrino Observatory (SNO) provide sensitivity to all three neutrino flavours.

SNO

The Sudbury Neutrino Observatory[39] is an imaging Čerenkov detector, which uses heavy water (D_2O) as both an interaction and observation medium. The detector consists of 1000 tonnes of heavy water, contained in a 12 m diameter transparent acrylic vessel. The Čerenkov light produced by neutrino interactions is detected by an array of 9456 photomultiplier tubes (PMTs), which are supported on a geodesic sphere. Seven kilotonnes of purified H_2O shield the heavy water from external radioactive backgrounds.

SNO was designed to provide direct evidence of solar neutrino flavour transitions by

measuring the ^{8}B neutrino interaction rates for three different processes:

$$\begin{aligned} \text{ES}: &\quad \nu_x + e^- \rightarrow \nu_x + e^- \\ \text{CC}: &\quad \nu_e + d \rightarrow p + p + e^- \\ \text{NC}: &\quad \nu_x + d \rightarrow p + n + \nu_x \end{aligned}$$

The first of these processes, elastic scattering (ES), is sensitive to all neutrino flavours, although the cross-section is largest for electron neutrinos (for which a CC contribution is possible, in addition to NC). In the second process, deuterium in the heavy water provides loosely bound neutron targets for exclusively CC interactions. At solar neutrino energies, this process is therefore sensitive to only electron neutrinos. The final reaction is a purely NC process, which is equally sensitive to all neutrino flavours and so provides a direct measurement of the total active neutrino flux.

For the ES and CC reactions, the recoil electron is detected directly through its production of Čerenkov light. For the NC reaction, the neutrons are not observed directly, but neutron capture by another deuteron releases a γ-ray that typically Compton scatters an electron or produces an e^+e^- pair. These secondary particles produce Čerenkov light that can be detected. In SNO phase II, NaCl was dissolved in the heavy water to increase the neutron capture efficiency. In phase III, the NaCl was removed and an array of ^{3}He proportional counters were installed to detect neutrons directly.

The measured rates for the three interaction types can be combined to produce separate measurements of the ν_e and ν_μ/ν_τ fluxes. During phase II of the experiment, the following fluxes were measured[20]:

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

Figure 2.3 illustrates the SNO phase II results by plotting the non-ν_e component of the flux ($\phi_{\mu\tau}$) against the ν_e component (ϕ_e). Possible 1σ bands are shown for the CC, NC and ES measurements, in addition to the total ^{8}B neutrino flux predicted by the SSM. The intersection of these bands shows that the measurements are consistent, and allows the value of $\phi_{\mu\tau}$ to be resolved.

The results indicate a non-zero value of $\phi_{\mu\tau}$, which implies that a non-ν_e component exists in the solar neutrino flux. When combined with the ν_e component, this produces a total neutrino flux in good agreement with the SSM prediction. SNO therefore provides

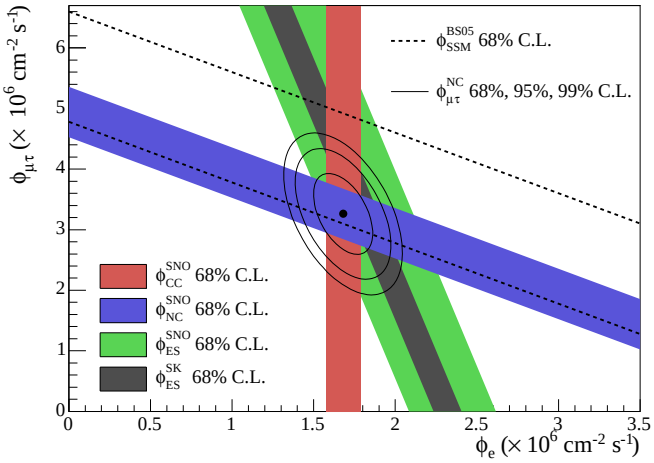

Figure 2.3: SNO phase II results, showing the flux of μ and τ neutrinos versus the flux of electron neutrinos[20]. Separate filled bands are indicated for the CC, NC and ES measurements and the $\pm 1\sigma$ uncertainties are represented by the intercepts of these bands with the axes. The Super-Kamiokande measurement in [40] is also indicated, as is the SSM prediction for the total ^{8}B neutrino flux.

strong evidence for neutrino flavour transitions.

The observation of a suppressed ν_e flux, alongside the total neutrino flux predicted by the SSM, can be analysed in the context of neutrino oscillations. For solar neutrinos, this analysis is complicated by the Mikheyev-Smirnov-Wolfenstein (MSW) effect[41, 42]. As neutrinos propagate through an electron-rich medium, such as the Sun, coherent forward scattering occurs. NC scatters occur for all three neutrino flavours, but electron neutrinos are additionally affected by CC scatters with electrons. This affects the transition probabilities.

Figure 2.4 shows the results of a two-flavour oscillations analysis of the SNO phase II data, accounting for the MSW propagation of neutrinos through dense matter in the Sun and Earth. The Figure displays confidence limits for the parameters $\tan^2 2\theta_{12}$ and Δm^2_{21}. The overall best fit point is $\Delta m^2_{21} = 5.0^{+6.2}_{-1.8} \times 10^{-5}\,\mathrm{eV}^2$, $\tan^2 2\theta_{12} = 0.45^{+0.11}_{-0.10}$[20].

KamLAND

The Kamioka Liquid scintillator Anti-Neutrino Detector (KamLAND)[43] is located in the Kamioka observatory, beneath the Japanese Alps. It investigates neutrino oscillations by observing electron anti-neutrinos emitted from distant nuclear reactors. Assuming $\mathcal{CPT}$ invariance, these $\bar{\nu}_e$ measurements can be combined with solar ν_e measurements to constrain the solar neutrino oscillation parameters.

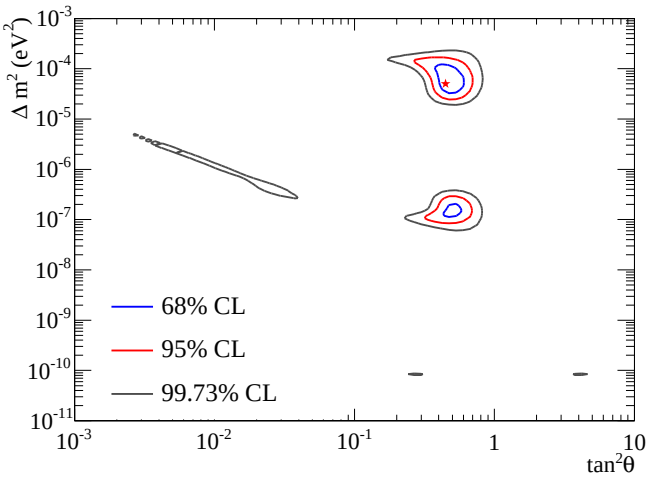

Figure 2.4: SNO phase II results, showing the confidence limits and best fit value for the oscillation parameters Δm^2_{21} and $\tan^2 2\theta_{12}$. Taken from [20].

KamLAND consists of 1 kton of highly purified liquid scintillator contained in a 13 m diameter transparent nylon-based balloon. This balloon is suspended in non-scintillating oil and is surrounded by 1879 PMTs. Electron anti-neutrinos are detected by the inverse β-decay reaction $\bar{\nu}_e + p \rightarrow e^+ + n$, which has an energy threshold of 1.8 MeV. The prompt scintillation light produced by the e^+ gives a measure of the incident $\bar{\nu}_e$ energy, via:

$$E_{\bar{\nu}_e} = E_{\text{prompt}} + \bar{E}_n$$

where E_{prompt} includes the positron kinetic and annihilation energy, and $\bar{E}_n$ is the average neutron recoil energy. The delayed production of a 2.2 MeV γ-ray (due to neutron capture after approximately 200 μs) helps to considerably reduce backgrounds.

The anti-neutrinos are produced by 55 Japanese nuclear power reactor units, which are distributed around KamLAND. Each reactor unit acts as an isotropic $\bar{\nu}_e$ source. The expected $\bar{\nu}_e$ flux and energy spectrum at the detector can be calculated from detailed information about reactor operations, such as regular thermal power measurements. Comparison of the expected number of events and energies with the experimental observations allows the presence of oscillations to be inferred.

Recently published results from the KamLAND collaboration[23] are based on data collected between March 2002 and May 2007. In the absence of $\bar{\nu}_e$ disappearance, $2179 \pm 89\,(syst.)$ events were expected from the reactors, whilst the background event expectation was 276.1 ± 23.5. Only 1609 events were observed, representing a significant

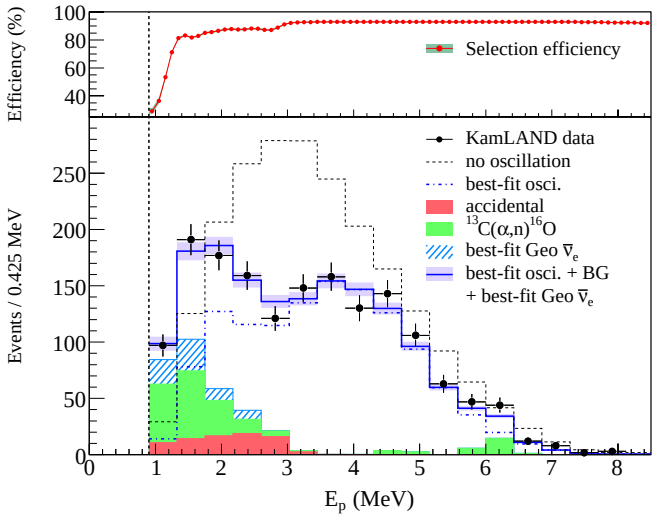

Figure 2.5: Prompt energy spectrum of selected $\bar{\nu}_e$ events at KamLAND and the fitted backgrounds. The no oscillation prediction is shown, in addition to the best fit prediction from a two-flavour oscillation analysis. Taken from [23].

deficit in the $\bar{\nu}_e$ flux. Figure 2.5 shows the prompt event energy spectrum of the $\bar{\nu}_e$ candidate events.

The measured spectrum has been analysed with a maximum likelihood fit to two-flavour neutrino oscillations. This fit accounts for variations in the reactor flux and includes Earth-matter effects. The best fit distribution is shown in Figure 2.5, whilst the allowed contours in neutrino oscillation parameter space are shown in Figure 2.6. The best fit oscillation parameters are $\Delta m^2_{21} = 7.58^{+0.14}_{-0.13}\,(stat.)^{+0.15}_{-0.15}\,(syst.) \times 10^{-5}\,\mathrm{eV}^2$ and $\tan^2 2\theta_{12} = 0.56^{+0.10}_{-0.07}\,(stat.)^{+0.10}_{-0.06}\,(syst.)$. Incorporating the results of solar neutrino experiments, particularly SNO, further constrains the parameter space and produces best fit parameters of $\Delta m^2_{21} = 7.59^{+0.21}_{-0.21} \times 10^{-5}\mathrm{eV}^2$ and $\tan^2 2\theta_{12} = 0.47^{+0.06}_{-0.05}$. This is the most precise measurement of Δm^2_{21} to date.

Figure 2.7 shows the ratio of the background-subtracted $\bar{\nu}_e$ candidate events to the expectation assuming no neutrino oscillations, plotted as a function of L_0/E. L_0 is the effective baseline for the KamLAND experiment and is taken as a flux-weighted average of the distances to the reactor units. This distribution decisively confirms the signal of neutrino oscillations, displaying almost two full cycles of the expected periodic variations with L_0/E.

The oscillatory structure observed in Figure 2.7 also provides strong evidence against the pure decoherence and decay models of neutrino propagation. These models produce survival probabilities that are exponential functions of L_0/E; only neutrino oscillations

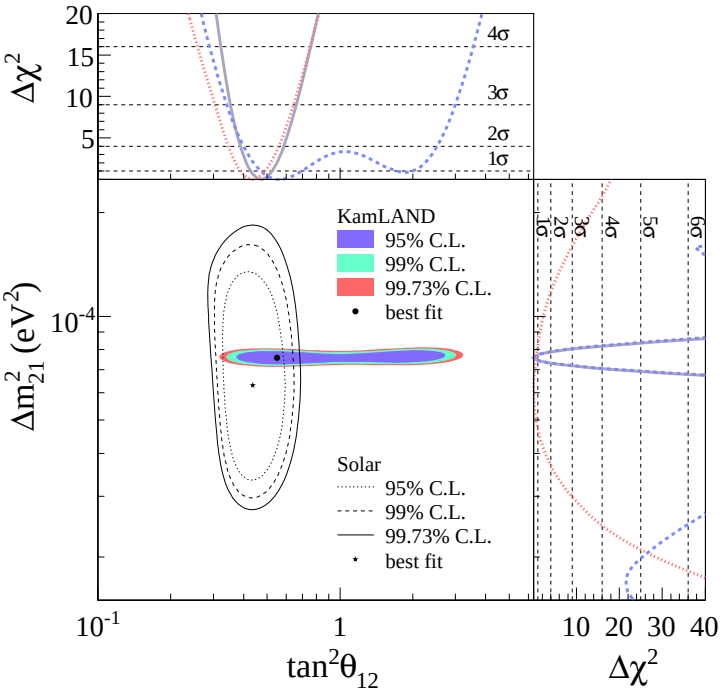

Figure 2.6: Allowed region for neutrino oscillation parameters, as obtained in the KamLAND oscillation analysis and in an analysis of solar neutrino data. Taken from [23].

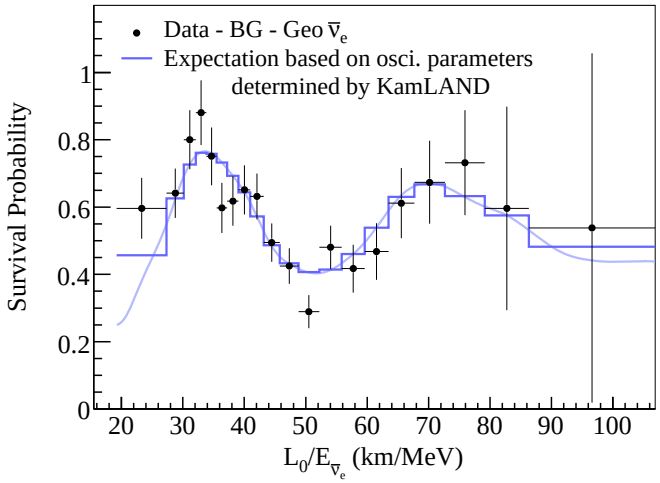

Figure 2.7: Ratio of the KamLAND background-subtracted $\bar{\nu}_e$ spectrum to the unoscillated prediction, plotted as a function of L_0/E, where L_0 is the effective baseline. Taken from [23].

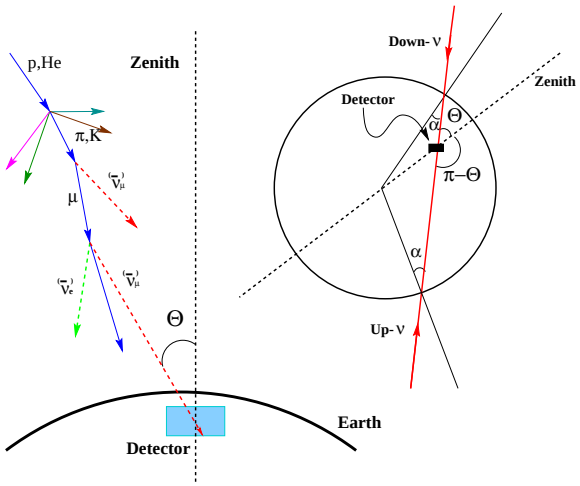

Figure 2.8: The production of atmospheric neutrinos in the collisions between cosmic rays and nuclei in the upper atmosphere. A definition of the zenith angle is also shown. Taken from [46].

can produce the observed periodic variations.

2.2.2 Short Range Oscillations

Atmospheric neutrinos are produced as a result of collisions between cosmic rays (typically protons) and nuclei in the Earth's upper atmosphere. These collisions produce cascades of secondary pions and kaons, which subsequently decay to produce large numbers of electrons and muons, with energies in the range 0.1 GeV to 100 GeV. Calculations of the absolute flux of atmospheric neutrinos[44, 45] have large uncertainties (∼20%), due to modelling of the primary cosmic ray flux and of the Earth's magnetic field and atmosphere. However, the ratio of muon to electron flavour content is known to ∼5%, as it arises from the following decay chain (also illustrated in Figure 2.8):

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

The decay chain predicts that the ratio of neutrino flavours, $R = \frac{N(\nu_{\mu}+\bar{\nu}_{\mu})}{N(\nu_{e}+\bar{\nu}_{e})}$, should be about 2. The ratio should gradually rise at energies above 5 GeV, as the intermediate muons are more likely to reach the Earth's surface before decaying. Experimental measurements of this ratio involve observation of CC atmospheric neutrino interactions. In these interactions, the charged lepton tags the flavour of the incident neutrino and

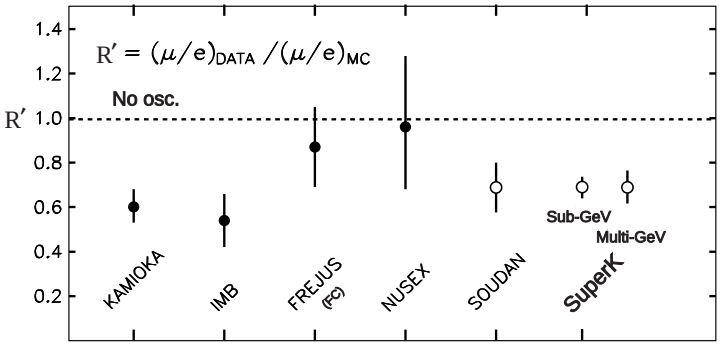

Figure 2.9: Summary of measurements of the double ratio, R', from different atmospheric neutrino experiments. Taken from [49].

allows estimation of its direction and energy. Experimental data are usually presented as a double ratio between the observed and predicted flavour ratios:

$$R' = \frac{R_{\mathrm{Data}}}{R_{\mathrm{MC}}} \tag{2.19}$$

The first measurements of this double ratio were performed by the IMB[47] and Kamiokande[48] experiments, which both used water Čerenkov detectors. Electrons and muons were separated by examining the observed Čerenkov rings; muons create 'sharp-edged' rings, whilst electrons induce electromagnetic showers and so produce 'fuzzy-edged' rings. Both experiments observed a significant deficit in the ν_μ flux relative to the expectation, with IMB measuring $R' = 0.54^{+0.05}_{-0.05}\,(stat.)^{+0.12}_{-0.12}\,(syst.)$ and Kamiokande finding $R' = 0.60^{+0.07}_{-0.06}\,(stat.)^{+0.05}_{-0.05}\,(syst.)$. This deficit was known as the 'atmospheric neutrino anomaly' and a summary of measurements of R' is shown in Figure 2.9.

Neutrino oscillations provide a possible explanation of the atmospheric neutrino anomaly, suggesting that the deficit in the ν_μ flux is due to transitions between ν_μ and other neutrino flavours. A means of testing this hypothesis is to measure the variation of the ν_μ flux with the zenith angle. This is the angle between the neutrino direction and the vertical, as illustrated earlier in Figure 2.8.

As the attenuation of neutrinos in the Earth is so small, neutrinos produced all around the atmosphere can be detected. The pathlengths for these neutrinos to arrive at a detector near the surface of the Earth therefore varies from about 15 km at small zenith angles to 13,000 km for neutrinos produced on the opposite side of the Earth. If neutrino oscillations occur within this range of distances, then deficits in the ν_μ flux should be observed at zenith angles for which the propagation distance produces a large oscillation probability.

Super-Kamiokande

The results of the Super-Kamiokande experiment[50] provide the most compelling evidence that the atmospheric neutrino anomaly is caused by neutrino oscillations. Super-Kamiokande is a 50 kton water Čerenkov detector. It is instrumented by 11,146 PMTs, which face an inner 22.5 kton fiducial volume of highly purified water. The inner volume is surrounded by a 2 m thick outer detector, which is instrumented by 1885 outward-facing PMTs.

Atmospheric neutrinos are observed via their CC interactions with nuclei in the detector fiducial volume. Interactions are classified as fully-contained (FC) if all of the energy is deposited inside the inner detector, or as partially-contained (PC) if one or more particles exit the inner detector and deposit energy in the outer veto region. The neutrino energies that produce PC events are typically ten times higher than those that produce FC events. FC events are divided into 'sub-GeV' and 'multi-GeV' samples, based on the observed amount of Čerenkov light. In multi-GeV events, the path of the charged lepton closely follows the direction of the incident neutrino, so the neutrino propagation distance can be determined accurately.

Neutrinos are also detected via their interactions with the surrounding rock. CC ν_μ interactions with the rock produce muons, which can enter the detector. If these muons travel downwards through the detector, they cannot be distinguished from the constant flux of cosmic ray muons. However, those muons that travel in an upward direction through the detector must be neutrino induced. Upward-going muons are divided into those that stop in the detector and those that pass through the entire detector volume. Upward stopping muons have typical parent neutrino energies of 10 GeV, whilst upward through-going muons have larger parent neutrino energies of around 100 GeV.

Super-Kamiokande has published an analysis of atmospheric neutrino data collected between April 1996 and July 2001[21]. Figure 2.10 shows the measured number of events and the expectation in the absence of oscillations, plotted as a function of the cosine of the zenith angle. Separate distributions are shown for the different classes of event and for different energy samples. While the measured zenith angles for electron-like events are consistent with expectations, significant zenith angle dependent deficits are observed for muon-like events. In particular, the distribution of multi-GeV muon-like events shows agreement with the expectation at low zenith angles, but falls to almost half the expectation at large zenith angles.

The interpretation of the data is that the atmospheric neutrinos undergo $\nu_\mu \leftrightarrow \nu_\tau$ oscillations with a large mixing angle. Such oscillations are mostly characterised by ν_μ disappearance, rather than ν_τ appearance, as the majority of the neutrinos are below the

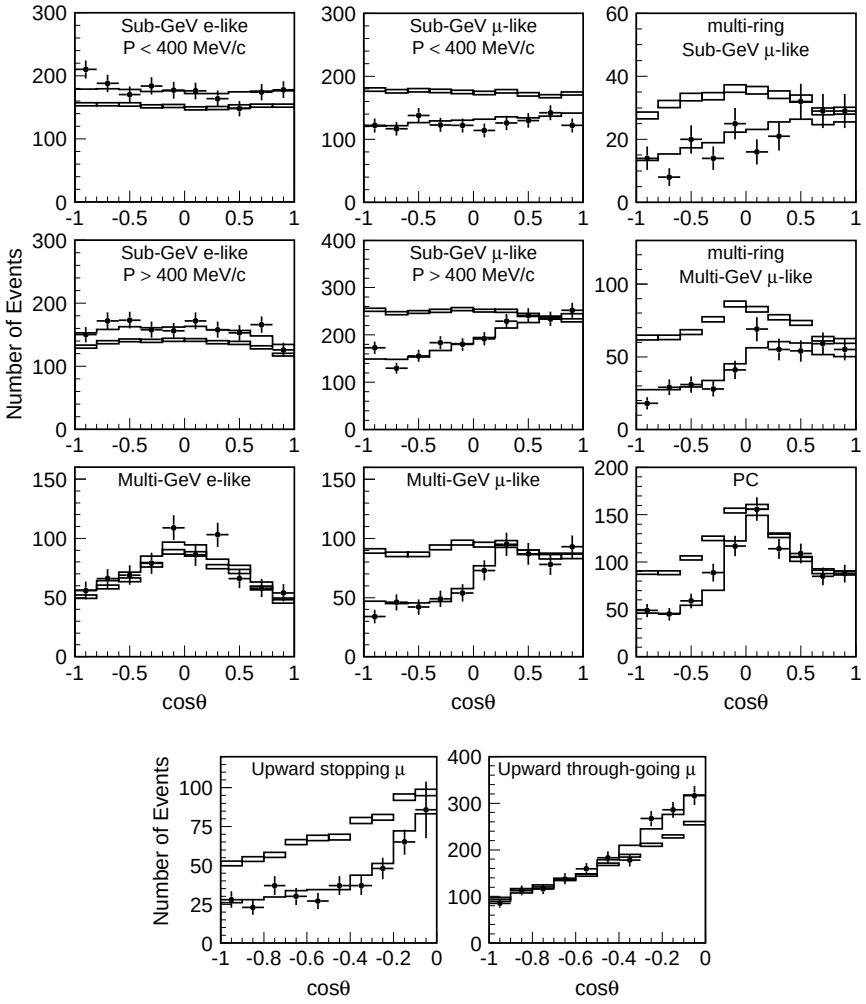

Figure 2.10: Super-Kamiokande zenith angle distributions for different classes of atmospheric neutrino events. The points show the data, the boxes show the unoscillated predictions and the lines show the best fit expectations for $\nu_\mu \leftrightarrow \nu_\tau$ oscillations with $\sin^2 2\theta_{23} = 1.0$ and $\Delta m^2_{32} = 2.1 \times 10^{-3}\,\mathrm{eV}^2$. Taken from [21].

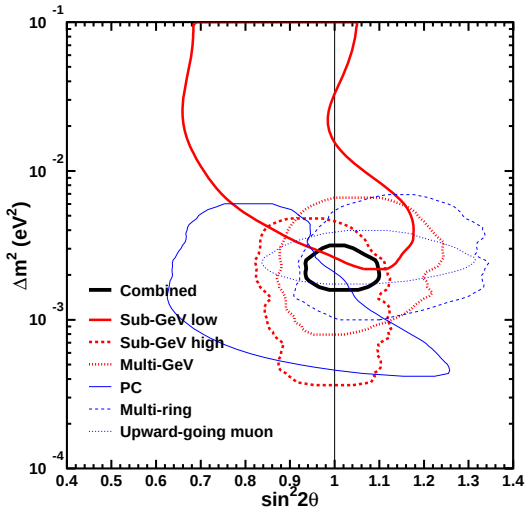

Figure 2.11: 90% confidence level allowed regions for $\nu_\mu \leftrightarrow \nu_\tau$ oscillations from different samples of atmospheric neutrino events in the Super-Kamiokande zenith angle analysis. Taken from [21].

energy threshold for CC tau production. A two-flavour oscillation interpretation of the data yields 90% confidence limits for the parameters controlling short-range oscillations: $1.5 \times 10^{-3} < \Delta m^2_{32} < 3.4 \times 10^{-3}\,\mathrm{eV}^2$, $\sin^2 2\theta_{23} > 0.92$. The best fit parameters are $\Delta m^2_{32} = 2.1 \times 10^{-3}\,\mathrm{eV}^2$, $\sin^2 2\theta_{23} = 1.0$. Figure 2.11 shows the 90% confidence level contours obtained in oscillation fits to the different classes of event.

In addition to the zenith angle study, Super-Kamiokande has also used the same dataset to perform an analysis of ν_μ disappearance as a function of L/E[22]. Only data with an expected L/E resolution of below 70% was used. This removed events near the horizon, for which L varies rapidly with zenith angle, and also removed low energy events, for which the mean angle between the neutrino and muon is large. Figure 2.12 shows the ratio between the measured L/E distribution and the expectation in the absence of oscillations. A dip is observed in this ratio at around 500 km/GeV.

A two-flavour oscillation fit to the L/E distribution yields 90% confidence limits of $1.9 \times 10^{-3} < \Delta m^2_{32} < 3.0 \times 10^{-3}\,\mathrm{eV}^2$, $\sin^2 2\theta_{23} > 0.90$ and identifies the best fit parameters as $\Delta m^2_{32} = 2.4 \times 10^{-3}\,\mathrm{eV}^2$, $\sin^2 2\theta_{23} = 1.0$. These parameters are consistent with those obtained in the zenith angle analysis. Figure 2.13 shows the final 68%, 90% and 99% confidence level contours for the L/E and zenith angle analyses.

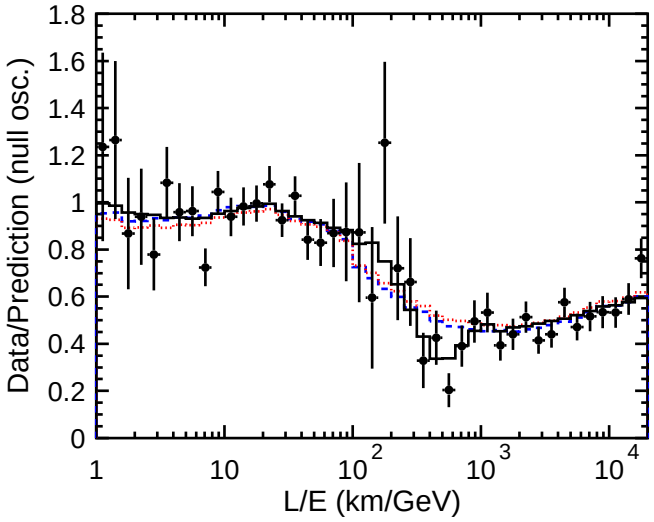

Figure 2.12: Ratio of Super-Kamiokande data to the unoscillated prediction as a function of L/E. The solid line shows the best fit expectation for two-flavour $\nu_\mu \leftrightarrow \nu_\tau$ oscillations. Also shown are the best fit expectations for neutrino decay (dashed line) and neutrino decoherence (dotted line). Taken from [22].

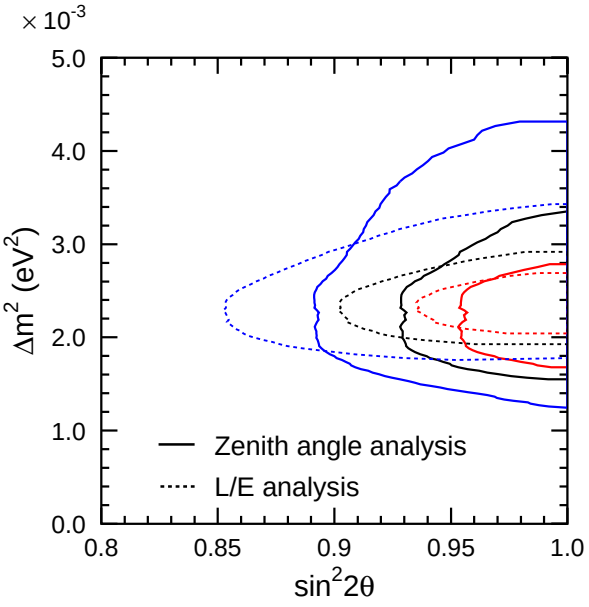

Figure 2.13: 68%, 90% and 99% confidence level contours obtained in the Super-Kamiokande L/E and zenith angle analyses. Taken from [21].

K2K

The signal of $\nu_\mu \leftrightarrow \nu_\tau$ oscillations has been confirmed by the K2K (KEK-to-Kamioka) experiment[51]. K2K used an accelerator ν_μ beam with a neutrino flight distance of 250 km to probe the same region of parameter space as the atmospheric neutrino experiments. The neutrino beam was produced by firing 12 GeV protons (from the KEK proton-synchrotron) onto an aluminium target, creating large numbers of pions and kaons. The positively charged pions/kaons were focussed by a pair of magnetic horns and directed into a 200 m decay pipe, where they subsequently decayed to produce a beam of 98% ν_μ with a mean energy of 1.3 GeV.

The initial unoscillated ν_μ spectrum was measured by a near detector, located 300 m from the target. The near detector complex consisted of a 1 kton water Čerenkov detector and a fine grained detector system. The neutrino energy spectrum was then measured again by the Super-Kamiokande detector at a distance of 250 km from the target. The energy spectrum and flux normalisation measured at the near detector were used to accurately predict the far detector spectrum for different models of neutrino propagation.

K2K has published an analysis of data collected between June 1999 and November 2004[24]. In the absence of oscillations, $158^{+9.2}_{-8.6}$ beam events were expected in the fiducial volume of the Super-Kamiokande detector. Only 112 such events were observed, representing the deficit expected from neutrino oscillations. The spectral distortion associated with oscillations was also observed in the 58 single-ring muon-like events for which the neutrino energy was reconstructed. Figure 2.14 shows the observed far detector energy spectrum, together with the distribution expected in the absence of oscillations and the best fit oscillation prediction.

Figure 2.15 shows the confidence level contours obtained in the two-flavour K2K oscillations analysis. These contours are consistent with the results of the Super-Kamiokande atmospheric L/E analysis. With maximal mixing, the 90% confidence limits for Δm^2_{32} are: $1.9 \times 10^{-3} < \Delta m^2_{32} < 3.5 \times 10^{-3}\,\mathrm{eV}^2$. The best fit point was $\Delta m^2_{32} = 2.8 \times 10^{-3}\,\mathrm{eV}^2$, $\sin^2 2\theta_{23} = 1.0$.

2.2.3 Subdominant Short Range Oscillations

In addition to solar and atmospheric neutrino oscillations, there is a third sector of the neutrino mixing matrix, which is controlled by the angle θ_{13}. Measurement of this mixing angle is important in order to understand the structure of the PMNS matrix. Accurate measurement is also vital for the study of subdominant oscillation effects and $\mathcal{CP}$ violation in the lepton sector, which are only observable with non-zero θ_{13} values.

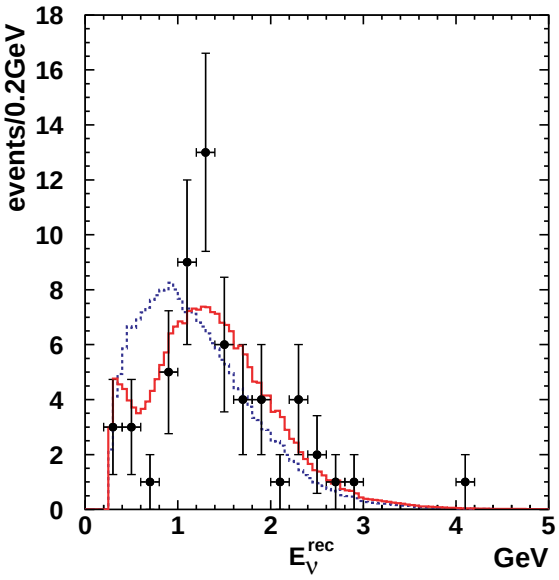

Figure 2.14: The reconstructed E_ν distribution for the 58 single-ring muon-like events in the K2K data. Points with error bars show the data, whilst the dashed line shows the expectation without oscillations. The solid line shows the best fit spectrum with neutrino oscillations. The histograms are normalised to the number of observed events. Taken from [24].

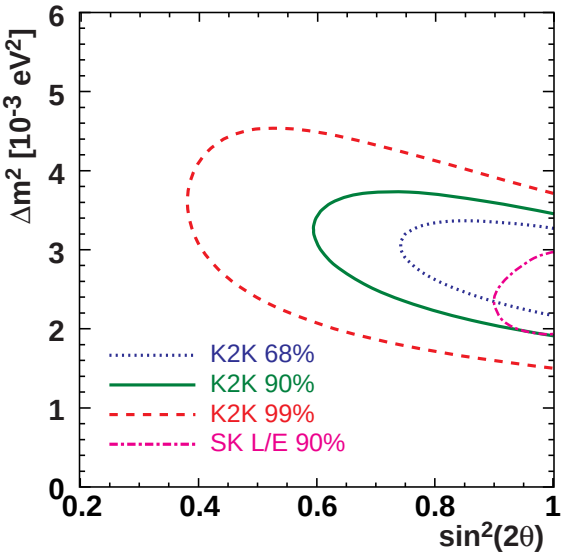

Figure 2.15: The results of the K2K two-flavour oscillation analysis, showing the allowed regions for oscillation parameters Δm^2_{32} and $\sin^2 2\theta_{23}$. For comparison, the findings of the Super-Kamiokande atmospheric L/E analysis are also displayed. Taken from [24].

However, whilst the solar and atmospheric mixing angles are large, θ_{13} has been identified as very small and current experiments have only managed to set upper limits to its value. The angle is currently constrained by the results of the CHOOZ reactor neutrino experiment[52].

CHOOZ

The CHOOZ experiment searched for the subdominant mode of $\bar{\nu}_e \leftrightarrow \bar{\nu}_\mu/\bar{\nu}_\tau$ oscillations in a reactor flux of electron anti-neutrinos. The experiment used a 5 ton liquid scintillator detector, located a distance of 1 km away from the two reactors at the CHOOZ nuclear power plant in northern France. These reactors produced a pure and isotropic flux of $\bar{\nu}_e$, with a mean energy of 3 MeV and an intensity known to better than 2%.

As in KamLAND, reactor $\bar{\nu}_e$ interactions were detected via the inverse β-decay reaction $\bar{\nu}_e + p \rightarrow e^+ + n$. The signature of a $\bar{\nu}_e$ interaction was prompt scintillation light from the e^+, followed by photon emissions associated with capture of the neutron. In order to increase the neutron capture cross-section and the energy of the resulting photons, the scintillator was loaded with Gadolinium.

The signal for oscillations at CHOOZ was a deficit of $\bar{\nu}_e$ events and a distorted $\bar{\nu}_e$ energy spectrum with respect to the unoscillated predictions. Figure 2.16 shows the observed prompt e^+ spectrum superimposed on the no oscillation prediction; no deficit of events or spectral distortion was observed. Figure 2.17 shows the confidence limits obtained in a two-flavour oscillation analysis. Subdominant oscillations $\bar{\nu}_e \rightarrow \bar{\nu}_x$ were excluded at the 90% confidence level for $\Delta m^2_{32} \geq 8 \times 10^{-4}\,\mathrm{eV}^2$ at maximal mixing and for $\sin^2 2\theta_{13} \geq 0.17$ at large Δm^2_{32} values[52].

2.2.4 Summary

The SNO and KamLAND experiments have provided evidence to validate the SSM and identify the solar neutrino anomaly as due to $\nu_e \leftrightarrow \nu_\mu/\nu_\tau$ flavour oscillations. Meanwhile, the Super-Kamiokande and K2K experiments strongly suggest that the atmospheric neutrino anomaly is caused by $\nu_\mu \leftrightarrow \nu_\tau$ oscillations. These solar and atmospheric oscillation modes are coupled by the mixing angle θ_{13}, which the CHOOZ experiment has identified as very small. The next information will be provided by the MINOS experiment, which is described in detail in Chapter 3. MINOS (and the analysis described in this thesis) aims to confirm the existence of oscillations in the region of parameter space suggested by Super-Kamiokande and to measure the oscillation parameters Δm^2_{32} and $\sin^2 2\theta_{23}$ with high precision.

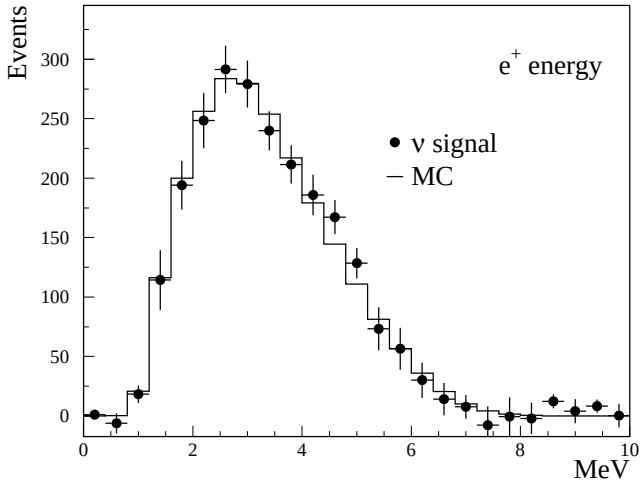

Figure 2.16: Prompt e^+ spectrum measured at CHOOZ, superimposed on the no oscillation prediction (solid line). Taken from [52].

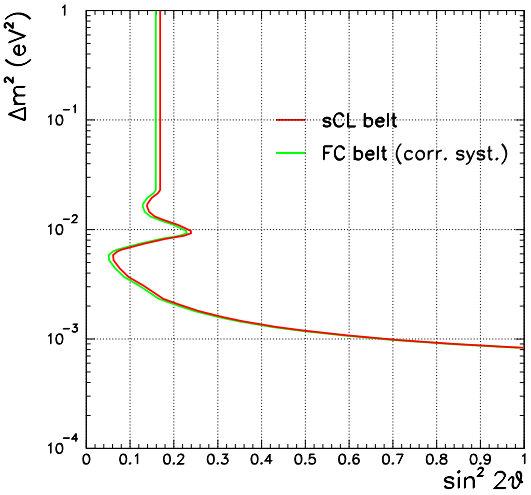

Figure 2.17: CHOOZ confidence limits for the parameters $\sin^2 2\theta_{13}$ and Δm^2_{32}. The region to the right of the contour is excluded at the 90% confidence level. Contours obtained using the strong confidence level (sCL) and Feldman-Cousins (FC) methods are shown. Taken from [52].

Chapter 3

The MINOS Experiment

The Main Injector Neutrino Oscillation Search (MINOS) is a long baseline, two-detector neutrino oscillation experiment. It uses an accelerator beam of muon neutrinos in order to look for evidence of neutrino oscillations in the region of parameter space indicated by atmospheric neutrino experiments. The experiment aims to confirm the existence of neutrino oscillations and perform a precise measurement of the oscillation parameters.

The neutrino beam is manufactured by the Neutrinos at the Main Injector (NuMI) facility at the Fermi National Accelerator Laboratory in Illinois, USA. As illustrated in Figure 3.1, the neutrinos are observed by functionally identical detectors at two different points in the beamline. The Near Detector is located close to the beam source and provides information about the beam before oscillations can develop, whilst the Far Detector is located 735 km away in the Soudan Underground Laboratory in Minnesota.

To search for evidence of oscillations, the neutrino energy spectrum observed at the Far Detector is compared with the spectrum expected in the absence of oscillations. To incorporate understanding of the beam and the detector response, and so address many sources of systematic errors, the expected Far Detector spectrum is produced by extrapolation of the measured Near Detector data spectrum. The signature for oscillations is a deficit of muon neutrinos in the Far Detector data, together with an associated distortion of the neutrino energy distribution.

In this Chapter, the stages required to obtain information for an oscillation analysis are described. This includes the production of the neutrino beam and details of the

Figure 3.1: The MINOS neutrino beamline. The neutrinos are produced at Fermilab, Illinois, and travel 735 km through the Earth to the Soudan Underground Laboratory in Minnesota.

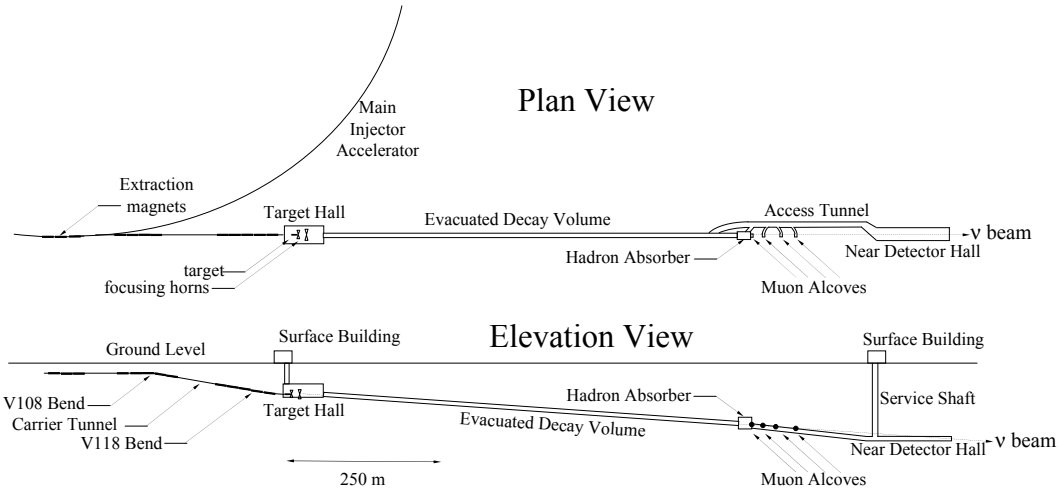

Figure 3.2: The NuMI beam facility. A beam of protons is directed onto a graphite target, producing pions and kaons that are focussed into an evacuated decay pipe by a system of magnetic horns. The neutrino beam is produced by the subsequent decay of the pions and kaons. Taken from [53].

MINOS detectors. Also described are detector calibration and, very briefly, the development of an accurate beam, event and detector Monte Carlo (MC) simulation. Finally, the reconstruction software used to identify high level objects, such as muon tracks and hadronic showers, is discussed.

3.1 The NuMI Beam

Figure 3.2 shows a diagram of the NuMI beamline. The production of the neutrino beam begins with the acceleration of protons to 120 GeV in the Main Injector accelerator. Three kicker magnets are then used to extract batches of these protons during a 10 μs 'spill'. The spills are spaced 2.2−2.4 s apart and typically contain 2.2×10^{13} protons. After extraction, the protons are directed downward by 58 mrad, to point towards the Far Detector. The protons are transported along a carrier tunnel and directed onto the hadron production target. This target is a rectangular graphite rod, which is divided longitudinally into 47 segments. It is 6.4 mm wide, 15 mm high and has a total length of 940 mm. A diagram of the target is shown in Figure 3.3, which also indicates its various cooling mechanisms.

When it reaches the target, the proton beam has a width of 1.2−1.5 mm. The protons interact in the target and produce pions and kaons with a wide range of longitudinal and transverse momenta. The long and narrow shape of the target helps to maximise the

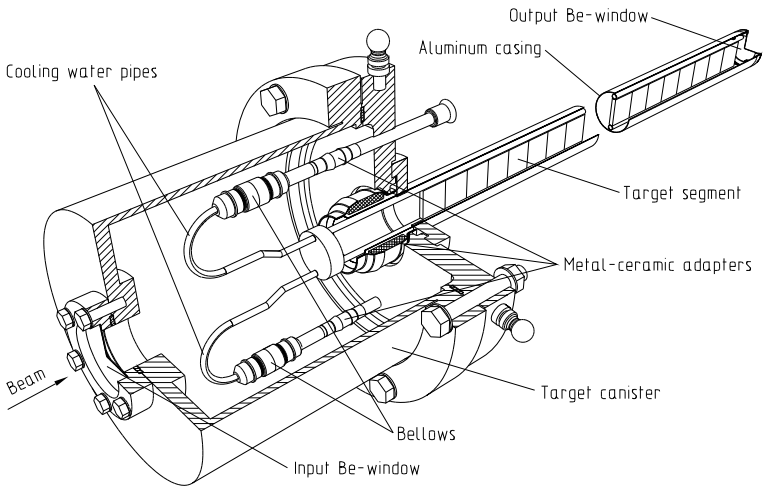

Figure 3.3: Schematic of the NuMI target, which consists of 47 water-cooled graphite segments. Taken from [54].

number of proton interactions, whilst allowing the pions and kaons to escape through the sides of the target (thereby reducing the chance of re-absorption).

A system of two magnetic horns, illustrated in Figure 3.4, is used to focus hadrons of a particular charge sign onto the beam axis. The two horns are 10 m apart and each consists of an outer cylindrical conductor and an inner conductor, which is parabolic in shape. During beam spills, pulses of up to 200 kA are applied to the conductors, generating a toroidal magnetic field. The magnetised horn system acts as a lens, focussing the charged particles that pass between the conductors. The direction of the horn current is normally chosen to focus positively charged hadrons.

The focal length of each horn is approximately proportional to the momentum of the incoming hadrons, so this means that the momentum of the focussed pions and kaons is determined by the relative positions of the target and horns. As the neutrino beam is produced through decay of the focussed hadrons, moving the target closer to the horns will shift the typical neutrino energy to lower values.

The target can be moved up to 2.5 m along the beam axis, allowing a large range of energy spectra to be obtained. Figure 3.5 shows the expected rates and energies of charged current ν_μ interactions in the Near Detector for different configurations of horn current and target position. The results of the Super-Kamiokande experiment suggest that the low energy configuration is optimal, as this maximises the event rate in the energy range where the largest oscillation signal is expected. In the low energy configuration, the target is placed 10 cm upstream from its default position inside the

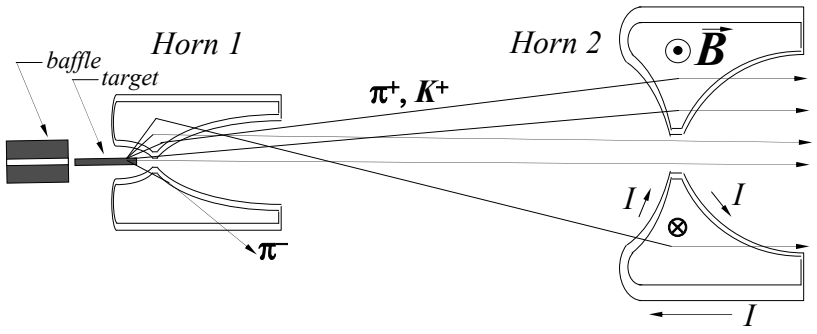

Figure 3.4: Hadrons produced in the NuMI target are focussed by a system of two magnetic horns. The horns are separated by 10 m and each consists of an outer cylindrical conductor and an inner parabolic conductor. Taken from [53].

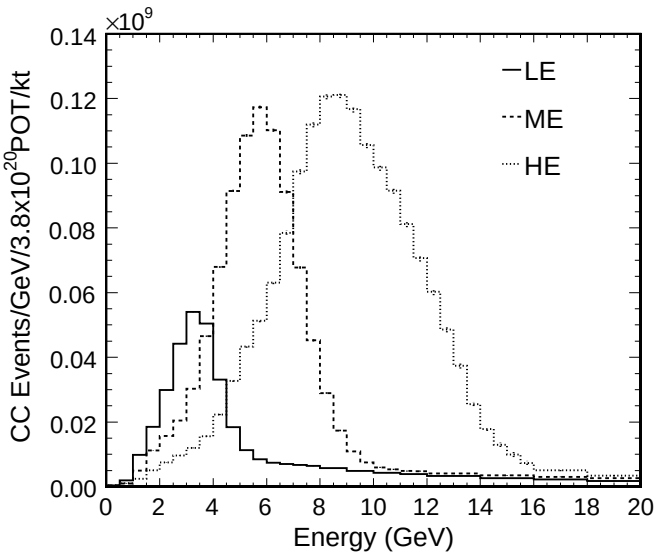

Figure 3.5: Expected rate of CC ν_μ interactions in the Near Detector for three different target/horn configurations. The configurations shown are Low Energy (LE: target 10 cm upstream of first horn, 185 kA horn current), Medium Energy (ME: 100 cm/200 kA) and High Energy (HE: 250 cm/200 kA). Taken from [55].

first horn and the horn current is 185 kA. This produces a Near Detector spectrum that peaks at approximately 3 GeV.

After the hadrons have been focussed, they pass into an evacuated decay pipe. This is 2 m in diameter and 675 m in length, making it slightly longer than the decay length of a 10 GeV pion. In the decay pipe, the pions and kaons undergo $\pi^+ \to \mu^+\nu_\mu$ and $K^+ \to \mu^+\nu_\mu$ decays, producing the beam of ν_μ.

At the end of the decay pipe, the remaining hadrons are removed by an absorber. This consists of a water-cooled steel and aluminium core, followed by steel and concrete blocks. After the absorber, the beam passes through 240 m of rock before reaching the Near Detector, thus removing any decay pipe muons. When it reaches the Near Detector, the beam composition is 92.9% ν_μ, 5.8% $\bar{\nu}_\mu$ and 1.3% $\nu_e/\bar{\nu}_e$[56]. The backgrounds in the ν_μ beam are due to kaon decays via $K^+ \to \pi^0 e^+ \nu_e$, muon decays via $\mu^+ \to e^+ \bar{\nu}_\mu \nu_e$ and the decays of any remaining negatively charged hadrons.

During beam operations, important monitoring information is collected. This includes details of the beam position and spot size (measured using secondary emission monitors) and the beam intensity (measured using toroidal beam current transformers). Arrays of ionisation detectors also monitor the profiles of the hadrons and muons at the end of the decay pipe, providing useful information about the beam alignment and focussing.

3.2 The MINOS Detectors

The two MINOS detectors are designed to be as similar as possible. By ensuring that the detectors respond to neutrino interactions in the same way, the systematic uncertainties in the neutrino flux, cross-sections and detector acceptance have a reduced impact on the oscillation analysis. Both detectors are steel/scintillator sampling calorimeters, which are capable of observing neutrino interactions with visible energy greater than 500 MeV. Each detector has a toroidal magnetic field, allowing reconstruction of the charge sign and momentum of muons based on their curvature. This provides a means of distinguishing between ν_μ and $\bar{\nu}_\mu$ CC interactions.

A comprehensive description of the detectors is provided in [57], whilst [53] provides a useful summary. In this Section, the important features of the detectors are highlighted. The common features of the two detectors are discussed first, before an examination of the details specific to each detector.

The active medium in the MINOS detectors consists of strips of plastic scintillator. These are 4.1 cm wide, 1.0 cm thick and can be up to 8 m long. By placing these strips side by side and encasing them in aluminium sheets, light-tight modules of typically 20 or 28 strips are created. These are arranged to form scintillator planes, which are attached to 2.54 cm thick plates of steel.

The detectors consist of arrangements of these steel/scintillator planes, which are aligned vertically with a spacing of 5.94 cm. As illustrated in Figure 3.6, the orientation

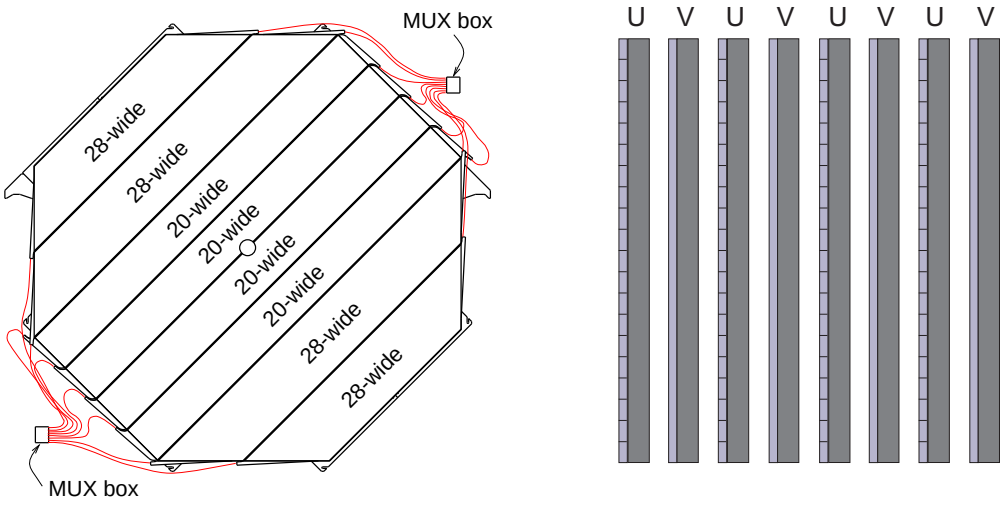

(a) (b)

Figure 3.6: (*a*) Construction of a Far Detector scintillator plane from modules of 20 and 28 scintillator strips. (*b*) Arrangement of steel/scintillator planes in the detector, illustrating the two different strip orientations. Adapted from [57].

of the strips alternates from $+45°$ to $-45°$ between successive planes. These orthogonal strip orientations allow three-dimensional event reconstruction and are referred to as the u and v 'views'.

A cross-section through a scintillator strip is shown in Figure 3.7. The strips contain a 2.0 mm wide and 2.0 mm deep groove and, aside from the groove, the outside of the strips are covered in a 0.25 mm reflective layer. This reflective layer consists of polystyrene doped with TiO_2. A 1.2 mm diameter wavelength-shifting (WLS) fibre is glued into the groove, which is then sealed with reflective tape.

When a charged particle passes through a strip, excitations of the scintillator material produce scintillation light. This light is collected by the WLS fibre and is transported to the end of the strip without significant re-absorption. At the end of the WLS fibre, the light is transferred to highly transparent clear fibres and then to multi-anode photomultipliers (PMTs). This optical readout system is illustrated in Figure 3.8.

The PMTs convert the light signals from the scintillator into photoelectrons, which are multiplied by secondary emission in the dynode chain. The output from the PMTs is then digitised by sensitive readout electronics, which are controlled by the data acqui-

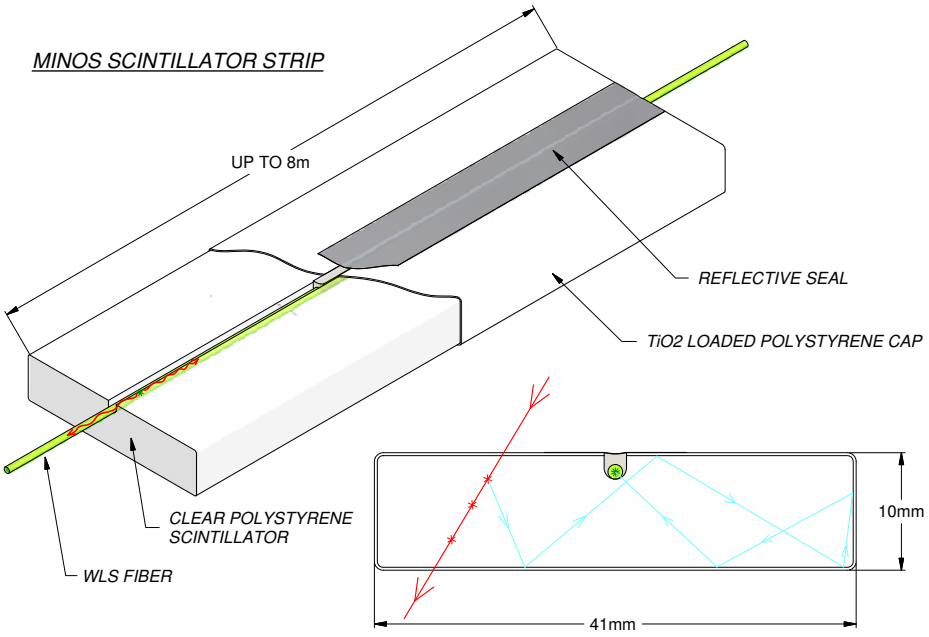

Figure 3.7: Cross section through a scintillator strip. Scintillation light produced by ionising particles is reflected by the outer layer of the strip and collected by an embedded wavelength-shifting fibre. Taken from [57].

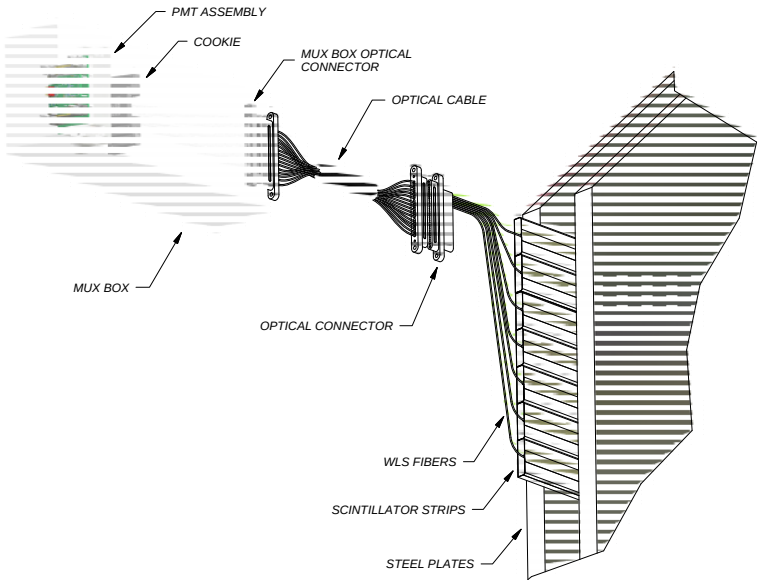

Figure 3.8: The optical readout system in the MINOS detectors. Scintillation light captured by the wavelength-shifting fibres is transferred to clear fibres and then to multi-anode photomultipliers. Taken from [57].

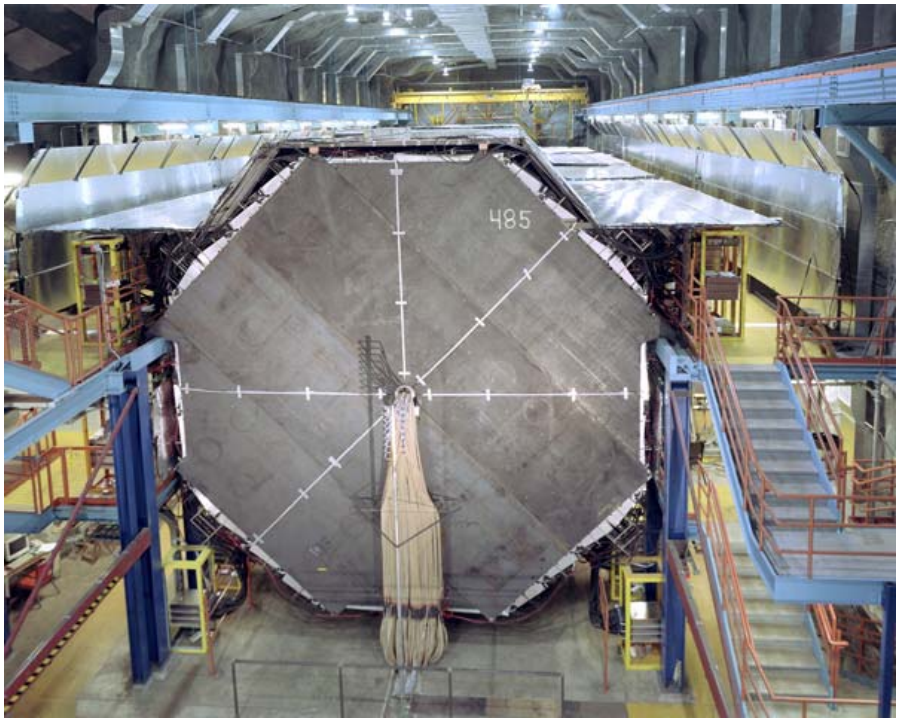

Figure 3.9: The MINOS Far Detector.

sition (DAQ) and timing systems[58]. The different rate of neutrino interactions at the Near and Far Detectors means that the two detectors have different readout electronics.

3.2.1 Far Detector

The MINOS Far Detector is shown in Figure 3.9. It is located in the Soudan Underground Laboratory at a depth of 705 m and a distance of 735 km from the NuMI target. The detector has a mass of 5.4 kton and consists of 486 octagonal steel planes, which are 8 m wide. The planes are divided into two 'supermodules' (of 249 and 237 planes) and, aside from the first plane in each supermodule, every plane is fully instrumented with scintillator. The supermodules are separated by a 1.1 m air gap, so the detector has a total length of approximately 30 m.

The two supermodules are independently magnetised by current-carrying coils that pass through a hole at the centre of the planes. The coils provide 15 kA-turns and induce an average field of 1.27 T in the steel. The direction of the coil current is normally chosen so as to focus negatively charged particles. By focussing negatively charged muons, the fraction of beam ν_μ interactions that are fully contained in the detector is increased.

The scintillator planes consist of 192 strips, which are up to 8 m in length and cover

Figure 3.10: The MINOS Near Detector.

the steel octagons. The strips are read out at both ends by Hamamatsu 16-anode PMTs. However, to reduce the detector instrumentation, the signals from eight different strip ends (separated by 1 m and in the same detector plane) are multiplexed and read out by the same PMT pixel. The multiplexing pattern is deliberately different for the two sides of the detector, allowing the resulting ambiguities to be resolved.

Details of the Far Detector front-end electronics can be found in [59], whilst a description of the DAQ is provided in [58].

3.2.2 Near Detector

The MINOS Near Detector, shown in Figure 3.10, is located 104 m underground at Fermilab, 1040 m downstream of the NuMI target. The detector has a mass of 0.98 kton and consists of 282 steel planes arranged in a single magnetised module. The steel planes are 3.8 m high, 4.8 m wide and have an elongated octagonal shape. The magnetic field is provided by a 40 kA-turns coil that passes through a square hole in each plane. The Near Detector coordinate system is centred on this coil hole, which is offset horizontally 0.56 m from the centre of the detector.

The neutrino beam has a diameter of approximately 50 cm when it reaches the Near Detector, and the centre of the beam strikes the detector a horizontal distance of 1.49 m

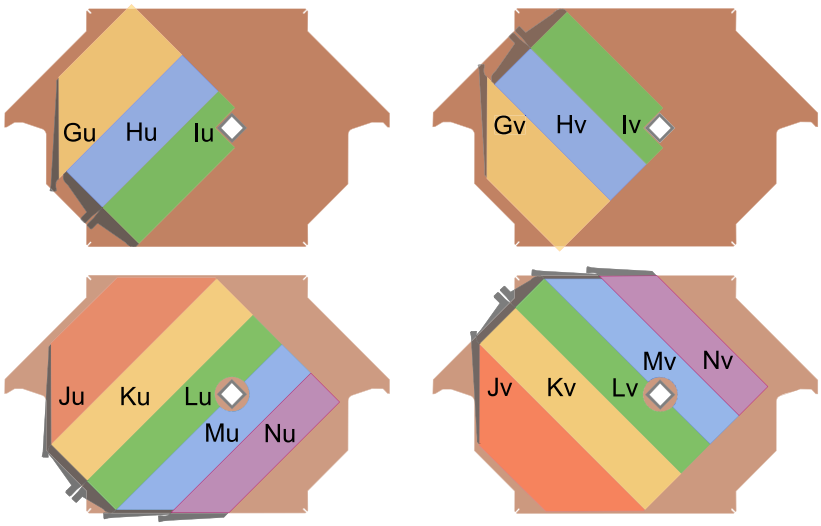

Figure 3.11: The four different configurations of scintillator plane used in the Near Detector. The upper two figures show partially instrumented u and v planes, whilst the lower two figures show fully instrumented u and v planes. The labels G-N denote the different shapes of the scintillator modules.

from the coil hole. The detector is designed to contain neutrino interactions around the beam spot and the magnetic field induced in this region is 1.17 T. The instrumentation in this region is quite uniform and most closely resembles that at the Far Detector. Away from the beam spot, however, financial constraints introduced differences in the instrumentation.

The Near Detector contains two different types of scintillator plane: fully and partially instrumented. These different scintillator planes are illustrated in Figure 3.11. The fully instrumented planes consist of 96 scintillator strips and cover 13.2 m^2 of the steel planes. Every fifth plane along the entire length of the detector is fully instrumented. The partially instrumented planes consist of only 64 scintillator strips and cover a reduced area of 6.0 m^2. Of the first 121 steel planes, four out of every five are partially instrumented. At higher plane numbers there are no partially instrumented planes and four out of five planes remain uninstrumented. This arrangement of scintillator planes is illustrated in Figure 3.12.

In order to reduce the required electronics, Near Detector scintillator strips are only read out at one strip end, using Hamamatsu 64-anode PMTs. This is possible because the Near Detector scintillator planes are smaller than those in the Far Detector. The strips are therefore shorter and the attenuation of scintillation light is reduced. To

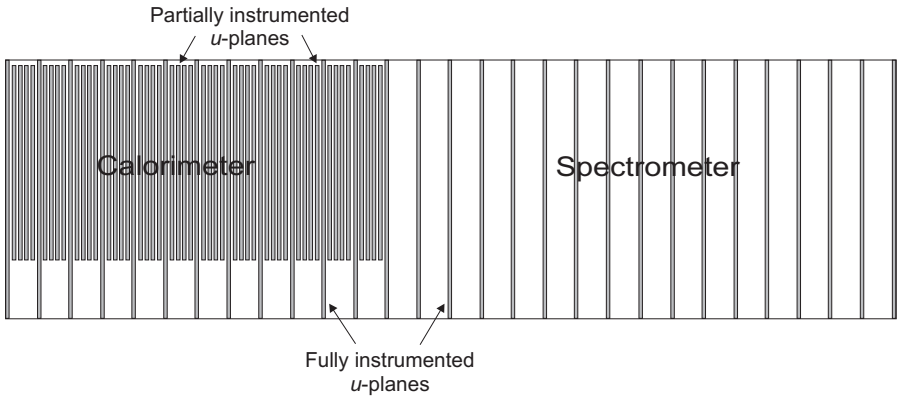

Figure 3.12: The arrangement of scintillator planes in the Near Detector, illustrated for u-type planes.

increase the light yield, the uninstrumented strip end is covered with reflective tape.

The Near Detector is arranged in two sections. The first 121 steel planes form the calorimeter section, whilst the remaining 161 steel planes form the muon spectrometer. The steel planes in the calorimeter sections are all instrumented with scintillator and there is no multiplexing of the strip ends in the readout electronics. The calorimeter is intended to allow identification of neutrino interaction vertices and to contain any associated hadronic showers. It also allows accurate reconstruction of the upstream sections of muon tracks.

The muon spectrometer is intended to track the high energy muons produced by neutrino interactions in the calorimeter. Only one in every five steel planes are instrumented in this region, and the signals from four strip ends (widely separated in position, but from the same plane) are multiplexed and read out by the same PMT pixel. The resulting ambiguity can only be resolved by propagating the upstream sections of muon tracks into the spectrometer region.

The divergence of the neutrino beam means that the interaction rate at the Near Detector is much higher than that at the Far Detector. With the low energy target/horn configuration and a typical beam intensity of 2.2×10^{13} PoT/spill, 16 beam neutrino interactions per spill are expected in the entire length of the Near Detector. Of these, about half will occur in the calorimeter region.

To allow accurate reconstruction of the multiple events recorded during each spill, the Near Detector uses specially designed readout electronics [60]. These high-speed electronics are capable of continuous (no dead-time) processing during the $10\,\mu$s beam spills. The charge from the PMTs is divided into 19 ns timing 'buckets', and this allows for an efficient separation of events occurring at similar times.

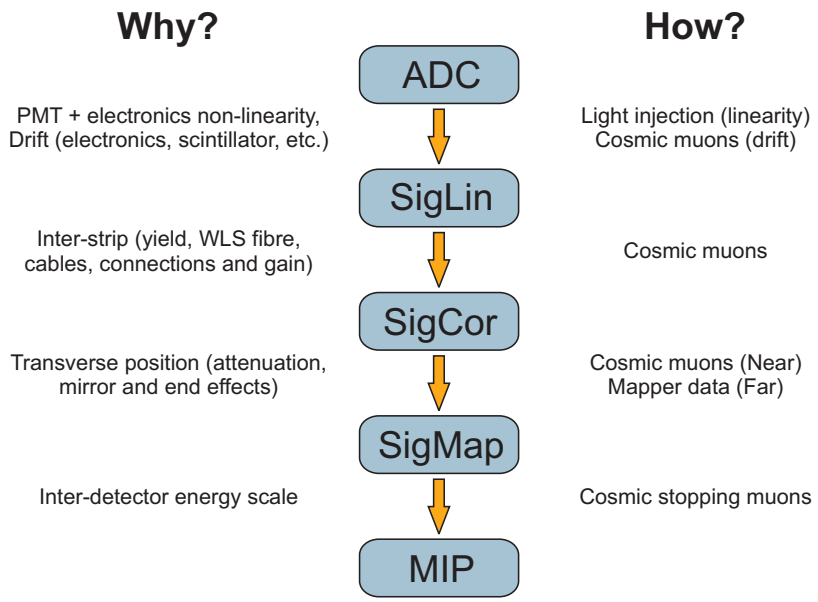

Figure 3.13: Summary of the different stages in the detector calibration scheme. Details of the motivation and the implementation for each stage are indicated. Adapted from [62].

3.3 Detector Calibration

In order to make a precise measurement of the neutrino oscillation parameters, the MINOS experiment must accurately reconstruct the neutrino energy spectra in the Near and Far Detectors. This requires measurements of the absolute muon energies and absolute hadronic shower energies in the observed neutrino interactions. These interactions are collected over long periods of time and are located in different regions of two different detectors, so the energy reconstruction is complicated by large variations in detector response.

To allow accurate energy reconstruction, a calibration process is required to remove the temporal and spatial variations in the detector responses and to relate these responses to the true energy deposits. An overview of the calibration procedure is illustrated in Figure 3.13, whilst detailed descriptions are provided in [61, 62].

The different steps in the calibration procedure are summarised below:

- **Drift calibration.** This corrects for shifts in the detector response over time. Such shifts may be due to temperature variations, ageing of the scintillator and PMTs or replacement of components in the electronics. The drifts are monitored by recording the energy deposited per plane by through-going cosmic muons. The median energy deposited per plane is calculated each day and this value tracks

any changes in the detector response.

- **Linearity calibration.** In order to correctly measure large energy deposits, calibration is required to ensure linear response of the PMTs and electronics. The electronics systems are linearised using charge injection circuits, whilst the PMT responses are linearised using a light injection (LI) system. The LI system uses UV LEDs to inject controlled pulses of light into the WLS fibres, in order to mimic the signals produced by energy deposits in the scintillator.

  - The **Far Detector** PMTs are expected to have linear responses up to about 100 photoelectrons, but become non-linear at higher pulse heights due to space charge effects. To perform a calibration[63], the linearity at low pulse heights is first established by fitting data recorded when light is injected at the opposite end of the fibre to the readout. These data then provide a linear scale to correct the PMT responses when light is injected at the same end of the fibre to the readout.
  - At the **Near Detector**, the intensity of light from each LED in the LI system is monitored using Positive-Intrinsic-Negative (PIN) photodiodes. These PIN diodes are sufficiently linear to allow PMT linearity corrections to be obtained through comparison of PMT data with the corresponding PIN diode data. The relationship between the PMT responses and the PIN diode responses is closely approximated by a second order polynomial[64].

- **Strip-to-strip calibration.** This normalises the response of all the scintillator strips in a detector, addressing issues such as differences in light output between strips, different PMT or electronics gains and differences in the WLS fibres. The calibration is performed by monitoring the response to through-going cosmic muons. After calibration, the response of each strip (of a particular length) is the same for a given energy deposit at the strip centre.

- **Attenuation correction.** This correction addresses the distance of an energy deposit from the centre of the strip. Scintillation light that travels for a longer distance in the WLS fibre is subject to greater attenuation. At the Far Detector, the attenuation along each WLS fibre was evaluated during construction by scanning a radioactive source along the length of each strip, allowing a correction to be parameterised. At the Near Detector, sufficient cosmic muon data is collected to allow the attenuation to be evaluated.

- **Inter-detector calibration.** This determines the relative calibrations of the Near and Far Detectors, allowing signals observed at the two detectors to be compared. The calibration is performed by measuring the detector responses to stopping cosmic muons, using only the portion of the track in which the muon energy should be between 0.5 GeV and 1.1 GeV. Using this region of track avoids the rapid changes in $\mathrm{d}E/\mathrm{d}x$ that occur at lower energies. The mean response is calculated for each strip and the median value is used to calculate a signal scale representing the entire detector.

An absolute energy calibration is possible through use of the data collected by a special MINOS calibration detector [65]. This was exposed to particle beams of known energy at the CERN PS accelerator and the data recorded were used to verify a MC simulation of hadronic and electromagnetic showers[61,66,67]. After the inter-detector calibration, the fully corrected PMT signals are expressed in muon equivalent units (MEU, also referred to as minimum ionising particles or MIP). The verified MC simulation is used to extract a MIP to GeV conversion, so that reconstructed shower energy measurements represent the real energy deposits.

3.4 Monte Carlo Simulation

The MINOS oscillation analysis makes extensive use of simulated data, which are used to produce predictions of the Far Detector neutrino energy spectrum. It is only through comparison of these predicted spectra with the real data spectrum that evidence for ν_μ disappearance and spectral distortion can be observed and a measurement of the oscillation parameters extracted. The simulated data are also used in the development of reconstruction software and event selection techniques, whilst further uses include the investigation of systematic uncertainties and the evaluation of sensitivities for different measurements.

To provide an accurate simulation, Monte Carlo methods are employed alongside a detailed model of the entire experiment; from hadron production in the target to the responses of the detectors to energy deposits from the final state particles.

Hadron production in the NuMI target is calculated using FLUKA[68], which is a general purpose tool used for the modelling of particle transport and interactions with matter. The particles that leave the target are recorded and are propagated in a GEANT3[69] simulation of the NuMI beamline. This beamline simulation includes the magnetic focussing horns, surrounding shielding and the decay pipe. The GEANT3/FLUKA code

is also responsible for modelling the hadronic interactions and decays that occur within the decay pipe. Information is stored for any of the decays that produce a neutrino, completing the simulation of the beam neutrino flux. The simulated neutrinos can be selected later, for input to event generation software.

The detector simulation, GMINOS, is based on GEANT3. It randomly samples neutrinos from the predicted beam flux and propagates them through the Near and Far Detector halls. Events are generated inside the detectors and also within the surrounding structures and rock. Details of the neutrino interactions are provided by interfacing with NEUGEN[70].

NEUGEN is a neutrino event generator and cross-section library that simulates neutrino-nucleus interactions in the energy range between 100 MeV and 100 GeV. It provides a list of the particles produced in the simulated interaction and this list is passed to GEANT3. The particles are then propagated through a detailed model of the detector geometry and a list of detector 'truth hits' is generated. The simulation of the hadronic interactions of the final state particles is provided by GCALOR[71].

The final step in the simulation is the transformation of the truth hits to a realistic detector output. This is provided by the 'PhotonTransport' and 'DetSim' software packages. PhotonTransport simulates the scintillator response to the true energy deposits and tracks the scintillation light through the WLS fibres and clear fibres to the PMTs. It then simulates the initial conversion of the light to photoelectrons. DetSim simulates the amplification of the photoelectron signal in the PMT dynode stages. It then models the front end electronics and DAQ for the detectors, outputting a RawDigitDataBlock exactly as would be produced by real data. This data block can be processed by the reconstruction software. DetSim is also responsible for the simulation of detector effects such as noise, readout non-linearities and PMT cross-talk (produced by the leakage of light onto adjacent PMT pixels).

3.5 Reconstruction Software

The raw data recorded by the detectors (or produced by the MC simulation) are processed with the MINOS event reconstruction software. This software examines the topology and timing of the low level detector hits in order to reconstruct higher level objects, such as muon tracks and hadronic showers. The tracks and showers are then used to identify individual neutrino interactions and to determine event kinematics, including the neutrino energy.

The reconstruction software consists of an ordered chain of algorithms, each of which

builds a set of candidate reconstructed objects (digits, strips, slices, tracks, showers and events) that are passed to the next algorithm in the chain. The different links in this chain are summarised below:

- **Digit formation.** The input to the reconstruction consists of the raw data collected during a particular block of time, such as a beam spill. This input is converted into a list of digits, each of which contains a digitised measurement of the pulse height recorded in an electronics channel, a digitised timing measurement and a list of possible associated strip ends. In the Far Detector, separate digits are formed for the scintillation light observed at the two ends of each strip, whilst the multiplexing means that every digit has eight possible associated strip ends.

- **Demultiplexing.** This stage of the reconstruction attempts to identify which of the eight strip ends associated with each Far Detector digit is really responsible for the scintillation light. The demultiplexing algorithm[72] compares digits from the two sides of the detector in order to produce a list of the possible solutions in each plane. Unambiguous digit combinations are identified whenever only one pairing of strip ends can place two digits on the same strip. These unambiguous combinations, together with timing information, are used to constrain the possible event region.

- **Strip formation.** The reconstruction software analyses the list of digits in order to form strip objects, which each represent a single energy deposit in a scintillator strip. At the Far Detector, strip objects are created by simply combining the digits recorded at the two ends of each scintillator strip. At the Near Detector, the scintillator strips are only read out at one end, but multiple digits can be produced for each energy deposit. Groups of digits originating at similar times from a single scintillator strip are combined to form a strip object.

- **Slicing.** The high event rate at the Near Detector means that multiple neutrino interactions are expected per spill. In order to simplify the reconstruction of individual events, the list of strip objects is divided into slices. Each slice contains strips that are closely associated in space and time and which are likely to originate from a single neutrino interaction. At the Far Detector, the event rate is so low that only one slice will typically be formed. The remaining reconstruction processes are all performed separately for each slice.

- **Track finding.** The primary signature of a charged current (CC) ν_μ interaction in the MINOS detectors is the presence of a muon track with a contained vertex.

It is therefore very important to identify the strips produced by muon tracks and to reconstruct the kinematic properties of the muons. The track reconstruction is performed in two stages: track finding and track fitting. The track finding algorithm analyses the topology of the strips in each slice in order to identify 'seed tracks' (which specify the basic structure of a muon track). The algorithm is described in detail in Chapter 4.

- **Track fitting.** In the second stage of track reconstruction, each seed track is used as the input to a Kalman filter. The Kalman filter uses knowledge of muon propagation in order to refine the track strip selection, extrapolate tracks into the Near Detector spectrometer and calculate many important track properties. The muon charge sign and momentum are estimated from the curvature of the track in the magnetic field, whilst a further momentum estimate is obtained from the track range. The track fitting algorithm is described in detail in Chapter 5.

- **Shower finding.** Within each slice, any strips that are clustered together in time and space are combined to form a shower. The energy of the shower can be calculated from the total pulse height deposited in the shower strips, after the subtraction of pulse height contributions from any tracks with which strips are shared. For accurate energy reconstruction in CC ν_μ interactions, the most important showers are the vertex hadronic showers associated with the interaction between the neutrino and nucleus.

- **Event formation.** The final step in the reconstruction process is the combination of tracks and showers to form events. The optimal combinations are identified by examining the spatial and temporal separations of the different tracks and showers in each slice. The reconstructed events represent the best estimates of the energy deposits associated with individual neutrino interactions. For a reconstructed CC ν_μ event, the neutrino energy is specified as the sum of the vertex shower energy and the energy of the longest track.

After the formation of events, the reconstruction process is complete. The details of each reconstructed event, including the estimated neutrino energy and properties of the tracks and showers, are recorded for use in an oscillation analysis.

Chapter 4

Track Finding

Interactions in the MINOS detectors consist of showers (localised clusters of hits) and tracks. The tracks that can be most clearly resolved in the detectors are those produced by muons. Figure 4.1 shows an example event display for a charged current (CC) ν_μ interaction, in which the hadronic shower and the muon track can be clearly distinguished. In order to perform an oscillation analysis, it is vital that the muon track strips in an event are identified; the successful reconstruction of a track is used as the basis for selecting a sample of CC ν_μ interactions and there are many important kinematic properties that can be derived from the track. The track reconstruction must therefore be both accurate and efficient.

As mentioned in Section 3.5, track reconstruction in MINOS is performed in two distinct stages: track finding and track fitting. Both these stages have been fully implemented as part of this thesis. The new software developed is now the standard MINOS track reconstruction, used by the collaboration in all its physics analyses. The software represents a significant upgrade to the previous MINOS code, providing a more accurate and more efficient reconstruction of important physics quantities (such as the muon momentum), whilst reducing reconstruction failures.

The new track finding software uses the topology of the reconstructed strips in order to identify muon 'seed tracks', which specify the basic structure of a muon trajectory. This track finding software is described in detail throughout this Chapter. The new track fitting software is described in Chapter 5, which also includes a brief performance comparison with the previous MINOS reconstruction.

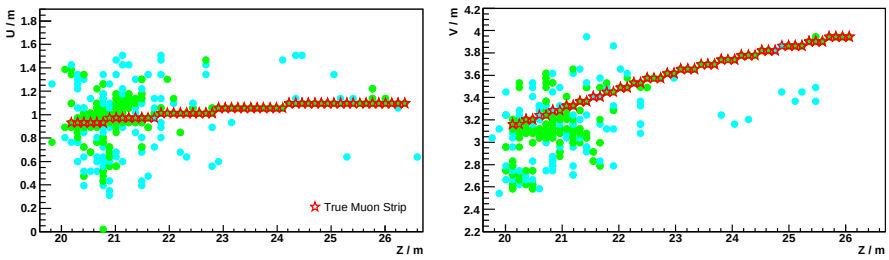

Figure 4.1: Event display for a typical CC beam ν_μ interaction. The track reconstruction software attempts to identify the true muon strips.

4.1 How the Algorithm Works

4.1.1 The General Aim

Many different track finding algorithms exist and the first task in writing a new reconstruction package was to select an algorithm, based on its effectiveness, simplicity and ease of application to the MINOS detectors. The aim was to engineer software that would be effective for all classes of event in both the Near and Far detectors. The same code should cope with steep cosmic muon events and with shallow beam neutrino events, where the muon can disappear into the coil hole region for many planes.

The algorithm selected has previously been implemented to detect atmospheric neutrino interactions at the Far Detector [28] and works effectively for this class of events. Whilst this first implementation does not work for beam neutrino interactions or with the different instrumentation of the Near Detector, the algorithm itself is simple and flexible. The algorithm operates by identifying the combinations of reconstructed strips that could represent small segments of track. It then looks for possible 'associations' between the small track segments and chooses the best segments to join together, gradually constructing the track.

The basic steps in the algorithm are shown in Figure 4.2. These steps are described in detail in this Chapter. Before implementing the algorithm, a decision was made to optimise the track finding efficiency. This means that the software should identify the muon tracks in even very low energy events and high y events (where the majority of the neutrino energy goes into the hadronic shower). The disadvantage of this high efficiency is that the algorithm will also tend to find the most plausible track in neutral current (NC) events. The approach is to accept a large number of tracks in NC events and to use the techniques discussed in Chapter 6 to separate CC and NC events by identifying the differences between muon tracks and 'fake' tracks.

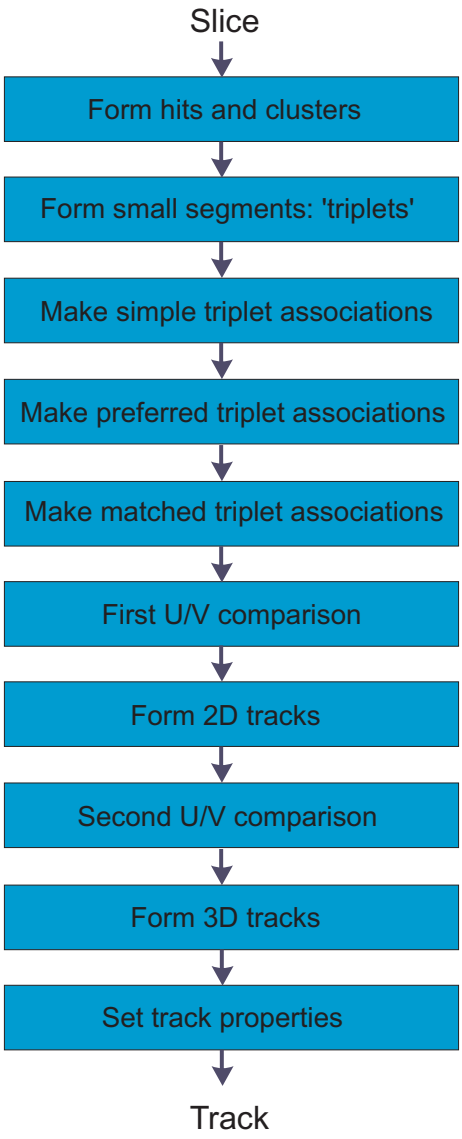

Figure 4.2: Summary flow diagram, showing the different steps in the track finding algorithm.

4.1.2 Hit and Cluster Formation

The input to the track finding algorithm is the list of strips in a slice. When a new slice is received, the first task is the formation of 'Hits'. These are basically equivalent to strips, and they are the smallest building blocks for the track. Hits are created for all strips in the slice, except those below a low pulse height threshold (two photoelectrons) or those flagged as cross-talk by the Far Detector demultiplexing software. Hits are not created for strips in the Near Detector spectrometer, as the track finding software does not consider this region. The identification and demultiplexing of spectrometer track strips is instead performed by the track fitting software, as described in Section 5.3.

After Hit formation, groups of adjacent Hits on a detector plane are combined to form 'Clusters'. The subsequent steps in the track finding algorithm use these Clusters, rather than the Hits, in order to reduce the complexity of events. After the final track Clusters have been identified, linear fits can be used to select the best hits from each track Cluster. Once all the Clusters have been formed, an estimate is made as to whether each Cluster is track-like or shower-like. This estimate is based on the Cluster width and the fraction of the plane's charge that it contains. The proximity of Clusters on nearby planes is also examined; densely packed Clusters are flagged as shower-like.

4.1.3 Triplet Formation

The track finding begins in earnest with the formation of 'Triplets', which are small segments of track containing three Clusters on separate planes. Triplets are formed separately for the u and v views, and the different possible configurations are illustrated in Figure 4.3. The algorithm firstly tries to form the most simple Triplets, which have Clusters on adjacent planes in the same view. For a given Cluster on plane i, all the combinations with Clusters on planes $i-2$ and $i+2$ are evaluated; for each combination, the beginning and end transverse positions of the Clusters are compared. If the extremal Clusters are within a few strip widths of the central Cluster, the Clusters could plausibly be part of a muon track and a Triplet is created.

Once the simple Triplets have been formed, the code starts to look for more complicated configurations and searches for Triplets that have one or more gap planes. This identification proceeds largely as for the simple Triplets, but important checks are made to ensure that any gap planes are genuine; Triplets spanning smaller numbers of planes are preferred and the algorithm must avoid creating Triplets that overlap. The Triplets are important objects and the next sections of the algorithm are concerned with the associations between nearby Triplets. Triplets are examples of TrackSegment objects,

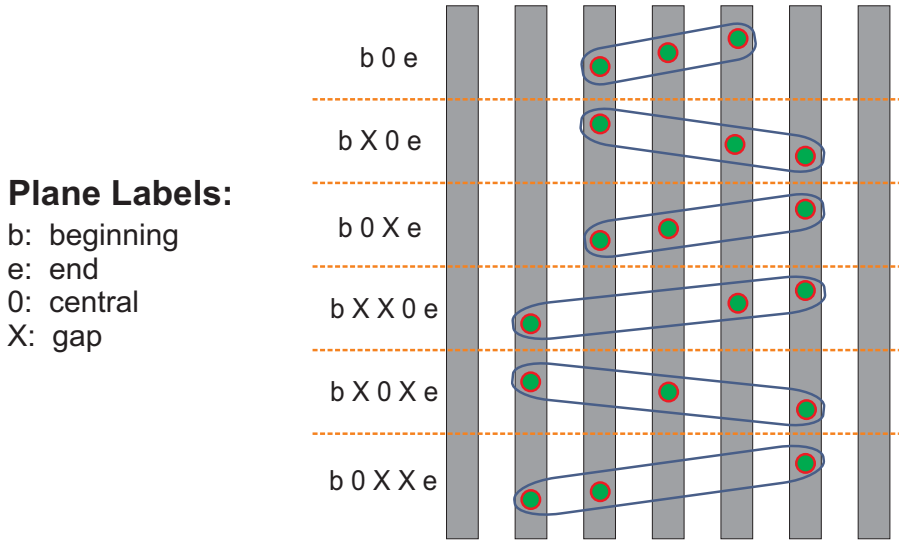

Figure 4.3: The different configurations of Clusters that can be used to form Triplets.

which can be associated with each other or joined together, and for which properties such as beginning and end directions can be calculated.

4.1.4 Simple Associations

To help choose which Triplets to join together, there are three levels of association that can be made between TrackSegment objects. The possible criteria for a simple association between two Triplets are summarised in Figure 4.4. The most obvious association is for the Triplets to overlap and share two Clusters, with the unshared Cluster in each Triplet being that which is furthest from the other Triplet. The second possible method of association is for the two Triplets to share a beginning/end Cluster and the remaining Clusters to obey transverse position criteria that would indicate that the Triplets are compatible (and not oriented in very different directions).

The third method of association is for the two Triplets to have compatible beginning/end directions. This method requires no overlap for the Triplets and is useful, for example, in associating TrackSegments separated by the coil hole. To test for this compatibility the beginning/end directions are used to extrapolate a straight line from each Triplet. If the line from the end (beginning) of one Triplet passes suitably close to the beginning (end) of the other Triplet, the Triplets can be declared associated (the proximity required depends on the difference in z position between the Triplets, but is of the order of $1-2$ strip widths).

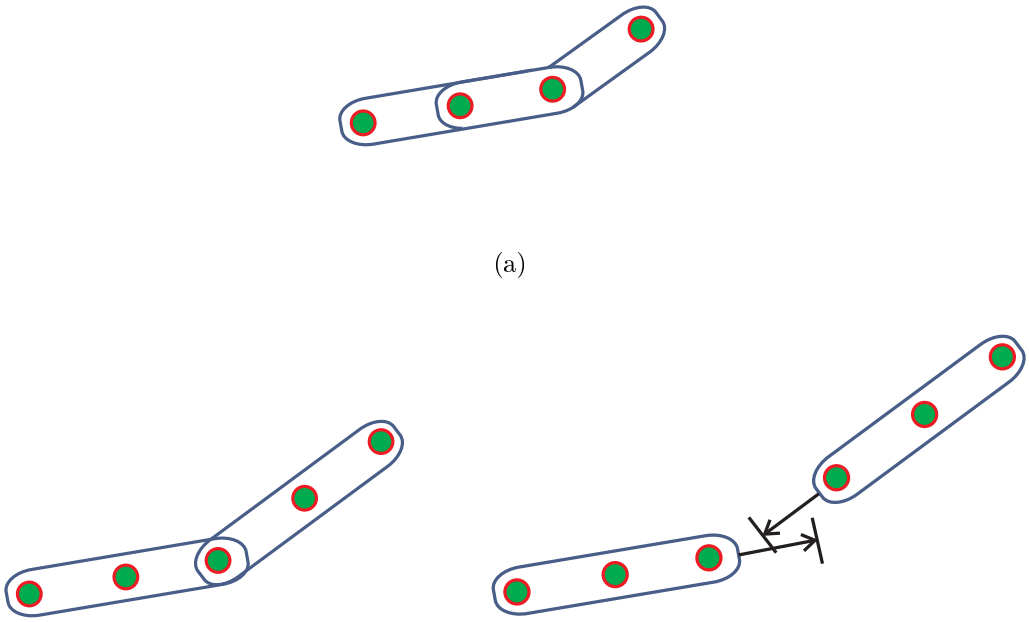

(a)

(b)

(c)

Figure 4.4: The criteria used in the identification of associations between Triplets: (*a*) the Triplets share two Clusters, (*b*) the Triplets share one Cluster and the remaining Clusters are sufficiently close, (*c*) the Triplets have compatible beginning/end directions.

4.1.5 Preferred Associations

From the list of possible associations between Triplets, the algorithm next tries to identify the associations that are most track-like and hence 'preferred'. It proceeds to search for chains of associated Triplets. For a given Triplet, it is known which Triplets are associated with its beginning and which are associated with its end. If these beginning and end Triplets are themselves associated, a chain of track-like segments has most likely been found. A preferred association can therefore be made, as illustrated in Figure 4.5.

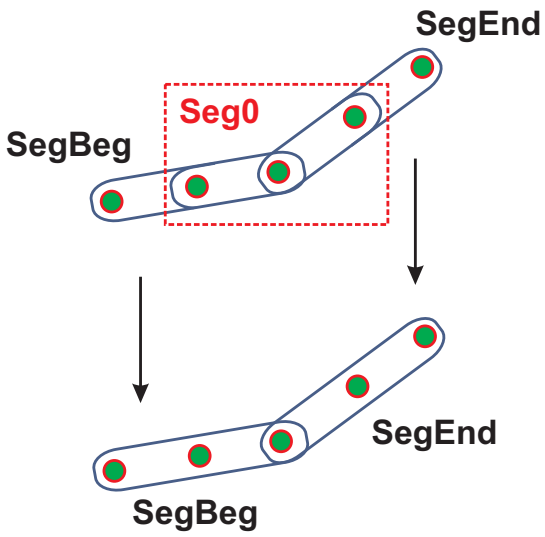

Figure 4.5: The identification of preferred associations between Triplets. In this example, SegBeg has an end association with Seg0, whilst SegEnd has a beginning association with Seg0. SegBeg and SegEnd are themselves associated, so a preferred association is formed between SegBeg and Seg0.

4.1.6 Matched Associations

The next procedure in the algorithm is to identify any long chains of Triplets with preferred associations. Such chains are strong candidates for muon tracks. If the Triplets in these chains have only one preferred beginning association and one preferred end association (i.e. there are no branches in the chain), they can be safely joined together to form a longer TrackSegment. If there are possible branches, 'matched' associations are formed between the relevant TrackSegments, as illustrated in Figure 4.6. Matched associations are also made between TrackSegments believed to be associated across the detector coil hole.

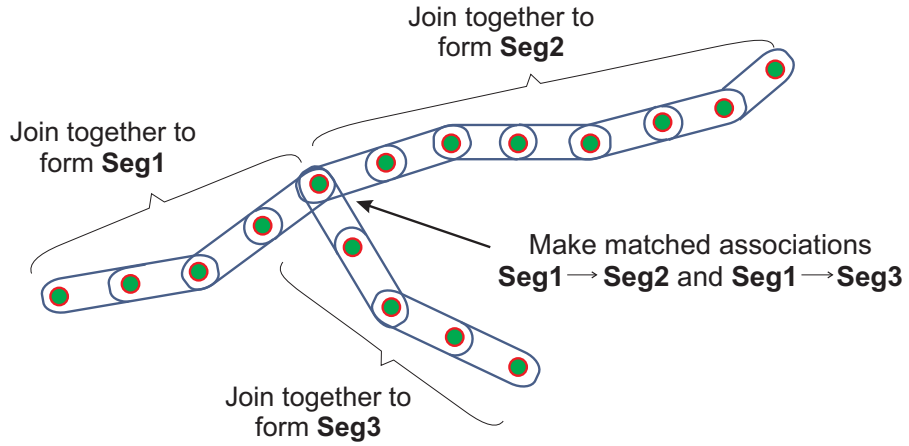

Figure 4.6: Triplets with single preferred beginning and end associations are joined together, creating chains of TrackSegments. Where these chains branch, matched associations are formed between the TrackSegments.

4.1.7 Seed Segment Selection

By this stage in the algorithm, a large amount of information has been generated about the associations between TrackSegments. The algorithm now focusses on finding the longest chains of TrackSegments linked by matched associations. These chains are the most likely candidates for '2D Tracks', which specify the track Clusters in a particular view. The approach to identifying the 2D Tracks is to locate 'seed segments', which are TrackSegments from which it is possible to move back and forth across a large number of planes, *along only paths of matched associations*.

The ideal candidates for seed segments are the 'clean' parts of the muon track, where there is little shower activity. To find these segments, the initial flagging of the Clusters as track-like or shower-like is used. If a segment is identified as particularly track-like, containing narrow Clusters in regions of low Cluster density, the segment is marked as an ideal seed. If an ideal seed segment is found in one view, the other view is examined to see if there is a corresponding segment that overlaps in z position and which could also be used as a seed. This is the first point at which information from one view can influence reconstruction in the other.

An ordered list of possible seed segments is therefore created, with the ideal seed segments placed first in the list. As described in Section 4.1.8, the algorithm considers each seed segment in turn and investigates the extent of the possible paths of matched associations. The longest paths are flagged as possible 2D Tracks. Once a segment has been considered as part of a path of matched associations, it is marked so that it is not considered again. Thus, by ordering the seed segments, a good 2D Track is typically found very quickly and, unless there is a promising second track, the remaining seed

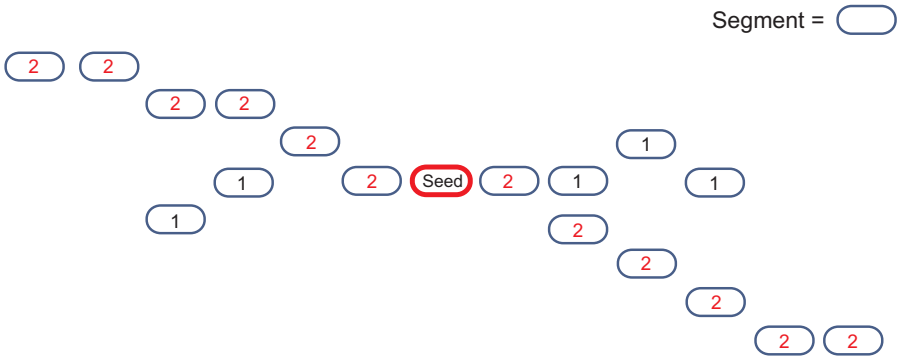

Figure 4.7: The identification of 2D Tracks from the possible chains of Track-Segments. By propagating back and forth from a seed segment, along paths of matched association, the segments in the longest 2D Track are flagged with a '2'.

segments quickly come to nothing; the same 2D Tracks are not found by multiple seeds. To increase the tracking efficiency, all segments have a chance to act as a seed if no ideal seed segments can be found.

4.1.8 2D Track Formation

For each seed segment, the code identifies the longest possible chains of matched associations. It does this by finding all the segments connected to the seed via matched associations; these connected segments are flagged with a '1'. The segments that are far away from the seed, and which could represent the possible beginning/end points of 2D Tracks, are then identified. The straightest paths (via matched associations) from these extremal segments back to the seed are then calculated and the segments in these paths are flagged with a '2'. These segments represent the constituents of the most plausible 2D Tracks, as illustrated in Figure 4.7.

The algorithm typically finds a selection of possible 2D Tracks, often originating from multiple different seed segments. The best possible 2D Track for each seed segment is identified by assessing the number of constituent Clusters and the 'straightness' of each 2D Track candidate, as illustrated in Figure 4.8. The number of Clusters is the most important quantity; the straightness score is calculated using a 'sliding' linear fit and is intended only to distinguish between 2D Tracks of equal size. For each seed segment, only the best 2D Track is stored for later use in the algorithm.

After the 2D Tracks have been created and stored, the algorithm investigates the opportunities for joining 2D Tracks together. The possible combinations of 2D Tracks in each view are considered, and the combination with the smallest beginning/end separation is identified. If these 2D Tracks have compatible beginning/end directions, the tracks are joined. The combinations of 2D Tracks are then recomputed and the pro-

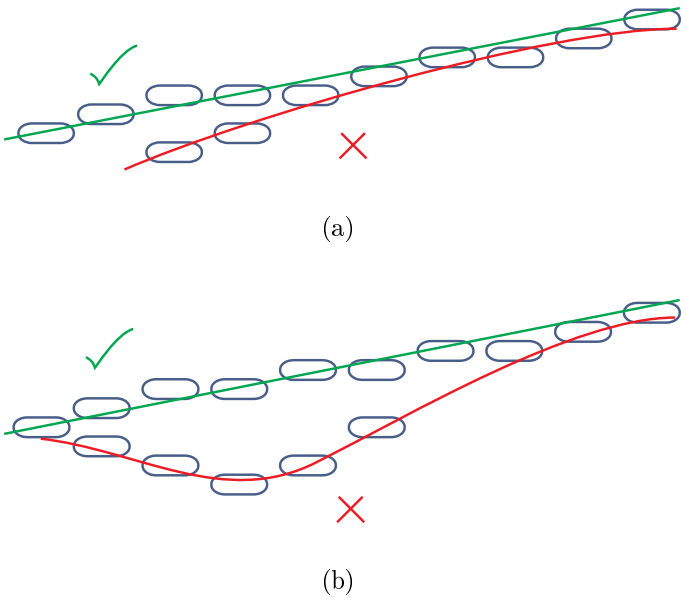

(a)

(b)

Figure 4.8: The best 2D Track in each view is identified by considering, (a), the number of clusters in each 2D Track and, (b), how straight each 2D Track is.

cess continues. Care is taken to ensure that distinct tracks are not mistakenly joined together; the algorithm uses information from both the u and v views to assess whether there is genuinely more than one track before it will make a join.

4.1.9 3D Track Formation

By this stage, the algorithm has created a list of complete 2D Tracks in the u and v views. The next procedure is to find the best pairings of the 2D Tracks between the two views. The optimal combinations are selected by evaluating the degree of plane overlap. To quantify this overlap, the u track and v tracks are examined and the internal and external beginning/end planes are identified, as shown in Figure 4.9. The difference between the internal beginning/end planes is labelled ΔPlane1, whilst the difference between the external planes is labelled ΔPlane2. A figure of merit is defined, $(\Delta\text{Plane1})^2/\Delta\text{Plane2}$, which is greater for better overlap.

With this u/v matching, the identification of the track Clusters is complete. Multiple tracks may have been identified in the slice, and each will consist of a list of Clusters in the u and v views. The tracks therefore have an appearance resembling that in Figure 4.10. The remaining steps in the algorithm must now address the selection of track Hits from the Clusters. The Clusters can be quite wide in shower-like regions, so this is often a complex task.

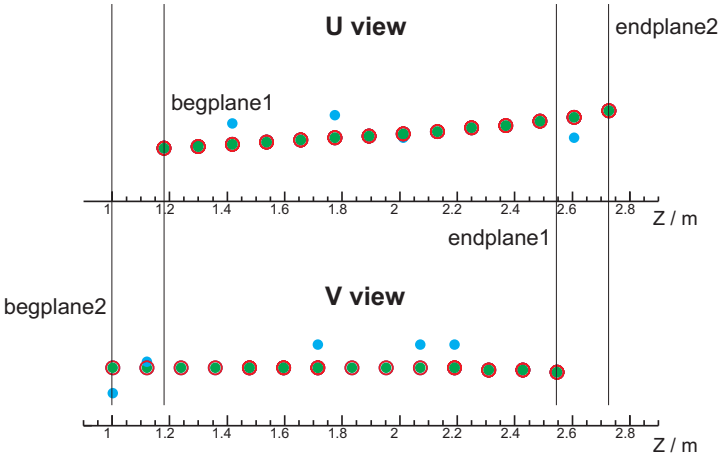

Figure 4.9: The plane definitions used to evaluate the overlap between the 2D Tracks in the u and v views.

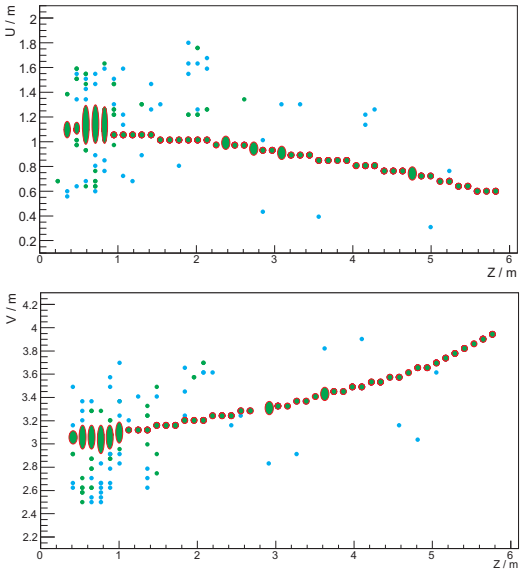

Figure 4.10: Event display for typical 2D Tracks in the u and v views. At this stage in the algorithm, the tracks consist of Clusters, which may contain multiple Hits in shower-like regions.

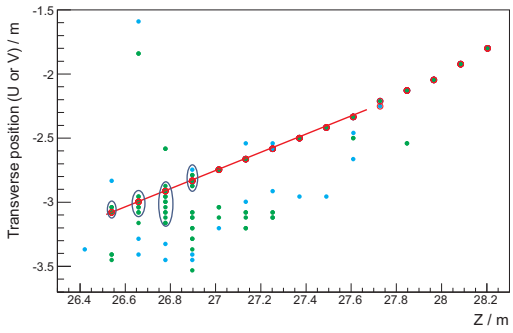

Figure 4.11: The selection of Hits from track Clusters, using linear fits and information from the 'clean' sections of the track.

4.1.10 The Final Track

The last section of the algorithm selects the best track Hits from the newly identified track Clusters. Its approach is simply to use linear fits through these Clusters. However, before attempting to find the final seed track strips, the algorithm tries to refine any track Clusters that are very wide and shower-like. It does this by using information from nearby track-like Clusters:

1. If there are track-like Clusters on either side of a shower-like Cluster, a linear fit through the shower region is attempted.

2. If there are track-like Clusters only on one side of a shower-like Cluster, a linear fit through the track-like Clusters is propagated into the shower region.

3. If the track in the given view consists entirely of shower-like Clusters, a linear fit through all the shower-like Clusters is attempted.

If one of these linear fits intercepts Hits in the shower-like Cluster, a new Cluster is made from these Hits. This 'refined' Cluster overrides the existing track Cluster on the plane. The algorithm is then ready to find the final seed track Hits, using only refined Clusters or the Clusters on track-like planes. This final selection again uses linear fits. However, all the Hits within a small plane range are now used, and each Hit is weighted by the fraction of its Cluster's charge that it represents. The intercepted Hits are the final seed track Hits. An example of selecting the best Hits from track Clusters is illustrated in Figure 4.11.

The seed track is completed by attempts to fill any obvious gaps in the track (using linear fits over the gap region) and by attempts to extend the beginning or end of the

track. These extensions are attempted using linear fits and are useful for including isolated Hits that could not form part of a Triplet. Small groups of Hits that were missed across the Far Detector supermodule gap can also be included in this way. The extension uses pulse height information to help avoid adding stray noise hits.

4.1.11 Track Properties

After identification of the tracks in a slice, a number of useful properties are calculated for each track. Firstly, interpolation between the strip positions in each view is used to specify u and v coordinates for the muon at every track plane. These coordinates are then used to calculate values for the pathlength and range (the path integral of the density) of the muon at each plane. Knowledge of the longitudinal positions of the muon in the strips also allows corrections to be made to the recorded strip times, providing useful muon timing information.

The variation of the strip times with the muon pathlength can be used to identify the direction of the track in the detector. Two linear timing fits are performed, as developed in [28]. In the first fit, the muon is constrained to travel at speed c in a direction such that the detector z coordinate increases with time. In the second timing fit, the muon is constrained to travel at speed c in the opposite direction. For each direction, an RMS value is calculated to describe how closely the direction hypothesis matches the measurements. The smallest RMS value identifies the most likely track direction and allows the vertex and end coordinates and direction cosines to be assigned.

Finally, an initial estimate of the muon momentum is obtained, using a simple parameterisation of the track range. With the calculation of these properties, the seed tracks are completed and are ready for use as the input to the Kalman filter track fitter, described in Chapter 5.

4.2 Near Detector Modifications

For the purpose of track reconstruction, the Near Detector can be divided into two sections, each of which must be treated separately:

1. **The 'forward' section** of the detector consists of the first 120 planes. In this region, only every fifth plane is fully instrumented. The other planes are partially instrumented, with only enough scintillator to cover the region around the beam spot. In this region of the detector, the reconstruction therefore needs to deal with tracks that pass through areas covered only by the fully instrumented planes. The

resulting track strips are then only observed every ±10 planes within a given view, rather than ±2 planes as expected by the standard track finding algorithm.

2. **The muon spectrometer**, comprising planes 121-281, is fully instrumented every five planes, with the remaining planes uninstrumented. The electronics channels in this region read out four different strips on the same plane, reducing the required electronics, but introducing an ambiguity in the strip positions. The track finding algorithm does not attempt to find spectrometer tracks. However, during the track fitting process, the seed tracks identified in the forward region of the detector can be extrapolated into the spectrometer and the most likely strip positions identified. This is described in more detail in Section 5.3.

The track finding algorithm therefore only needs to consider the forward region of the Near Detector, in which modifications are only required to address the combination of fully and partially instrumented planes. The modified algorithm is summarised in Figure 4.12. The approach is to make ±10 plane Triplets to aid the track finding in the parts of the detector covered only by the fully instrumented planes. These are identical to the normal ±2 plane Triplets, except that the basic plane gap between Clusters is now ten planes, rather than two.

The algorithm proceeds as normal until after the seed segment identification; a number of Near Detector specific operations are then performed. The first of these operations is the identification of Clusters available for the construction of ±10 plane Triplets. This identification is extremely simple; for each Cluster on plane i, the code looks to see if there are Clusters on either plane $i + 2$ or $i - 2$ with similar transverse positions. If a nearby Cluster is found, both are flagged as unavailable for forming ±10 plane triplets. Any Clusters that are part of an ideal ±2 plane seed segment are also flagged as unavailable. An example of this identification is shown in Figure 4.13.

The ±10 plane Triplets are then formed from the available Clusters, and any simple associations between these Triplets are identified. As there are rarely many ±10 plane Triplets, there is no search for preferred associations, but chains of simple associations are identified. If chains of Triplets with only one beginning association and one end association are found, the Triplets are joined together. The ±10 plane TrackSegments are then added back into the main algorithm by searching for associations with any existing ±2 plane TrackSegments. If compatible TrackSegments are found, matched associations are formed. If no associations are found, the ±10 plane TrackSegments are still added to the list of TrackSegments, so that they may act as seed segments.

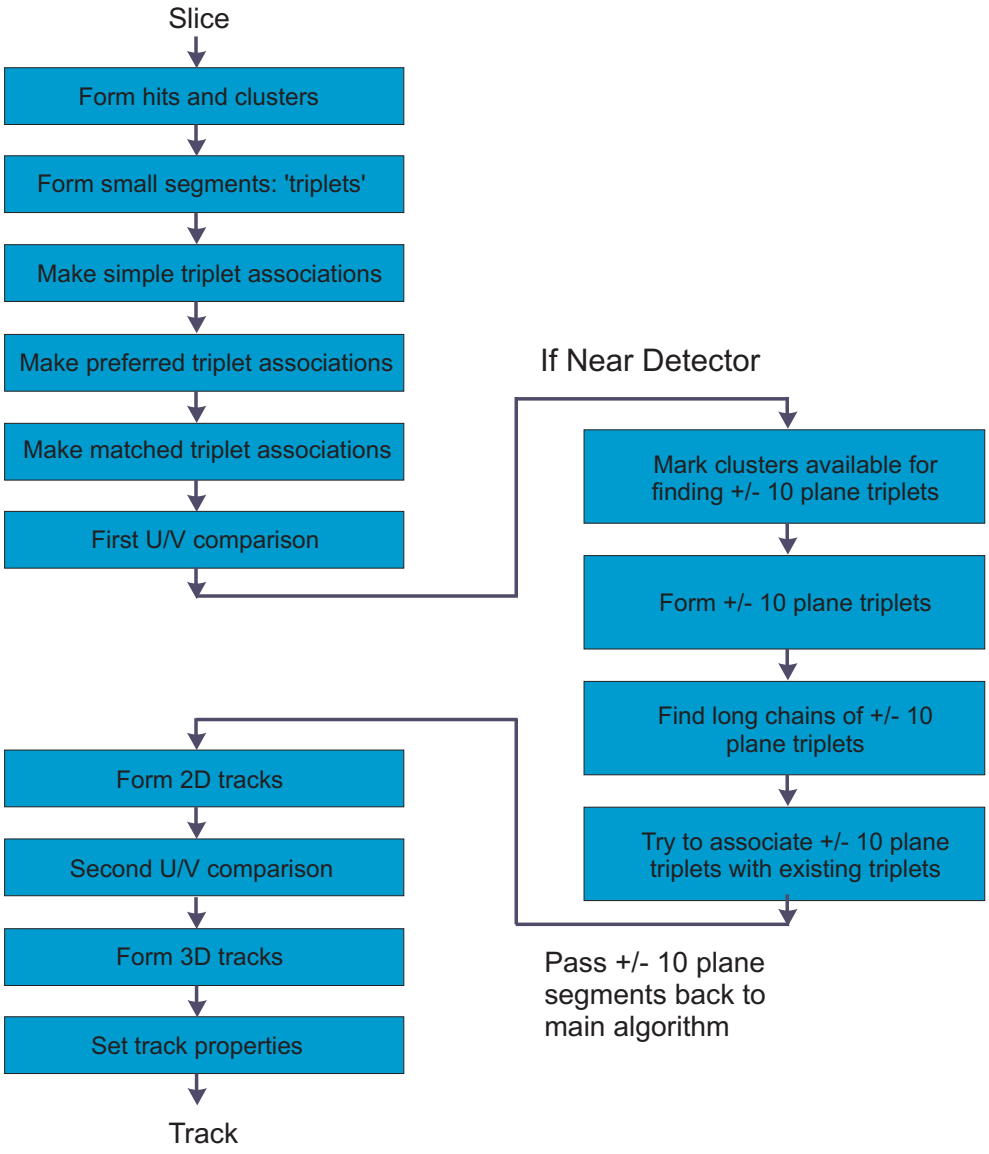

Figure 4.12: The modifications required for use of the track finding algorithm in the Near Detector.

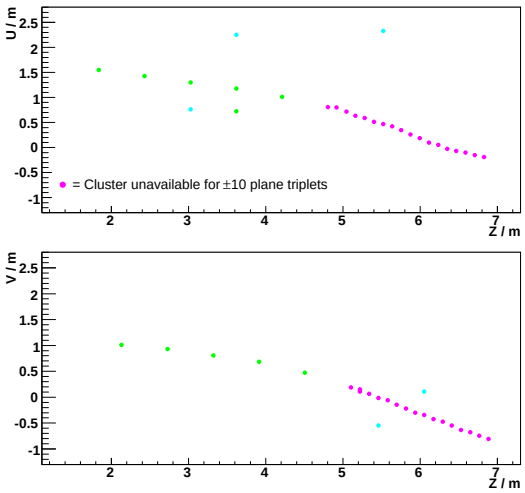

Figure 4.13: The Clusters used in the formation of ±10 plane Triplets are those which are 'isolated', with no nearby Clusters on adjacent planes.

4.3 Results

The aim of the track finding software is to analyse the topology of an event and find a seed track, which identifies the most important features of the muon track. The performance of the software can be evaluated using the MC simulation. In this Section, results are presented showing the performance for samples of Far and Near Detector beam MC.

The first quantity to investigate is the tracking efficiency. This is a measure of how likely the software is to identify true muon tracks. For the purposes of this study, events with a true muon track are defined as true CC ν_μ events which contain at least six true muon strips and which have a true interaction vertex in the detector fiducial volume (described in Section 6.3.2). The tracking efficiency is then calculated as:

$$\text{Tracking Efficiency} = \frac{\#\text{ events with true muon track and reconstructed track}}{\#\text{ events with true muon track}} \quad (4.1)$$

It is also important to investigate the accuracy of the track strip identification. This is best characterised by the purity and the completeness of the reconstructed tracks. The purity describes the fraction of the strips in the reconstructed track that are really true muon strips, as defined in Equation 4.2. The completeness describes the fraction of the true muon hits in an event that are included in the reconstructed track, as defined in Equation 4.3. Purity and completeness values are only calculated for events that contain

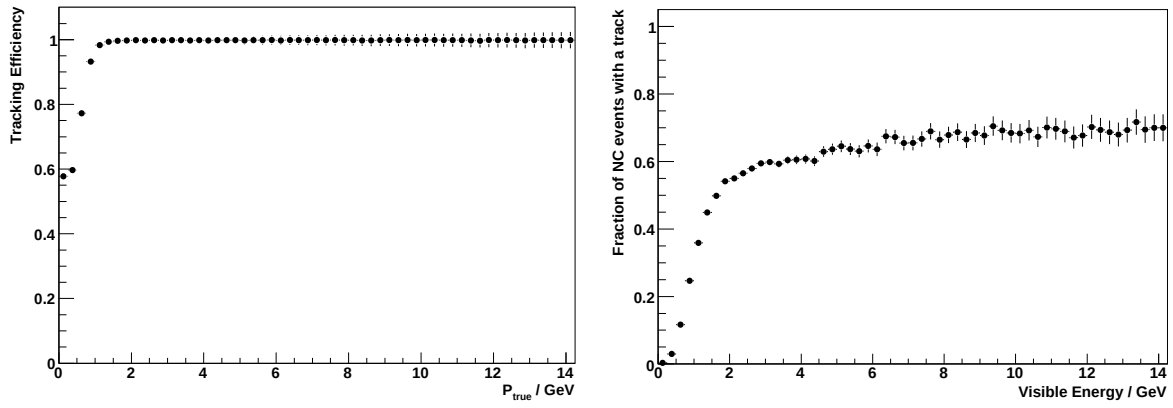

(a) (b)

Figure 4.14: (*a*) The tracking efficiency for true CC ν_μ Far Detector MC events as a function of the true muon momentum. (*b*) The fraction of true NC events that contain a reconstructed track as a function of the true visible energy.

a true muon track:

$$\text{Track Purity} = \frac{\#\text{ true muon strips in track}}{\#\text{ strips in track}} \tag{4.2}$$

$$\text{Track Completeness} = \frac{\#\text{ true muon strips in track}}{\#\text{ true muon strips in event}} \tag{4.3}$$

It should be remembered that the results presented in this Section are simply for the seed tracks. In the track fitting procedure, an understanding of muon propagation is used to improve the strip selection. The track fitting procedure also calculates final values for the kinematic properties of the muon, and the reconstruction of these properties is assessed in Chapter 5.

4.3.1 Far Detector Beam MC

Figure 4.14(a) shows the tracking efficiency for Far Detector beam MC as a function of true muon momentum. The efficiency is extremely high, with a track being reconstructed in essentially all events containing a true muon track with momentum greater than 1 GeV. Below 1 GeV, the efficiency remains high. This is important for an oscillation analysis, as the events of interest are low energy CC ν_μ interactions.

The high tracking efficiency means that many true NC events also contain a recon-

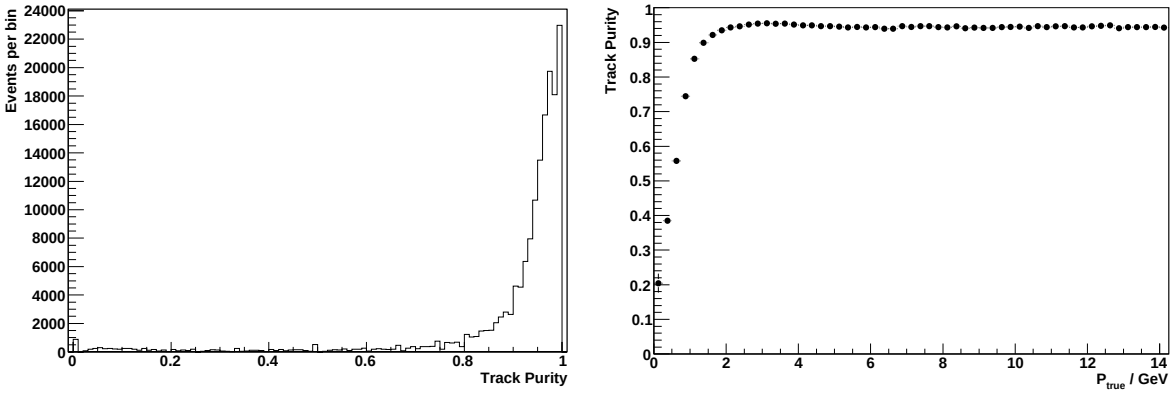

(a) (b)

Figure 4.15: (*a*) The distribution of seed track purity for Far Detector MC events with a true muon track. (*b*) The variation of seed track purity with true muon momentum.

structed track. Figure 4.14(b) shows the fraction of NC events that contain a track, displayed as a function of the true visible energy. NC events are less likely to contain reconstructed tracks than CC events, but more than 60% of NC events with large visible energy will nevertheless contain a track. However, as described in Chapter 6, the fake tracks in NC events can be efficiently distinguished from true muon tracks.

Figure 4.15 shows the distribution of track purity and the variation of the track purity with true muon momentum. High purities are obtained, indicating that the algorithm correctly identifies true muon strips. At low muon momentum, the muon tracks are short and difficult to resolve. This means that the reconstructed track may not accurately represent the muon and will have a low purity. The purity then quickly rises with muon momentum, as the tracks become longer and easier to resolve.

At low neutrino energies, below 3 GeV, events are largely quasi-elastic (QEL) and resonance (RES) interactions, with relatively few deep-inelastic-scattering (DIS) interactions. The purity reaches a maximum when most events are QEL and RES interactions with sufficient energy to produce long tracks. At higher muon momenta, most events are DIS interactions and the purity remains constant; both the tracks and showers increase in size, but the fraction of the track emerging from the shower changes very little.

Figure 4.16 shows the distribution of track completeness and the variation of the completeness with true muon momentum. The completeness of the seed tracks does not peak at one because of the removal of low pulse height strips from the track finding

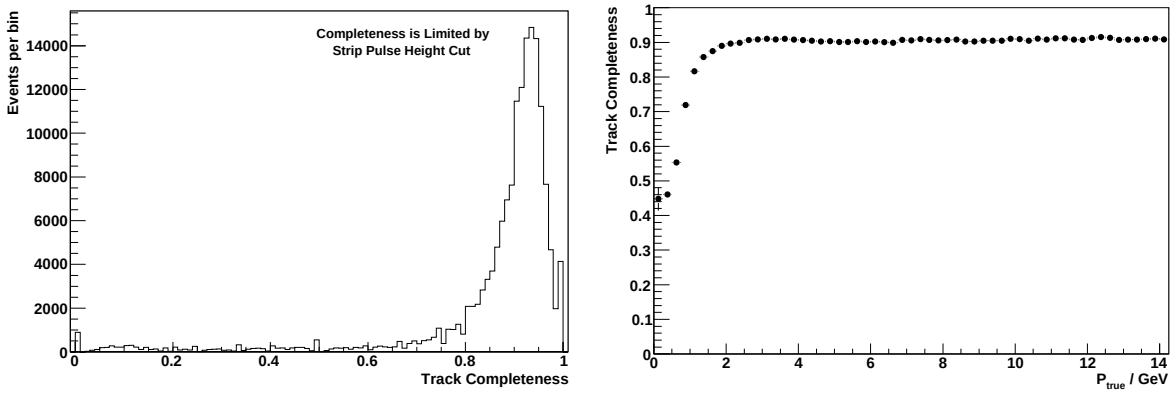

(a) (b)

Figure 4.16: (*a*) The distribution of seed track completeness for Far Detector MC events with a true muon track. (*b*) The variation of seed track completeness with true muon momentum.

algorithm. The seed track is intended to simply identify the basic topology of the muon track and is not expected to find every single muon hit; the pulse height cut is used to reduce the chance of stray noise hits disrupting the pattern recognition.

The track completeness and purity are controlled by the same factors, which determine the topology of the event. The completeness and purity therefore display very similar variations with true muon momentum, with the same important features produced by the same underlying effects.

4.3.2 Near Detector Beam MC

The track finding software does not consider the spectrometer region of the Near Detector. This means that the performance of the software can only be evaluated by considering events with true muon tracks in the forward section of the detector. The tracking efficiency, track purity and track completeness are calculated using only these events, with true muons strips in the spectrometer not contributing to the track completeness. It should also be remembered that the track finding in the Near Detector is strongly affected by the performance of the slicing software. The track finding relies on the slicing algorithm to correctly separate individual events, without splitting the strips on a single track into multiple slices.

The tracking efficiency for Near Detector beam MC is shown as a function of true muon momentum in Figure 4.17(a), whilst the fraction of true NC events that contain a

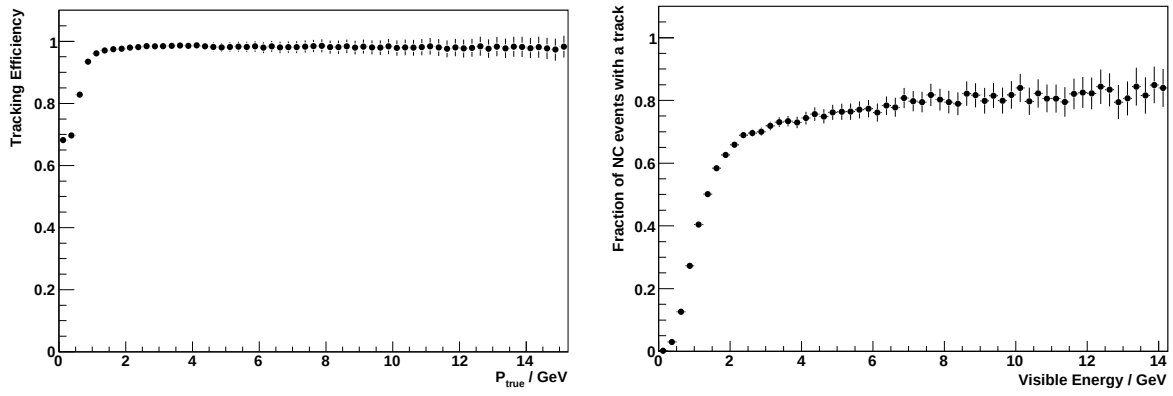

(a) (b)

Figure 4.17: (*a*) The tracking efficiency for true CC ν_μ Near Detector MC events as a function of the true muon momentum. (*b*) The fraction of true NC events that contain a reconstructed track as a function of the true visible energy.

reconstructed track is shown as a function of true visible energy in Figure 4.17(b). The distributions are very similar to those obtained for the Far Detector, demonstrating the same high tracking efficiency. However, the fraction of NC events with reconstructed tracks is a little higher than at the Far Detector.

Figures 4.18(a) and 4.19(a) show the distributions of track purity and track completeness respectively. High purity and completeness values are obtained, although the distributions are a little broader than those for the Far Detector. This is simply a reflection of the increased difficulty of finding tracks at the Near Detector, due to the non-uniform instrumentation and the splitting of single events into multiple slices.

The variations in the track purity and completeness with true muon momentum are shown in Figures 4.18(b) and 4.19(b). The observed behaviour is similar to that at the Far Detector. Initially the purity and completeness rise with true muon momentum, as the muon tracks become longer and easier to resolve. A maximum is then obtained, corresponding to the point at which the greatest proportion of events are QEL and RES interactions with clear muon tracks. However, unlike at the Far Detector, the purity and completeness then decrease a little as the true muon momentum increases. This is because of the increasing fraction of events that are DIS interactions. The DIS interactions fill the small Near Detector fiducial volume with large hadronic showers, whilst the clean sections of the muon track disappear into the spectrometer, making the tracking finding more difficult.

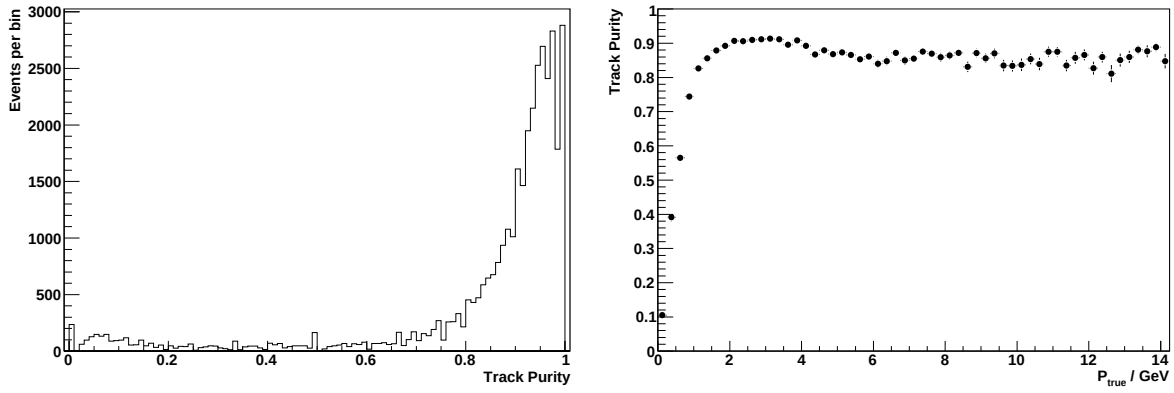

(a) (b)

Figure 4.18: (*a*) The distribution of seed track purity for Near Detector MC events with a true muon track. (*b*) The variation of seed track purity with true muon momentum.

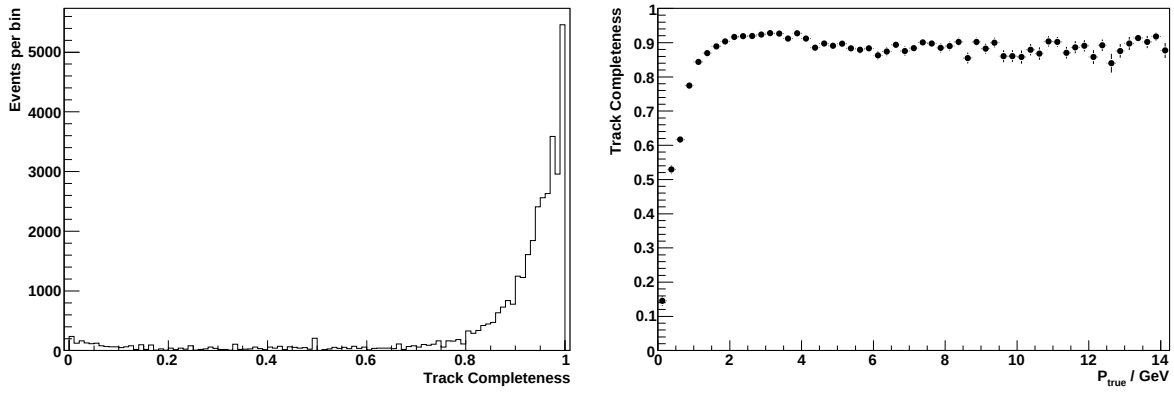

(a) (b)

Figure 4.19: (*a*) The distribution of seed track completeness for Near Detector MC events with a true muon track. (*b*) The variation of seed track completeness with true muon momentum.

Chapter 5

Track Fitting

In the first stage of muon track reconstruction, the track finding algorithm identifies a 'seed track'. This seed track is intended to indicate the most obvious features of the muon trajectory and is identified using only the topology of the reconstructed strips.

The second stage of muon track reconstruction is the track fitting process. Instead of relying purely on pattern recognition, the track fitting algorithm uses information from the seed track in conjunction with knowledge of the propagation and energy loss of muons. This allows accurate track strip identification, even in large vertex showers, and allows the properties of the muon to be fully specified at every track plane. The track fitting process is ideally suited to a Kalman filter implementation.

The Kalman filter [73] is a set of simple recursive equations that enables the state of a dynamic system to be estimated from a series of incomplete and noisy measurements. In the context of track fitting, the filter uses information about muon propagation in order to remove the effects of noise and multiple scattering and converge efficiently to a complete and accurate description of the muon.

This Chapter describes a new software package that has been implemented to fit the seed tracks identified in the MINOS detectors. The package uses a Kalman filter and details are provided as to how track fitting is performed within the framework of the filter. Also included are details of how the Kalman filter can be used to control track strip identification in large vertex showers and in the Near Detector spectrometer. Finally, the performance of the software is evaluated using MC events.

5.1 How the Algorithm Works

5.1.1 The Kalman Filter

The first requirement in a Kalman filter implementation is the definition of a state vector. In the track fitting package, this state vector must fully specify the properties of the muon at a particular point on the track. As suggested in [74], a five element state vector is used at each track plane:

$$\mathbf{x}_k = (u, v, \mathrm{d}u/\mathrm{d}z, \mathrm{d}v/\mathrm{d}z, q/p) \tag{5.1}$$

At detector plane k, the state vector specifies the transverse positions of the muon, u and v, its transverse direction, $\mathrm{d}u/\mathrm{d}z$ and $\mathrm{d}v/\mathrm{d}z$, and its ratio of charge sign to momentum, q/p. The track fitting process aims to obtain accurate estimates of the state vector at every track plane, providing a complete description of the muon trajectory.

In order to carry out the fit, the Kalman filter requires a series of measurements of the muon trajectory. In the first instance, these measurements can simply be the strips in the seed track. To illustrate the mechanics of the fit, consider a situation where the seed track identifies a muon strip on a plane labelled $k-1$.

To begin the fit at plane $k-1$, estimates must be provided for the initial state vector at this plane and for the initial state covariance. The state covariance indicates the expected error in the state vector and is represented by a 5×5 matrix, $\mathbf{C}_{k-1}$. The initial estimate for the state vector is provided using the properties of the seed track, whilst the covariance matrix is initialised with non-zero values only along its diagonal. These entries along the diagonal each indicate the typical uncertainty in the associated state vector element.

The fit proceeds by predicting the state vector at the next plane with a track strip, plane k, then using the actual measured track strip position in order to update the prediction and produce a 'filtered' state vector. The prediction of the state vector at plane k is produced by using knowledge of muon propagation to produce a 5×5 propagator matrix, $\mathbf{F}_{k-1}$. The propagator matrix is described in more detail in Section 5.1.2.

The covariance matrix must also be updated to account for the increased uncertainty in the predicted state vector due to multiple scattering and energy loss of the muon in its motion between the planes. This increased uncertainty is represented by a 5×5 'noise covariance' matrix, $\mathbf{Q}_{k-1}$, which is described in Section 5.1.3.

The final input to the filter is an estimate of the measurement precision of the track strip position at plane k, V_k. The filtered state vector at plane k is then calculated by a single use of the Kalman update equations shown in Equations 5.2–5.5. In the update equations, $\mathbf{H}_k$ is the measurement function, equal to (1,0,0,0,0) for u planes and (0,1,0,0,0) for v planes, whilst $\mathbf{1}$ is the 5×5 identity matrix and m_k is the measured track strip position.

$$\mathbf{C}_k^{k-1} = \mathbf{F}_{k-1}\mathbf{C}_{k-1}\mathbf{F}_{k-1}^T + \mathbf{Q}_{k-1} \tag{5.2}$$

$$\mathbf{K}_k = \mathbf{C}_k^{k-1}\mathbf{H}_k^T(V_k + \mathbf{H}_k\mathbf{C}_k^{k-1}\mathbf{H}_k^T)^{-1} \tag{5.3}$$

$$\mathbf{x}_k = \mathbf{F}_{k-1}\mathbf{x}_{k-1} + \mathbf{K}_k(m_k - \mathbf{H}_k\mathbf{F}_{k-1}\mathbf{x}_{k-1}) \tag{5.4}$$

$$\mathbf{C}_k = (\mathbf{1} - \mathbf{K}_k\mathbf{H}_k)\mathbf{C}_k^{k-1} \tag{5.5}$$

In the update equations, the difference between the predicted and filtered state vectors at plane k is determined by two important factors. The first of these is the difference between the measured strip position and the predicted position. The second is a five element vector known as the Kalman gain, $\mathbf{K}_k$. The Kalman gain uses the state covariance and the uncertainty in the track strip position to determine how much a single measurement can influence the filtered state vector. If the agreement between the predictions and measurements is good, the implication is that the state vector closely matches the true muon state vector. The reduced uncertainty is reflected in the covariance matrix, which is updated in the final step of the update equations.

After application of the Kalman update equations, plane k can be relabelled as plane $k-1$ and a new plane k can be identified. The process described above can then be repeated for the new pair of planes, taking advantage of the newly calculated filtered state vector and covariance matrix as input to the Kalman update equations. The fit can begin at any plane and can propagate either backwards or forwards along the track, accounting for the direction of motion in the modelling of the propagator and noise matrices. Moving between adjacent track planes in this way, and moving back and forth along the entire track multiple times, allows convergence towards the true state vectors at each track plane. The final state vectors at each track plane specify the reconstructed properties of the muon, whilst the final covariance matrix at each plane indicates the estimated uncertainty in these properties.

Of particular interest are the contents of the state vector at the track vertex, which provide the best estimates of the track vertex position and direction. This state vector also specifies the reconstructed values of the muon charge sign and momentum, as determined from the curvature of the track in the detector magnetic field.

5.1.2 The Propagator Matrix

An important step in the track fitting process is the calculation of the propagator matrix. The aim of this matrix is to transform the state vector at plane $k-1$ to produce a prediction of the state vector at plane k, as shown in Equation 5.6. The propagator matrix is the means by which information about muon propagation and the detector magnetic field is incorporated in the fit.

$$\mathbf{x}_{k,predicted} = \mathbf{F}_{k-1}\mathbf{x}_{k-1} \tag{5.6}$$

The propagator matrix must be explicitly calculated for every pair of track planes considered in the fit. The matrix depends on the difference in z coordinate between the two planes, δz, and on the components of the magnetic field, B_u and B_v. An approximation to the forward propagator matrix can be written in functional form, as shown in Equation 5.7.

$$\mathbf{F}_{k-1} = \begin{pmatrix} 1 & 0 & \delta z & 0 & \frac{1}{2}B_v\delta z^2 \\ 0 & 1 & 0 & \delta z & -\frac{1}{2}B_u\delta z^2 \\ 0 & 0 & 1 & 0 & B_v\delta z \\ 0 & 0 & 0 & 1 & -B_u\delta z \\ 0 & 0 & 0 & 0 & 1+\epsilon \end{pmatrix} \tag{5.7}$$

Each row of the propagator matrix is responsible for the transformation of one element of the state vector. In Equation 5.7, the u and v coordinates and gradients are transformed using a straight line approximation for the path of the muon, together with the leading term to account for the effects of the magnetic field. The transformation of the q/p element of the state vector is hugely simplified; a small perturbation (ϵ is a small positive constant) is applied to the q/p value in order to represent energy loss when moving between the planes.

The simple approximation to the propagator matrix is quick to calculate during the track fit. For the positions and gradients in the state vector, the simple propagator was found to perform well. However, a more sophisticated propagator matrix is required in order to correctly model the variations in q/p between the planes and to provide a more complete description of bending in the magnetic field. For this reason, in the new track fitting software, the first four columns in the matrix are calculated using the simple approximation, whilst a more accurate (and time consuming) approach is used

to calculate the final column of the matrix.

The propagator matrix is completed by using a muon 'Swimmer' package [75], which accurately models muon motion and energy loss in the specific geometry of the MINOS detectors. Once provided with an initial state vector and z coordinate, the Swimmer can numerically calculate the new state vector at any requested z coordinate. As shown in Equation 5.8, the final propagator matrix column, $\mathbf{A}$, can then be determined by linearising the propagator function about the current state vector.

$$\mathbf{A} = \frac{\mathbf{S}(q/p + \delta(q/p)) - \mathbf{S}(q/p - \delta(q/p))}{2\delta(q/p)} \tag{5.8}$$

In Equation 5.8, $\mathbf{S}(q/p+\delta(q/p))$ is the Swimmer prediction for the new state vector, after the q/p value in the input state vector has been incremented by a small amount $\delta(q/p)$. To linearise the propagator function, new state vectors are evaluated for a small increment to the q/p value and for a small decrement to the q/p value. The difference between the two new state vectors indicates the importance of the q/p change for each state vector element. Dividing the difference by twice the incremental value produces the final propagator matrix column.

5.1.3 The Noise Covariance Matrix

At each step in the track fit, a noise covariance matrix must be calculated. This matrix accounts for the increase in state vector uncertainties due to multiple scattering and energy loss of the muon in its motion between the two track planes under consideration.

The noise covariance matrix can be divided into two distinct parts, as shown in Equation 5.9. The first part is a 4×4 matrix, $\mathbf{Q}_k^{MS}$, which addresses uncertainties due to multiple scattering. The second part is a single term, $Q_k^{\delta E}$, which addresses the uncertainties due to energy loss through ionisation. The remaining terms in the matrix prove to be very small and are neglected.

$$\mathbf{Q}_k = \left(\begin{array}{c|c} \mathbf{Q}_k^{MS} & 0 \\ \hline 0 & Q_k^{\delta E} \end{array}\right) \tag{5.9}$$

The multiple scattering section of the noise covariance matrix is derived in [76]. In this derivation, the scattering is modelled as though it occurs at a discrete point, taken

to be the centre of the next steel plane the muon will encounter. The final result is shown in Equation 5.10.

$$\mathbf{Q}_k^{MS} = \begin{pmatrix} z_0^2\sigma_{33}^2 & z_0^2\sigma_{34}^2 & -z_0\sigma_{33}^2 & -z_0\sigma_{34}^2 \\ z_0^2\sigma_{34}^2 & z_0^2\sigma_{44}^2 & -z_0\sigma_{34}^2 & -z_0\sigma_{44}^2 \\ -z_0\sigma_{33}^2 & -z_0\sigma_{34}^2 & \sigma_{33}^2 & \sigma_{34}^2 \\ -z_0\sigma_{34}^2 & -z_0\sigma_{44}^2 & \sigma_{34}^2 & \sigma_{44}^2 \end{pmatrix} \tag{5.10}$$

The terms in Equation 5.10 are simple to calculate, requiring only knowledge of the current Kalman state vector and the detector geometry. In Equation 5.10, z_0 is the difference in z coordinate between the scattering point and the next track plane, whilst the remaining terms are:

$$\sigma_{33}^2 = \frac{1}{2}\sigma_{MS}^2 \left(\frac{\mathrm{d}s}{\mathrm{d}z}\right)^2 \left[1 + \left(\frac{\mathrm{d}u}{\mathrm{d}z}\right)^2\right] \tag{5.11}$$

$$\sigma_{34}^2 = \frac{1}{2}\sigma_{MS}^2 \left(\frac{\mathrm{d}s}{\mathrm{d}z}\right)^2 \left[\left(\frac{\mathrm{d}u}{\mathrm{d}z}\right) \cdot \left(\frac{\mathrm{d}v}{\mathrm{d}z}\right)\right] \tag{5.12}$$

$$\sigma_{44}^2 = \frac{1}{2}\sigma_{MS}^2 \left(\frac{\mathrm{d}s}{\mathrm{d}z}\right)^2 \left[1 + \left(\frac{\mathrm{d}v}{\mathrm{d}z}\right)^2\right] \tag{5.13}$$

In Equations 5.11–5.13, the factors of $\mathrm{d}s/\mathrm{d}z$ account for the path length of non-perpendicular muons through the detector planes, whilst σ_{MS}^2 is the variance of the multiple scattering angle. This variance depends on the thickness of the scattering medium, x, measured in radiation lengths, X_0, as shown in Equation 5.15. For a perpendicular muon, the thickness is calculated as 1.47 radiation lengths per steel plane.

$$\frac{\mathrm{d}s}{\mathrm{d}z} = \sqrt{1 + \left(\frac{\mathrm{d}u}{\mathrm{d}z}\right)^2 + \left(\frac{\mathrm{d}v}{\mathrm{d}z}\right)^2} \tag{5.14}$$

$$\sigma_{MS} = \frac{13.6\,\mathrm{MeV}}{\beta cp}\sqrt{\frac{x}{X_0}}\left[1 + 0.038\ln\left(\frac{x}{X_0}\right)\right] \tag{5.15}$$

The final term in the noise covariance matrix accounts for energy loss due to ionisation. This term is described in [74], which suggests that the width of the energy loss distribution can be characterised by 25% of the total mean energy loss. The q/p covariance is then as shown in Equation 5.16, in which ΔE is the mean energy loss for

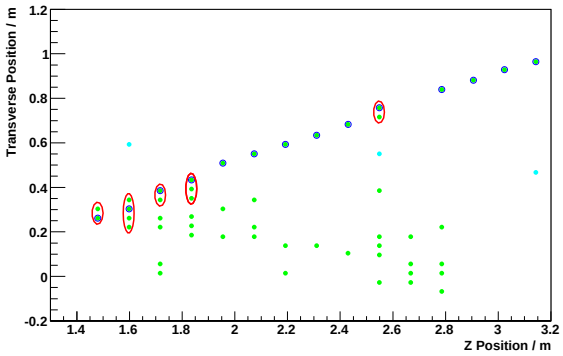

Figure 5.1: Information must be extracted from the seed track strips in order for use in the framework of the Kalman filter. The measured track strip positions and errors are determined by examining the clusters of strips around the seed track strips.

a perpendicular muon; approximately 40 MeV per steel plane in the MINOS detectors.

$$Q_k^{\delta E} = \left(0.25\frac{\Delta E}{p^2}\cdot\frac{ds}{dz}\right)^2 \tag{5.16}$$

5.1.4 Interface to the Kalman Filter

Filter Input

The initial input to the track fitting package is the list of strips in the seed track. Information must be extracted from these strips in order for use in the framework of the Kalman filter. The first step in the extraction of this information is to store the seed track strips in plane order. Also stored are details of all the strips in the same slice as the seed track.

At each plane with a seed track strip, the Kalman filter needs to know the view (u or v) and the z coordinate of the plane. It also needs to know the measured transverse position of the track strip and the associated precision of this measurement. This information is provided by investigating the clustering of the slice strips around the seed track strip. If a seed track strip is part of a cluster of adjacent strips, the measured track strip position that is recorded for use in the Kalman filter is actually the charge weighted transverse position of the cluster. The charge weighting proves useful at removing the impact of strips produced by noise. The recorded error in the track strip measurement is simply the transverse width of the cluster, divided by $\sqrt{12}$. This input information, reduced to a few numbers per seed track plane, is stored in plane order for use in the fit. An illustration of how the information is extracted is shown in Figure 5.1.

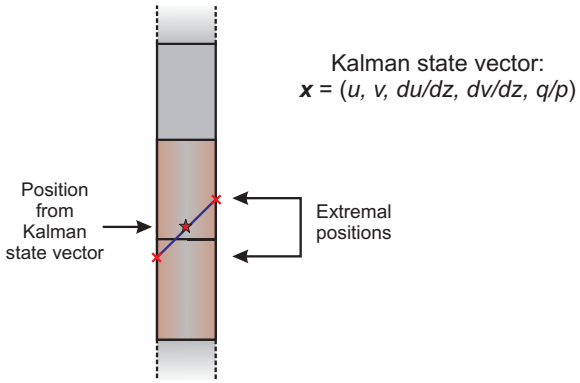

Figure 5.2: The strips that most closely match the Kalman state vectors are determined by using the positions and gradients from the state vector to estimate the muon trajectory through the scintillator.

The fit begins at the vertex of the seed track and then moves between the planes containing seed track strips, making use of the Kalman update equations. The fit continues with this motion until it reaches the end of the track. At this point, the direction is reversed and the fit moves back towards the vertex. Upon reaching the vertex, one iteration is complete.

Filter Output

The input to the Kalman filter is obtained by examining the seed track strips in order to produce the set of numbers required by the framework of the filter. At the end of each filter iteration, after state vectors have been calculated for each track plane, this process can be reversed. The Kalman state vectors at each plane are compared to the strips in the slice and the strips that are most consistent with the state vectors are identified. The new list of strips can be used as seed track strips for the next filter iteration. In this way, the track strip identification improves with each iteration of the Kalman filter in a very natural and elegant manner.

The comparison between the slice strips and state vectors is performed by using the positions and gradients from the state vector. As shown in Figure 5.2, these positions and gradients can be used to estimate the muon trajectory through the 1 cm of scintillator material. The strips intersected by the passage of the muon are those identified as seed track strips for the next filter iteration. If no strips are intersected by the line projected through the scintillator, the nearest slice strip within a certain tolerance can be selected. This technique for refining the track strip identification addresses most track finding problems. It is now only in very large vertex showers, where track finding is particularly difficult, that a different approach is required.

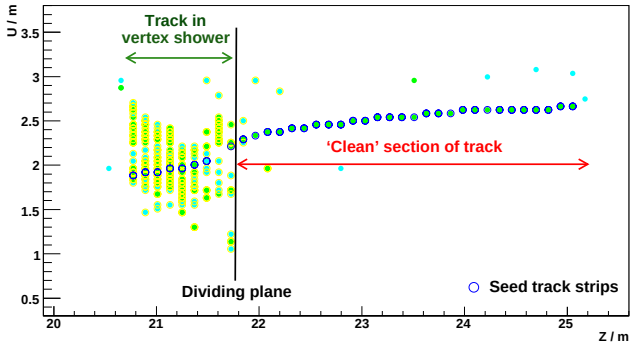

Figure 5.3: The division of a typical CC beam ν_μ interaction into a 'clean' section of track and a vertex shower. The strips in the seed track are likely to be accurate in the clean region, but will not be so accurate within the shower.

5.1.5 Track Strip Identification in Vertex Showers

Large vertex showers in an event can make the track finding process difficult and can lead to incorrect track strip identification within the shower. The track fitting software is able to correct for this and can provide accurate identification of the muon track strips in even very large vertex showers.

The approach to improving the track reconstruction in vertex showers is to identify the 'clean' section of the seed track, as illustrated in Figure 5.3. The Kalman filter can use this clean section of track in order to build up information about the muon and accurately calculate the state vectors just outside the shower. These state vectors can then be carefully projected back into the shower region.

The Kalman fit begins at the vertex of the seed track. It then moves between the seed track planes until it reaches the end of the track. At this point, the algorithm attempts to identify whether or not there is a significant vertex shower in the event. If a large vertex shower is found, the algorithm tries to locate the plane that divides the clean section of the track from the shower.

The planes in a vertex shower are identified by counting the number of slice strips on each plane and comparing the results with the expectation for a clean muon track with no shower. Only slice strips above a certain pulse height are counted, and the clean track expectation accounts for the gradient of the seed track. If the region around the track vertex contains significantly more slice strips per plane than expected for a clean muon track, the event is declared to have a vertex shower. Further examination of the number of strips per plane allows the plane at the end of the vertex shower to be identified.

The number of seed track planes in the clean section of the track is then counted.

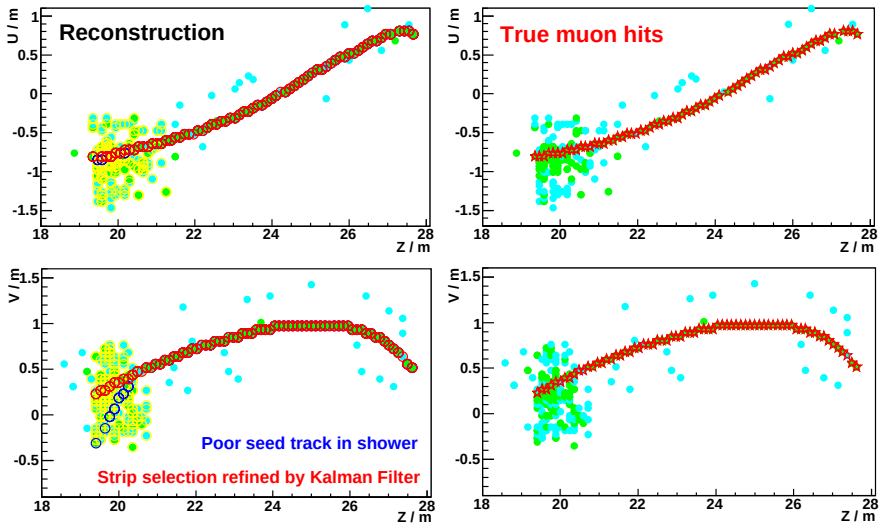

Figure 5.4: Example event display in which the track fitting process improves the track strip identification within a large vertex shower.

If the Kalman filter is to build up sufficient information about the muon, at least six clean seed track planes are required. If there are enough clean planes, the seed track information inside the vertex shower is simply discarded. The Kalman fit then proceeds back towards the vertex, stopping when it reaches the edge of the shower. At this point, the Swimmer is used to predict the state vector at the next plane, inside the shower.

The predicted state vector is compared to the strips in the slice, as described in Section 5.1.4. Any matching strips are then treated exactly as if they were provided by a track finding package; the information required by the Kalman filter is extracted and the Kalman update equations are applied. The swim back into the vertex shower continues until a plane window is reached in which the Swimmer finds no strips. Figure 5.4 shows an example event display in which the track fitter corrects the track strip identification in a large vertex shower.

During the Kalman fit, the state vectors converge very rapidly. By default, two full iterations along the track are performed. Only during the first iteration are strips in the vertex shower treated differently to strips in the clean section of the track. A complete summary of the track fitting process is illustrated in Figure 5.5.

5.1.6 Calculation of Track Properties

After the final Kalman filter iteration, filtered state vectors are available at every track plane. These state vectors provide a complete description of the muon at different points

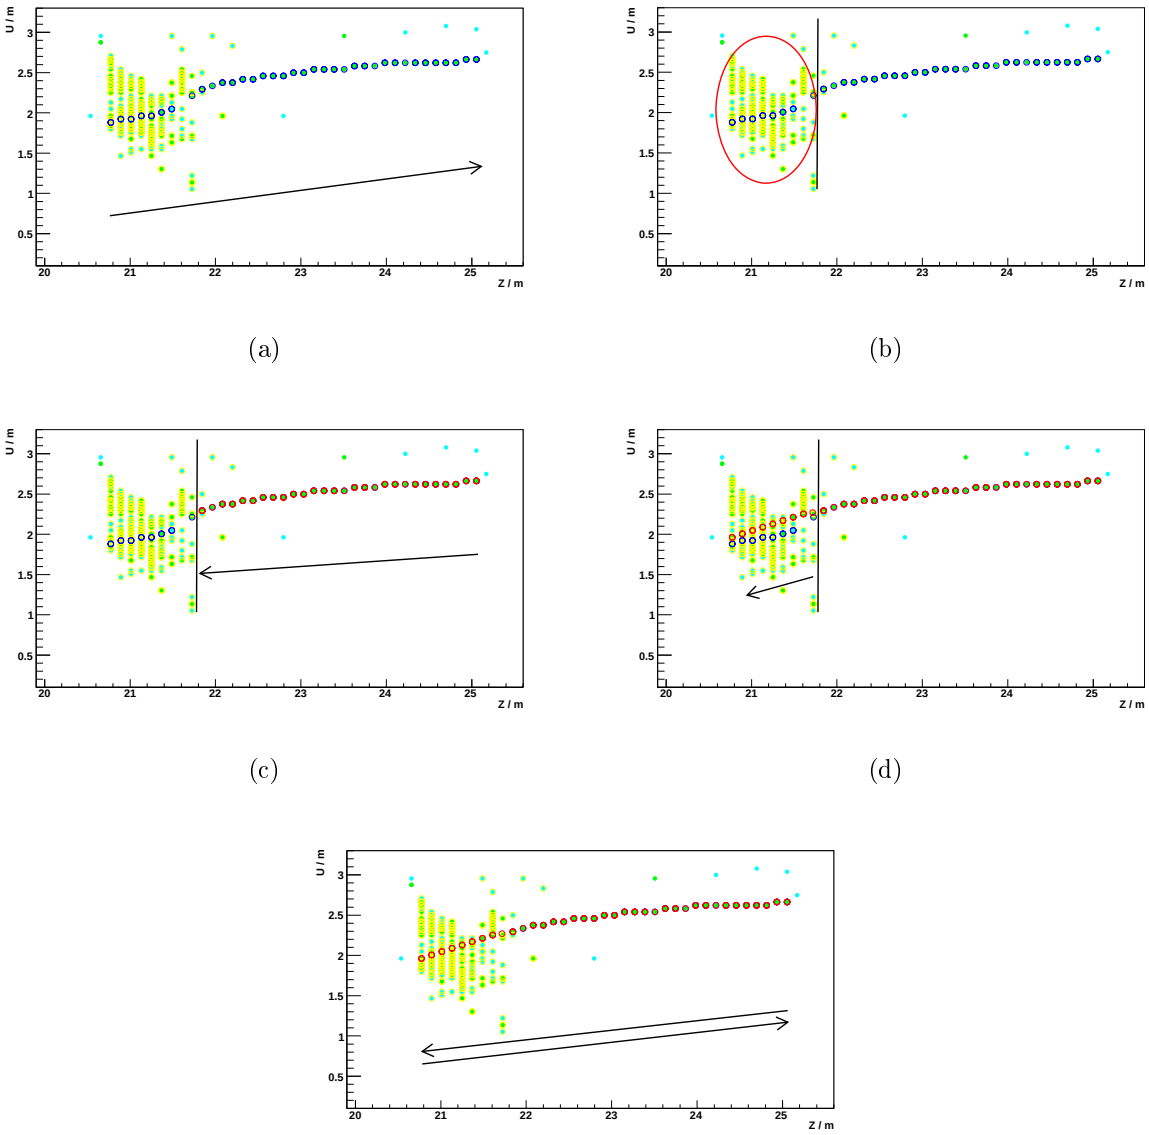

Figure 5.5: Summary of the track fitting process for a CC beam neutrino interaction with a large vertex shower. Blue markers represent strips in the seed track, whilst red markers represent strips in the fitted track. (*a*) Fit moves from seed track vertex to end. (*b*) Identification of vertex shower. (*c*) Fit moves back towards shower. (*d*) Swimmer identifies track strips in shower. (*e*) Remaining iterations completed.

along its trajectory. By extracting information from the state vectors, the properties of the muon can be specified for use in a physics analysis.

The first step is to identify the final track strips. This identification proceeds as detailed in Section 5.1.4, by using the final filtered state vectors to specify the muon trajectory through each scintillator plane. Any slice strips intersected by this trajectory are included in the final track.

The state vectors at each track plane are then stored. Also stored are the errors associated with each state vector, which are obtained from the covariance matrix at the relevant plane. This information provides a complete record of the position, direction and momentum of the muon in its motion through the detector.

Of particular importance is the state vector at the track vertex. This accurately specifies the track vertex position and direction, both of which are useful in a physics analysis. The vertex state vector also includes the q/p value at the track vertex, which specifies the reconstructed charge sign and momentum of the muon, as determined from the curvature of the track in the magnetic field. The uncertainty associated with the vertex q/p value, $\sigma_{q/p}$, is extracted from the final covariance matrix at the track vertex.

The muon momentum is also estimated from the range of the track in the detector. This estimate is obtained as described in [77]. Starting at the end of the track, the Swimmer is used to model the motion of the muon between the final state vectors, back to the track vertex. At each step in this smooth motion, the changes in muon energy are calculated and contribute towards the final muon momentum value. The smooth motion between filtered state vectors, together with the detailed detector model in the Swimmer, allows accurate momentum reconstruction.

5.2 Far Detector Results

The aim of the track fitting software is to provide precise muon track strip identification, even in large vertex showers, and to accurately reconstruct the properties of the muon. In this Section, the performance of the software is evaluated at the Far Detector by using the MC simulation. Results are presented for Far Detector beam MC events and a short comparison is made between the results obtained with the new track finding and fitting software and those obtained using the previous MINOS reconstruction. The application of the software to cosmic muon reconstruction is also briefly discussed.

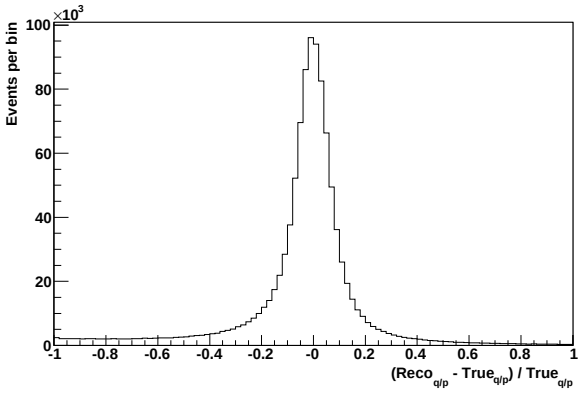

(a)

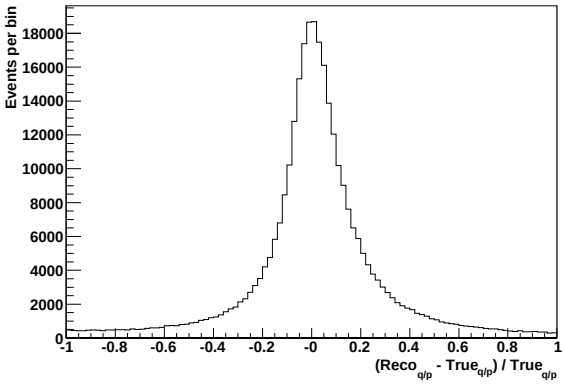

(b)

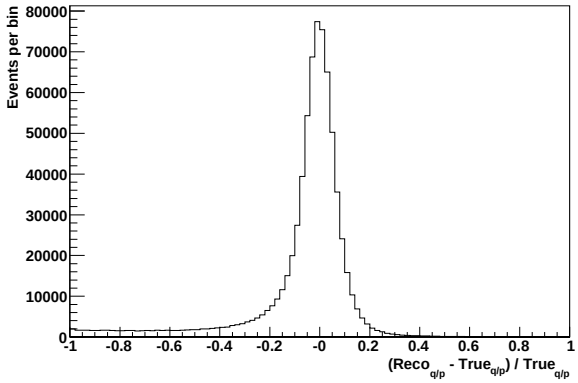

(c)

Figure 5.6: The fractional discrepancy observed between the reconstructed and true q/p values for Far Detector true CC events with a fiducial track vertex: (a) for all events, (b) for events with uncontained tracks and, (c) for events with fully contained tracks.

5.2.1 Far Detector Beam MC

Figure 5.6 shows the accuracy of the q/p reconstruction for Far Detector beam MC events. By considering only true CC events with a reconstructed track vertex in the detector fiducial volume, the fractional discrepancy between the reconstructed and true q/p values is investigated. Figure 5.6(a) shows this discrepancy for all the events, whilst Figure 5.6(b) considers only events with tracks that leave the detector and Figure 5.6(c) considers only events with fully contained tracks.

As expected, the q/p reconstruction is better for contained tracks than for uncontained tracks. This is because the contained tracks typically provide more information about their curvature in the magnetic field. However, the performance for both contained

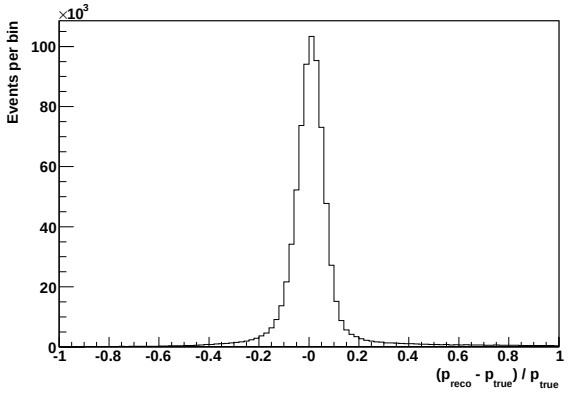

Figure 5.7: The fractional discrepancy between the reconstructed momentum from range and the true momentum, as observed for Far Detector true CC events with a fully contained track and fiducial track vertex.

and uncontained tracks is excellent, given the coarse instrumentation of the MINOS detectors. In a typical analysis, the momentum from curvature is only used for uncontained tracks, but the charge sign reconstruction is used for all tracks.

The fractional discrepancy between the reconstructed momentum from range and the true muon momentum is shown, for events with fully contained tracks, in Figure 5.7. Figure 5.8 shows the fractional offset in muon momentum from range and the percentage resolution, both plotted as a function of the true muon momentum. Only small fractional offsets are observed, but the mean offset is consistently greater than zero. This implies that reconstructed tracks are often slightly too long, probably due to tracking back too far into vertex showers. As the true momentum increases, the resolution quickly approaches 5%, as expected for range based estimates of muon momentum.

Figure 5.9 shows the fractional offset in muon momentum from curvature and the percentage resolution, both plotted as a function of true muon momentum. These distributions are satisfactory, demonstrating suitably small offsets and a resolution approaching 13% at muon momentum above 10 GeV.

The accuracy of the final track strip identification is evaluated using the purity and completeness variables that were defined in Section 4.3. Figure 5.10(a) shows the distribution of purity for the fitted tracks, whilst Figure 5.10(b) shows the variation of purity with true momentum. Despite the inclusion of low pulse height strips in the fitted tracks, the purity remains very high; essentially identical to that of the seed tracks.

Figure 5.11(a) shows the distribution of completeness for the fitted tracks. The variation of fitted track completeness with true momentum is shown in Figure 5.11(b), which also shows the variation observed for the seed tracks. The track fitting process clearly increases the completeness of the tracks. This is due to the careful selection of

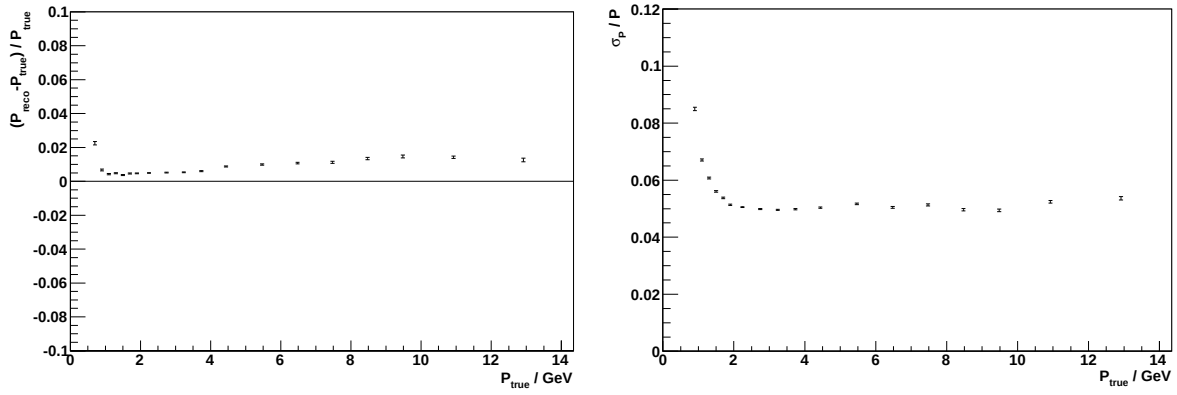

(a) (b)

Figure 5.8: (*a*) The fractional offset in muon momentum from range as a function of the true muon momentum. (*b*) The percentage resolution as a function of the true muon momentum. Both (*a*) and (*b*) consider only Far Detector true CC events with a fully contained track and fiducial track vertex.

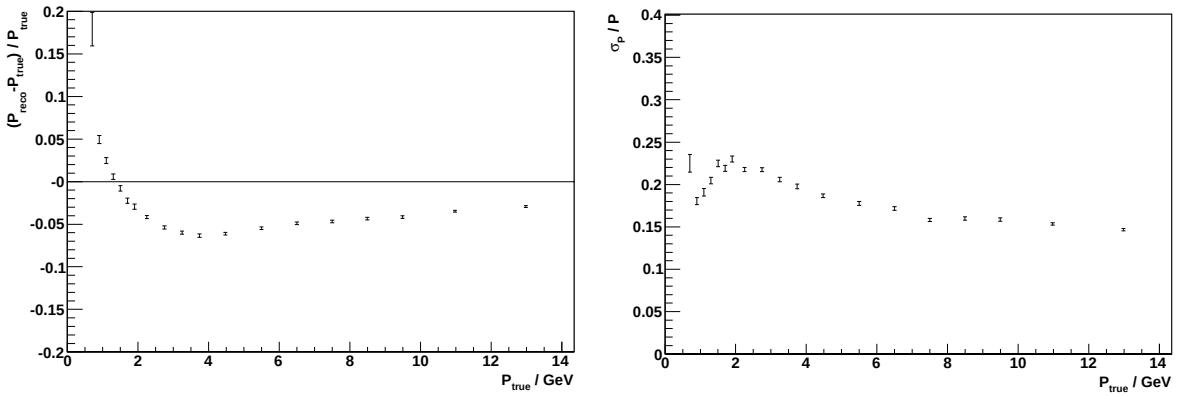

(a) (b)

Figure 5.9: (*a*) The fractional offset in muon momentum from curvature as a function of the true muon momentum. (*b*) The percentage resolution as a function of the true muon momentum. Both (*a*) and (*b*) consider only Far Detector true CC events with an uncontained track and fiducial track vertex.

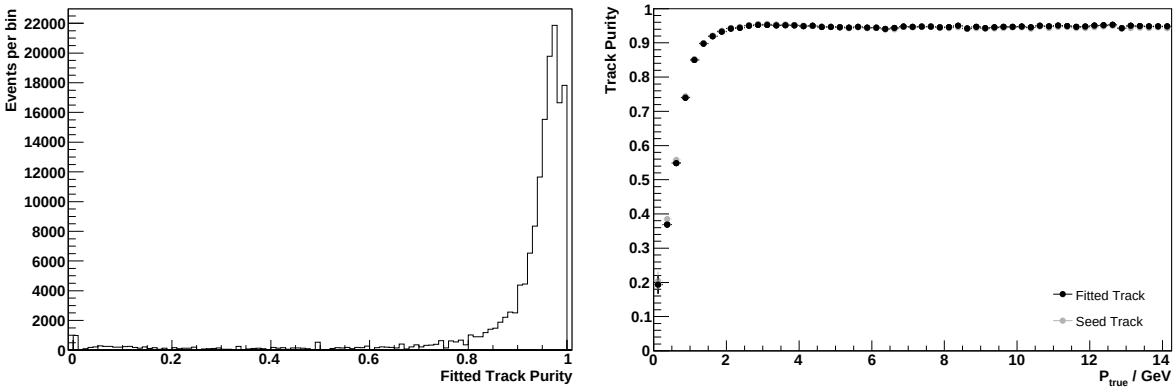

(a) (b)

Figure 5.10: (*a*) The distribution of fitted track purity for Far Detector events with a true muon track. (*b*) The variations of fitted track and seed track purity with true muon momentum.

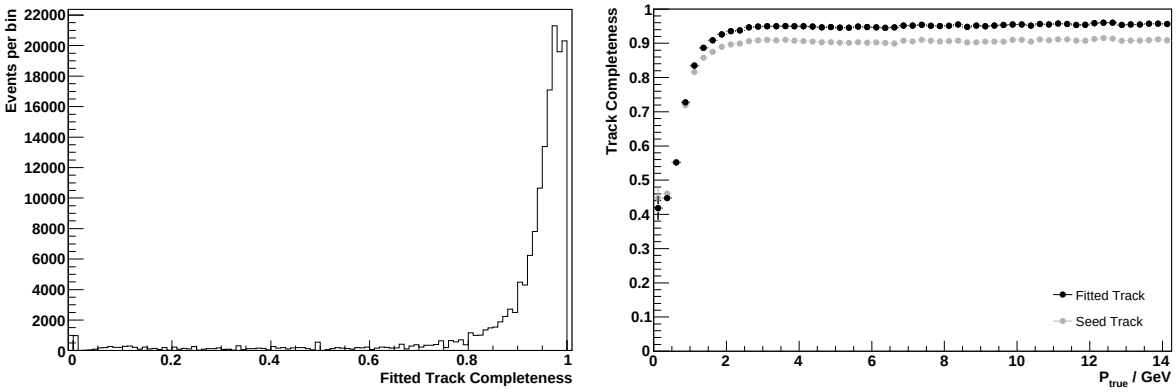

(a) (b)

Figure 5.11: (*a*) The distribution of fitted track completeness for Far Detector events with a true muon track. (*b*) The variations of fitted track and seed track completeness with true muon momentum.

low pulse height strips that were neglected in the track finding algorithm, together with the improved tracking in vertex showers.

A final point worthy of note in this Section is that some tracks can fail the Kalman fit. This typically happens because the pattern of track strips leads to an unphysical covariance matrix, with negative elements along the diagonal. Efforts are made to reduce the number of failures in the new software, and the failure fraction for this Far Detector beam MC sample was only 0.15%.

5.2.2 Comparison with Previous Reconstruction

The motivation for writing the new track finding and fitting software was to provide a more accurate and more efficient reconstruction than the previous software. The success of the new software can therefore be evaluated by processing a sample of Far Detector beam MC events with both versions of the reconstruction and comparing the results.

The first results examined are those indicating the performance of the q/p and momentum from range reconstruction. Figure 5.12 shows the fractional discrepancies observed between reconstructed and true values for a sample of true CC events with track vertices in the detector fiducial volume. Only events passing the Kalman fit contribute to these distributions and this immediately highlights a difference between the two reconstructions; the new software suffers from fewer fitting failures.

Figure 5.12(a) shows that the new software greatly improves the q/p reconstruction. In addition to the reduction in fitting failures, the number of poorly reconstructed tracks in the tails of the distribution is decreased, whilst the number of accurately reconstructed tracks in the peak is increased. Improvements are also observed for the momentum from range reconstruction; Figure 5.12(b) shows that the fractional discrepancy distribution contains more entries in the peak and is more closely centred on zero.

Figure 5.13 examines the track strip selection, showing distributions of fitted track purity and completeness. The distributions show that the new software maintains the high purity of the old reconstruction whilst introducing a large improvement in the completeness. The new software clearly identifies a greater fraction of the true muon strips, performing particularly well in situations where there are multiple true muon strips on a plane.

The reduction in track fitting failures is quantified in Figure 5.14(a). This Figure shows that the new software reduces the failure fraction from 5.1% to just 0.15%, allowing many perfectly good muon tracks to be reclaimed for use in a physics analysis. Finally, Figure 5.14(b) displays the CPU time required for the combined process of finding and

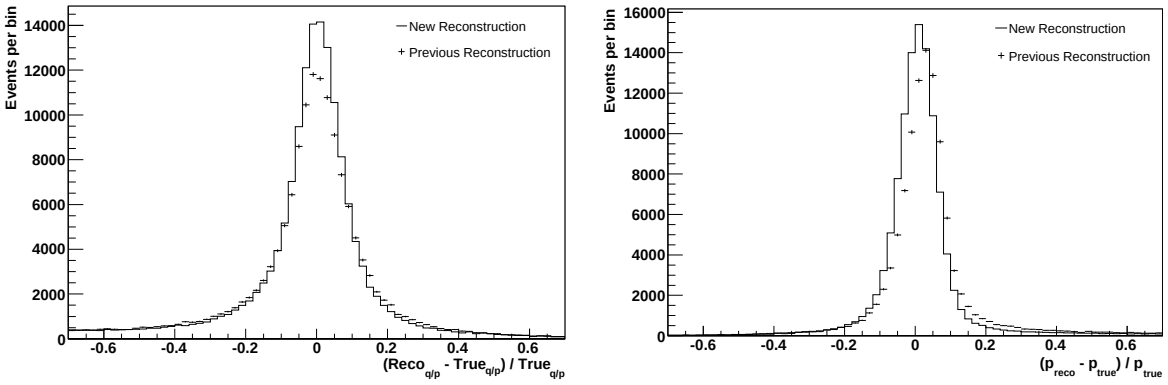

(a) (b)

Figure 5.12: Comparison of the performance of the new and old reconstruction software. (a) The fractional discrepancy between reconstructed and true q/p values. (b) The fractional discrepancy between the reconstructed momentum from range and the true muon momentum, for events with fully contained tracks.

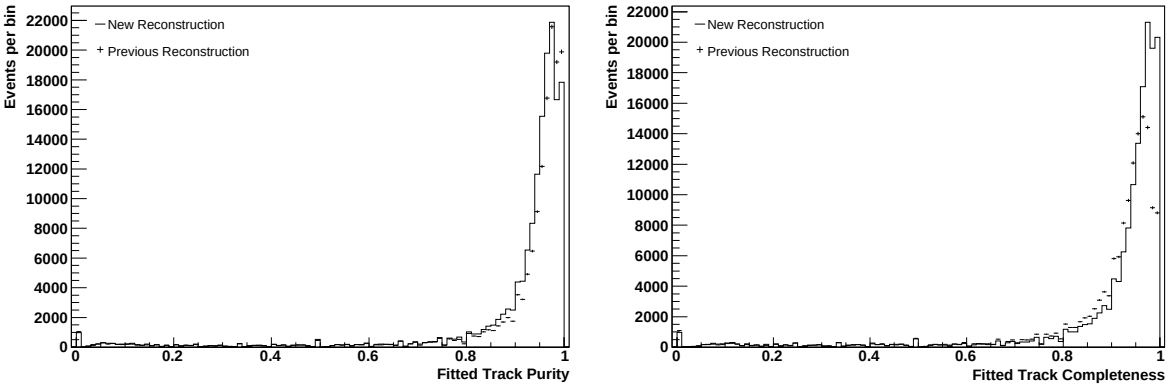

(a) (b)

Figure 5.13: Comparison of the strip selection with the new and old reconstruction software. (a) Distribution of fitted track purity. (b) Distribution of fitted track completeness.

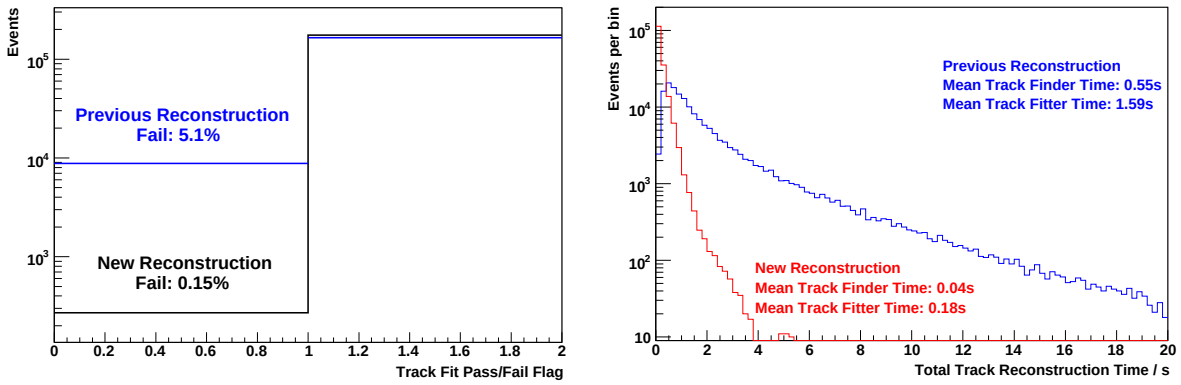

(a) (b)

Figure 5.14: (*a*) The reduction in track fitting failures obtained using the new reconstruction software. (*b*) The reduction in CPU time obtained using the new software.

fitting a track. The new software significantly reduces the required CPU time, cutting the mean track reconstruction time from 2.14 s to just 0.22 s.

5.2.3 Far Detector Cosmic Muons

The new track reconstruction software is intended for use with more than just beam neutrino interactions. The algorithms are designed to cope with steep tracks, where the muon can pass through a large number of strips on a small number of planes, and with muons that move either forwards or backwards along the detector z axis.

Processing a large sample of Far Detector cosmic muon MC events produced the results in Figure 5.15. Figure 5.15(a) shows the fractional discrepancy between the reconstructed and true q/p values for this sample. Despite the sample consisting of steep high energy muon tracks, which display little curvature in the magnetic field, successful q/p reconstruction is obtained. Figure 5.15(b) shows the reconstructed z direction cosines at the track vertex. The software successfully reconstructs muon tracks with a large range of gradients and deals effectively with muons moving either forwards or backwards along the detector z axis.

The track finding package is capable of identifying multiple tracks in a slice, and Figure 5.16 demonstrates successful reconstruction of a real multiple muon event in the Far Detector. The software ensures that the beginning and end z coordinates for each separate track match closely in the u and v views. It is also careful to avoid tracking through the low pulse height cross-talk strips, which the demultiplexing software often

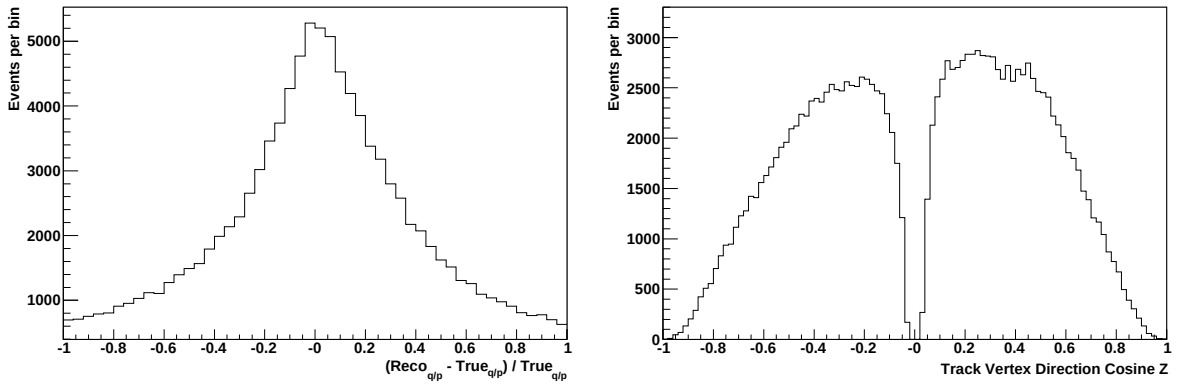

(a) (b)

Figure 5.15: Reconstruction performance for a sample of MC cosmic muons. (*a*) The fractional discrepancy between the reconstructed and true q/p values. (*b*) The distribution of reconstructed vertex z direction cosines.

incorrectly aligns with the tracks. Events with very high muon multiplicity can be accurately reconstructed.

The track reconstruction software is now used for a large number of different physics analyses in the MINOS experiment. For instance, Figure 5.17 illustrates how the new software helps with the identification of upward-travelling muons, induced by atmospheric neutrinos. The Figure shows an event display for a real cosmic muon with a large and complicated vertex shower. Poor tracking with the old reconstruction incorrectly identifies the track as upward-travelling, but the new reconstruction provides accurate tracking and so helps reduce background events in the analysis.

5.3 Near Detector Spectrometer Demultiplexing

The Near Detector spectrometer consists of planes 121−281 and, in this region, each electronics channel reads out four different strips on the same plane. This multiplexing reduces the required electronics, but introduces an ambiguity in the strip positions that must be resolved during the track reconstruction.

As detailed in Chapter 4, the track finding software does not consider the spectrometer region and identifies only seed tracks in the first 120 planes of the detector. To complete the track reconstruction and remove the ambiguity in the strip positions, the seed tracks need to be carefully extended into the spectrometer. As suggested in [78], this track extension can be controlled by the Kalman filter.

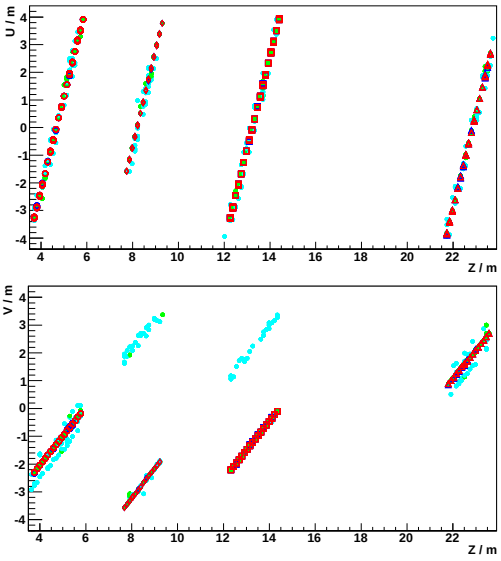

Figure 5.16: Event display for a typical multiple muon event recorded in the Far Detector. Red markers indicate reconstructed track strips, whilst green markers represent strips in the slice. Cyan markers indicate low pulse height slice strips (largely due to cross-talk).

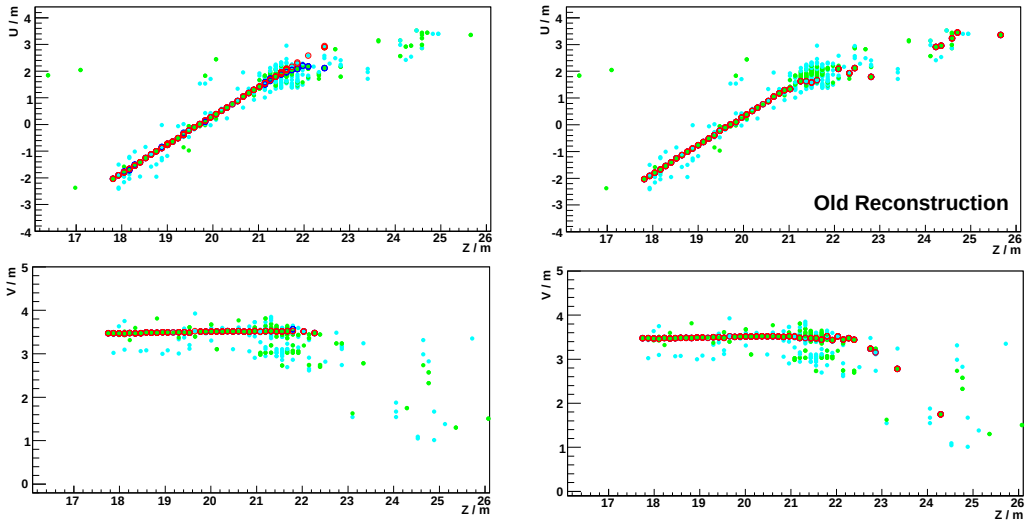

(a) (b)

Figure 5.17: The reconstruction of a complex Far Detector cosmic muon interaction, (a) with new reconstruction software and, (b) with previous software.

The procedure for demultiplexing the spectrometer region is summarised in Figure 5.18. Figure 5.18(a) shows the arrangement of strips in the slice, including all the possible strip positions in the spectrometer. The true muon path is simple to identify by eye, but the challenge is to develop software to successfully automate the process.

Figure 5.18(b) shows the seed track that is identified in the forward region of the detector. This seed track provides the input to the Kalman filter, which begins the track fit at the vertex and then moves towards the spectrometer. When the fit reaches the edge of the spectrometer, as shown in Figure 5.18(c), the Swimmer is used to predict the Kalman state vector at the next plane containing reconstructed strips. At this plane, the predicted state vector is compared to the list of possible slice strip positions. This list must account for the multiplexing of the electronics and must represent all the different possible strip patterns that could produce the observed output from the electronics. Any strips that closely match the Swimmer prediction are included in the Kalman fit, using an iteration of the Kalman update equations.

The extension of the seed track continues, as shown in 5.18(d), until a large plane window is encountered in which no strips are found. At this point, the situation resembles that shown in Figure 5.18(e). The Kalman fit then reverses its direction and moves back towards the track vertex. The remaining iterations are then completed, with no further special treatment for the spectrometer region. The final strips included in the fitted track are those that match the filtered state vectors after completion of the last iteration.

After the final track strips are identified, the strips representing the alternative demultiplexed positions can be safely removed. This produces the final output shown in Figure 5.18(f). Any spectrometer strips that are not included in a track are not demultiplexed. Figure 5.19 shows two further examples of successful spectrometer track strip identification and demultiplexing.

In order to obtain precise track strip identification in the spectrometer, the seed tracks need to provide the Kalman filter with sufficient information about the muon. For this reason, tracks are typically only used in an analysis if they cross a large enough number of planes in the forward region of the detector. This selection helps to remove the majority of events with poor spectrometer demultiplexing.

However, poor demultiplexing can also occur if there are several muon tracks in the slice, typically if there is a second track that is entirely contained in the spectrometer. The problem is that the seed track can be extended too far, by incorrectly adding the strips produced by the second muon. Figure 5.20 shows an example of this problem, with Figure 5.20(a) showing the reconstruction output and Figure 5.20(b) revealing the true strip locations.

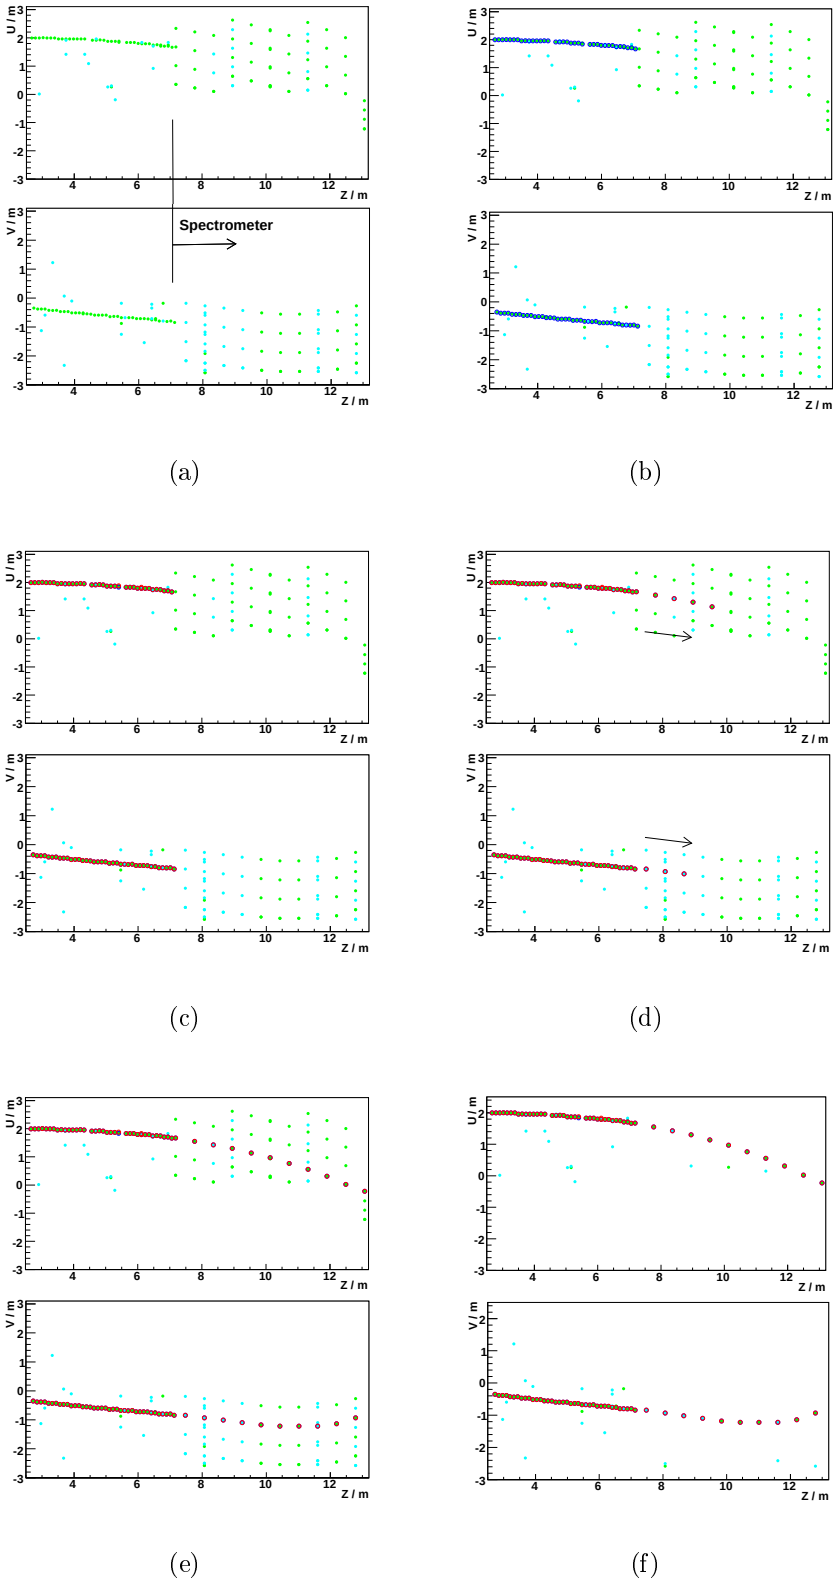

Figure 5.18: Summary of the demultiplexing procedure in the Near Detector spectrometer.

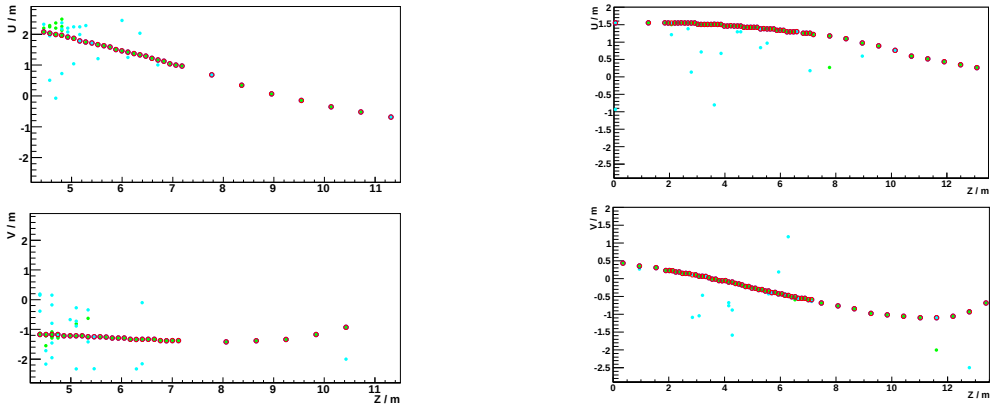

Figure 5.19: Event displays for typical Near Detector CC beam neutrino interactions, for which the track strips in the spectrometer region have been successfully demultiplexed.

To avoid including strips from a different muon track, the demultiplexing procedure can make use of timing information. Using nearby track strips, a sliding estimate of the track time is calculated during the Kalman fit. A spectrometer strip can then only be included in the fit if its strip time is within ± 40 ns of the current track time estimate. The Near Detector provides timestamps in 19 ns buckets, so a time window of 40 ns is sufficiently large to avoid causing any problems with the reconstruction of simple muon tracks in either data or MC. Its sole function is to avoid joining separate muon tracks together and, as shown in [79], it works very effectively.

With the use of timing information, the reconstruction of the example in Figure 5.20 is improved, producing the final output shown in Figure 5.20(c). In this final output, the seed track is extended correctly, whilst the track contained entirely in the spectrometer is neglected by the reconstruction.

5.4 Near Detector Results

Having implemented the spectrometer track finding and demultiplexing, the performance of the new reconstruction software can be evaluated for the Near Detector. In this Section, the results obtained using a sample of Near Detector beam MC are examined.

Figure 5.21 shows the fractional discrepancy between the reconstructed and true q/p values for true CC events with reconstructed track vertex in the detector fiducial volume. Figure 5.21(a) shows the discrepancy for all these events, whilst Figure 5.21(b) shows only events with uncontained tracks and Figure 5.21(c) shows only events with fully contained tracks.

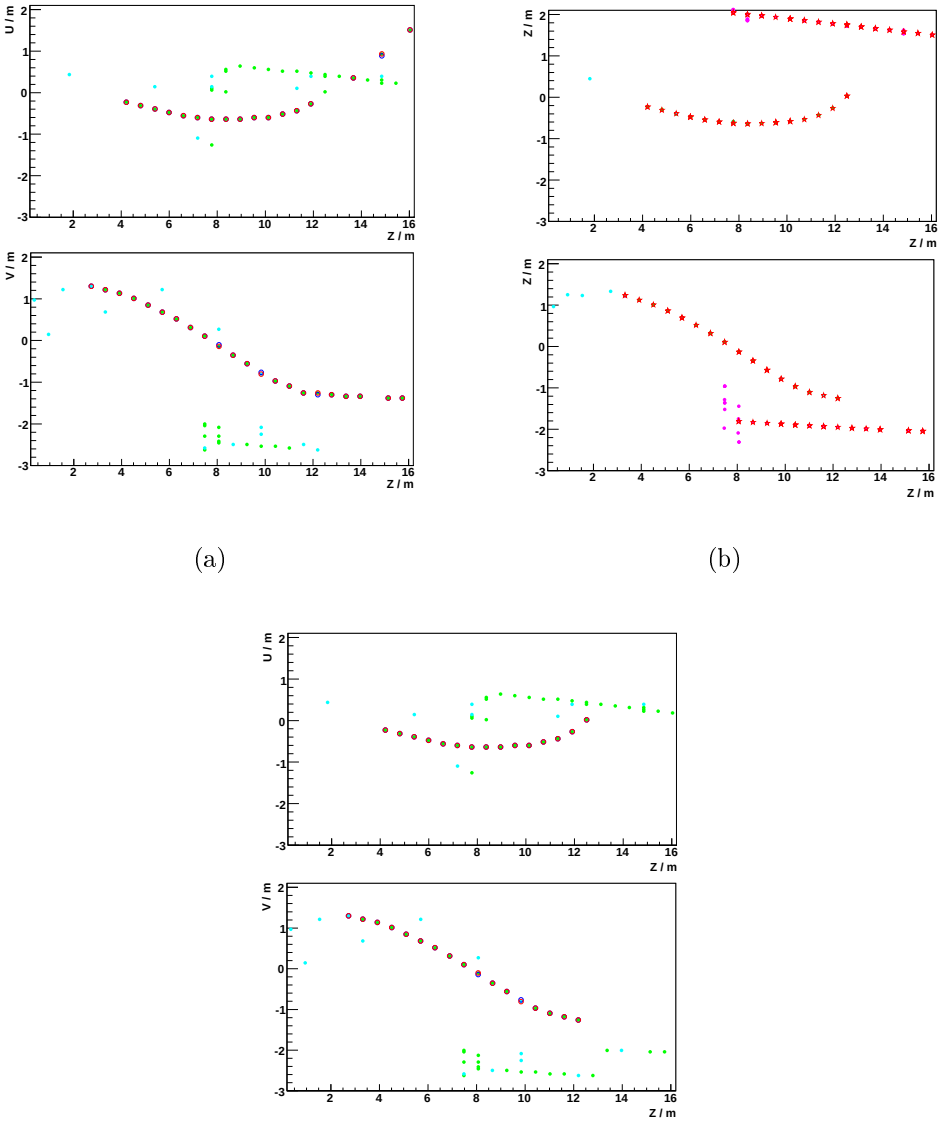

(a) (b)

(c)

Figure 5.20: Example event displays in which there are two true muon tracks, one of which is entirely contained in the spectrometer. (*a*) The default demultiplexing algorithm incorrectly extends the first track too far, using strips from the second track. (*b*) The true arrangement of strips for this slice. (*c*) The improved reconstruction, after use of timing information in the demultiplexing algorithm.

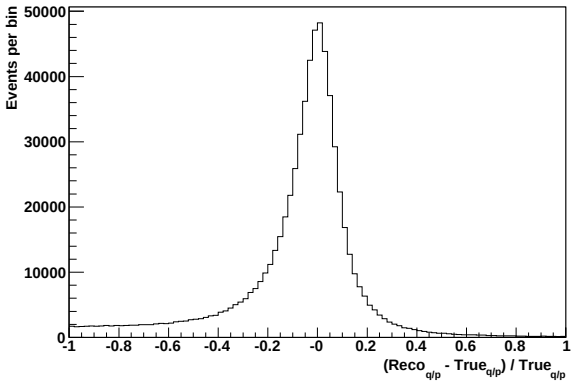

(a)

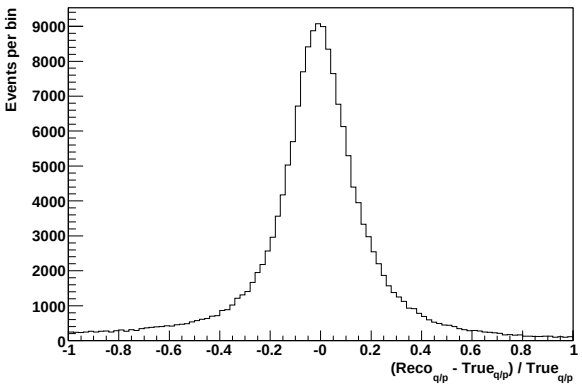

(b)

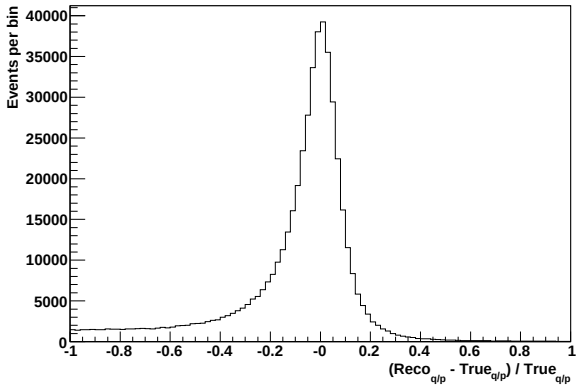

(c)

Figure 5.21: The fractional discrepancy observed between the reconstructed and true q/p values for Near Detector true CC events with a fiducial track vertex: (a) for all events, (b) for events with uncontained tracks and, (c) for events with fully contained tracks.

The fractional discrepancy between the reconstructed momentum from range and the true muon momentum is shown in Figure 5.22 for true CC events with fully contained tracks and fiducial track vertices. Figure 5.23 shows the fractional offset in muon momentum from range and the percentage resolution, both plotted as a function of the true muon momentum. The important features in these distributions can be explained by the different nature of the fully contained tracks at low and high muon momentum. At low momentum, the tracks are typically contained in the calorimeter section of the detector. Every plane is instrumented in this region of the detector, so the tracks have well measured momentum. The resolution approaches 5%, as observed at the Far Detector, and the fractional momentum offset is negligible.

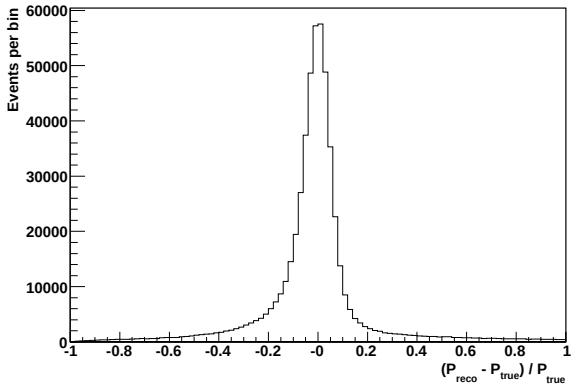

Figure 5.22: The fractional discrepancy between the reconstructed momentum from range and the true momentum, as observed for Near Detector true CC events with a fully contained track and fiducial track vertex.

At higher momentum, the tracks tend to enter the spectrometer. Only one plane in every five is instrumented in this region of the detector, and the multiplexing of the electronics makes track strip identification more difficult. These complications reduce the accuracy of the momentum from range reconstruction. The mean offset in the reconstructed values becomes increasingly negative, suggesting that the momentum is often underestimated. This observation is consistent with the increased difficulty in identifying the true end point of the track.

Figure 5.24 shows the fractional offset in the muon momentum from curvature and the percentage resolution, both plotted as a function of the true muon momentum. Two different regimes can be observed, particularly in the resolution plot, due to different classes of uncontained tracks. At low momentum, the uncontained tracks tend to be short and exit through the side of the detector. Momentum reconstruction is difficult for these short tracks, as they display little curvature. At higher momentum, the uncontained tracks tend to be those that pass through the entire spectrometer and often exit the detector through the final plane. These long tracks demonstrate clear curvature in the detector magnetic field and allow accurate calculation of the momentum.

The accuracy of the track strip identification is demonstrated in Figures 5.25 and 5.26, which show the purity and completeness of the fitted tracks. These purity and completeness variables are described in Section 4.3, but now account for strips both in the forward section of the detector and in the spectrometer region. The distributions indicate that the final fitted tracks have slightly higher purity and completeness than the seed tracks. However, the variation of both purity and completeness with muon momentum is unchanged and remains as described for the seed tracks in Section 4.3.

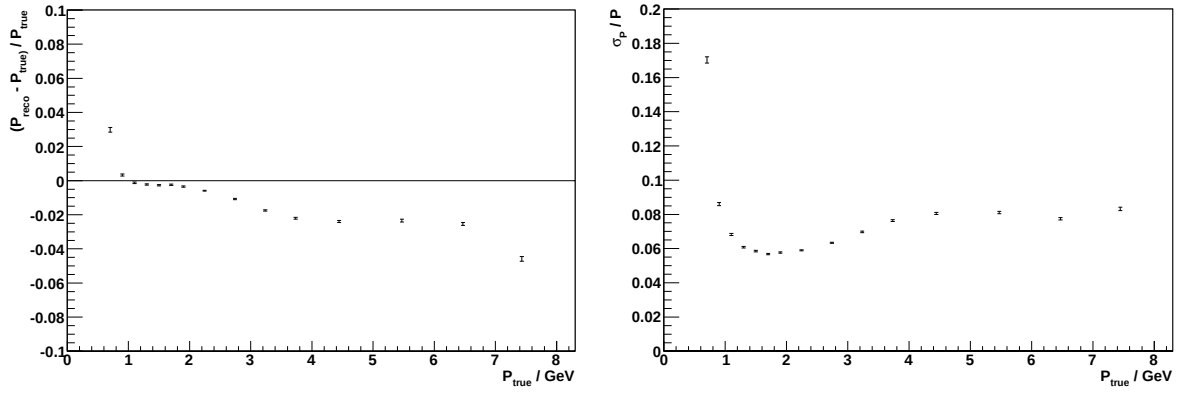

(a) (b)

Figure 5.23: (*a*) The fractional offset in muon momentum from range as a function of the true muon momentum. (*b*) The percentage resolution as a function of the true muon momentum. Both (*a*) and (*b*) consider only Near Detector true CC events with a fully contained track and fiducial track vertex.

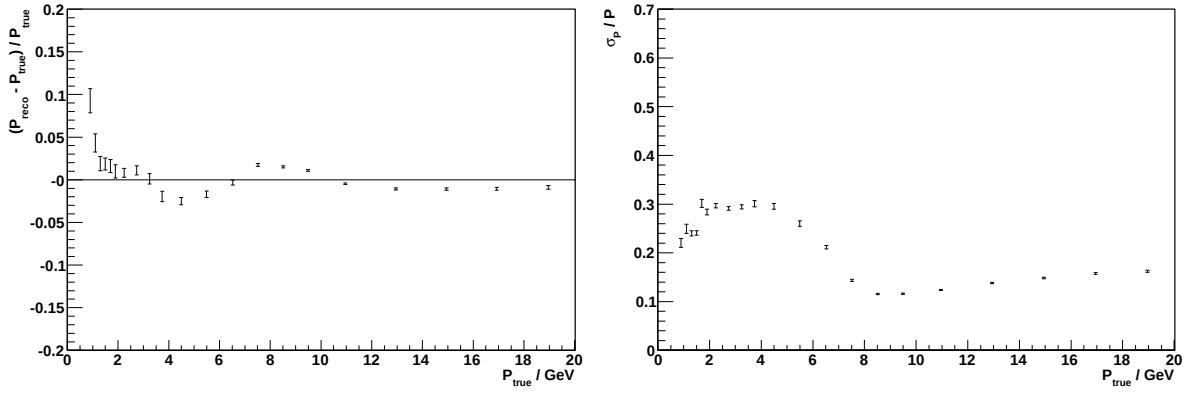

(a) (b)

Figure 5.24: (*a*) The fractional offset in muon momentum from curvature as a function of the true muon momentum. (*b*) The percentage resolution as a function of the true muon momentum. Both (*a*) and (*b*) consider only Near Detector true CC events with an uncontained track and fiducial track vertex.

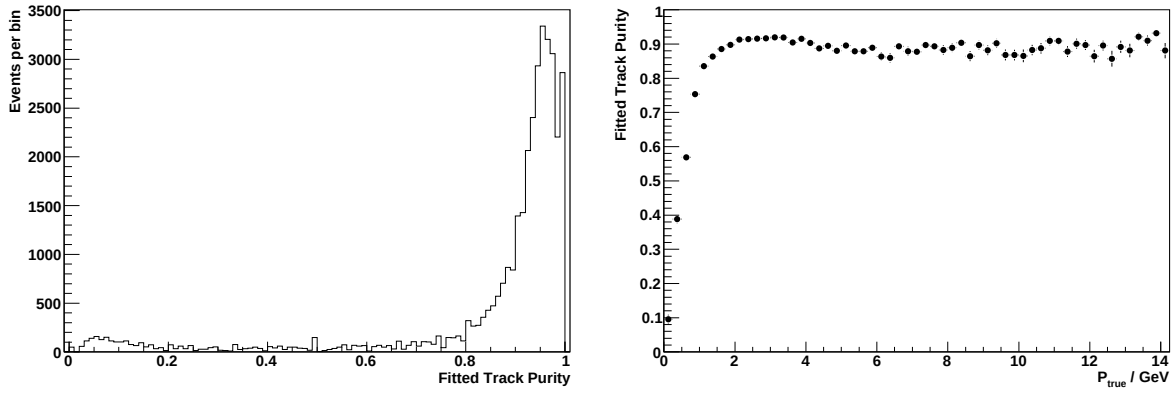

(a) (b)

Figure 5.25: (*a*) The distribution of fitted track purity for Near Detector events with a true muon track. (*b*) The variation of fitted track purity with true muon momentum.

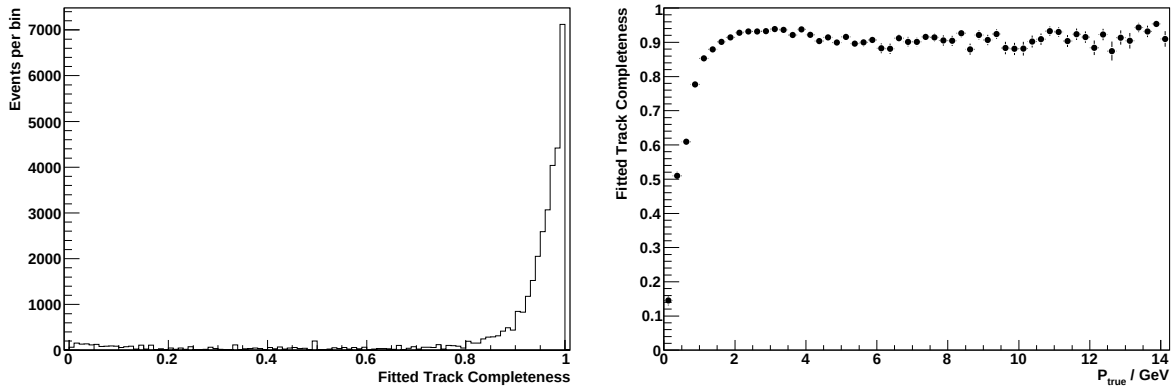

(a) (b)

Figure 5.26: (*a*) The distribution of fitted track completeness for Near Detector events with a true muon track. (*b*) The variation of fitted track completeness with true muon momentum.

The final result of interest concerns the fraction of tracks that fail the Kalman fit. This fraction is higher for the Near Detector than for the Far Detector, due primarily to the difficulties associated with non-uniform instrumentation and the spectrometer. However, the Near Detector failure rate remains low, at only 2.5%. This represents a sizeable decrease from the failure rate of 9% for the old reconstruction software.

Chapter 6

Event Selection

The aim of the analysis described in this thesis is to measure the neutrino oscillation parameters that control $\nu_\mu \leftrightarrow \nu_\tau$ flavour transitions. To make this measurement, pure samples of charged current (CC) beam ν_μ interactions must be selected in both the MINOS Near and Far Detectors. The Near Detector sample will provide information about the neutrino beam before oscillations have a chance to develop, whilst the Far Detector sample will fully incorporate the effects of oscillations.

The selection of beam CC ν_μ interactions is therefore of fundamental importance to the analysis and must be performed carefully and accurately to ensure the validity of the physics results. The selection procedure can be broadly divided into three stages:

1. Data quality checks to ensure that the beam and detectors were functioning correctly at the time of data collection.

2. Identification of CC beam neutrino events, based primarily on the presence of a reconstructed track with vertex inside the detector. At this stage, selection cuts are made to remove cosmic muons and anti-neutrino events, but a large neutral current (NC) background is expected.

3. Removal of the NC background from the sample, through use of a multivariate analysis, which compares the properties of events with the MC expectation for CC and NC interactions.

These stages in the selection procedure are described in detail in this Chapter. Also

described are the dataset used in the analysis and a data-driven technique to improve the MC modelling of the NC background.

6.1 Data Quality

6.1.1 Detector Status

The requirements for Far Detector physics data to be classified as 'good data' are described in detail in [80]. During detector operations, a large volume of detector monitoring information is stored and this information can be used to exclude bad data on an event-by-event basis. The aim is to use only data collected at times when the detector and all its subsystems were fully operational. The selection cuts made to identify good Far Detector data can be summarised as follows:

- All parts of the detector readout system must be enabled, with all 16 front end readout crates operational. If any readout crates were not operational, regions of the detector would be inactive and incomplete events may be recorded.
- The detector PMTs are powered by 8 high voltage (HV) mainframes, each providing power to approximately 200 PMTs. Any HV trips will create large regions of dead readout, so HV trips must be identified and data collected during a time of dead readout discarded. The HV trips are identified by monitoring the number of triggered 'singles' rates on each PMT. Singles are dynode triggers due to photons naturally incident on the PMT photocathode (such as those produced by the relaxation of mechanical stress in the optical fibres). A PMT whose singles rate has dropped significantly below the normal rate of 300 Hz is labelled as a 'cold chip'. Problems with the HV mainframes lead to a large number of cold chips, and a data quality cut of ≤ 20 cold chips is used to exclude data collected at the time of a HV trip.
- The detector magnetic field must be fully operational, with normal coil currents in the range 80 ± 1 A. The data for this analysis were all collected with the detector in forward field configuration, focussing μ^- arriving in the direction from Fermilab.
- The GPS timing system, used to correlate beam spills in the Near and Far Detectors, must report a timing error of less than 1000 ns.

- Data collected during Light Injection (LI) runs are excluded. LI is used to measure the gain of each PMT pixel and can be identified by the triggering of a special PMT, and by topological cuts.

For the Near Detector, a cut is made on the coil current to ensure correct operation of the detector magnetic field. However, further data quality cuts simply involved performing run selection by hand, using the distribution of physics quantities such as track vertices to reject bad runs. The volume of Near Detector data means that this simple data quality selection will suffice. However, a more rigorous approach to Near Detector data quality is currently being implemented for future analyses [81].

6.1.2 Beam Monitoring

During each beam spill, many beam quantities are measured and recorded. This allows the selection of only data associated with good quality spills. A list of beam quality criteria is presented in [82], including cuts on the position of the beam. These cuts were determined by finding the range of vertical and horizontal positions for which the hadron and muon monitor signals were stable.

The beam position cuts vary with time, in order to ensure an optimal selection across the entire dataset. Typical values for the position cuts are:

- $-2.0\,\text{mm} <$ Beam Spot mean x-position $< -0.01\,\text{mm}$
- $0.01\,\text{mm} <$ Beam Spot mean y-position $< 2.0\,\text{mm}$
- $0.1\,\text{mm} <$ RMS of Beam Spot x profile $< 1.5\,\text{mm}$
- $0.1\,\text{mm} <$ RMS of Beam Spot y profile $< 2.0\,\text{mm}$

In addition to these beam position cuts, the following selection cuts are also applied to ensure consistent and high quality beam conditions:

- $0.5 \times 10^{12} <$ PoT in spill $< 50 \times 10^{12}$, to reject spills reporting abnormal proton on target values.
- $-200\,\text{kA} <$ Horn current $< -155\,\text{kA}$, to ensure the horns are functioning correctly.
- $|t_{snarl} - t_{spill}| < 1\,\text{s}$, ensuring sensible values for the difference between the start of the window during which data is recorded (t_{snarl}) and the Main Injector prediction for the spill time (t_{spill}).

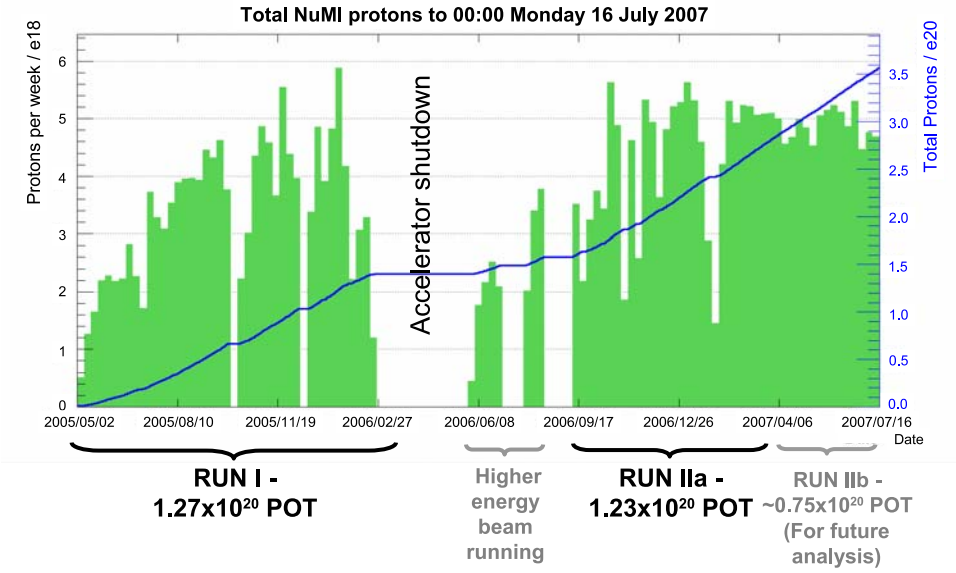

Figure 6.1: Summary of protons delivered to the NuMI target during the first two periods of operation of the MINOS experiment. Taken from [83].

- The NuMI target is required to be in the LE-10 position. This is the target position used for the majority of MINOS beam data collection. As will be described in Chapter 7, different beam conditions need to be considered separately when extrapolating Near Detector data to make predictions of Far Detector distributions. Some high energy (HE) beam data has been collected, totalling 1.5×10^{19} PoT, but this analysis uses only the LE-10 data.

6.2 The Analysis Dataset

This analysis uses data collected during the first two periods of operation of the MINOS experiment. Figure 6.1 shows a summary of the protons on target (PoT) during these two periods. Run I commenced in late May 2005 and ended with the scheduled accelerator shutdown in February 2006. During Run I, the NuMI target was in the LE-10 position and, after making all the data quality and beam quality cuts described above, the Run I Far Detector data sample represents 1.27×10^{20} PoT.

Immediately after the accelerator shutdown there was a period of high energy beam running, but the data collected with this beam configuration are not included in this analysis. During this time, problems with the NuMI target required a complete target replacement. A new target was installed and placed in the LE-10 position for the start

of Run II.

Run II commenced in mid September 2006 and a dataset labelled Run IIa was closed at the end of March 2007. The closure of this dataset enabled analysis in time for conference presentations in summer 2007. However, datataking continued and Run IIb consists of the data collected from April 2007 to mid July 2007. After making data quality cuts, the Run IIa Far Detector data sample represents 1.23×10^{20} PoT. The combined Run I and Run IIa Far Detector data samples therefore represent a total of 2.5×10^{20} PoT and the final dataset for this analysis. This final Far Detector analysis dataset is summarised in Table 6.1.

For the Near Detector, the large volume of data collected means that it is not necessary to use the entire Run I and Run IIa Near Detector datasets. In the analysis, all that is required is a Near Detector data sample that tracks any changes in the beam energy or composition. For this reason, approximately one complete Near Detector run (lasting about 24 hours) was selected per week. This resulted in final Near Detector data samples representing 2.37×10^{19} PoT from Run I and 2.21×10^{19} PoT from Run IIa.

In Section 6.5, the observed Near Detector neutrino energy distributions are examined month by month to check for consistency. It is observed that the change of target between the Run I and Run IIa samples resulted in a clear change in the neutrino energy spectrum. The change in beam conditions means that, in the analysis, the two samples need to be treated separately.

A summary of the Near Detector dataset used in this analysis is shown in Table 6.2.

Far Detector			
Beam Type	**Start (Run/Subrun)**	**End (Run/Subrun)**	**PoT / $\times10^{19}$**
LE-10 Run I	20/05/2005 (31720/0)	26/02/2006 (33791/23)	1.27
LE-10 Run IIa	12/09/2006 (36570/0)	31/03/2007 (37832/23)	1.23

Table 6.1: Summary of the Far Detector analysis dataset.

Near Detector			
Beam Type	**Start (Run/Subrun)**	**End (Run/Subrun)**	**PoT / $\times 10^{19}$**
LE-10 Run I	02/06/2005 (7876/0)	23/02/2006 (9857/0)	2.37
LE-10 Run IIa	13/09/2006 (10783/0)	30/03/2007 (11995/23)	2.21

Table 6.2: Summary of the Near Detector analysis dataset.

6.3 Selection of Charged Current ν_μ Interactions

6.3.1 Triggers

The Data Acquisition (DAQ) systems for the MINOS detectors only write out detector data if a trigger signal has been received. For the analysis of beam data, special spill triggers are used to ensure that the detectors are read out whenever beam neutrinos pass through. Not reading out detector data at other times has the advantage of reducing backgrounds from cosmic muons and noise. The spill triggers provide the first and lowest level cut for selecting a sample of beam CC ν_μ interactions.

The Near and Far detector spill triggers are prompted by a signal generated in the Main Injector. This signal is an instruction for the kicker magnet to fire and extract protons into the NuMI beamline. The Main Injector operates with a clock frequency of 53 MHz and the kicker signal is generated 20 cycles before the magnet actually fires. In the Near Detector, the signal prompts the continuous digitisation of every PMT pixel in a 13 μs period, starting about 1.5 μs before the arrival of the neutrinos. A process known as the SpillServer also sends the GPS timestamp of the spill to the Far Detector trigger processors.

The Far Detector DAQ has data buffering that is large enough to allow it to wait for the GPS timestamp of the spill to arrive from the Near Detector via the Internet. A simple time of flight calculation allows prediction of when the beam neutrinos should arrive at the Far Detector and this forms the basis of a bias-free trigger. This trigger prompts the recording of all detector hits within a time window of 100 μs around the predicted time of arrival at the Far Detector.

The splitting of events at the Far Detector is avoided by extending the spill trigger window so that there is a period of at least 156 ns without detector activity on either side of the window. Finally, a period of 30 μs before the spill trigger window is also read out. This is required to determine if there was any earlier detector activity that might cause electronics dead time during the spill trigger window.

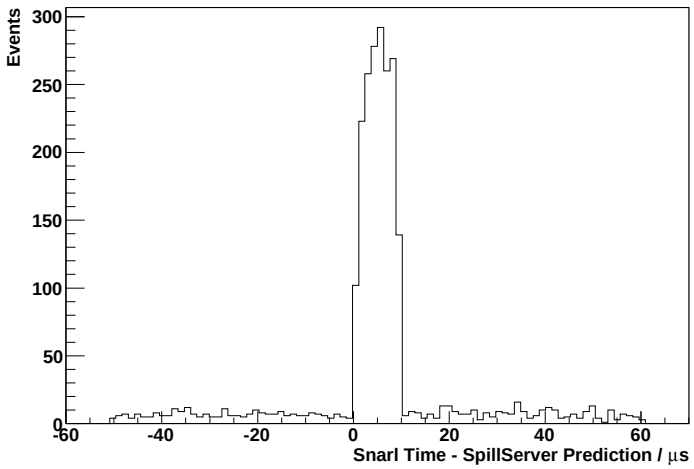

Figure 6.2: The difference between the Far Detector snarl time and the SpillServer prediction, for Far Detector data events satisfying data quality cuts.

Figure 6.2 shows the distribution of the difference between the Far Detector snarl time and the prediction made by the SpillServer. This distribution is made for all the events recorded in the Far Detector, after data quality cuts. Most of these events will be simply due to noise and will likely have no track and a very low reconstructed energy. To reduce the amount of noise and cosmic muons in the sample, a further Far Detector timing cut is implemented in the analysis. This cut requires that the difference between the snarl time and the SpillServer prediction be between -20 μs and 30 μs.

6.3.2 Muon Track and Detector Fiducial Volume

The primary signature of a CC ν_μ interaction in the MINOS detectors is a muon track, so events are only selected if they contain a reconstructed track. The efficiency of the track finding software is very high (greater than 99% for muons with momentum above 1.5 GeV), meaning that nearly all muon tracks will be identified. However, tracks due to pions or protons, or simply fake tracks, can also be found. The process of distinguishing the true muon tracks from the fake tracks is the primary task required for the removal of the NC background and is addressed in Section 6.4.

Reconstructed tracks have well measured vertex coordinates, specifying the location of the neutrino interaction in the detector. The position in the detector is important, as neutrino interactions in some areas of the detector can lead to incomplete reconstruction and so to poor energy measurement. For instance, interactions near the edge of the detector can lead to a large loss of shower pulse height outside the detector and a

likely underestimate of the neutrino energy. Similarly, events where the reconstructed track starts at the edge of the detector are probably due to neutrino interactions in the surrounding rock. To help ensure the events in the final sample are accurately reconstructed beam neutrino events with well measured energies, events are only selected if the track vertex is contained within a specified fiducial volume of the detector.

Near Detector

For the Near Detector, the large volume of data collected means that the fiducial volume can be relatively small, sacrificing statistics in order to achieve a sample of very well reconstructed events. A fiducial vertex is defined by:

- Track vertex plane > 16
- Track vertex plane < 85
- $r_{vtx} < 1.0\,\mathrm{m}$, where r_{vtx} is the radial distance of the reconstructed track vertex from the position of the beam spot (at $x = 1.48\,\mathrm{m}$ and $y = -0.24\,\mathrm{m}$).

The spectrometer region of the detector begins on plane 121 and the back plane cut at plane 85 is intended to ensure that showers are fully developed before the spectrometer. This is necessary because there is no identification of shower strips in the spectrometer. The cut also ensures that, for any track with fiducial vertex, there will be enough information to accurately extend the track into the spectrometer during the track fitting process.

The front plane cut at plane 16 and the cut on radial distance from the beam spot help to remove rock muons and cosmic muons from the sample. The radial cut helps define a narrow region of the detector in which the instrumentation is entirely uniform and helps to avoid reconstruction problems associated with the detector coil hole.

The fiducial volume represents only 4.7% of the total Near Detector mass, but helps avoid many of the problems associated with the non-uniform detector instrumentation. Figure 6.3 shows the distribution of data track vertices in the Near Detector. The different regions covered by the partially and fully instrumented planes can be clearly seen. The areas included in the fiducial volume are highlighted.

Far Detector

For the Far Detector, the low event rate means that a high selection efficiency is very important. There is a delicate balance between a small fiducial volume that discards too

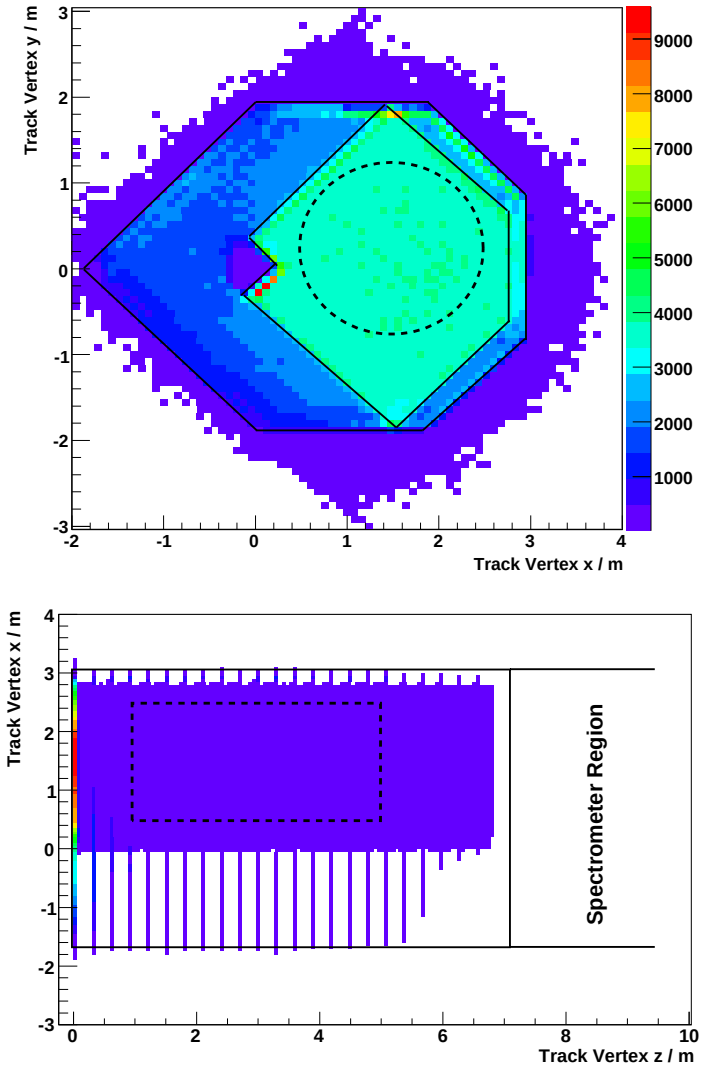

Figure 6.3: Distribution of reconstructed track vertices observed in Near Detector data. The detector fiducial volume is indicated by the dashed lines, whilst the solid lines show the outlines of the fully and partially instrumented regions.

many events and a larger volume that includes too many poorly reconstructed events. Both the exclusion of well reconstructed events and the inclusion of badly reconstructed events decrease the sensitivity to the oscillation parameters. The fiducial volume used should be that which maximises the sensitivity.

The Far Detector fiducial volume cuts were optimised in [84], resulting in the following definition for a fiducial vertex:

- Track vertex plane > 4
- Track vertex plane < 466
- The track vertex plane must not be in the region $240 <$ track vertex plane < 254, in order to remove interactions close to the supermodule gap.
- $r_{vtx} < \sqrt{14}\,\mathrm{m}$, where r_{vtx} is the radial distance of the reconstructed track vertex from the centre of the detector.

The front and back plane cuts help to remove rock muon and cosmic muon events. They also remove beam events for which the energy reconstruction is likely to be poor. The division of the Far Detector into two supermodules means that similar cuts must be placed around the supermodule gap. The radial cut removes the majority of cosmic muons and rock muons that enter the detector. The position of the cut provides the optimal balance between including as many events as possible and excluding events for which shower pulse height may be lost outside the detector. A cut around the coil hole, vetoing events with track vertex less than 0.4 m from the centre of the detector, is also advocated in [84]. This suggestion was rejected by the collaboration for use in the 'standard' fiducial volume, on the grounds of simplicity and increased selection efficiency. However, the idea of a coil hole cut is revisited in Chapter 9.

The optimised fiducial volume represents 78% of the total Far Detector mass. Figure 6.4 shows the distribution of track vertices for a large MC sample. The areas included in the fiducial volume are highlighted.

6.3.3 Track Properties

The remaining cuts that complete the basic selection of CC ν_μ events are all based on the properties of the reconstructed track. Firstly, $\bar{\nu}_\mu$ events are excluded by demanding that the track has a negative reconstructed charge sign. This identifies the track as being due to the passage of a μ^-, rather than a μ^+. Charge reconstruction requires that the track does not fail the Kalman fit, so events are only selected if they are flagged

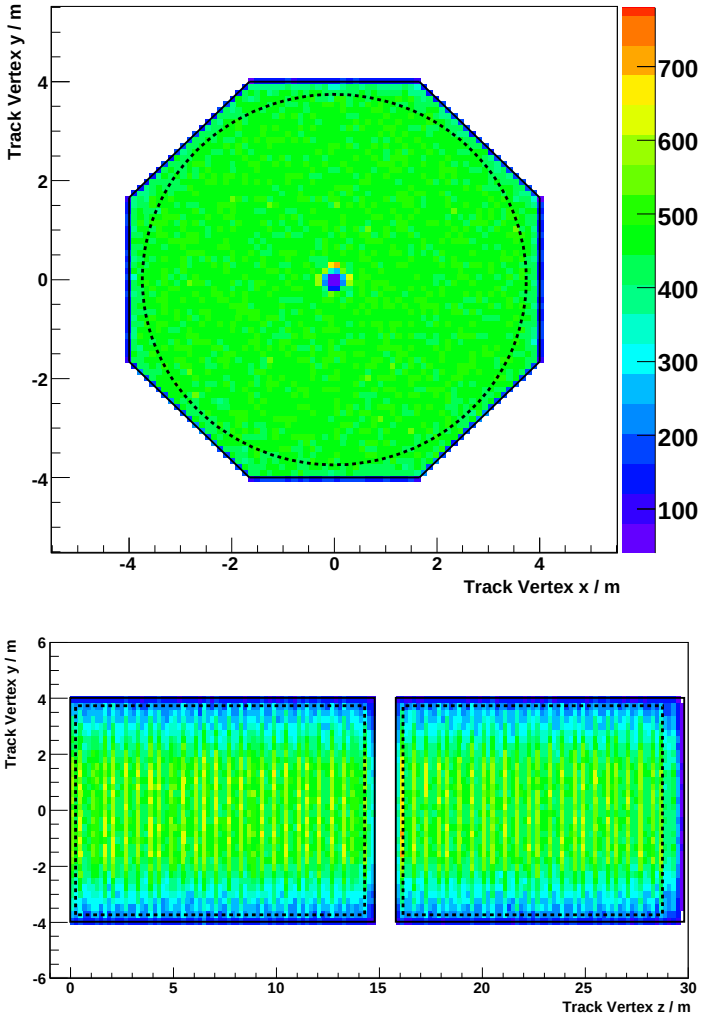

Figure 6.4: Distribution of reconstructed track vertices observed for a sample of Far Detector MC events. The detector fiducial volume is indicated by the dashed lines.

as having passed the fit. However, in the Near Detector (where the number of events failing the Kalman fit is of the order of a few percent), tracks without a reconstructed charge sign can be used in the analysis if they satisfy certain quality cuts. The quality cuts demand close agreement between the reconstructed track vertex planes in the u and v views and between the reconstructed end planes in the two views. At the Near Detector, the beam composition is 93% ν_μ, so it is assumed that these 'reclaimed' tracks are due to ν_μ interactions.

The final direct cut on a reconstructed track property affects only the Far Detector sample. In order to reduce the cosmic muon background at the Far Detector, a cut is made on the reconstructed direction of the muon track. Cosmic muon tracks are expected to be much steeper than the tracks induced by beam neutrinos. The track vertex direction cosine is calculated with respect to the beam direction and, for selection, this direction cosine must be greater than 0.6.

With the track direction cut and use of the fiducial volume, the cosmic muon background in the analysis sample is expected to be small. A detailed study [85] has reported acceptably small rock muon and cosmic muon background expectations for the 2.5×10^{20} PoT Far Detector data sample:

Rock muons:	1.7 ± 0.2 (stat.) events
Cosmic muons:	0.5 ± 0.1 (stat.) events

6.4 Removal of Neutral Current Background

The process of identifying and removing NC background events from the CC sample requires careful examination of the track reconstruction. In a signal CC event, the reconstructed track should represent the motion of a real muon. The track should be minimally ionising and should curve smoothly in the magnetic field. The muons produced in CC interactions lie in characteristic kinematic regions and this should be reflected in the reconstructed properties of the track and the neutrino event.

In a NC background event, however, the reconstructed track is not due to a muon. These tracks are termed 'fake muon' tracks. The fake muon tracks are often the result of tracking the motion of a pion or a proton, but they may simply be a series of hits selected from a NC shower. Fake tracks are to be expected if a high CC tracking efficiency is demanded at very low energies.

In order to separate the real muon tracks from the fake tracks, a multivariate analysis is used. This analysis is described fully in [86]. The basic idea is to select a number

of uncorrelated variables that characterise the topologies and kinematics of the tracks. For each event, the variables are compared with the MC expectation for CC and NC interactions and the information from these comparisons is incorporated into a 'Particle Identification' (PID) variable. NC background events can then be removed by making a selection cut on the value of the PID.

To calculate the PID, a set of CC and NC probability density functions (PDFs) must be generated for the chosen set of variables. These PDFs are created using true CC and true NC events from a large MC sample. Only events passing the standard CC selection cuts are used in the generation of the PDFs. For a given event, the CC and NC PDFs give a probability of the event being CC and NC respectively. The product of the probabilities for all the different variables gives the overall probability of the event being CC and NC:

$$P_{CC} = P(CC) \prod_i P(x_i|CC) \tag{6.1}$$

$$P_{NC} = P(NC) \prod_i P(x_i|NC) \tag{6.2}$$

In Equations 6.1 and 6.2, $P(CC)$ and $P(NC)$ are the overall normalisations of CC and NC events, x_i is the value of the i^{th} variable for the event and $P(x_i|A)$ is the value extracted from the PDF, i.e. the probability of obtaining this value of the variable given interaction type A. The overall CC and NC probabilities are combined together to create the final PID value:

$$\text{PID} = \frac{P_{CC}}{P_{CC} + P_{NC}} \tag{6.3}$$

Events with CC-like properties will have PID values close to one, whilst events with fake tracks will have PID values approaching zero. The use of the overall normalisations of CC and NC events means the PID represents the relative likelihood of an event being a CC interaction; the normalisations have no CC/NC separation effect.

The most important step in the PID analysis is the selection of the input variables. These need to provide strong discriminating power between the real and fake muon tracks. The following variables are used in the construction of the PID:

- **Track Topology.**
  - **Number of track-like planes.** The reconstructed track in most CC events will extend beyond the reconstructed vertex shower, whilst fake muon tracks are typically contained within the shower. The extent to which the track

extends beyond the shower is represented by the number of track-like planes: the number of planes containing strips that belong only to the track and not to any shower.

- **Pulse height per track-like plane.** The muons produced in CC interactions are approximately minimally ionising. The mean pulse height deposited in track-like planes will therefore occupy a sharp peak for real muons. The distribution will be much broader for the fake muon tracks in NC interactions.
- **Goodness of muon track fit.** The Kalman filter algorithm used in the track fitting process uses information based on the kinematics and energy loss of real muons. The relative error returned by the algorithm provides a measure of the goodness of the track fit and better fits are expected for real muon tracks than for fake muon tracks.
- **Reconstructed track charge.** The NuMI beam is composed primarily of neutrinos, rather than anti-neutrinos. Most CC interactions will therefore produce a μ^- and should be reconstructed with a negative track charge. The fake tracks reconstructed in NC events should have no preferred charge.

- **Event Kinematics.**
  - **Reconstructed y.** This is defined as the ratio between the reconstructed shower and neutrino energies. For CC interactions, the reconstructed y distribution will be peaked towards zero, particularly at low energies. A much flatter distribution is expected for NC interactions.

- **Relative CC/NC Spectrum.**
  - **Track length.** The NC spectrum is peaked towards low energies, whilst the CC spectrum covers the full range of energies from the neutrino beam. The distribution of track lengths reflects this difference without being highly correlated with neutrino energy. The variable used is actually the number of track hit planes, accounting for gaps in the tracks.

The input variables for the PID can be used in either 1D or 2D PDFs. In the case of 2D PDFs, the track topology variables are stored as a function of the track length, whilst reconstructed y is stored as a function of the reconstructed energy. Each slice in energy (or length) is normalised separately so that the PDFs can track variations with energy without directly depending on energy themselves. The PDFs used to generate the PID in this analysis are:

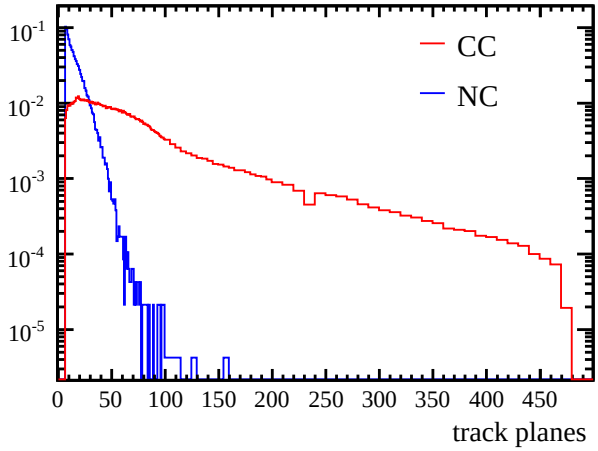

Figure 6.5: PDFs of reconstructed track length (number of hit planes) for selected CC events and NC background events. Taken from [86].

- Track length (number of hit planes),
- Reconstructed track charge,
- Number of track-like planes *vs.* track length,
- Pulse height per track-like plane *vs.* track length,
- Goodness of muon track fit *vs.* track length,
- Reconstructed y *vs.* reconstructed energy.

The PDFs used to calculate PID values at the Far Detector are shown in Figures 6.5 – 6.10. The PDFs all show the expected differences between CC and NC interactions that warranted their inclusion in the PID analysis.

The PID distributions generated for independent MC test samples are shown for the Near Detector in Figure 6.11(a) and for the Far Detector in Figure 6.12(a). The distributions obtained at the two detectors are very similar. In both detectors the CC distribution is strongly peaked at one, whilst the NC distribution is broadly peaked towards zero. A broad rise towards zero is also observed in the CC distribution. This is due to the CC events that have the most NC-like properties, with low muon energies and high shower energies. Similarly, a broad rise towards one is observed in the NC distribution, due to NC events with small showers and relatively long pion or proton tracks.

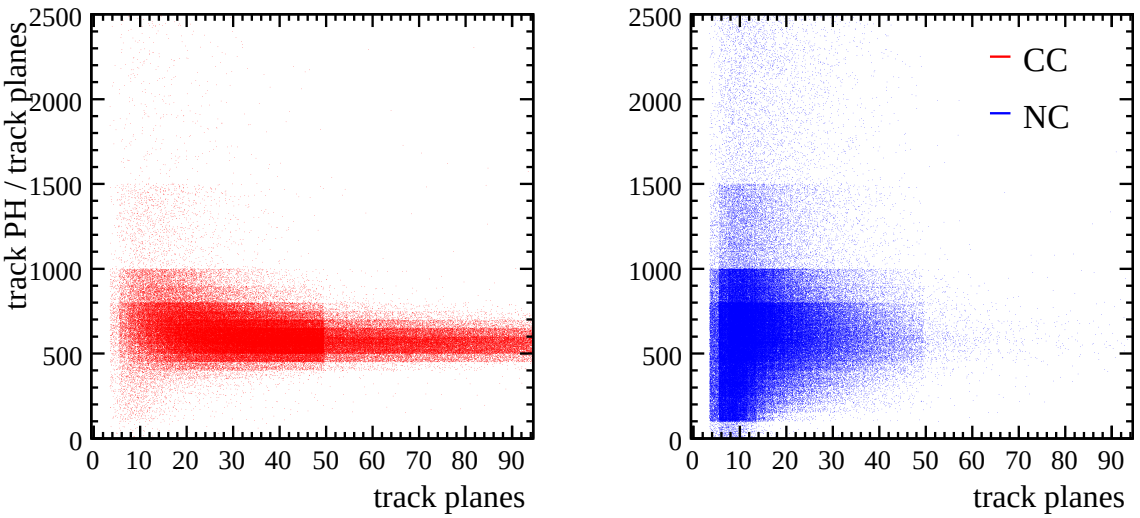

Figure 6.6: PDFs of pulse height per track-like plane *vs.* track length for selected CC events and NC background events. Taken from [86].

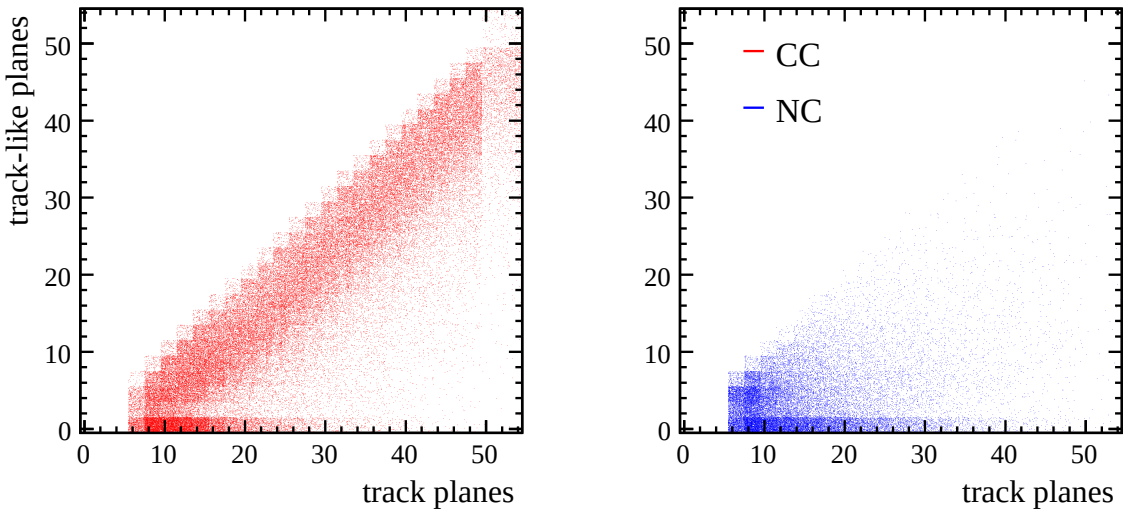

Figure 6.7: PDFs of number of track-like planes *vs.* track length for selected CC events and NC background events. Taken from [86].

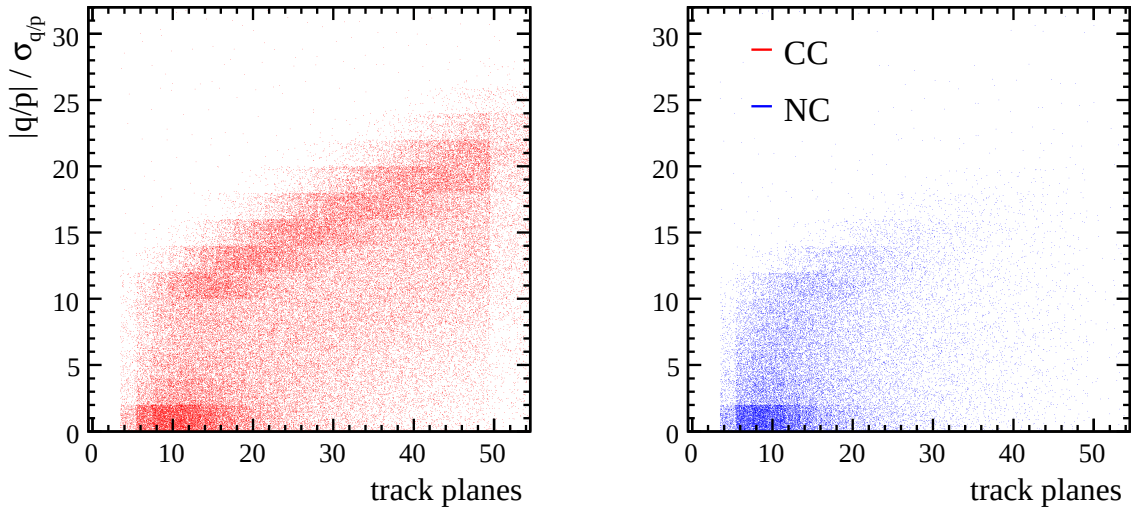

Figure 6.8: PDFs of goodness of muon track fit *vs.* track length for selected CC events and NC background events. Taken from [86].

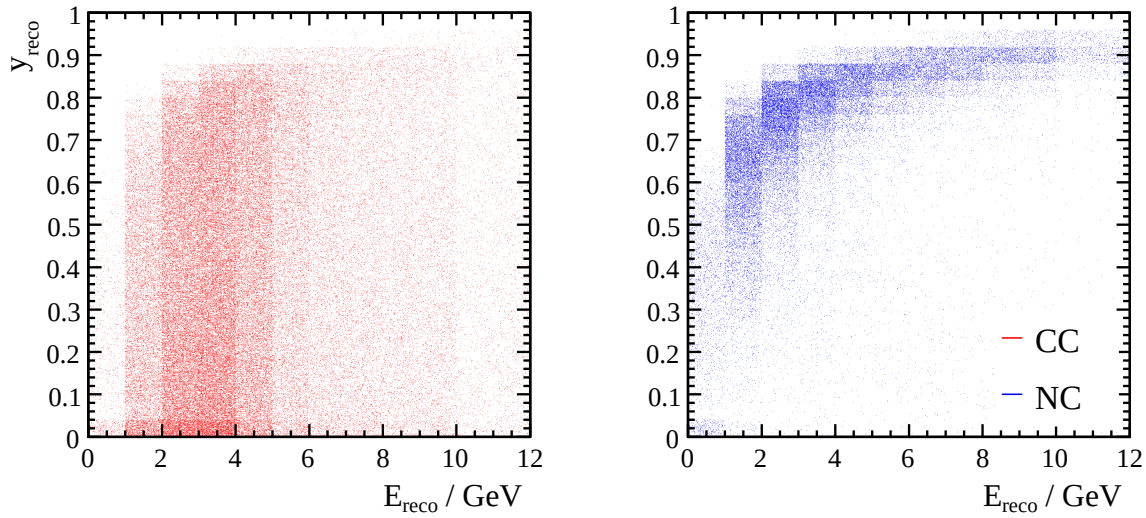

Figure 6.9: PDFs of reconstructed y *vs.* reconstructed neutrino energy for selected CC events and NC background events. Taken from [86].

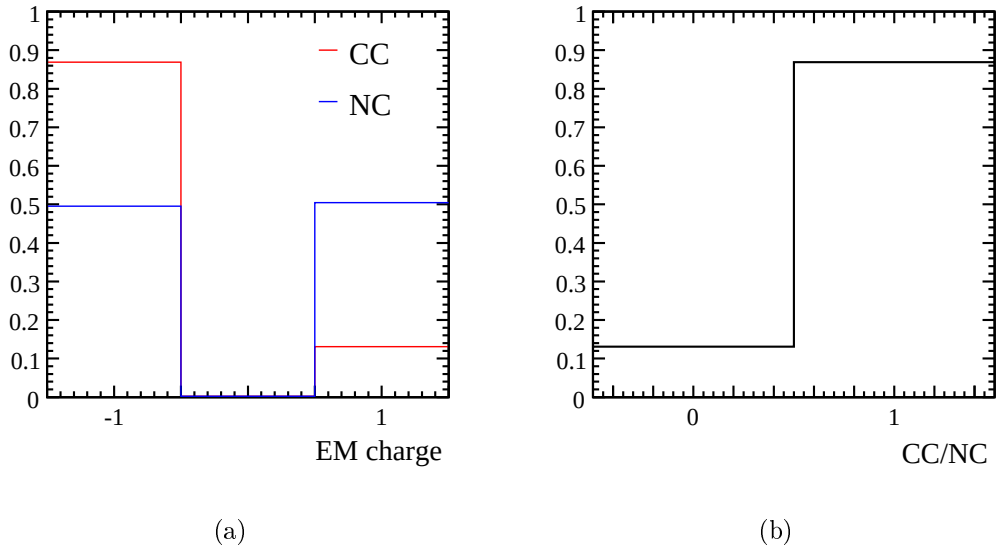

Figure 6.10: (*a*) PDFs of reconstructed track charge for selected CC events and NC background events. (*b*) The relative CC/NC normalisation. Both taken from [86].

Figures 6.11(b) and 6.12(b) show the Near Detector and Far Detector variations in the purity and efficiency of the CC selected sample for different PID cuts. Purity is defined as the fraction of the sample represented by true CC events, whilst efficiency is defined as the fraction of the original CC events that remain after the implementation of the PID cut. At both detectors, high purities and efficiencies are achieved.

The optimal value of the PID cut is that which provides the best balance between the purity of the sample and the efficiency. The value of the cut was optimised for maximal sensitivity to the oscillation parameters and reported in [87]. In both the Near and Far Detectors, events require a PID value greater than 0.85 in order to be selected. Applying this PID cut completes the selection of CC ν_μ events for use in this analysis.

6.5 Examination of Near Detector Data

Applying the CC selection cuts to the Near Detector dataset described in Section 6.2 allows examination of the properties of real beam CC ν_μ interactions. A detailed comparison of these data with the MC expectation is included in Chapter 8 of this thesis. In this Section, a few simple checks are made to ensure that the beam and Near Detector operated as expected during data collection.

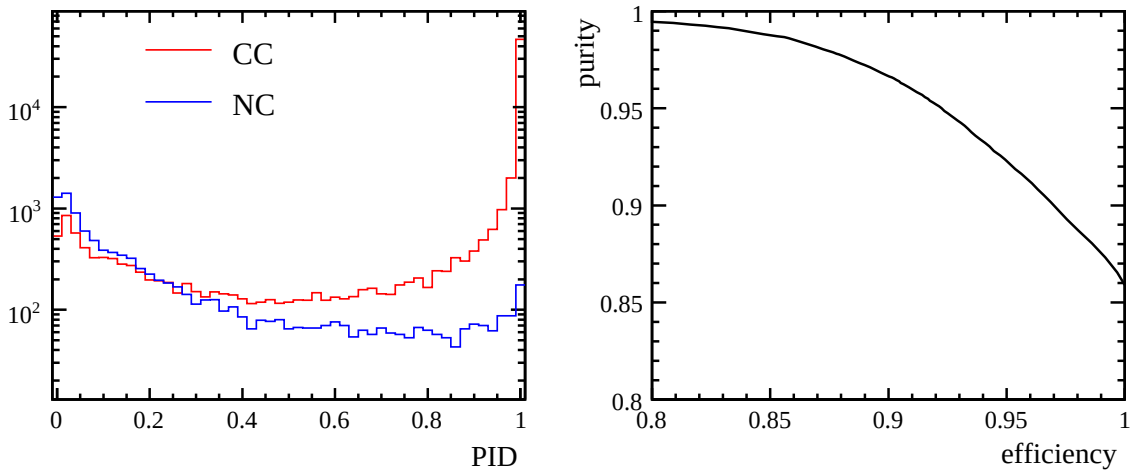

(a) (b)

Figure 6.11: (*a*) PID distributions obtained for selected CC events and NC background events in a 1.6×10^{19} PoT Near Detector MC sample. (*b*) The variations in purity and efficiency obtained by applying different PID cuts. Taken from [86].

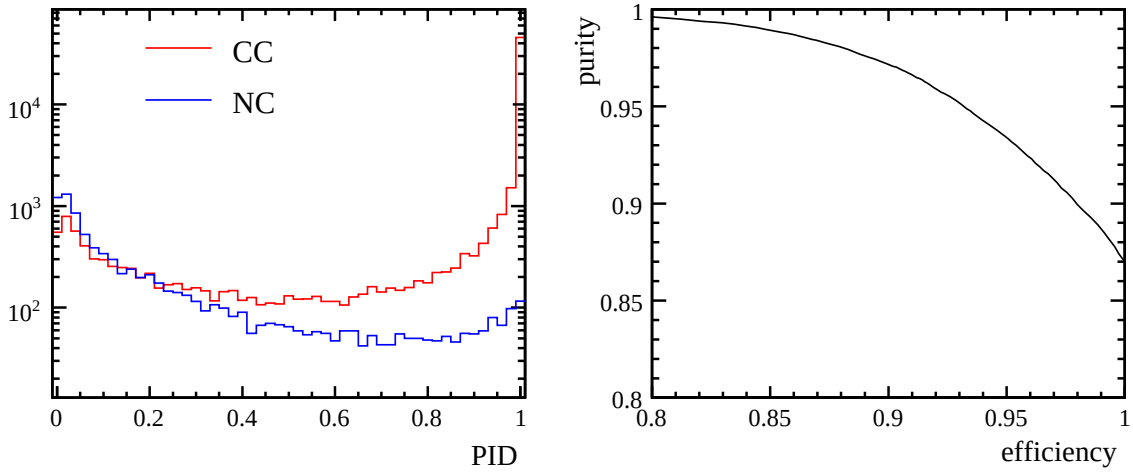

(a) (b)

Figure 6.12: (*a*) PID distributions obtained for selected CC events and NC background events in a 1.9×10^{20} PoT Far Detector MC sample. (*b*) The variations in purity and efficiency obtained by applying different PID cuts. Taken from [86].

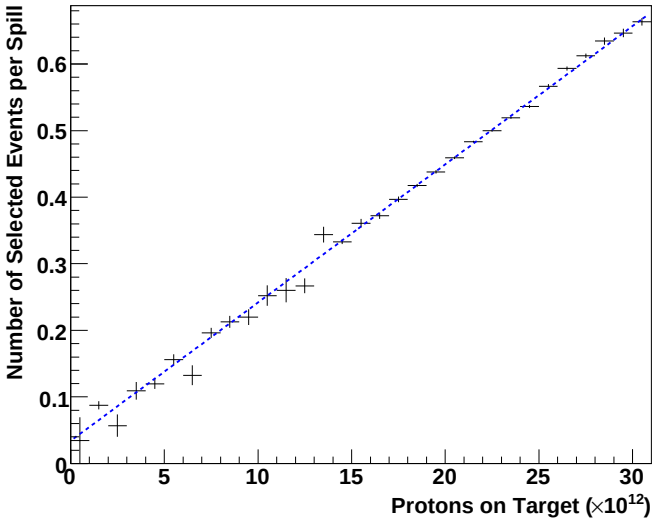

Figure 6.13: The number of selected Near Detector events per spill, shown as a function of the Protons on Target in the spill.

Figure 6.13 shows the number of selected events per beam spill as a function of the protons on target in the spill. A clear linear relationship is observed, suggesting that the reconstruction software can cope with the range of beam intensities in the dataset. This is an important observation, as high beam intensities can produce multiple overlapping events in the Near Detector and so make the reconstruction considerably more difficult.

Figure 6.14 shows the variation in the observed event rate throughout the dataset, indicating the number of selected events per 1×10^{16} PoT during each month. As expected, the event rate remains approximately constant within Run I and within Run IIa. However, a clear decrease in the event rate is observed upon moving from Run I to Run IIa. This can be attributed to the change in beam conditions between the two datasets; in particular, the replacement of the NuMI target will have had an important impact.

Figure 6.15 shows the neutrino energy spectrum observed month-by-month for selected events in the Run I and Run IIa datasets. Within each dataset, the monthly distributions display only small statistical fluctuations. However, Figure 6.16 suggests that there is a real difference between the energy spectra observed in Run I and Run IIa. This is again a reflection of the change in beam conditions, most likely due to the target replacement. The change in energy spectra of the selected Near Detector events emphasises the need to treat the Run I and Run IIa datasets separately in the analysis.

Aside from the changes between Run I and Run IIa, the examined Near Detector

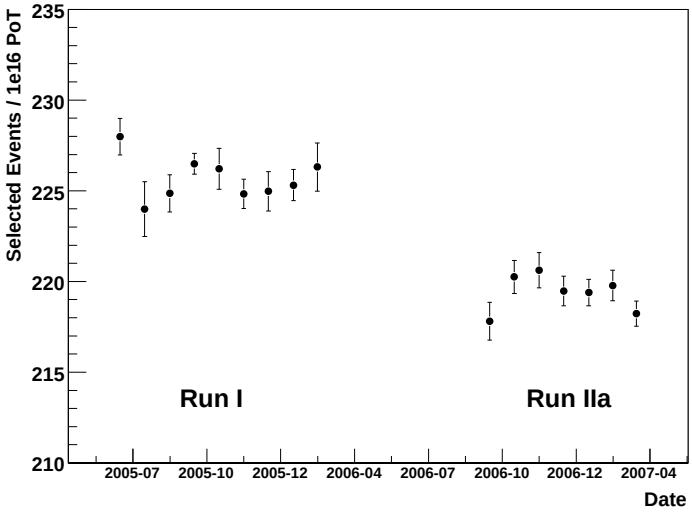

Figure 6.14: The number of selected Near Detector events per 1×10^{16} PoT, for each month in Run I and Run IIa.

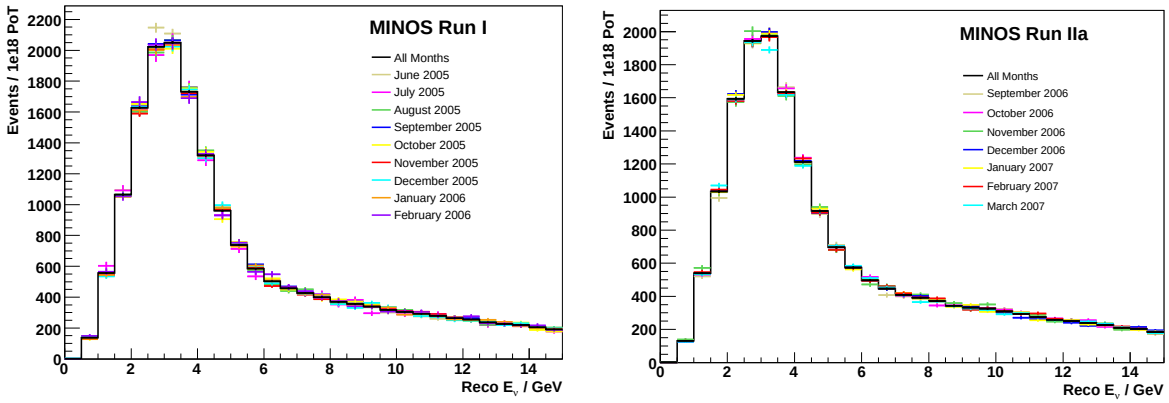

(a) (b)

Figure 6.15: The Near Detector neutrino energy spectrum observed each month during (*a*) Run I and (*b*) Run IIa.

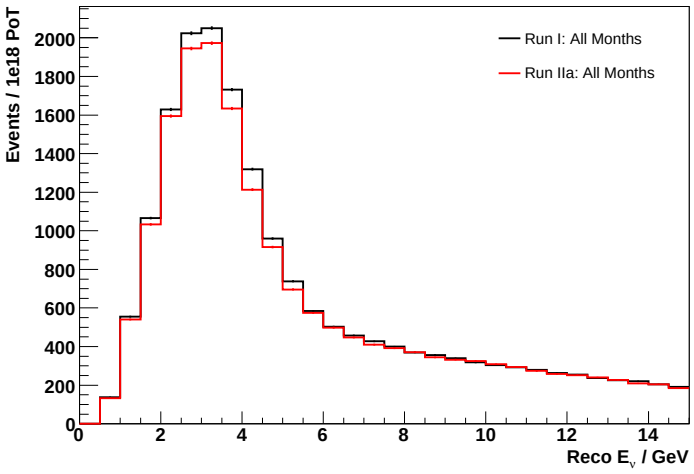

Figure 6.16: The difference between the mean Run I and Run IIa Near Detector neutrino energy spectra.

data distributions are as expected. There is no cause for concern about the stability of the beam or the consistency of the recorded data.

6.6 Neutral Current Background Systematic

The method for removing NC events from the CC sample has been developed and tested using the MC simulation. It is important to remember that, whilst MC tests indicate high CC selection efficiencies and small NC backgrounds, the nominal MC predictions of the NC background will not perfectly represent the data. This is due to the following factors:

- There are large uncertainties in the NC cross-sections, as reviewed in [88].
- The kinematic properties of NC interactions are not well known.
- There are significant hadronisation uncertainties, affecting the modelling of which hadrons are produced for the given interaction kinematics.
- There are large uncertainties concerning intra-nuclear effects. These uncertainties affect the modelling of which of the hadrons produced actually make it to the detector and how much energy is lost in traversing the nucleus.
- There are uncertainties associated with modelling the interactions of the hadrons in the detector.

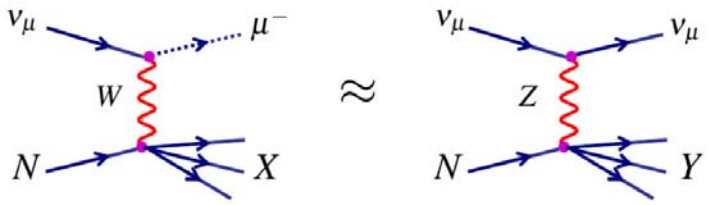

Figure 6.17: Muon removed CC events can be considered as 'fake' NC-like events, although kinematic differences mean the agreement will not be perfect. Taken from [90].

- The selected NC events represent only a small fraction of all NC events, from the tails of the distributions.

With these factors in mind, a systematic uncertainty of 50% in the NC background is suggested in [53]. It is therefore desirable to use the Near Detector data to correct the NC component of the MC. This Section describes techniques for obtaining just such a data-driven correction.

Firstly, 'fake' NC events are generated using the showers from CC data and CC MC events, enabling the different selection efficiencies for data and MC NC events to be examined. Secondly, the cross-section uncertainties are addressed by scaling the NC events in the MC in order to achieve the best agreement between the data and MC PID distributions.

6.6.1 Muon Removal Technique

An established technique for correcting the NC component of the MC involves removing the reconstructed muon tracks from selected CC events in data and MC. After muon removal, the remaining CC showers are passed back to the reconstruction software and new events are created. These can be considered as 'fake' NC events. Comparison of fake NC events in data and MC provides an insight into the required MC NC corrections. This technique has been used successfully in past MINOS analyses [53, 89].

The fake NC events will not perfectly resemble NC events, as there are kinematic differences in the production of NC and CC showers, illustrated in Figure 6.17. A comparison between MC true NC events and MC CC showers must therefore be performed to assess the level of agreement. However, even with some disagreement, a comparison of Muon Removed (MR) events in data and MC should provide a useful first order correction for MC NC events.

The track removal procedure requires considerable care. If the original muon strips are not removed effectively then the second reconstruction pass will tend to reconstruct a track along or close to the original muon path. However, the track removal cannot be too aggressive, as the core of the shower must be retained if the fake NC events are to resemble true NC events.

The algorithm for muon track removal is summarised in [91]. It simply examines the strips in an event in order to form a new Muon Removed DigitList. Its first action is to identify the 'maximum shower plane', which is defined as being six track-like planes from the track vertex. The parent digits of all strips beyond this plane are rejected and not included in the new DigitList. The remaining strips in the event are then examined. If a strip is included in a track, the track angle (θ_z, measured with respect to the z axis) and the deposited charge (Q, measured in MIP, as defined in Section 3.3) are used to judge whether the track is genuinely from a muon. There are then three different actions that the algorithm can take:

1. If the strip is not in a track identified as a muon track, its parent digits are retained for the next reconstruction pass.

2. If the strip is in a muon track and has $Q\cos\theta_z \geq 1.2\,\mathrm{MIP}$, its parent digits are included in the new DigitList, but their charges are scaled by $(1-1\,\mathrm{MIP}/Q\cos\theta_z)$. This reflects the expected muon charge deposition per plane of $Q \simeq 1\,\mathrm{MIP}/\cos\theta_z$.

3. If the strip is in a muon track and has $Q\cos\theta_z < 1.2\,\mathrm{MIP}$ then its parent digits are rejected and not included in the new DigitList.

The new DigitList is passed to the Strip making software and the other reconstruction packages, which should treat the CC shower in a very similar manner to a true NC event. The results of this procedure are illustrated by example in Figure 6.18.

The power of the muon removal technique is that it enables the different behaviour of fake NC events (and so hopefully real NC events) in the data and the MC to be characterised. For instance, the difference in selection efficiencies can be observed by taking a fixed number of MR data events and the same number of MR MC events. CC selection cuts (excluding PID) can be applied and the data/MC ratio plotted as a function of PID. This ratio then provides a correction to both the shape and the normalisation of NC events as a function of their PID value. The correction can be calculated for different ranges of reconstructed E_ν and applied to the MC NC events. The correction factor can also be analysed in the PID range above the PID cut, yielding information about the NC background in the CC selection.

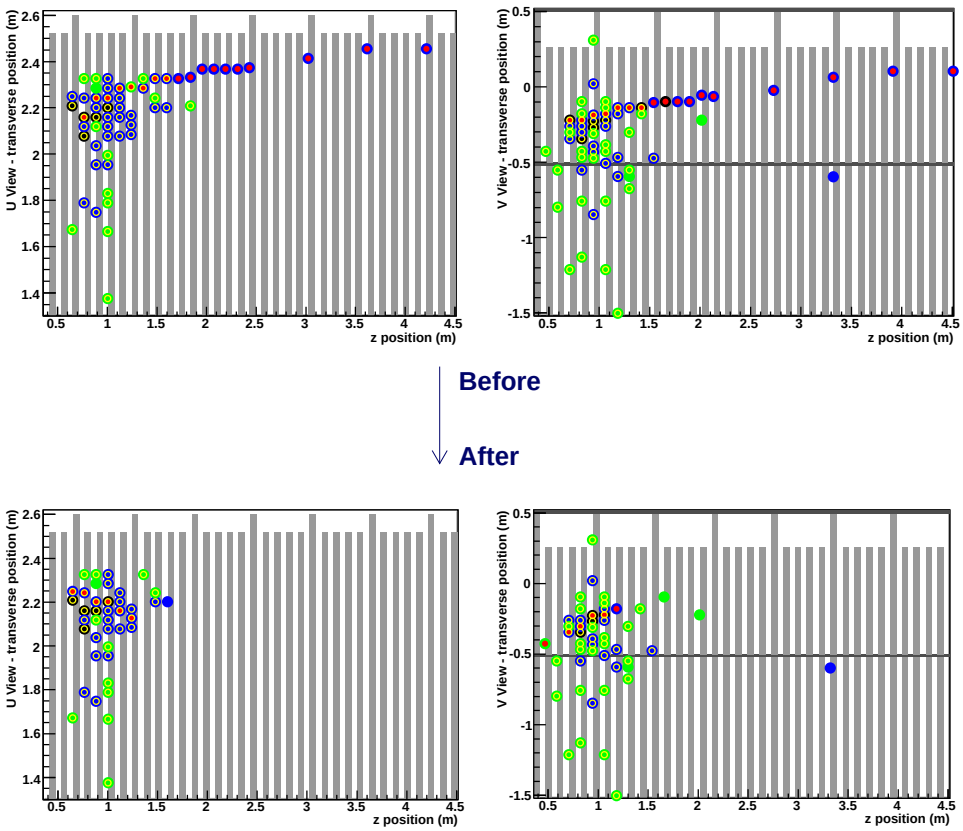

Figure 6.18: The removal of the reconstructed muon track from a typical Near Detector CC event, producing a NC-like event. The coloured markers represent strips in the slice, with red markers indicating strips in a reconstructed track.

6.6.2 Muon Removal Results

The data and MC used in this study are summarised in Table 6.3. In order to improve agreement between data and MC, the beam weighting scheme described in [92] is applied to the MC. In the oscillation analysis, a similar beam weighting technique (described in detail in Chapter 7) will be applied to the MC, so it is NC background discrepancies between the data and the reweighted MC that need to be investigated.

Description	PoT / $\times 10^{19}$	Comments
ND MC	2.91	
ND Data	5.07	Data from Jun, Oct, Nov, Dec 2005
MR ND MC	3.37	
MR ND Data	1.67	Data from Jun, Oct, Nov, Dec 2005

Table 6.3: Summary of the Near Detector data and MC used in the muon removal study.

An initial comparison of the data and MC PID distributions is shown in Figure 6.19. Poor agreement is observed between the data and MC, particularly in the low PID region containing most of the true NC events.

The first step in the muon removal study is to examine the level of agreement between MC true NC events and MC MR CC events (i.e. 'fake' NC events). Figure 6.20 shows the level of agreement for the PID and its input variables. The distributions are normalised to unit area to highlight the level of agreement in the shape of the distributions.

It can be seen that the agreement between the true and fake NC events is not good, with the PID distribution shifted towards lower values for true NC events. However, as already mentioned, differences between NC events and MR CC events are expected on kinematic grounds, as well as from pathologies in the muon removal and reconstruction software. The level of disagreement observed does not prevent the MR events being used to obtain a first order MC correction.

The next important comparison is between the MR events in MC and those in data. The PID distributions (scaled to the PoT exposure of the MR data) and the distributions of the PID input variables (normalised to unit area) are shown in Figure 6.21. For all the distributions, good agreement is observed between MR MC and MR data events. The events reconstructed in data CC showers clearly resemble the events reconstructed in MC CC showers, although there are differences when compared to the events reconstructed in MC NC showers.

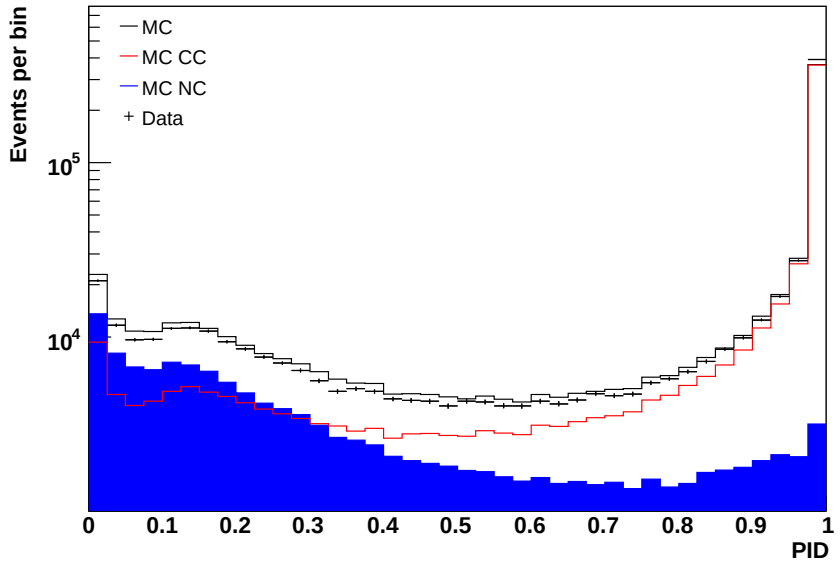

(a)

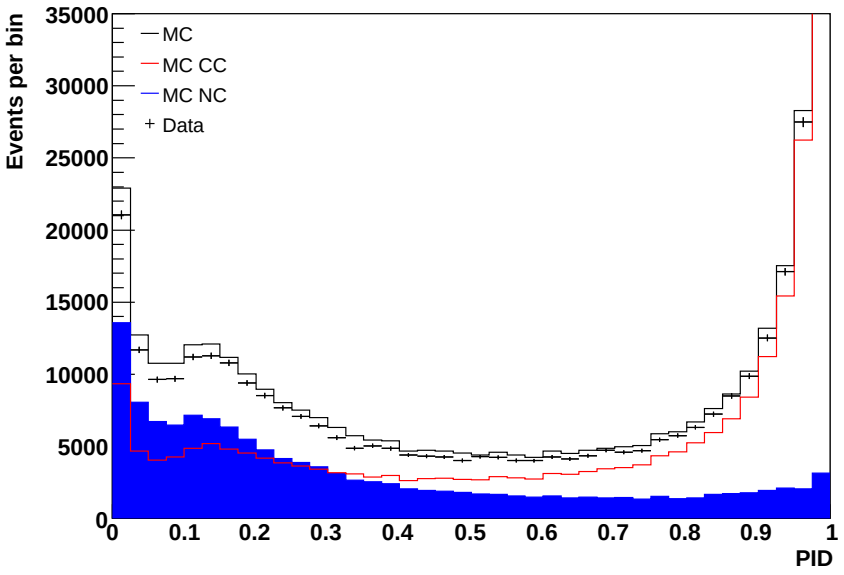

(b)

Figure 6.19: Comparison of the PID distributions for Near Detector data and MC. The MC distribution is divided into CC and NC contributions. The errors shown for the data distribution incorporate the statistical errors in both the data and MC samples. (*a*) With log scale. (*b*) Without log scale.

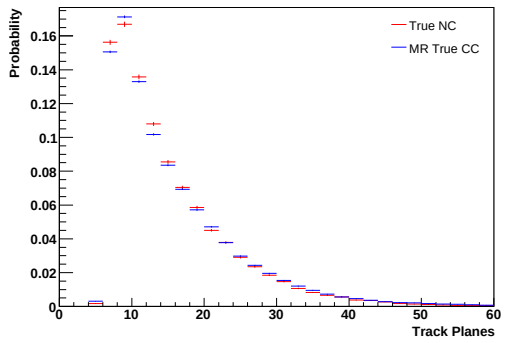

(a)

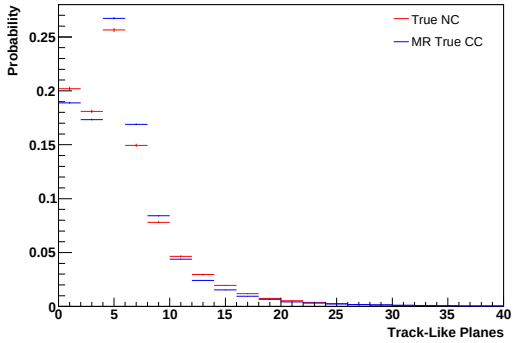

(b)

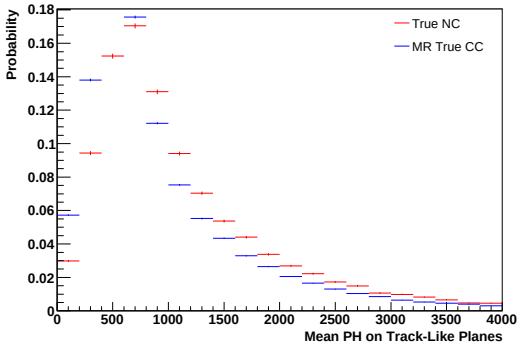

(c)

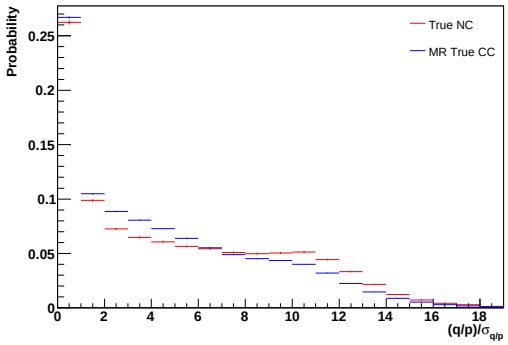

(d)

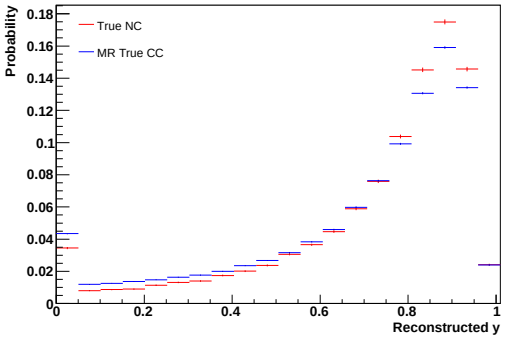

(e)

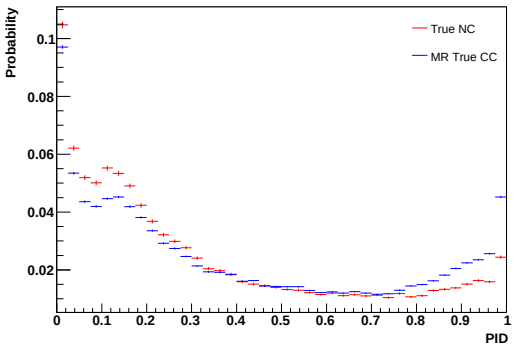

(f)

Figure 6.20: Comparison of the PID distribution, and the distributions of PID input variables, for MC muon removed true CC events (blue) and MC true NC events (red). (*a*) Track planes. (*b*) Track-like planes. (*c*) Mean pulse height on track-like planes. (*d*) Goodness of muon track fit. (*e*) Reconstructed y. (*f*) PID distribution. All distributions are normalised to unit area.

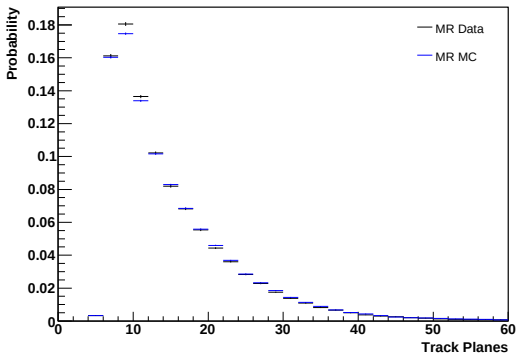

(a)

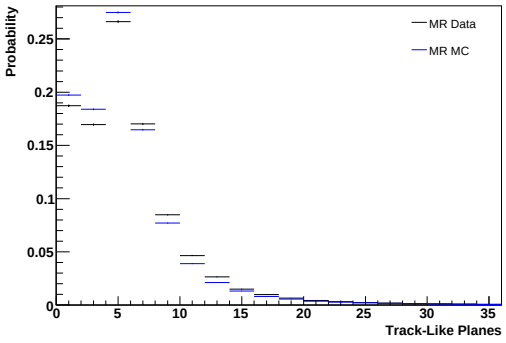

(b)

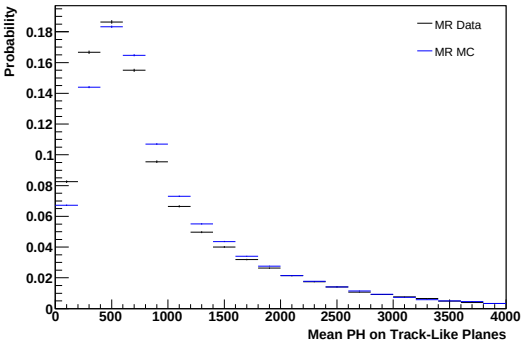

(c)

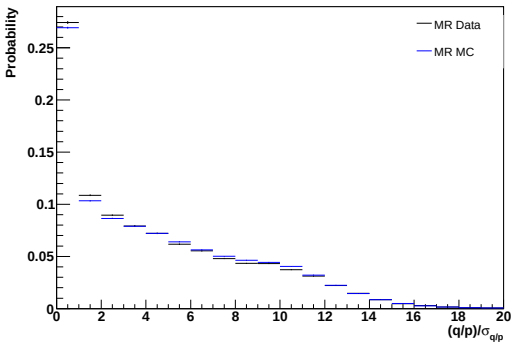

(d)

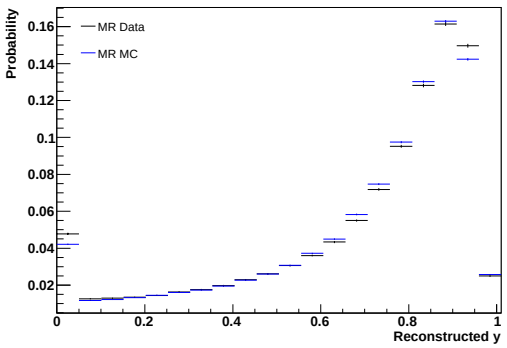

(e)

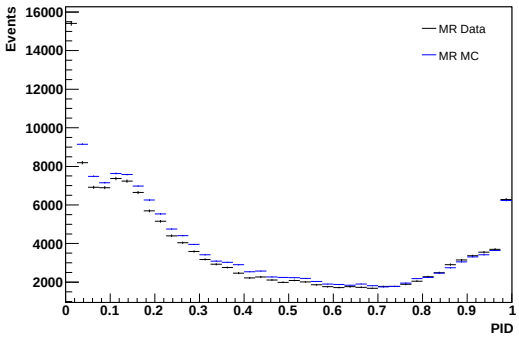

(f)

Figure 6.21: Comparison of the PID distribution, and the distributions of PID input variables, for muon removed events in MC (blue) and data (black). (*a*) Track planes. (*b*) Track-like planes. (*c*) Mean pulse height on track-like planes. (*d*) Goodness of muon track fit. (*e*) Reconstructed y. (*f*) PID distribution. The PID distribution is scaled to the PoT exposure of the muon removed data, whilst the remaining distributions are normalised to unit area.

To correct the MC for differences in the selection efficiency between data and MC NC events, CC selection cuts (excluding PID) are applied to a large sample of MR data events and to an identical number of MR MC events. Differences between MR data events and MR MC events (which should represent the differences between data and MC NC events) mean that a different number of MR data and MR MC events are selected. This results in different PID distributions for selected MR data events and selected MR MC events, as shown in Figure 6.22(a).

The ratio of selected MR data events to selected MR MC events is shown as a function of PID in Figure 6.22(b). This ratio represents the selection efficiency correction required for MC NC events and it shows no strong PID dependence. Corrections can be obtained for different ranges of reconstructed E_ν and constants can be fitted to the PID range representing the CC selected region. These constants help to quantify the scaling of MC NC background events required to account for the different selection efficiencies in data and MC. The values obtained are listed in Table 6.4.

Reco E_ν / GeV	**NC MC Correction (Efficiency)**
$0-1$	$-11.4\% \pm 6.3\%$
$1-2$	$-4.6\% \pm 2.4\%$
$2-4$	$-9.0\% \pm 1.5\%$
$4-6$	$-15.2\% \pm 1.8\%$
$6-12$	$-18.8\% \pm 1.5\%$
$12-100$	$-15.1\% \pm 2.1\%$

Table 6.4: Normalisation corrections for the MC NC background, addressing the different selection efficiencies of NC events in data and MC. The quoted errors are purely statistical.

By using the correction ratio from the relevant bin of reconstructed E_ν, NC events can be assigned weights as a function of their PID and reconstructed E_ν values. Application of these weights improves the data-MC PID agreement across the entire PID range. Figure 6.23(a) reproduces the original data-MC PID distributions (shown in Figure 6.19), whilst Figure 6.23(b) displays the improved agreement obtained when the NC events in the MC are weighted using the selection efficiency correction factors. Particularly large improvements are observed in the NC-rich region at low PID.

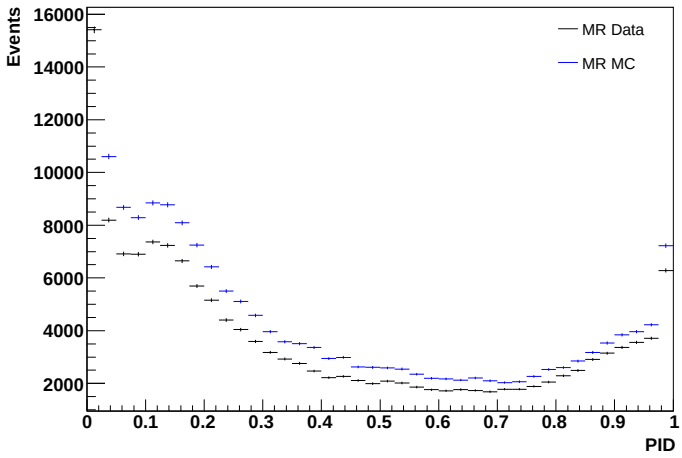

(a)

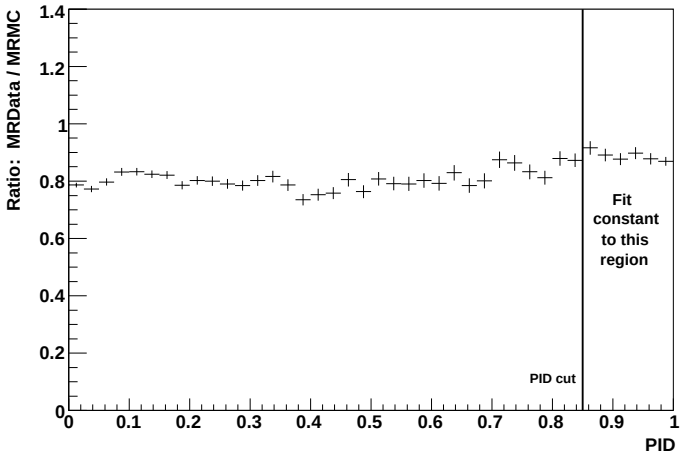

(b)

Figure 6.22: (*a*) PID distributions obtained after applying CC selection cuts to a large sample of muon removed MC events and an identical number of muon removed data events, (*b*) ratio of selected muon removed data events to selected muon removed MC events as a function of PID. The vertical line indicates the PID cut.

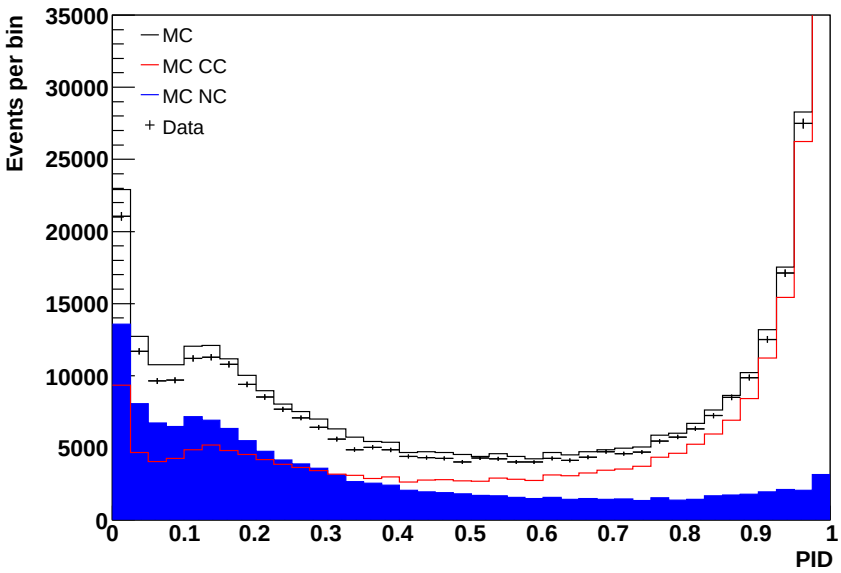

(a)

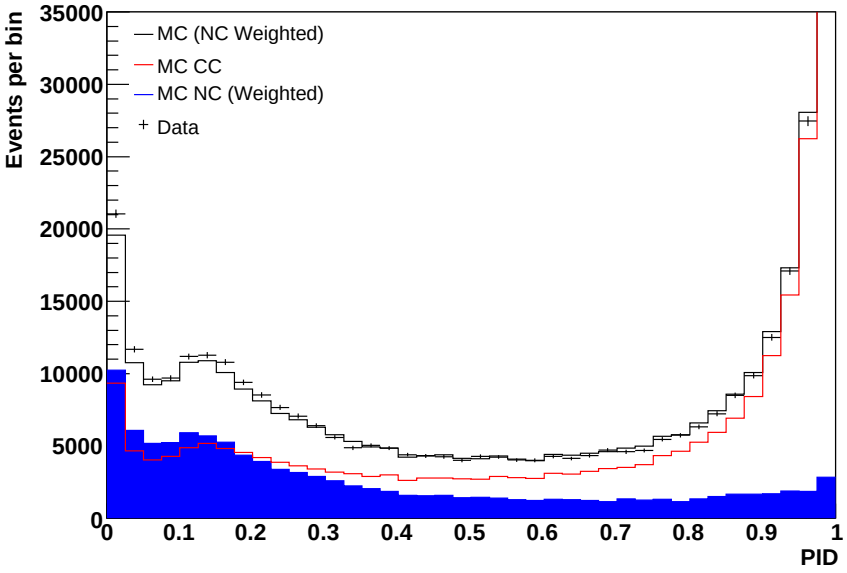

(b)

Figure 6.23: Comparison of the PID distributions for Near Detector data and MC. The error bars displayed for the data distributions incorporate the statistical errors in both the data and MC samples. (*a*) Before and, (*b*), after correction of the MC NC events to account for the different selection efficiencies of NC events in data and MC.

6.6.3 Final Results

After applying the efficiency corrections, the uncertainties in the NC cross-sections can be addressed. The approach is to correct the overall MC NC normalisation such that the agreement between the data and MC PID distributions is optimised. This involves scaling the normalisation of the MC NC PID distribution in such a way that a χ^2 value (comparing the overall agreement between the data and MC PID distributions) is minimised.

Figure 6.24 compares the data PID distribution with a final MC distribution, obtained with the inclusion of both the efficiency and normalisation corrections. The final agreement between data and MC is a significant improvement over that obtained using the nominal MC.

The efficiency and normalisation corrections have been combined together in Table 6.5 to give final values for data-driven NC corrections. The errors quoted are purely statistical and result from combining the efficiency and normalisation scale errors in quadrature. The overall value and statistical error for the entire energy range is $(-6.3 \pm 1.3)\%$.

Reco E_ν / GeV	**NC MC Correction (Eff & Norm)**
$0-1$	$-11.1\% \pm 6.9\%$
$1-2$	$+2.6\% \pm 3.1\%$
$2-4$	$+0.4\% \pm 2.4\%$
$4-6$	$-17.6\% \pm 3.3\%$
$6-12$	$-8.3\% \pm 2.9\%$
$12-100$	$-14.6\% \pm 3.4\%$

Table 6.5: Final corrections for the MC NC background. These corrections account for both MC normalisation uncertainties and the different selection efficiencies of NC events in data and MC. The quoted errors are purely statistical.

It is desirable to use the results of this study to obtain an estimate of the systematic uncertainty in the NC background. No precise estimate can be provided, but the efficiency and normalisation corrections are uncorrelated and roughly indicate the typical levels of agreement between data and MC. The efficiency corrections are typically less than 20%, whilst the normalisation corrections are typically less than 10%. Adding these typical efficiency and normalisation corrections in quadrature suggests an estimated systematic uncertainty of 25%.

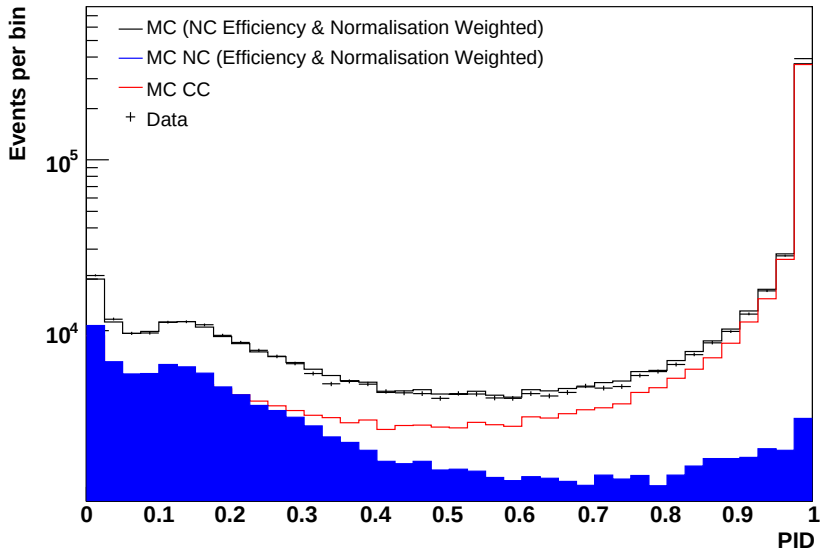

(a)

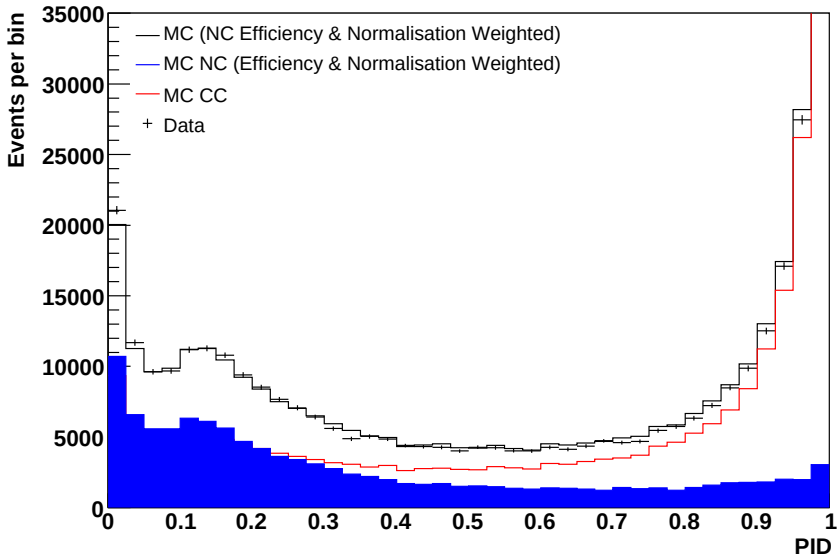

(b)

Figure 6.24: Comparison of the PID distributions for Near Detector data and MC, after correction of MC NC events to account for MC normalisation uncertainties and the different selection efficiencies of NC events in data and MC. The error bars displayed for the data distributions incorporate the statistical errors in both the data and MC samples. (*a*) With log scale. (*b*) Without log scale.

Chapter 7

Extrapolation and Fitting Far Detector Data

Of fundamental importance to a long baseline neutrino oscillation analysis is the ability to accurately predict the Far Detector reconstructed neutrino energy distribution. In an analysis, predictions are made for different oscillation parameters and compared with the observed data distribution. It is through these comparisons that sensitivity to the oscillation parameters is achieved.

In principle, the Far Detector predictions can be made using knowledge of the neutrino beam and the neutrino interaction cross-sections. However, absolute predictions of the neutrino energy spectrum are very difficult. Understanding of the neutrino beam is a major source of systematic error[93] and there are also sizeable uncertainties in the cross-sections[94]. For this reason, data collected at the Near Detector are used to establish the beam and cross-section characteristics. Once the response of the Near Detector to the neutrino beam is understood, accurate Far Detector predictions can be made by simply extrapolating this understanding to the Far Detector.

In this Chapter, a new extrapolation method is developed and described in detail. The most important systematic uncertainties in an oscillation analysis are also identified and an oscillation fit is developed and tested. This fit extracts a measurement of the oscillation parameters from the data, whilst fully accounting for the effects of the systematic uncertainties.

7.1 The Need for an Extrapolation

A survey of neutrino cross-section measurements at $\mathcal{O}(1\,\mathrm{GeV})$[88] reveals that most measurements have large uncertainties. These are due to low statistics and inconsistencies between the results of different experiments. There are also errors associated with the conversion of data from nuclear targets to free nucleon cross-sections. Even for quasi-elastic CC processes, with relatively high statistics, the uncertainties are typically 20%. This makes it difficult to produce direct Far Detector predictions. Poor understanding of the neutrino beam then introduces further complications. A discussion of beam uncertainties is included in [93], and the most important issues are summarised below.

Hadron production

At small angles to the beam, the main contribution to the neutrino flux is from the decay of pions produced in the target. There are large uncertainties in the energy distributions of these pions and further uncertainties are introduced by the non-zero target length, which allows the pions to be absorbed or to interact again. These uncertainties affect the prediction of the absolute neutrino flux by as much as 20%. Poor knowledge of the pion p_t distribution also affects understanding of the divergence of the hadron beam and the distribution of pion decay points in the decay pipe. This introduces uncertainties in both the shape and normalisation of the Far Detector prediction.

Finite target length

The target length is not negligible when compared to the distance between the target and the first focussing horn. The horn has a finite acceptance, so this means that hadrons focussed from the beginning of the target will have higher momentum than those focussed from the end of the target. The hadron spectrum is therefore shifted to lower energies along the target length. Proton attenuation in the target also introduces uncertainties by decreasing the number of hadrons produced along the target. Finally, the fraction of hadrons passing through the opening in the horn will increase along the target, raising the contribution to the high energy tail.

Beam divergence

The focussing horns operate by providing a p_t 'kick' to hadrons, proportional to the radius at which they cross the horn. However, the finite target length means that hadrons are produced both upstream and downstream of the focal point of the first horn.

This causes divergence of the hadron beam, as upstream particles are over-focussed and downstream particles are under-focussed. Importantly, the combination of this effect with the shifting of hadron momentum along the target, strongly correlates the energy and the divergence of the hadron beam. Other factors affecting the divergence of the hadron beam include the finite width of the initial proton beam and the fraction of hadrons that pass through the openings in the horns and so enter the decay pipe without focussing.

Hadron decays

The production of the beam neutrinos typically proceeds via the decay $\pi^+ \rightarrow \mu^+ + \nu_\mu$. For a given pion decay, the energy of the neutrino depends on the decay angle, θ, and the Lorentz boost of the pion, γ, as shown in Equation 7.1. The resulting neutrino flux through an area A a distance z from the decay point is given by Equation 7.2.

$$E_\nu = \frac{0.43\, E_\pi}{1 + \gamma^2\theta^2} \tag{7.1}$$

$$\text{Flux} = \left(\frac{2\gamma}{1 + \gamma^2\theta^2}\right)^2 \frac{A}{4\pi z^2} \tag{7.2}$$

The distribution of hadron decays throughout the decay pipe has a large impact on the neutrino energies and fluxes at the Near and Far Detectors and is a large source of systematic error:

- The divergence of the hadron beam and the finite size of the decay volume means that most hadrons will hit the walls of the decay pipe, where they will quickly interact or decay. Due to the correlation between hadron momentum and divergence, this leads to a relationship between the hadron energy (and hence neutrino energy) and the decay position.
- Unlike the Far Detector, the solid angle subtended by the Near Detector varies considerably along the decay pipe. This means that neutrinos produced towards the end of the decay pipe represent a larger fraction of the neutrino flux at the Near Detector than at the Far Detector. The relationship between hadron energy and decay position then introduces differences between the energy distributions at the two detectors.

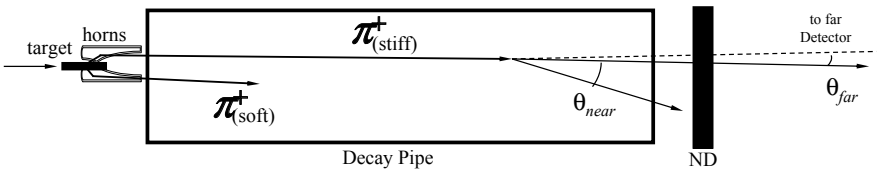

Figure 7.1: The different decay angles required to produce neutrinos that interact at the Near and Far Detectors. Taken from [53].

- Neutrinos produced in decays at large radii must have different decay angles in order to reach the Near and Far Detectors, as shown in Figure 7.1. The neutrino energy varies with the decay angle, so the energy of the neutrinos observed at the two detectors will be different, with the Far Detector energy typically being higher.

It is only through examination of Near Detector data that the distribution of hadron decays in the decay pipe (and the other uncertainties described in this Section) can be understood. To generate accurate predictions of Far Detector distributions, this understanding must then be extrapolated from the Near Detector to the Far Detector.

7.2 How the Extrapolation Works

The principle behind the extrapolation method is to relate the neutrinos that interact in the Near and Far Detectors via their parent pions and kaons. These pions and kaons are the 'beam particles', which decay to produce the neutrino beam. The extrapolation considers a large MC sample of beam particles. For each beam particle, details are required of a Near Detector event *and* a Far Detector event that could occur after the beam particle decays to produce a neutrino.

To address the uncertainties in the beam and cross-sections, the beam particles are assigned weights so that the distributions of their associated Near Detector MC events closely match the distributions of Near Detector data events. Accurate Far Detector predictions are then possible by simply considering the distributions of the Far Detector MC events associated with the weighted beam particles. The steps in this extrapolation method are summarised in Figure 7.2 and are described in detail in this Section.

7.2.1 Reweighting the Near Detector MC

The process of predicting the Far Detector energy distribution begins with large libraries of selected Near Detector data and MC events. The MC events form the basis of the

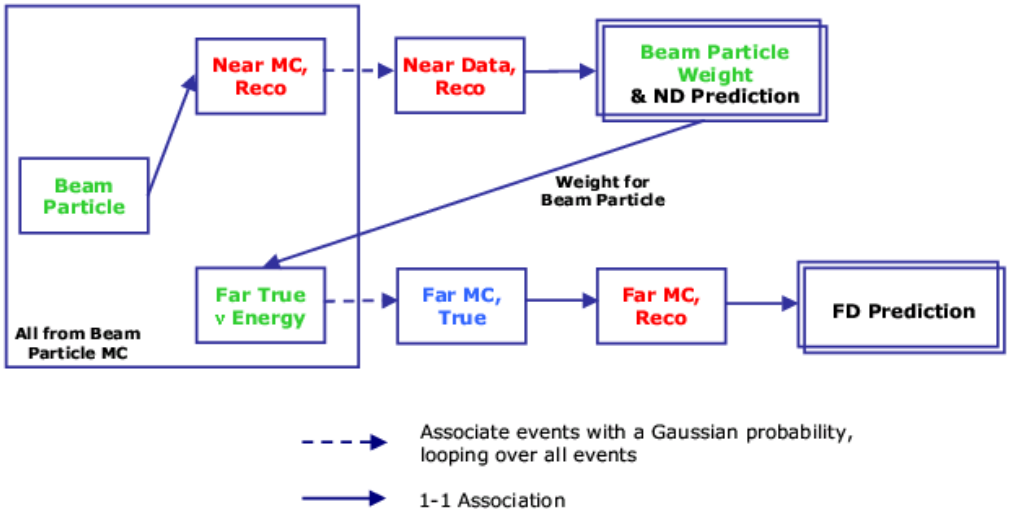

Figure 7.2: Summary flow diagram for the extrapolation method.

'Beam Particle MC' used by the extrapolation; each entry contains full details of a reconstructed Near Detector interaction, together with some details of an associated Far Detector interaction (the interaction that would occur if the same parent pion/kaon decayed to produce a neutrino that interacted in the Far Detector). These details include the energy of the Far Detector neutrino interaction and other kinematic properties.

Weights are calculated for each beam particle, to account for differences between the real neutrino beam and cross-sections and those in the MC simulation. To calculate these weights, several 'matching properties' are selected; physical properties that characterise an event. The weights are then constructed so that the distributions of these properties for the beam particle Near Detector events match those for the Near Detector data.

The choice of matching properties is very important; in the event of a disagreement between the data and the nominal MC, the extrapolation method needs to detect the discrepancy in the distributions of the matching properties and correct the beam particle weights accordingly. It is therefore desirable to use matching properties that help distinguish between different classes of event (e.g. quasi-elastic interactions and NC background events).

Unfortunately, the required MC statistics increase rapidly with the number of matching properties. Two properties is the maximum number that can realistically be used without encountering problems due to low statistics in a bin of matching properties. The matching properties chosen for use in the extrapolation are E_ν and y (these are reconstructed properties, allowing comparison with the real Near Detector data). These properties help to associate beam particle Near Detector MC events and data events that

have similar energies and similar divisions of energy between the track and the shower. Different interaction types typically populate different areas of the reconstructed $E_\nu - y$ distribution, so this choice of properties should help the extrapolation to respond to cross-section discrepancies between data and MC.

If a beam particle's Near Detector event has a certain set of reconstructed properties, E^R_{MC} and y^R_{MC}, the desired beam particle tuning weight is:

$$W(E^R_{MC}; y^R_{MC}) \equiv \frac{\text{Number of ND data events with these properties}}{\text{Number of beam particle ND events with these properties}} \tag{7.3}$$

$$= \frac{D_{Data}(E^R_{MC}; y^R_{MC})}{D_{MC}(E^R_{MC}; y^R_{MC})} \tag{7.4}$$

This weight could be obtained by creating histograms of the matching properties and then examining the ratios of bin contents. However, the Near Detector data can be sampled without smearing the distributions by performing an integral/sum over the properties. Each beam particle Near Detector event is compared to all the events in the data sample. In each comparison, narrow normalised Gaussian distributions (approximating to δ-functions) are used to assign weights as a function of the differences between the matching properties.

Calculating weights in this way, the beam particle Near Detector events sample the data distributions of the matching properties, providing the numerator in Equation 7.3. This sampling is discussed further in Section 7.2.2. The denominator in Equation 7.3 is provided by ensuring that the MC has a flat $E_\nu - y$ distribution, as discussed in Section 7.2.3. This calculation of the beam particle weight is summarised in Figure 7.3.

7.2.2 Gaussian Matching of Events

Using the reconstructed neutrino energy and reconstructed y as matching properties, the numerator for the beam particle weight in Equation 7.3 can be calculated as:

$$W_{Gauss}(E^R_{MC}; y^R_{MC}) = \sum_{\text{Data Events}} \left(\frac{1}{\sqrt{2\pi\sigma^2_E}} e^{-(E^R_{Data} - E^R_{MC})^2/2\sigma^2_E} \times \frac{1}{\sqrt{2\pi\sigma^2_y}} e^{-(y^R_{Data} - y^R_{MC})^2/2\sigma^2_y} \right) \tag{7.5}$$

For each pairing of beam particle Near Detector event and data event, the contribution to this weight is the product of a normalised Gaussian distribution for each matching property. The Gaussian distributions have specified narrow widths, σ_E and σ_y, which are

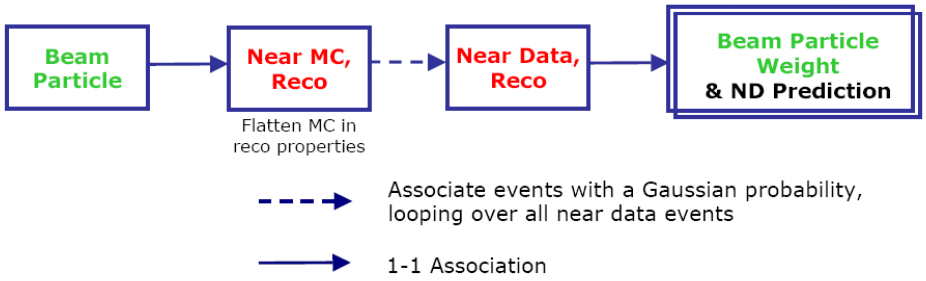

Figure 7.3: The steps required to obtain the beam particle weights.

constant below 20 GeV, but then increase smoothly with energy to combat the reduced MC statistics in the high energy tail. The value returned by each Gaussian distribution represents how closely the matching property agrees between the data event and beam particle MC event.

By summing over the data events, calculating weights of this form, the beam particle Near Detector events actually sample the Near Detector data distributions at the given matching properties. To demonstrate this, it is helpful to consider the limit of infinitely narrow Gaussian distributions (δ-functions) and infinitely large data and MC samples. For simplicity, the neutrino energy is considered as the only matching property in this demonstration. The beam particle Near Detector reconstructed energy distribution can be defined as:

$$D_{MC}(E^R_{MC}) = \sum_{\text{MC Events}} p(E^R_{MC}; E^T_{MC}) \times \epsilon(E^T_{MC}) \tag{7.6}$$

where $p(E^R_{MC}; E^T_{MC})$ is the probability of obtaining a reconstructed energy value, E^R_{MC}, given the true energy, E^T_{MC}, and ϵ is the selection efficiency, which is a function of the true energy. A Near Detector data prediction is obtained by considering the Near Detector energy distribution of the reweighted beam particles, as follows:

$$D_{Pred}(E^R_{MC}) = \sum_{\text{MC Events}} \left(p(E^R_{MC}; E^T_{MC}) \times \epsilon(E^T_{MC}) \times F(E^R_{MC}) \times \sum_{\text{Data Events}} \delta(E^R_{Data} - E^R_{MC}) \right) \tag{7.7}$$

$$= \sum_{\text{MC Events}} \left(p(E^R_{MC}; E^T_{MC}) \times \epsilon(E^T_{MC}) \times F(E^R_{MC}) \times D_{Data}(E^R_{MC}) \right) \tag{7.8}$$

$$= \sum_{\text{MC Events}} \left(p(E^R_{MC}; E^T_{MC}) \times \epsilon(E^T_{MC}) \times W(E^R_{MC}) \right) \tag{7.9}$$

$$= D_{Data}(E^R_{MC}) \tag{7.10}$$

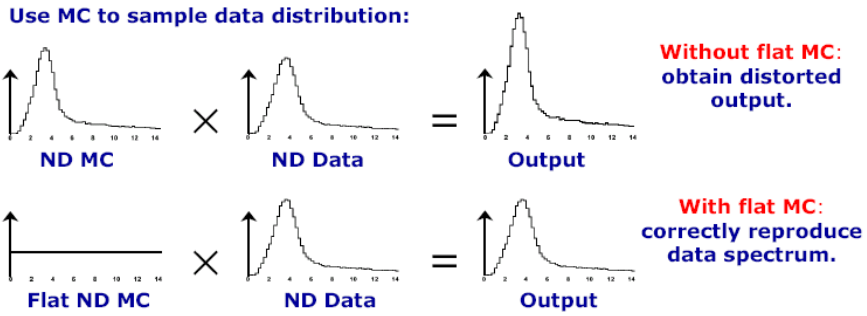

Figure 7.4: The shape of the Near Detector MC distribution must not influence the sampling of the Near Detector data, or the final predictions will be distorted.

In the sum over the data events, the δ-functions simply return the value of the data distribution at the given matching properties. This provides the numerator in Equation 7.3. The denominator in this Equation is then provided by the flattening function, F, which accounts for the number of beam particle Near Detector events with the same properties. The flattening function is discussed further in Section 7.2.3.

This exercise provides a proof of concept for the idea of matching events using Gaussian weights. In the limit of infinitely narrow Gaussian distributions and infinite statistics, the reweighted beam particle Near Detector events will exactly reproduce the data distributions for the chosen matching properties.

7.2.3 Flattening the MC

The Near Detector events are matched in reconstructed E_ν and y, so the beam particle Near Detector events must have flat distributions in these quantities in order to uniformly sample the data distribution. The effect of sampling without flattening would be to distort the final Far Detector and Near Detector predictions by the product of the Near Detector data and beam particle Near Detector distributions, as shown in Figure 7.4.

The flattening weights can be obtained by applying the same Gaussian weighting technique to the beam particle Near Detector events. However, for simplicity, and because sufficient MC statistics are available, finely binned distributions of the matching properties are used instead. The flattening weight for a given set of matching properties is then the bin width, divided by the bin contents. Examples of flattening weights for the beam particle Near Detector events are shown in Figure 7.5. Larger bin sizes are used in the regions of parameter space where the distributions are quite flat, whilst the binning must be fine wherever the distributions vary rapidly.

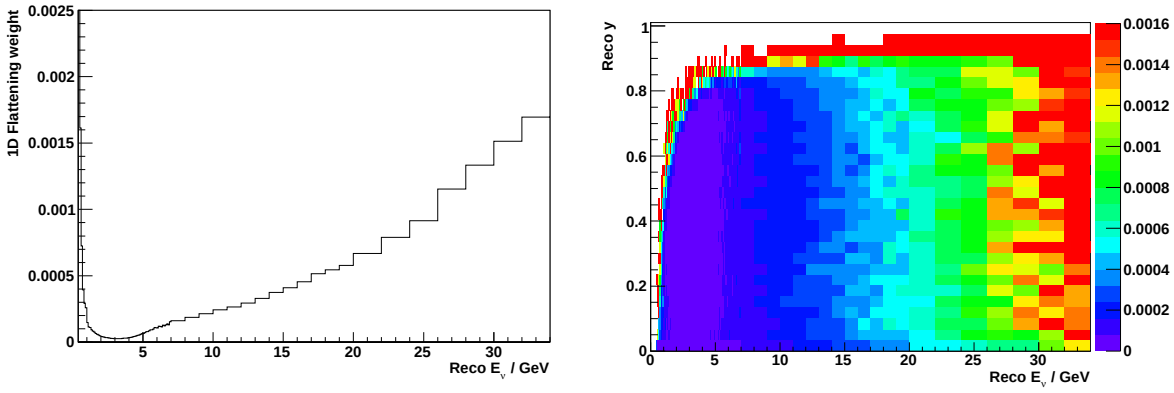

(a) (b)

Figure 7.5: MC flattening weights as a function of the matching properties, (a) using E_ν as the only matching property, (b) using both E_ν and y.

7.2.4 Transport to the Far Detector

After the beam particle weights have been obtained, the weighted beam particle Near Detector distributions should accurately match those of the Near Detector data. The extrapolation to the Far Detector can then be considered. In principle, the Far Detector E_ν prediction can be made by simply considering the reconstructed energy distribution of the weighted beam particle Far Detector events. However, instead of a complete and idealised beam particle MC, only Near Detector MC has been used. Although this contains full details of the reconstructed Near Detector events, it specifies only the true values of E_ν and y for the associated beam particle Far Detector events.

To complete the extrapolation and associate a fully reconstructed Far Detector event with each beam particle, a large library of selected Far Detector MC events is required. The technique of Gaussian matching is used again to assign weights to these Far Detector MC events and identify the events that are most representative of the beam particle Far Detector event. Unlike the previous use of Gaussian matching to reweight the beam particles, this process is only required in order to complete the construction of the beam particle MC.

The only quantities available for this matching are the true values of E_ν and y for the beam particle Far Detector events. These can be compared with the truth quantities in the Far Detector MC library. The truth information also allows true CC and NC background events to be treated separately. In the transport to the Far Detector, beam

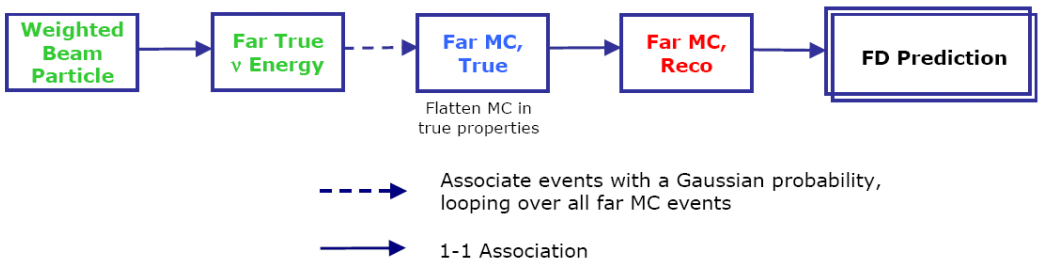

Figure 7.6: The steps required to propagate understanding of the beam from the Near Detector to the Far Detector.

particles producing true CC events in the Near Detector are matched only to true CC events in the Far Detector. True E_ν and true y are used as matching properties, requiring the true CC Far Detector MC to have a flat $E_\nu - y$ distribution (so that the Gaussian weights pick out contributions summing to one event for each beam particle). Beam particles producing true NC events in the Near Detector are only matched to true NC background events in the Far Detector. This matching uses true E_{vis} (true $E_\nu \times y$), which is a more natural quantity for NC events.

Individual Gaussian weights are calculated for every pairing of beam particle Far Detector event and MC event. Each Gaussian weight (divided by the flattening weight) indicates how closely the MC event represents the beam particle Far Detector event. With these weights, the Far Detector MC events can make a contribution to the final Far Detector prediction. Contributions are made for every possible combination of beam particle and Far Detector MC event. The basic form for a contribution is the reconstructed E_ν value for the MC event, weighted by its Far Detector Gaussian weight, the Far Detector flattening weight and the beam particle weight. This transport of the beam particle weights to the Far Detector is summarised in Figure 7.6.

7.2.5 Event Selection and Efficiencies

The extrapolation method uses only events that pass CC selection cuts on reconstructed quantities, as detailed in Chapter 6. This helps ensure that the events matched together are similar and well reconstructed. However, the use of selected events introduces a new complication, due to the different selection efficiencies for the two detectors. The Far Detector prediction, extrapolated from the Near Detector, is influenced by Near Detector selection efficiencies, whilst the Far Detector data is influenced by Far Detector selection efficiencies. This ratio of efficiencies must be addressed.

Look-up-tables are made of the selection efficiency as a function of the true matching properties. In the construction of the Far Detector prediction, every combination of beam particle and Far MC event then receives an extra weight, equal to the selection efficiency of the Far MC event, divided by the selection efficiency of the beam particle Near Detector event. True CC events and NC background events are treated separately. The weights for the CC events are ratios of the selection efficiency of true CC events as a function of true E_ν and true y. The weights for the NC events are ratios of the (much smaller) selection efficiencies of true NC events using the CC selection cuts. The NC efficiency look-up-tables are made as a function of true E_{vis}.

Construction of the efficiency look-up-tables involves using high statistics MC samples to plot the distributions of the true matching properties. Distributions are made using the standard selection cuts on reconstructed quantities and then made again using the equivalent cuts on truth quantities. The ratio of the two distributions (reconstructed selection, divided by true selection) characterises the variations in efficiency. The CC selection efficiencies are shown in Figure 7.7 and it can be seen that the CC selection efficiency for the Near Detector is typically lower than for the Far Detector. Without an efficiency correction, the Far Detector prediction would therefore be too low. There would also be some small distortion of the shape of the Far Detector prediction.

7.2.6 Final Far Detector Prediction

Only a few factors remain in order to obtain the final Far Detector predicted neutrino energy distribution:

- The contribution to the Far Detector prediction from each beam particle must be weighted by the probability of the beam particle decaying to produce a neutrino that passes through the Far Detector, divided by the probability of the beam particle decaying to produce a neutrino that passes through the Near Detector. This weight can be considered as a ratio of the neutrino fluxes.

- The Near and Far Detector interactions associated with each beam particle usually have different energies. The change in cross-section associated with these different energies must therefore be addressed. This produces only a very small correction.

- The final prediction must be scaled by the ratio of the exposures (PoT) for Far Detector data to Near Detector data.

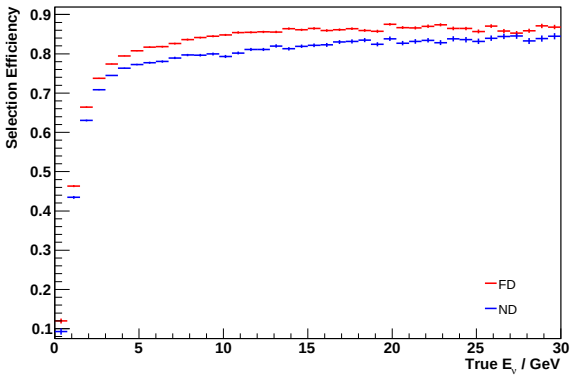

(a)

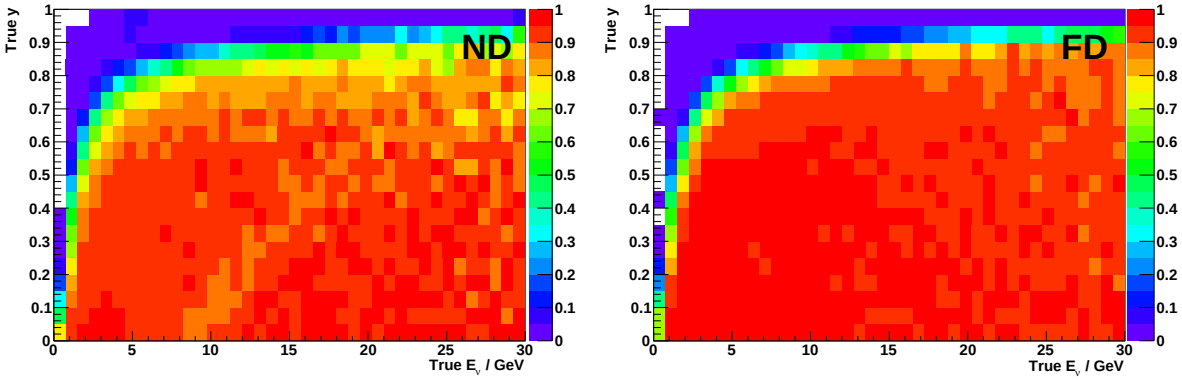

(b)

Figure 7.7: The selection efficiency of true CC events as a function of the true matching properties. Near and Far Detector efficiencies are shown, for use with, (a), E_ν as the only matching property and, (b), both E_ν and y as matching properties.

- The final prediction must also be scaled by the ratio of the Far Detector fiducial mass to the Near Detector fiducial mass.

The final Far Detector prediction is therefore obtained by plotting the reconstructed energy of the far MC events, each with their own individual weights:

$$\begin{aligned}\text{FinalWeight} = &(\text{NearFlatWeight} \times \text{NearGaussWeight}) \\ &\times (\text{FarFlatWeight} \times \text{FarGaussWeight}) \\ &\times (\text{FluxWeight} \times \text{PoTWeight} \times \text{FidVolWeight}) \\ &\times \text{EfficiencyWeight} \times \text{CrossSectionWeight} \end{aligned} \quad (7.11)$$

This prediction is the 'no oscillation' prediction of the Far Detector reconstructed energy distribution and assumes simple neutrino propagation. By including a propagation weight for all true CC ν_μ Far Detector MC events, predictions can be made for different oscillation parameters. The propagation weight is simply the survival probability of Equation 2.14.

Contribution from ν_τ Interactions

One small correction to the Far Detector prediction is to include the effects of ν_τ appearance (from the oscillation of the ν_μ). This is a simple process and requires a MC file containing only Far Detector ν_τ events. These events are weighted by the appearance probability (one minus the ν_μ survival probability) and standard Far Detector selection cuts are applied. The number of selected ν_τ events is expected to be small, about one event for an exposure of 2.5×10^{20} PoT. The spectrum of the contribution is illustrated for typical oscillation parameters in Figure 7.8.

Event Splitting and Reconstruction Thresholds

The final correction to the Far Detector predicted spectrum is required to address event reconstruction problems. In the Near Detector, a single event can be split into multiple reconstructed events, usually because the strips are incorrectly divided into multiple slices. This effect is rare at the Far Detector, but a related effect does occur in both detectors. This is the failure to reconstruct low energy true events if they do not produce enough detector activity. To correct for both effects, the number of true fiducial events in the MC input files (a true record of all events generated) can be compared to the

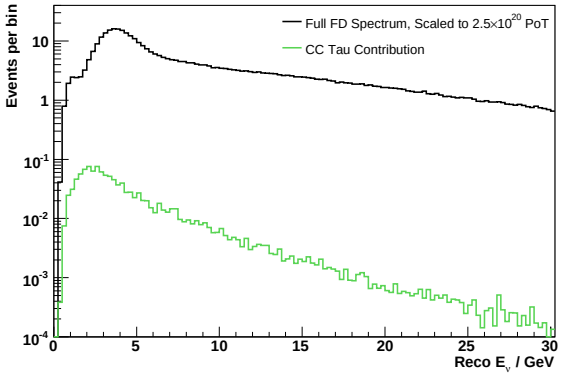

Figure 7.8: The small contribution that CC ν_τ events make to a $2.5{\times}10^{20}$ PoT Far Detector prediction, shown for $\Delta m^2_{32} = 3.0 \times 10^{-3}\,\mathrm{eV}^2$ and $\sin^2(2\theta_{23}) = 0.9$

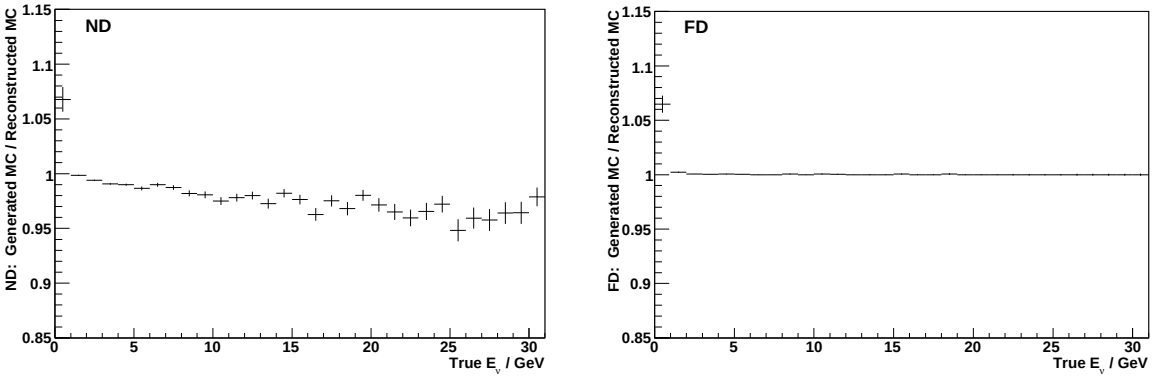

(a) (b)

Figure 7.9: The ratio of generated MC events to reconstructed events as a function of true E_ν in (a) the Near Detector and, (b) the Far Detector. The ratio indicates the effects of reconstruction thresholds and event splitting.

number of true fiducial events in the output of the reconstruction. This ratio has been evaluated as a function of true E_ν for the two detectors and the results are shown in Figure 7.9. The corrections are very small; slightly greater than one at low energies due to threshold effects, and slightly below one at higher energies in the Near Detector, due to event splitting.

7.2.7 Basic Oscillation Fit

The aim of an oscillation fit is to obtain a measurement of the neutrino oscillation parameters by comparing the Far Detector predicted energy distributions with the data. In the fit, each comparison between a predicted distribution and the data is assigned a likelihood value, which expresses how closely the prediction and data agree. Every predicted distribution represents different oscillation parameters and the best fit parameters are those for which the data and prediction agree most closely.

In this analysis, a negative log likelihood value is calculated for each comparison between a predicted distribution and the data. The negative log likelihood values are calculated using the binned likelihood function shown in Equation 7.12.

$$-\ln \mathcal{L} = \sum_{i}^{n_{Ebins}} (N_i^{MC} - N_i^{Data} \ln N_i^{MC}) \tag{7.12}$$

where N_i^{MC} and N_i^{Data} are respectively the predicted content and the actual data content of reconstructed energy bin i and n_{Ebins} is the total number of bins in the reconstructed energy distributions.

Comparison of the negative log likelihood values obtained for different oscillation parameters with that obtained at the best fit point allows confidence intervals to be constructed. The s-standard-deviation confidence interval is described by the values of the oscillation parameters which satisfy the relationship in Equation 7.13.

$$-\ln \mathcal{L}(\Delta m^2_{32}, \sin^2(2\theta_{23})) = -\ln \mathcal{L}_{max} + s^2/2 \tag{7.13}$$

$$\Delta \ln \mathcal{L} = s^2/2 \tag{7.14}$$

Neutrino oscillations are described by two parameters, so there are two degrees of freedom in the fit. The 68% confidence interval is therefore described by the contour for which $\Delta \ln \mathcal{L} = 2.279/2$, whilst the 90% confidence interval is described by the contour for which $\Delta \ln \mathcal{L} = 4.605/2$. The appropriate 1σ confidence intervals for the individual

oscillation parameters are given by the extreme limits of the $\Delta \ln \mathcal{L} = 1/2$ contour when it is projected onto the Δm^2_{32} and $\sin^2(2\theta_{23})$ axes.

Both the best fit oscillation parameters and the confidence level contours can be obtained using the CERN numerical minimisation program MINUIT [95]. MINUIT controls the variation of the oscillation parameters in the fit, so as to efficiently and accurately identify the global likelihood minimum and any specified contours of the likelihood surface.

7.3 Performance Tests

7.3.1 Simple MC Tests

After implementation, the extrapolation needed to be tested and demonstrated to work effectively. The first, simple tests, made use of nominal MC files for 'fake data'. The high statistics MC samples used are summarised in Table 7.1. The Near Detector MC represents the basis of the beam particle MC discussed in Section 7.2. It should be remembered that, from the files, only events passing CC ν_μ selection cuts are used.

Dataset	**PoT Details**		
Near 'Data'	500 files	2.42×10^{13} PoT per snarl	**9.68×10^{18} PoT total**
Near MC	502 files	2.42×10^{13} PoT per snarl	**9.72×10^{18} PoT total**
Far MC	524 files	6.5×10^{20} PoT per file	**3.41×10^{23} PoT total**
Far 'Data'	49 files	6.5×10^{20} PoT per file	**3.19×10^{22} PoT total**

Table 7.1: Summary of the MC used in the extrapolation method and the MC 'fake' datasets used to test the performance of the extrapolation.

In a simple MC test, the best way to evaluate the success of the extrapolation is to compare the Far Detector no oscillation prediction with the Far Detector no oscillation 'data'. If the prediction without oscillations is correct, it is simple to extend the predictions to include the effects of oscillations. A further required check is a comparison of the Near Detector 'data' with the Near Detector prediction, made using the reweighted beam particle Near Detector events.

The first test removed the effects of efficiencies and of smearing of the E_ν and y distributions by using only true quantities and true selections. This test was then the simplest possible proof of principle for the extrapolation. The results are shown in

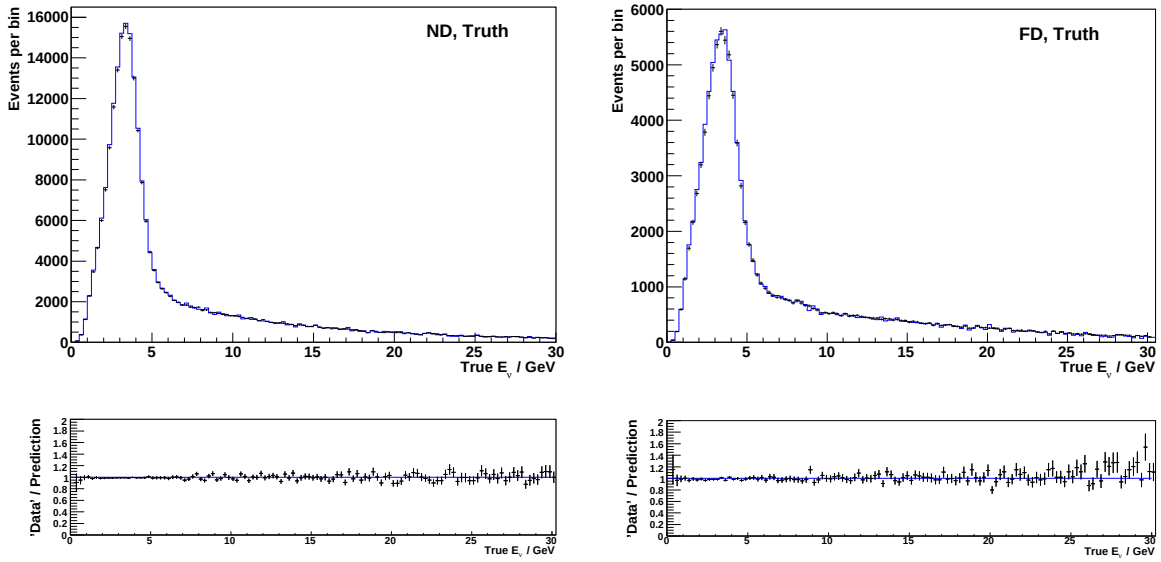

(a) (b)

Figure 7.10: Comparison of fake data (points) with the Near Detector prediction and Far Detector no oscillation prediction (solid lines). The results are for a simple test in which truth quantities and selections were used. Ratios of the data distributions to the predicted distributions are also shown.

Figure 7.10, and both the Near Detector and Far Detector no oscillation predictions closely match the 'data' distributions.

There is some evidence that the predictions are about 1% too high in the peak. This is because the widths of the Gaussian distributions used in the matching process have a small impact on the height of the peak. There is no specific value required for these Gaussian widths (the Gaussians simply need to be very narrow), so the widths were tuned to produce optimal agreement for extrapolations using reconstructed quantities, rather than true quantities. In Section 7.4, the effects of making large changes to the Gaussian widths are investigated. The impact on an oscillation fit is very small.

Figure 7.11 shows the results of tests using reconstructed quantities, and making selection cuts on reconstructed quantities. There is excellent agreement between predictions and the data for both the Near and Far Detectors. The agreement between the predicted high energy tail and the test data is also excellent, as shown in Figure 7.12(a). This is despite the reduced statistics at high energies and can be partly attributed to the slow increase in the width of the Gaussian distributions used above 20 GeV.

The nominal MC 'fake data' can also be used to test that the treatment of the NC background is satisfactory. Figure 7.12(b) shows that the Far Detector NC background

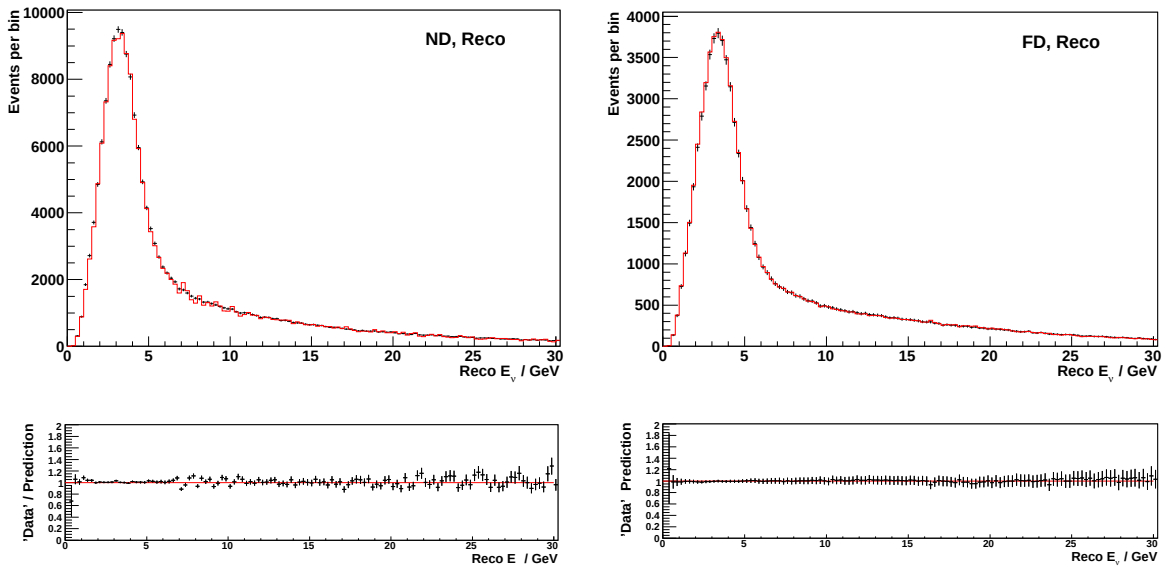

(a) (b)

Figure 7.11: Comparison of fake data (points) with the Near Detector prediction and Far Detector no oscillation prediction (solid lines). The results are for a realistic test of the method, in which reconstructed quantities were used and event selections were made by placing cuts on reconstructed quantities.

prediction is in close agreement with the actual NC contamination of the CC selected Far Detector data distribution.

A final test is to introduce the effects of oscillations into the fake data. Predicted spectra can then be made for different values of the oscillation parameters and a basic oscillation fit used to determine the parameters that best reproduce the data. If the method is successful, the best fit parameters should be consistent with the parameters used to introduce oscillations into the fake data.

The results of such a test are shown in Figure 7.13. This Figure shows that excellent agreement is obtained between the Far Detector 'data' spectrum and the best fit prediction, and that the true input oscillation parameters of $\Delta m^2_{32} = 3.0 \times 10^{-3}\,\mathrm{eV}^2$ and $\sin^2(2\theta_{23}) = 0.9$ are accurately reconstructed as:

$$\Delta m^2_{32} = (3.00^{+0.02}_{-0.02}) \times 10^{-3}\,\mathrm{eV}^2,\ \sin^2(2\theta_{23}) = (0.899^{+0.006}_{-0.007})$$

The errors quoted for this result are purely statistical; the effects of systematic uncertainties are discussed in detail in Sections 7.4 and 7.5.

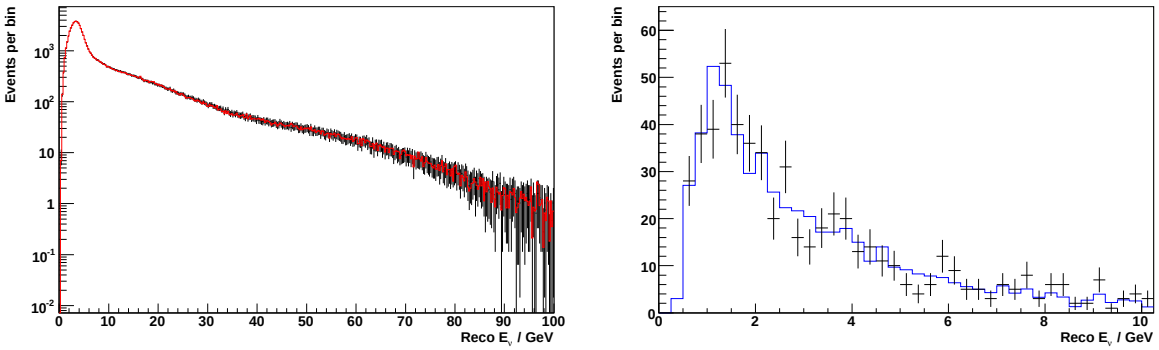

(a) (b)

Figure 7.12: (*a*) Comparison of Far Detector no oscillation prediction at high energies (red) and nominal Far Detector MC distribution (black). (*b*) comparison of Far Detector NC background prediction (solid line) and the true NC background in the fake data (points).

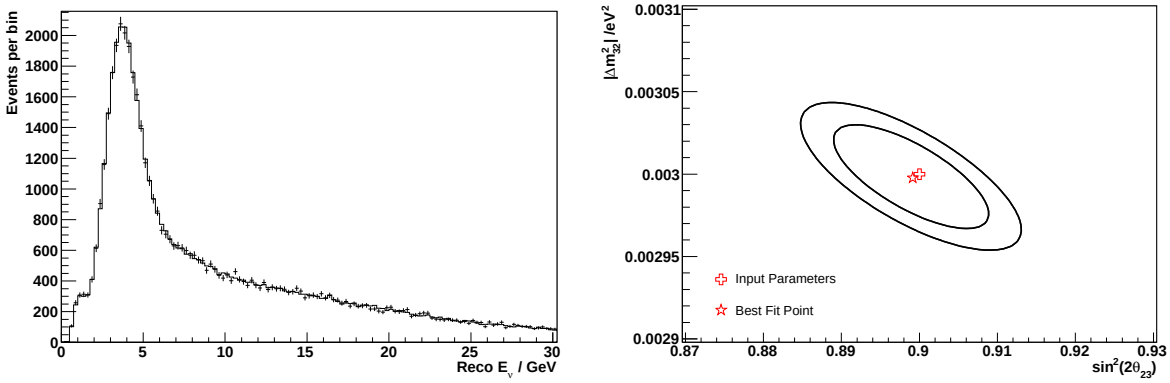

(a) (b)

Figure 7.13: Results of a basic oscillation fit to MC fake data. (*a*) Comparison of the fake data distribution (points) and the best fit prediction (solid line). (*b*) The 68% and 90% confidence level contours, and the best fit oscillation parameters.

7.3.2 Mock Data Challenges

After testing with nominal MC 'fake data', the next step is to test the ability of the extrapolation method to adapt to changes in the beam conditions. The MC required for such tests is available in the form of the MINOS mock data challenges[96].

The challenge datasets consist of fake Near Detector and Far Detector data samples. Both samples incorporate the effects of a tweaked hadron production spectrum, obtained using the parameterisation described in [97]. Other small distortions are included, representing a shift in the position of one of the horns, mis-calibration of the horn current and a change in the horn current distribution. The Far Detector challenge dataset also includes the effects of oscillations, for some specific values of the oscillation parameters. By using the Near Detector data to reweight the beam particle MC, the extrapolation method should successfully adapt to the change in beam conditions and allow accurate determination of the oscillation parameters.

The extrapolation was tested using two different mock datasets, labelled mock data challenge (MDC) II and III (MDC I can be considered as the test using nominal MC). The two datasets use the same tweaks to the beam conditions, but represent different oscillation parameters.

MDC II

The MDC II Far Detector dataset was divided into 100 small 'experiments' and each experiment, representing 1.032×10^{20} PoT, was fitted individually. Figure 7.14 shows the distribution of the best fit points and shows the mean 68% and 90% confidence level contours obtained for the 100 experiments. The Figure also shows the overall best fit point for the entire dataset, which is in good agreement with the MDC II input parameters.

Examining the positions of the best fit points for the 100 experiments shows that the number falling outside the two confidence level contours is approximately as expected from statistical fluctuations. The last feature of interest is the pile-up of best fit points at the boundary of the physical region, $\sin^2(2\theta_{23}) = 1$. This is expected and the distribution would be more uniform if the best fit was allowed to move into the unphysical region.

Figure 7.15 shows the results of the overall fit to the entire challenge data sample. There is close agreement between the best fit predicted energy distribution and the data distribution. The true input oscillation parameters also lie comfortably within the 68% confidence level contour. The best fit oscillation parameters are determined as:

$$\Delta m^2_{32} = (2.36^{+0.04}_{-0.04}) \times 10^{-3}\,\mathrm{eV}^2,\ \sin^2(2\theta_{23}) = (0.941^{+0.020}_{-0.020})$$

The errors quoted are purely statistical and the best fit values are to be compared with the MDC input values of $\Delta m^2_{32} = 2.38 \times 10^{-3}\,\mathrm{eV}^2$ and $\sin^2(2\theta_{23}) = 0.93$.

MDC III

MDC II and III share the same tweaks to the beam conditions, differing only in the input oscillation parameters. Any extrapolation method that is successful with MDC II should therefore also work effectively with the MDC III dataset. As with MDC II, the MDC III dataset was split into 100 small 'experiments', each representing 1.032×10^{20} PoT. Figure 7.16 shows the distribution of the best fit points for the 100 experiments, the mean 68% and 90% confidence level contours obtained and the overall best fit point for the entire dataset. The best fit points are clustered around the true input parameters, with the numbers outside the two contours consistent with the expectation from statistical fluctuations. The clustering of best fit points at the physical boundary is smaller than for MDC II; a reflection of the lower input value for $\sin^2(2\theta_{23})$.

Figure 7.17 shows the result of the fit to the entire 1.032×10^{22} PoT challenge dataset. There is good agreement between the best fit predicted energy distribution and the data energy distribution, but the true MDC III input parameters lie outside the 90% confidence level contour. This initially appears to be a cause for concern. However, an investigation was performed[98] in which the extrapolation was bypassed by applying the beam weights used to construct the mock dataset directly to a Far Detector MC sample. As indicated in Figure 7.17(b), it was found that the true MDC III input parameters are not precisely identified even if such a 'truth extrapolation' is used. The discrepancy between the input and best fit parameters must therefore be attributed to statistical fluctuations.

The best fit oscillation parameters for MDC III are determined as:

$$\Delta m^2_{32} = (2.86^{+0.05}_{-0.05}) \times 10^{-3}\,\mathrm{eV}^2,\ \sin^2(2\theta_{23}) = (0.800^{+0.015}_{-0.015})$$

The errors quoted are purely statistical and the best fit values are to be compared with the MDC input values of $\Delta m^2_{32} = 2.845 \times 10^{-3}\,\mathrm{eV}^2$ and $\sin^2(2\theta_{23}) = 0.772$.

7.4 Systematic Uncertainties

It is important for the extrapolation method to be robust against systematic uncertainties and to be relatively insensitive to the precise modelling in the nominal MC. With a robust extrapolation method, a difference between the data and the MC will not propagate to have a large impact on measurement of the oscillation parameters. However,

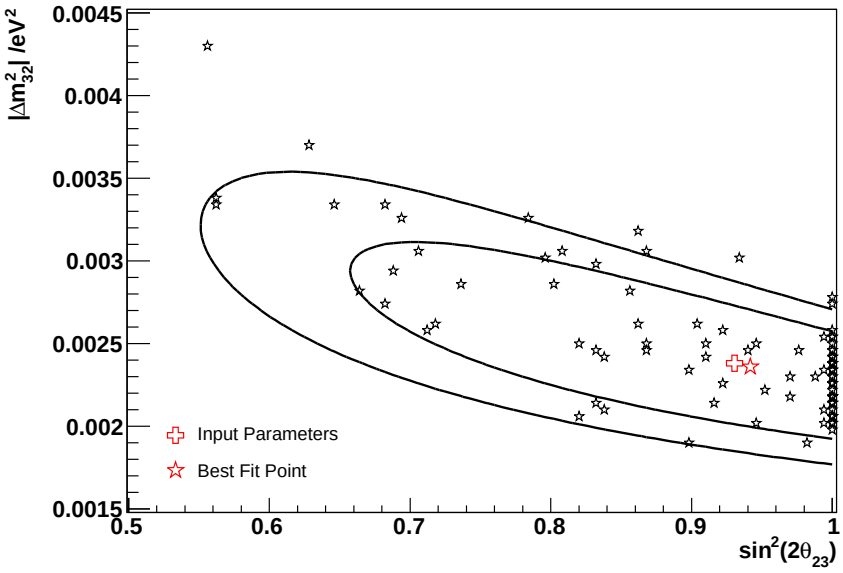

Figure 7.14: Distributions of best fit points obtained by fitting the 100 MDC II 'experiments'. The overall best fit point and the mean 68% and 90% confidence level contours are also shown.

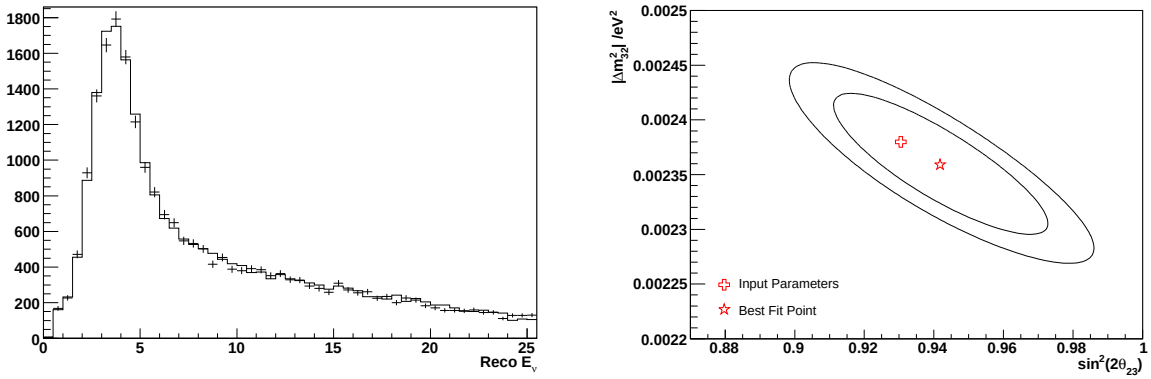

(a) (b)

Figure 7.15: Results of a basic oscillation fit to the full MDC II sample. (*a*) Comparison of the MDC II data distribution (points) and the best fit prediction (solid line). (*b*) The 68% and 90% confidence level contours, and comparison of the best fit oscillation parameters with the true input parameters.

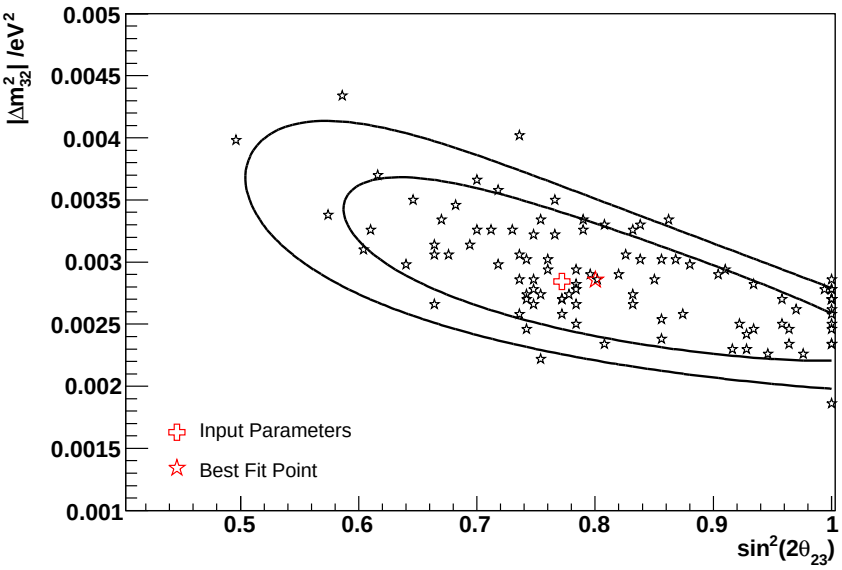

Figure 7.16: Distributions of best fit points obtained by fitting the 100 MDC III 'experiments'. The overall best fit point and the mean 68% and 90% confidence level contours are also shown.

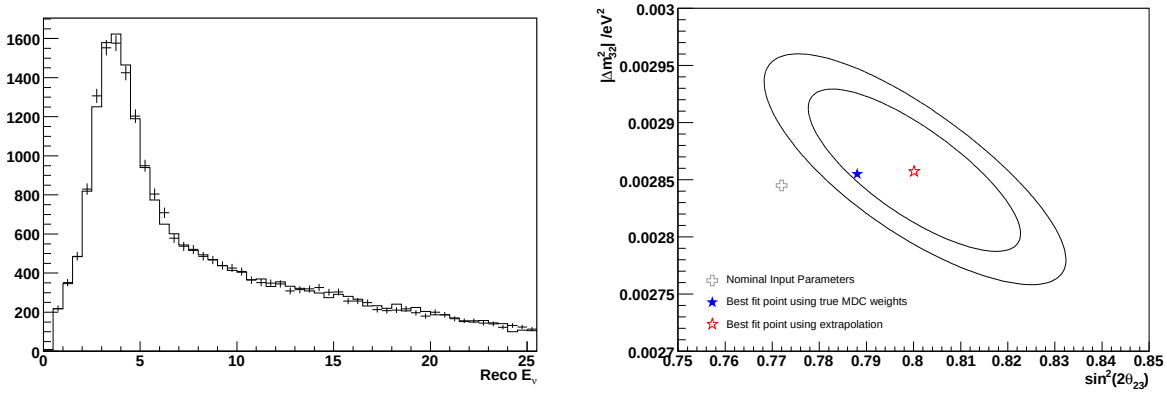

(a) (b)

Figure 7.17: Results of a basic oscillation fit to the full MDC III sample. (*a*) Comparison of the MDC III data distribution (points) and the best fit prediction (solid line). (*b*) The 68% and 90% confidence level contours, and comparison of the best fit oscillation parameters with the true input parameters and the results obtained using the true MDC weights instead of an extrapolation.

systematic uncertainties will unavoidably have some impact on the measurement, so it is important to characterise the response of the extrapolation to different systematic shifts. This enables the most important systematic uncertainties to be identified, so that they can be fully addressed in the fit to the oscillation parameters.

7.4.1 Approach and Results

The procedure to evaluate the impact of a shift in a poorly understood parameter is as follows:

1. Create some 'fake data' by oscillating nominal MC. This can then be extrapolated and the best fit oscillation parameters found by using MINUIT to control a basic oscillation fit.

2. Introduce a systematic shift into one of the poorly understood quantities in the fake data, leaving all the other MC unchanged.

3. Run the extrapolation again and find the new best fit oscillation parameters with this systematic shift.

4. The change in the best fit oscillation parameters then characterises the importance of the systematic uncertainty.

The use of MINUIT enables effectively continuous sampling in the oscillation parameter space, finding accurate values for the changes in best fit parameters.

To characterise the extrapolation method, fake data was generated with oscillation parameters $\Delta m^2_{32} = 3.0 \times 10^{-3}\,\text{eV}^2$ and $\sin^2(2\theta_{23}) = 0.9$. The fake data was scaled to $2.5{\times}10^{20}$ PoT and shifts of $\pm 1\sigma$ were considered for each of the following quantities:

- **NC level,** addressing uncertainties in the number of true NC events in the CC selected sample. A 50% uncertainty in the NC background is suggested in [53] and this is the initial estimate used in the analysis. However, in the optimised analysis presented in Chapter 10, the results from Chapter 6 are considered and a more realistic estimate of 25% is used.

- **Interaction cross-sections,** addressing uncertainties in the number (and relative fraction) of events due to quasi-elastic (QEL), resonance (RES) and deep-inelastic-scattering (DIS) interactions. The impact of 10% shifts in QEL and RES cross-sections and 20% in the DIS cross-section are considered, as suggested in [99].

- **$\bar{\nu}$ level,** addressing an important uncertainty in the beam composition. As suggested in [99], an uncertainty of 20% is considered for the number of $\bar{\nu}$ interactions.

- **Muon energy from range,** addressing systematic differences between the track reconstruction in data and MC, in addition to uncertainties in muon energy loss. The study in [100] reports that data and MC measurements agree to within 2%.

- **Muon energy from curvature,** addressing uncertainties in the modelling of the detector magnetic fields. The study in [100] suggests an uncertainty of 10%.

- **Normalisation**, addressing uncertainties in PoT counting, in event selection efficiencies and in the detector fiducial masses. These factors combine to introduce an uncertainty of 4% in the predicted Far Detector event rate.

- **Shower energy scale,** addressing uncertainties in the absolute hadronic energy scale and in the effects of intra-nuclear rescattering. An overall uncertainty of 10% is suggested in [101].

- **Relative Near/Far Detector shower energy scales,** addressing uncertainties in the different detector responses and calibrations. An uncertainty of approximately 3% is suggested in [62].

- **Relative Near/Far Detector muon energy from curvature,** addressing understanding of the differences between the magnetic fields at the two detectors. A plausible and conservative estimate of this uncertainty is 5%.

- **Gaussian widths,** addressing the freedom to vary the width of the Gaussian distributions used in the extrapolation. The actual widths used in the method were tuned for optimal performance in MC tests. Here, the impact of 10% shifts in the widths are considered.

- **Different beam conditions,** addressing the effects of altering the beam configuration. The effects were evaluated by using the beam reweighting scheme described in [92].

The results of this investigation to identify the most important systematic uncertainties are shown in Table 7.2 and illustrated in Figure 7.18.

Uncertainty	**Shift**	$\mathbf{\Delta(\Delta m^2_{32})/10^{-4}\,eV^2}$	$\mathbf{\Delta(\sin^2(2\theta_{23}))}$
NC Level	+50%	+0.165	−0.0075
	−50%	−0.165	+0.0077
QEL cross-section	+10%	+0.116	−0.0034
	−10%	−0.113	+0.0033
RES cross-section	+10%	−0.009	+0.0005
	−10%	+0.020	−0.0008
DIS cross-section	+20%	−0.177	+0.0048
	−20%	+0.216	−0.0058
$\bar{\nu}$ level	+20%	−0.0006	−0.000002
	−20%	−0.0001	+0.000049
E_μ range scale	+2%	+0.446	−0.0044
	−2%	−0.431	+0.0043
E_μ curvature scale	+10%	+0.102	−0.0068
	−10%	+0.023	+0.0117
Normalisation	+4%	−0.552	−0.0122
	−4%	+0.582	+0.0110
E_{shw} scale	+10%	+0.672	+0.0090
	−10%	−0.623	−0.0132
FD E_{shw} scale	+3%	+0.087	+0.0084
	−3%	−0.068	−0.0094
ND E_{shw} scale	+3%	+0.125	−0.0063
	−3%	−0.108	+0.0057
FD E_μ curvature scale	+5%	+0.063	+0.0047
	−5%	+0.002	−0.0071
ND E_μ curvature scale	+5%	+0.005	−0.0014
	−5%	+0.030	+0.0010
Gaussian width	+10%	−0.100	−0.0027
	−10%	+0.135	+0.0026
Beam reweighting	Run I weights	+0.274	+0.0007
	Run IIa weights	+0.284	−0.0004

Table 7.2: The shifts in best fit oscillation parameters, $\Delta(\Delta m^2_{32})$ and $\Delta(\sin^2(2\theta_{23}))$, induced by different systematic shifts in a fake data sample.

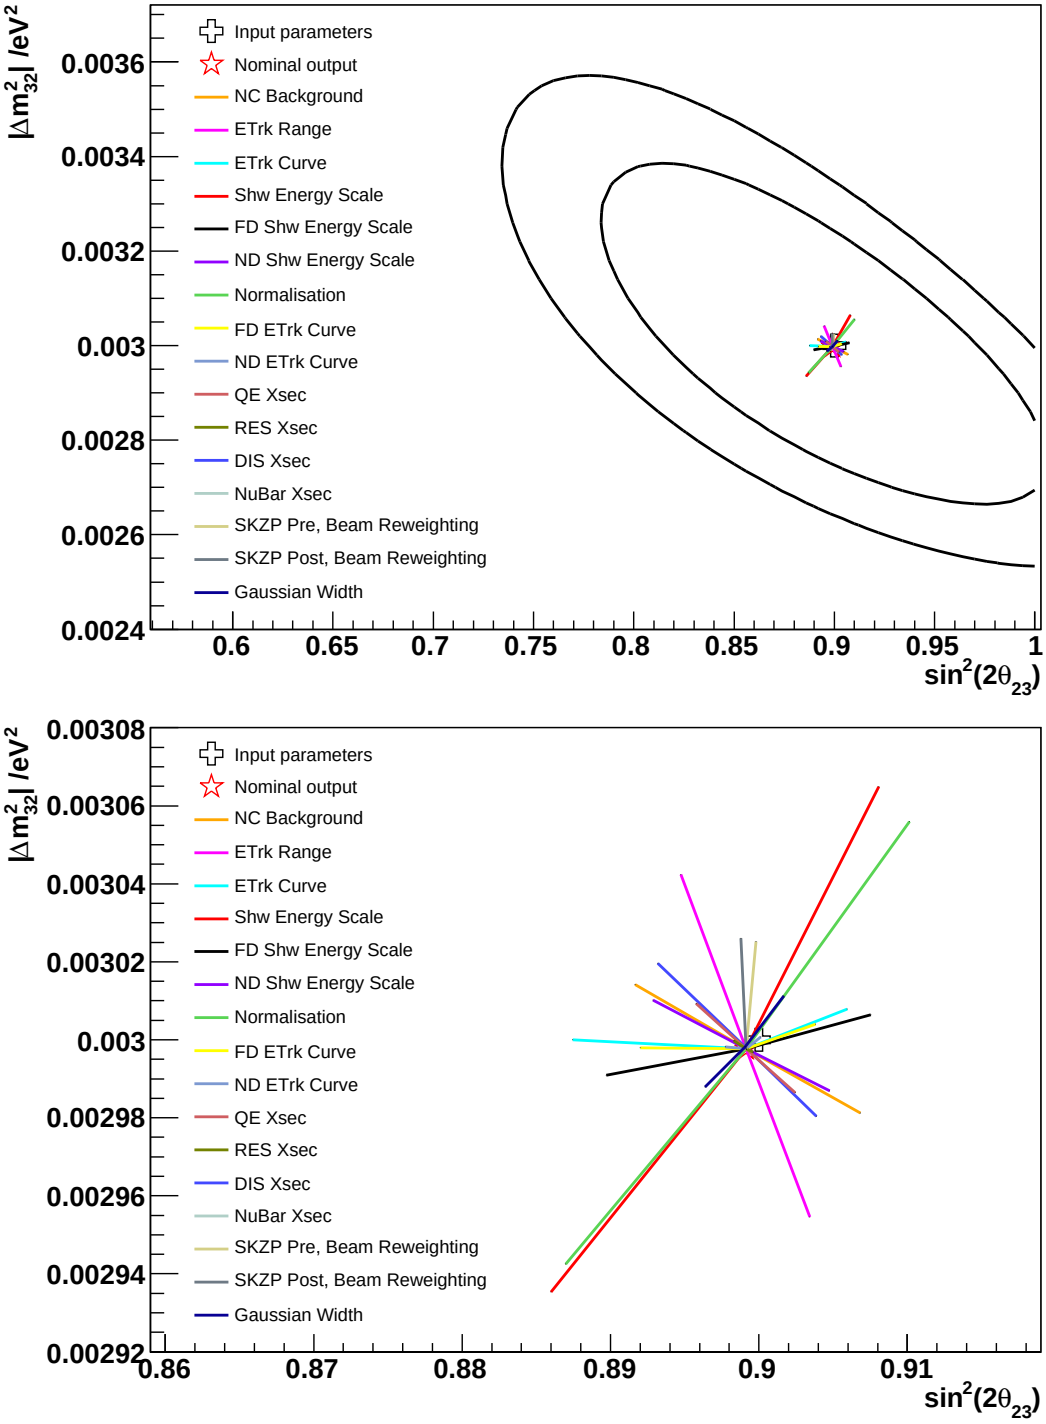

Figure 7.18: Summary diagrams of the shifts to the best fit oscillation parameters induced by different systematic shifts in a fake data sample. The 68% and 90% confidence level contours shown were produced using nominal MC as fake data and are scaled to 2.5×10^{20} PoT. The two diagrams show the same shifts, but on axes with different scales.

Figure 7.18 shows that the systematic uncertainties are relatively unimportant for an exposure of 2.5×10^{20} PoT, with the shifts contained well within the 68% confidence level contour. The Figure also shows that the two most important uncertainties are the shower energy scale and the normalisation. Also important are the muon range energy scale, the muon curvature energy scale, the Far Detector shower energy scale and the NC background. These six systematic uncertainties are accounted for in the full oscillation fit described in Section 7.5. The remaining uncertainties may prove to be important in future analyses with larger Far Detector datasets, but can be considered negligible at the current level of statistics.

In this study, distortions of the cross-sections and changes in the NC background only induced small shifts in the best fit point. The reasons why the extrapolation is so robust against these distortions are investigated in the remainder of this Section.

7.4.2 Cross-section Distortions

The impact of cross-section uncertainties on the extrapolation method can be investigated in further detail. The extrapolation needs to cope with distortions to the cross-sections so that it can produce accurate predictions even if the true cross-sections differ from those in the MC. One extreme set of tests is to set the cross-sections for the different interaction types to zero in the fake data. The MC libraries used in the extrapolation remain unchanged.

The three cases examined were removal of QEL events from the Near and Far Detector data, removal of RES events and removal of DIS events. In each case, the no oscillation Far Detector prediction from the extrapolation was compared to the actual no oscillation fake data. Oscillations were then introduced into the fake data and the shifts in the best fit parameters (from the parameters obtained using nominal MC) were evaluated.

Figure 7.19 shows that, for all these extreme cross-section distortions, the Far Detector no oscillation prediction remained in good agreement with the fake Far Detector data. Figure 7.20 shows that the shifts in best fit oscillation parameters are relatively small and unimportant for a data exposure of 2.5×10^{20} PoT. The extrapolation is therefore very robust against cross-section distortions and this can be attributed to the use of both E_ν and y as matching properties.

In an extrapolation, the only information available about the removal of a cross-section is the change in the reconstructed $E_\nu - y$ distribution for the Near Detector data. However, the different interaction types populate different areas of the recon-

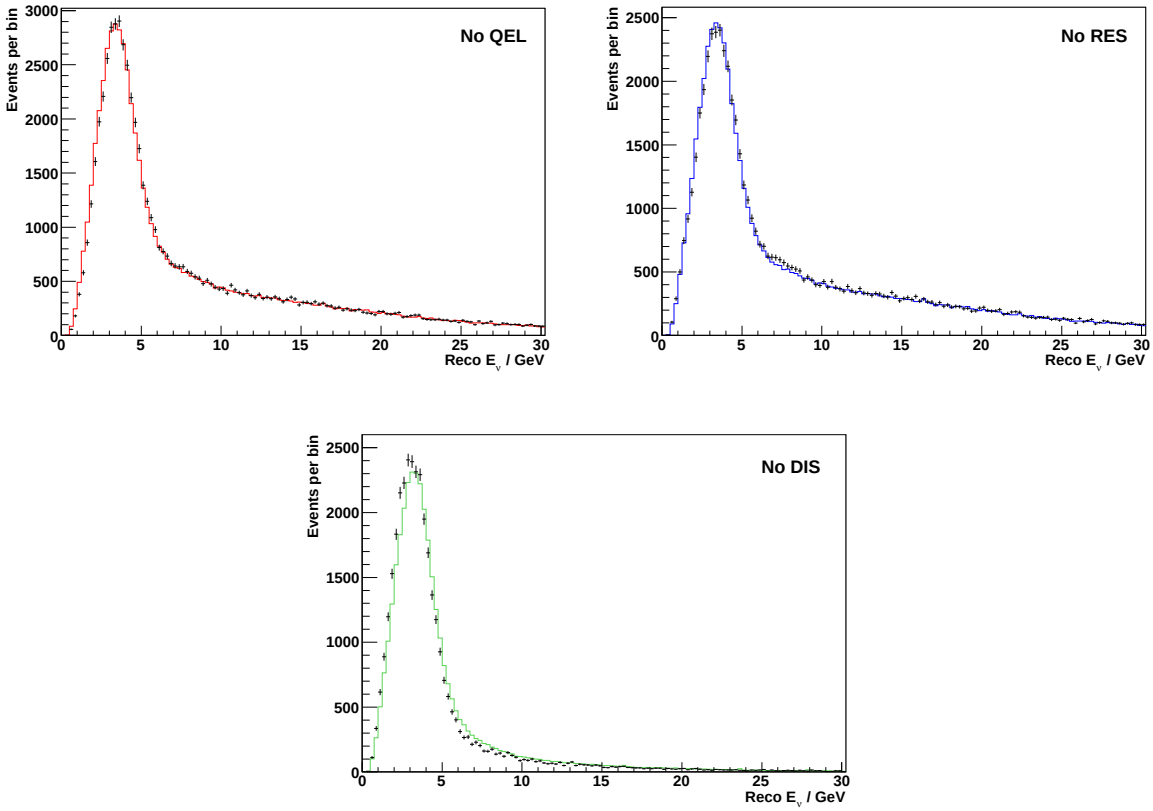

Figure 7.19: The Far Detector no oscillation predictions (solid lines) and no oscillation data (points), obtained after the removal of different interaction types from the Near and Far Detector fake data.

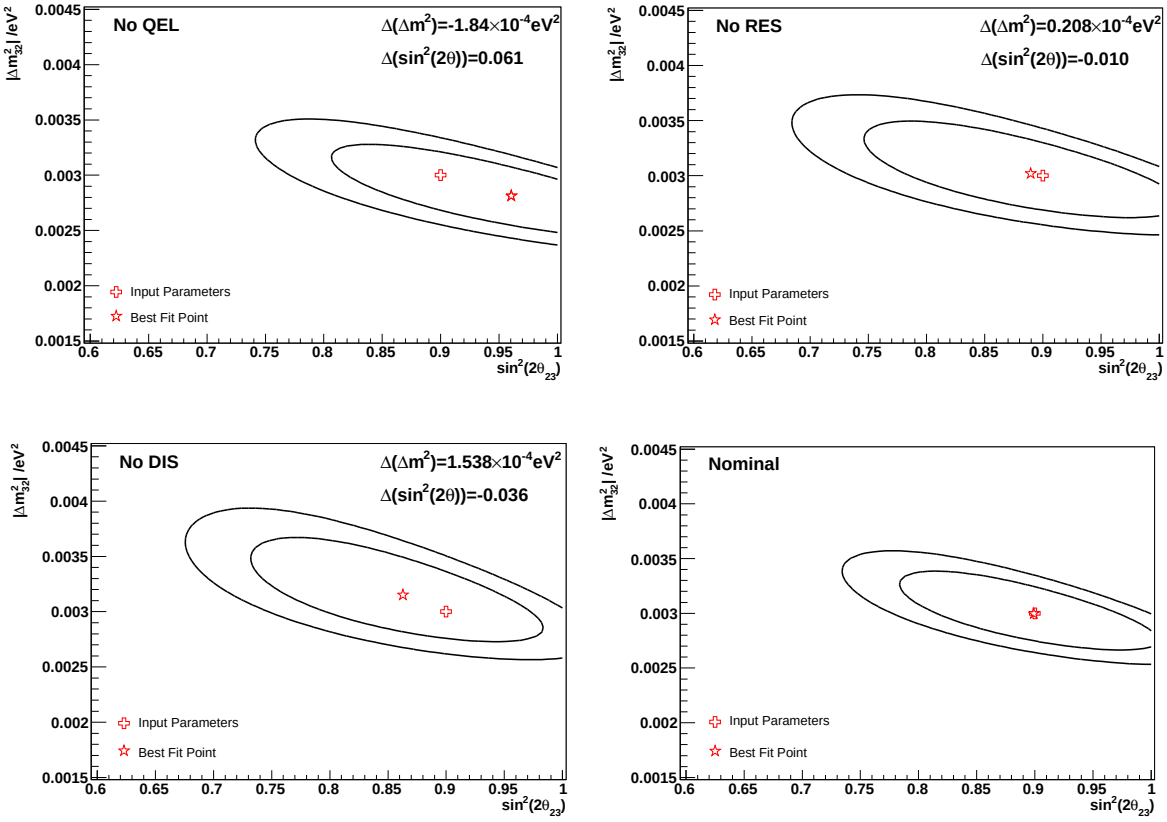

Figure 7.20: The shifts in best fit points and the new 68% and 90% confidence level contours (scaled to 2.5×10^{20} PoT) obtained after removal of different interaction types from the Near and Far Detector fake data. For comparison, the results obtained with nominal MC fake data are shown.

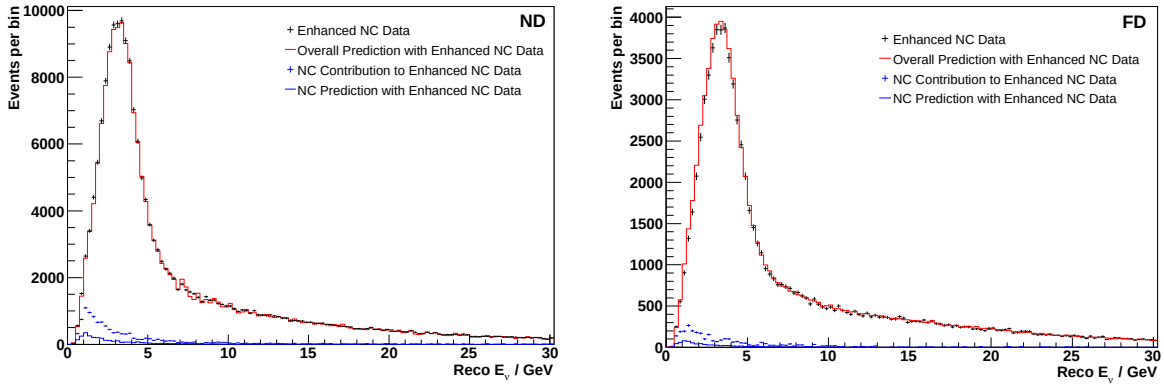

Figure 7.21: The predicted distributions and data distributions obtained when the NC fraction in the fake data is enhanced by a factor of five. Of particular interest are the predicted NC background distributions for the enhanced data.

structed $E_\nu - y$ space, so this is reflected in the beam particle weights obtained by Gaussian matching. These beam particle weights are then propagated to the Far Detector, allowing an accurate Far Detector prediction to be made, despite the extreme distortions that have been applied. If one were to use only E_ν as a matching parameter, the separation between the different interaction channels is not so apparent in the Near Detector data, increasing the sensitivity to cross-section variations.

7.4.3 Enhancement of NC

A further extreme test for the extrapolation is to greatly increase the number of NC events in the fake data samples and observe the impact on the Far Detector predictions. Figure 7.21 shows the results of increasing the number of NC events by a factor of five. The extrapolation is only aware of the increase because of changes in the $E_\nu - y$ distribution for the Near Detector data. However, these changes are enough to trigger a small increase in the beam particle weights for NC background events, increasing the NC background prediction.

The response of the extrapolation to the NC enhanced fake data is illustrated more clearly in Figure 7.22. Whilst the increase in beam particle weights for NC background events is far too small to allow accurate prediction of the NC background, the NC background prediction is larger than that obtained with nominal MC fake data. Although small, this response helps reduce the sensitivity of the extrapolation method to changes in the NC background.

Introducing oscillations to the NC enhanced data and performing an oscillation fit

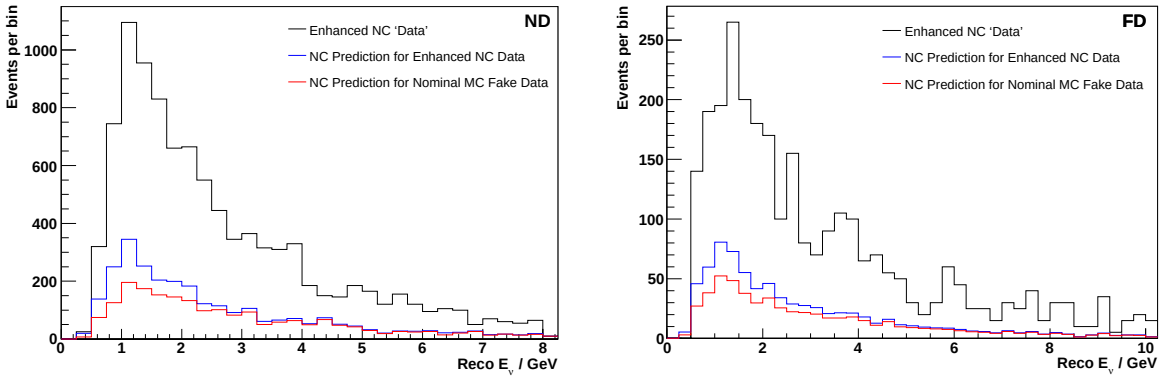

Figure 7.22: Comparison of the NC background predictions for NC enhanced fake data and for nominal MC fake data.

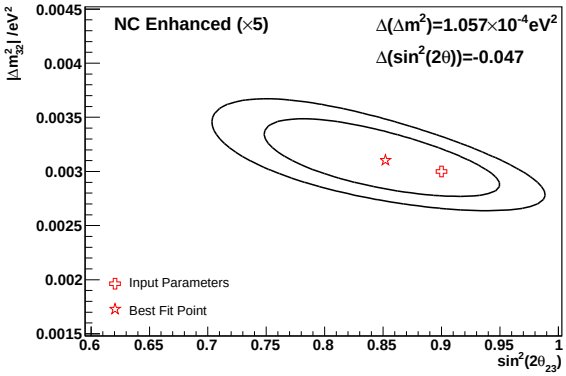

Figure 7.23: The best fit parameters and 68% and 90% confidence level contours (scaled to 2.5×10^{20} PoT) obtained by fitting fake data that has five times the nominal NC background.

produces the 68% and 90% confidence level contours (scaled to 2.5×10^{20} PoT) shown in Figure 7.23. As expected for a NC enhanced data sample, the best fit value of $\sin^2(2\theta_{23})$ is noticeably lower than the true input value.

7.4.4 Extrapolation without PID cut

As a final test of the extrapolation method, an extrapolation is attempted without use of a PID cut in the event selection. The Near and Far Detector MC libraries and the Near Detector fake dataset therefore consist of events passing CC selection cuts excluding PID and so contain many more NC events than usual. New selection efficiencies and flattening weights must be created, but the extrapolation is otherwise performed as normal. Only

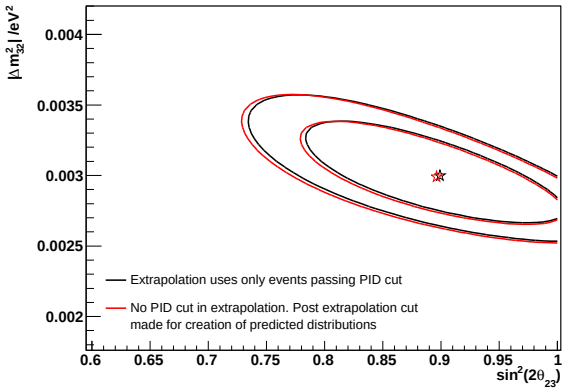

Figure 7.24: The best fit oscillation parameters and 68% and 90% confidence level contours (scaled to 2.5×10^{20} PoT) for two different MC tests. In the first test (black), only events passing the PID cut in the Near and Far Detectors are considered. In the second test (red), no PID cut is implemented in the Near Detector. Instead, a post-extrapolation PID cut is made on Far Detector events after the extrapolation weights have been assigned.

once the final extrapolation weights have been assigned is a PID cut implemented; the weighted Far Detector MC events only contribute to the final Far Detector predictions if they pass a post-extrapolation PID cut. At no point is a PID cut implemented at the Near Detector.

Once the Far Detector predictions have been obtained, they can be compared to the CC selected Far Detector fake data sample in the standard manner. Figure 7.24 compares the results of this extrapolation without PID cut to the standard results of extrapolating only CC selected events. It can be seen that there is little difference in the best fit oscillation parameters and in the 68% and 90% confidence level contours. This test demonstrates that the extrapolation method is insensitive to the precise details of the Near Detector event selection, providing further proof that the method is robust and suitable for use in an oscillation analysis.

7.5 Full Fit to Oscillation and Nuisance Parameters

In Section 7.4, the most important systematic uncertainties were identified as:

- Shower energy scale, $\sigma = 10\%$
- Relative normalisation, $\sigma = 4\%$
- Muon energy from range scale, $\sigma = 2\%$

- Muon energy from curvature scale, $\sigma = 10\%$
- Far Detector shower energy scale, $\sigma = 3\%$
- NC background, $\sigma = 50\%$ (σ is later reduced to 25%)

The effects of these uncertainties can be incorporated into the fit to the oscillation parameters by considering the uncertainties as 'nuisance parameters'. The fit proceeds largely as before, by producing predicted distributions for different values of the oscillation parameters, comparing the predictions to the data and calculating negative log likelihood values. However, at each sampling point in oscillation parameter space, the nuisance parameters are now allowed to vary; MINUIT can be used to alter these parameters to obtain the minimum negative log likelihood value possible for the given oscillation parameters. Scanning across oscillation parameter space, allowing the nuisance parameters to vary at each sampling point, the likelihood surface can be constructed and the best fit oscillation parameters identified.

The likelihood function of Equation 7.15 is used. This differs from the basic likelihood function in that it includes penalty terms. These increase the negative log likelihood value if a nuisance parameter moves from its nominal MC value. If a shift in a nuisance parameter is to produce the minimum negative log likelihood, the agreement between the prediction and data must improve enough to outweigh the penalty term.

$$-\ln \mathcal{L} = \sum_{i}^{n_{Ebins}} (N_i^{MC} - N_i^{Data} \ln N_i^{MC}) + \sum_{j}^{n_{syst}} \frac{a_j^2}{2\sigma_j^2} \tag{7.15}$$

where N_i^{MC} and N_i^{Data} are respectively the predicted content and the data content of reconstructed energy bin i. n_{syst} is the number of nuisance parameters included in the fit, a_j is the magnitude of the shift of nuisance parameter j from its nominal MC value and σ_j is the typical uncertainty in this parameter.

To perform the fit, Far Detector predictions need to be generated for different values of the oscillation parameters *and* any combination of the nuisance parameters. The predictions for a given shift in the nuisance parameters are produced by applying the systematic changes to the MC library files and extrapolating from the same Near Detector data. This is to be contrasted with the previous study, in which the most important systematic uncertainties were identified by applying shifts to fake data and leaving the MC unchanged.

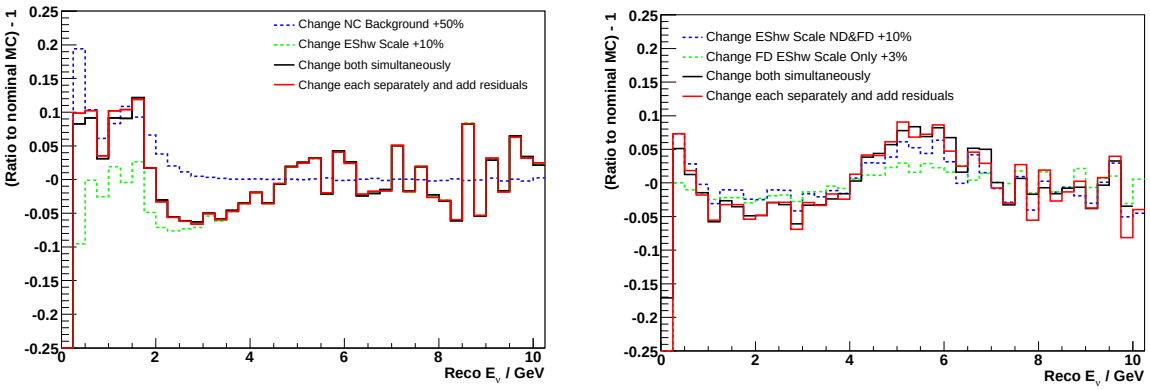

(a) (b)

Figure 7.25: Comparison of nominal Far Detector prediction with the predictions for different systematic shifts. Comparisons are shown for both independent and simultaneous shifts in different pairs of systematics. If two systematics are uncorrelated, the results of the simultaneous change can be reconstructed from the results of the independent changes. (*a*) Shower energy scale and NC background. (*b*) Shower energy scale and Far Detector shower energy scale.

The need to produce a prediction for any combination of oscillation and nuisance parameters appears to make the full fit prohibitively slow. However, there are some very good approximations that can be used:

1. It can be shown that the different systematic uncertainties are uncorrelated in the extrapolation. This means that the Far Detector prediction for a given set of nuisance parameters can be obtained by making separate predictions for each nuisance parameter and examining how each differs from the nominal prediction. The residuals from the nominal prediction can simply be added together to calculate the overall residual and so construct the final prediction.

   Figure 7.25 compares the nominal Far Detector prediction with the predictions for different combinations of nuisance parameters. It shows that the prediction for shifts in a pair of nuisance parameters can be reproduced by making separate changes and calculating the residuals from the nominal prediction. Figure 7.25(a) demonstrates that the shower energy scale and NC background systematics are uncorrelated, whilst Figure 7.25(b) provides the same demonstration for the shower energy scale and Far Detector shower energy scale.

2. For each extrapolation, a two dimensional histogram (a 'source histogram') can be

saved. This links every final extrapolation weight to a reconstructed E_ν value and to a true E_ν value. The reconstructed E_ν values should be binned in the same way as the energy distributions used in the oscillation fit, but the true energy values are used to calculate survival probabilities and so must be very finely binned. The source histogram provides all the information needed to produce Far Detector predicted energy spectra for any choice of oscillation parameters. Source histograms must be produced separately for CC and NC events and each source histogram is unique to the combination of nuisance parameters used in its production.

3. Source histograms need only be produced for a few values of each nuisance parameter, typically the nominal case and $\pm 1\sigma, \pm 2\sigma$. A quadratic fit between the bin contents then allows a prediction to be obtained for any reasonable value of the parameter.

When combined, these approximations enable accurate Far Detector predictions to be produced very quickly for any sensible combination of the oscillation parameters and the nuisance parameters. This then allows an accurate fit to be performed in timescales of a few hours.

Figure 7.26 shows example MC results for exposures of 1.27×10^{20} PoT, 2.5×10^{20} PoT and 12.7×10^{20} PoT. With only one year's data (1.27×10^{20} PoT), the various systematic uncertainties cause only modest broadening of the 68% and 90% confidence level contours and the error is largely statistical. With more data, the uncertainties begin to limit the sensitivity to the oscillation parameters, broadening the contours appreciably.

7.6 Data Predictions

In Chapter 6, the MINOS Run I and Run IIa datasets were introduced and it was explained that, due to the replacement of the NuMI target, the two Near Detector datasets must be extrapolated separately in an oscillation analysis. Applying the extrapolation method to the Near Detector Run I and Run IIa datasets produces the Near Detector data predictions shown in Figure 7.27 and the Far Detector no oscillation data predictions shown in Figure 7.28.

Figure 7.27 shows that the Near Detector data predictions are in close agreement with the Near Detector data. Also shown are the predictions made using nominal Near Detector MC as fake data (essentially just the nominal MC predictions). Weighting of the beam particles is clearly required in order to match the real data; the number of predicted events is enhanced by about 15% in the peak and by about 30% in the

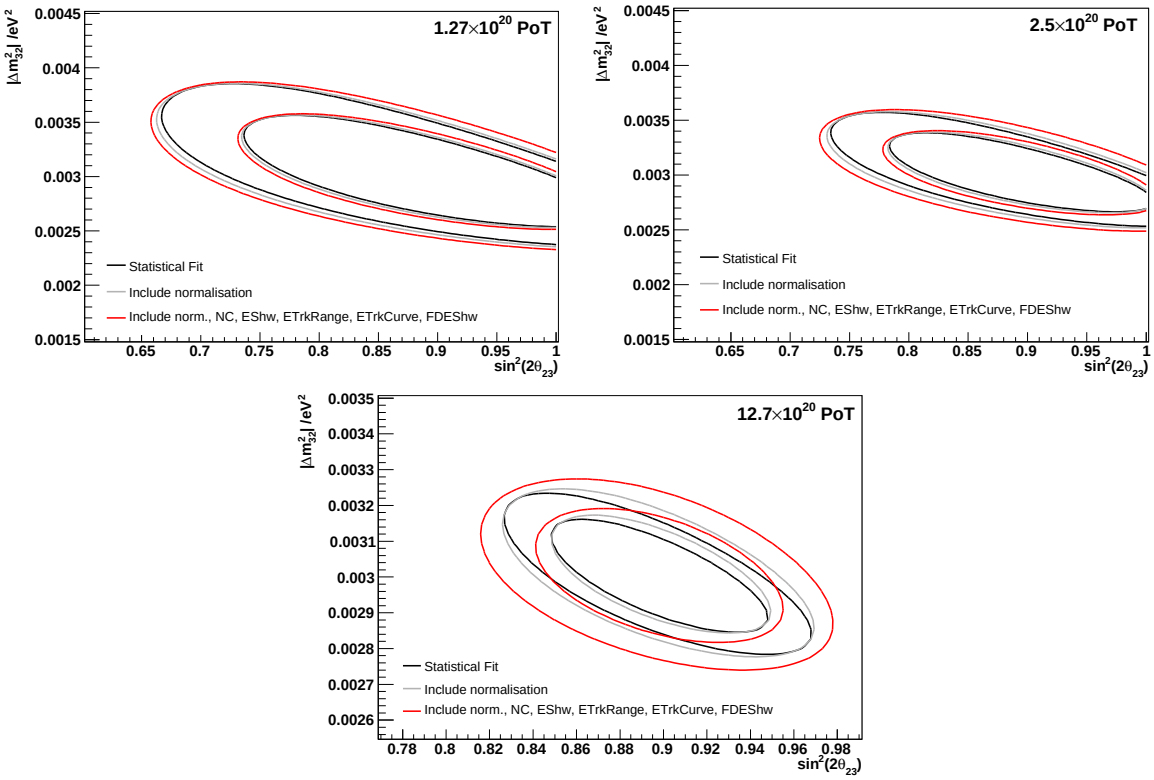

Figure 7.26: The broadening of the 68% and 90% confidence level contours, resulting from the inclusion of systematic uncertainties in the oscillation fit. Example MC results are shown for exposures of 1.27×10^{20} PoT, 2.5×10^{20} PoT and 12.7×10^{20} PoT.

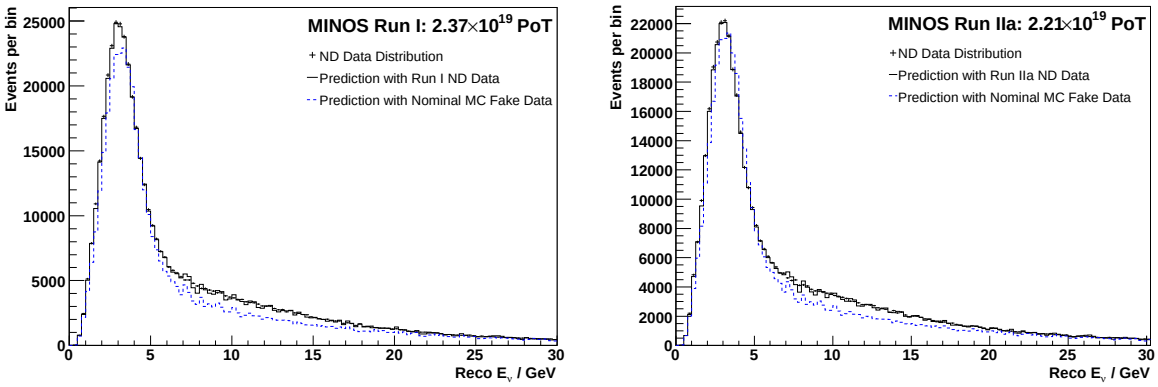

Figure 7.27: Near Detector data and predictions for Run I and Run IIa. For comparison, the predictions obtained using nominal MC fake data are displayed.

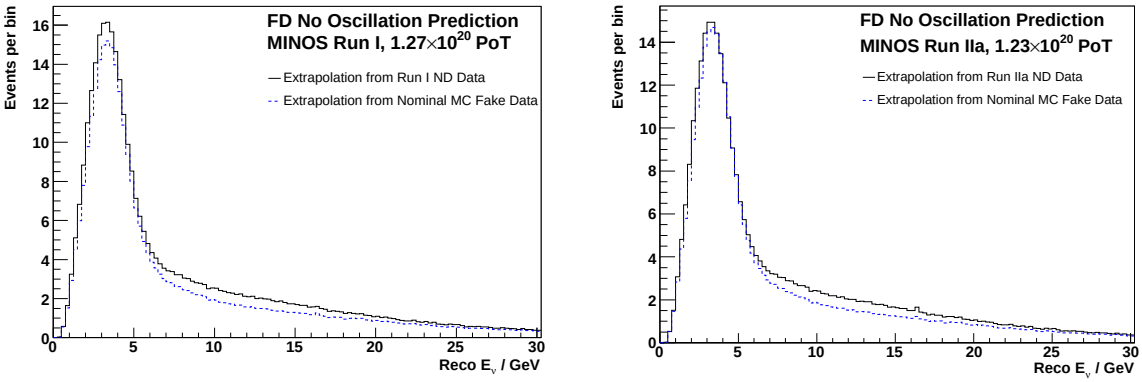

Figure 7.28: Far Detector no oscillation data predictions, obtained by extrapolating the Run I and Run IIa Near Detector data. For comparison, the predictions obtained by extrapolating nominal Near Detector MC are displayed.

$5-25\,\text{GeV}$ region. There is also evidence for a small shift to lower energies in the peak region. The corresponding Far Detector predictions in the absence of oscillations are shown in Figure 7.28. These predictions are also higher in the peak and $5-25\,\text{GeV}$ region than the extrapolations from nominal Near Detector MC.

To determine the best fit oscillation parameters and confidence limits from the data, all that remains is to compare the Far Detector predictions for different oscillation parameters with the actual Far Detector data, performing a fit as described in Section 7.5. The results of this fit to the $2.5\times10^{20}\,\text{PoT}$ Far Detector data sample are presented in detail in Chapter 10.

Chapter 8

Data/MC Comparisons

In order to extract a measurement of the neutrino oscillation parameters from the MINOS data, a selection procedure has been developed to identify CC ν_μ interactions and an extrapolation method has been implemented to predict the Far Detector neutrino energy spectrum. These analysis techniques have been developed using the MC simulation; the extrapolation method, in particular, makes extensive use of the MC. It is therefore very important to examine the agreement between the simulation and the real data and to understand the reasons for any discrepancies.

This Chapter presents a comparison of the data and MC in the Near and Far Detectors. The comparison focusses on the quantities used to separate NC and CC interactions, together with those used to determine the energies of the muon tracks and hadronic showers in CC interactions. The aim is to identify any major discrepancies that would not be addressed by the treatment of systematic uncertainties developed in Chapter 7. The agreement between the data and MC is not used to identify the most important uncertainties, nor to quantify the level of the uncertainties. Instead, the comparison is a cross-check to ensure that no major problems have been overlooked and that the analysis of the data can proceed.

In the process of the extrapolation, weights are calculated for each Near Detector and Far Detector MC event, as described in Chapter 7. These weights incorporate information about the neutrino beam and cross-sections from the Near Detector data and, as intended, applying the weights to MC events greatly improves agreement with the data. In this comparison of data and MC, extrapolation weights are calculated and

applied to each MC event. The issues that could affect an oscillation analysis are those data/MC disagreements that remain after this MC reweighting.

8.1 Near Detector

The large volume of data collected at the Near Detector enables high statistics comparisons to be made between data and MC. For this study, Near Detector data were selected as described in Chapter 6, resulting in a sample representing 2.37×10^{19} PoT (534,453 selected events) from Run I and 2.21×10^{19} PoT (483,772 selected events) from Run IIa. These data were compared with a MC sample equivalent to 2.91×10^{19} PoT. In all the distributions shown in this Section, the data and MC are both scaled to 1×10^{18} PoT. Separate Run I and Run IIa extrapolation weights were calculated for each MC event and the final MC weight was the PoT weighted mean of these two extrapolation weights.

8.1.1 PID related quantities

Figure 8.1 shows the comparison between data and MC for the five variables used in the generation of the PID parameter. No PID cut was made to produce these distributions, but all the other selection cuts detailed in Chapter 6 were made. There is reasonable agreement between the data and MC for these input variables, but there are some important features worth noting.

- There is evidence for poor modelling of the spectrometer region of the detector. This can be seen in Figure 8.1(a), which shows disagreement between data and MC at high values of the track planes variable. This variable is defined as the number of hit planes on the track, rather than the overall track length. As there is good agreement for the track length variable, see Figure 8.5(b), this indicates that tracks in data are more likely to have 'gaps' in the spectrometer region.

- More tracks in the data fail the Kalman fit compared to the MC expectation. The tracks failing the fit are not assigned a q/p or $\sigma_{q/p}$ value and are all found in the first bin of the $|q/p| \,/\, \sigma_{q/p}$ distribution, shown in Figure 8.1(c). This Figure shows a large excess of data events in the first bin.

- There is evidence for poor simulation of low energy showers. This can be seen from the excess of low y data events in Figure 8.1(e), which is due to a large excess of low energy showers in the data. Event displays reveal that the extra showers are

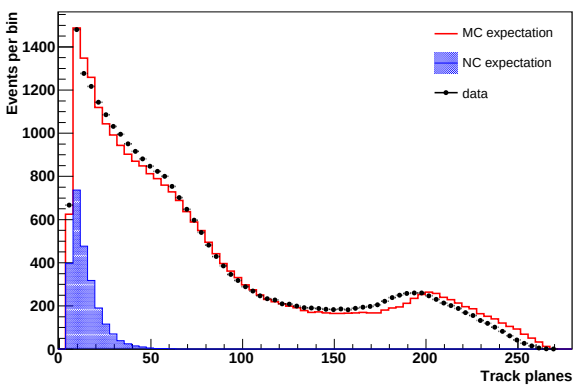

(a)

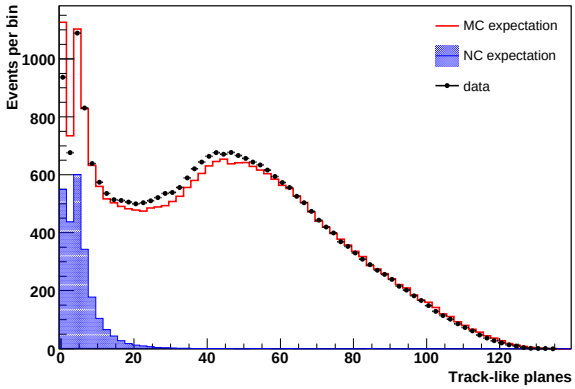

(b)

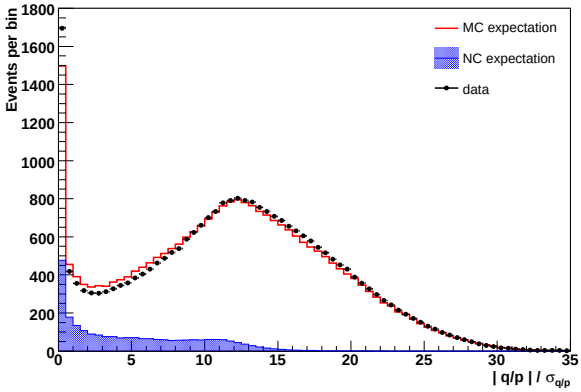

(c)

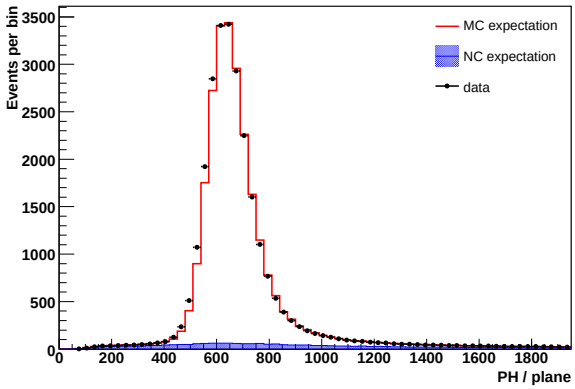

(d)

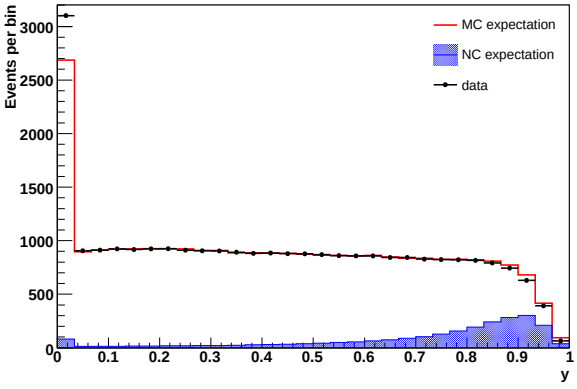

(e)

Figure 8.1: Comparison between Near Detector data and MC for the variables used in the generation of the PID parameter: (a) track planes, (b) track-like planes, (c) goodness of muon track fit, (d) mean pulse height on track-like planes, and (e) reconstructed y.

typically in close proximity to the track, so they also have the effect of reducing the number of track-like planes, as shown in Figure 8.1(b).

- One plausible explanation for the excess of low energy showers is the effect of the 'threshold' (minimum strip charge) for observation of a strip in the detector. This may be poorly modelled in the MC simulation. Further thresholds in the reconstruction packages, all determined from the MC, control how a strip is used in the algorithms (for instance, whether it can be included in a shower). Also relevant are the uncertainties in the modelling of intra-nuclear rescattering. These uncertainties affect the simulation of which hadrons (produced in a neutrino interaction) actually reach the detector and how much energy is lost in traversing the nucleus.

- The close data/MC agreement for the y distribution, shown in Figure 8.1(e), is primarily due to the use of the extrapolation weights. These weights try to fix the MC E_ν and y distributions so that they match the data as closely as possible. After the MC reweighting, the only remaining discrepancy in the y distribution is the excess of low y data events, discussed above.

The data and MC distributions for the resulting PID parameter are shown in Figure 8.2. The agreement is generally good, with the largest discrepancy being an excess of MC events in the NC rich region. This is in agreement with the findings of Chapter 6, where it was suggested that the NC component of the MC be adjusted by $(-6.3 \pm 1.3)\%$.

The overall effect of the discrepancies in the PID input variables can be evaluated by performing a MC test extrapolation and oscillation fit with two different Near Detector fake data samples. The first Near Detector fake data sample simply has standard values of the PID input variables. However, the input variables in the second sample are distorted so that they match those of the true Near Detector data. Performing this test with input oscillation parameters of $\Delta m^2_{32} = 3.0 \times 10^{-3}\,\mathrm{eV}^2$ and $\sin^2(2\theta_{23}) = 0.9$ produced shifts in best fit parameters of $\Delta(\Delta m^2_{32}) = 0.012 \times 10^{-3}\,\mathrm{eV}^2$ and $\Delta(\sin^2(2\theta_{23})) = 0.009$. These shifts are smaller than most of those observed in the systematics study of Chapter 7, indicating that the data/MC disagreement in the PID input variables does not result in a significant systematic error.

8.1.2 Track related quantities

Figure 8.3 shows the number of tracks observed in selected events; i.e. after the PID selection cut. As expected, this distribution is strongly peaked at one track. A 2.5%

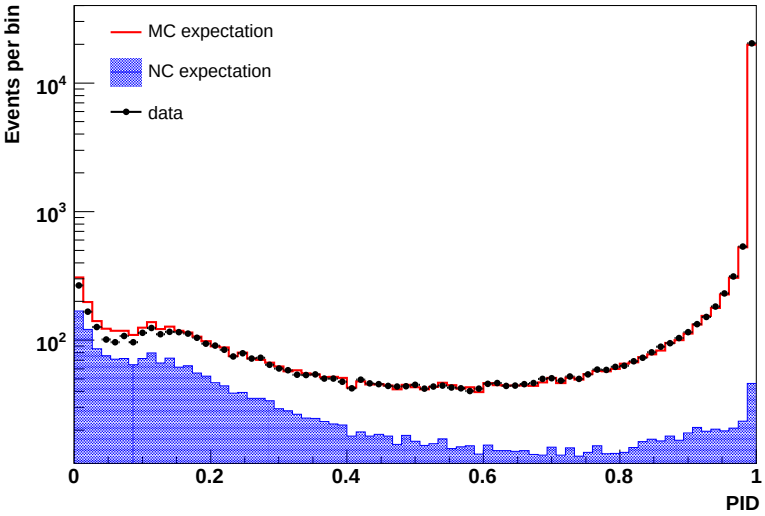

(a)

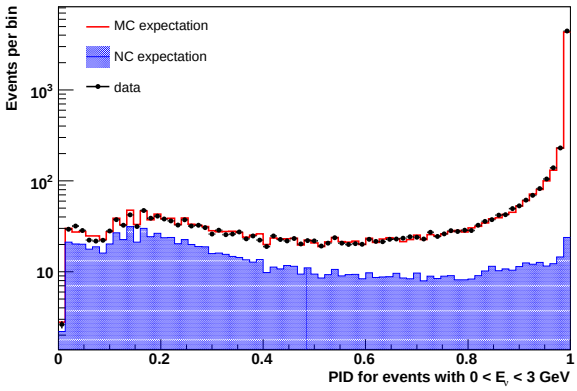

(b)

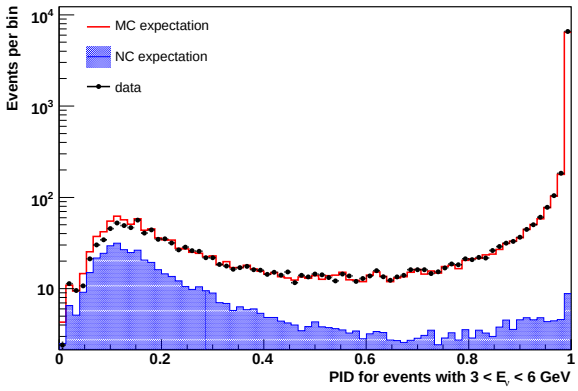

(c)

Figure 8.2: Comparison of PID distributions for Near Detector data and MC: (*a*) for all events, (*b*) for events with reconstructed E_ν below 3 GeV, and (*c*) for events with reconstructed E_ν between 3 GeV and 6 GeV.

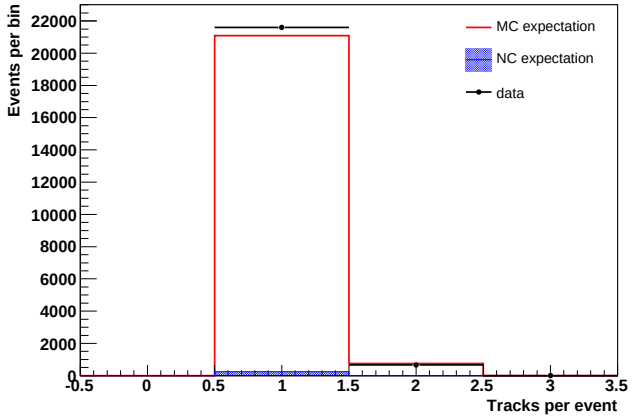

Figure 8.3: The number of reconstructed tracks observed in selected Near Detector data and MC events.

excess of one track events is observed in the data and this excess needs to be considered when interpreting the other Near Detector distributions. For the oscillation analysis, this excess is covered by the 4% systematic uncertainty in the normalisation.

The distributions of track vertex and end coordinates are shown in Figure 8.4. The excess of data events is visible in the track vertex distributions, which are otherwise well modelled. However, there are clear discrepancies between the track end distributions in data and MC, with a 10% excess of MC events at large x. This discrepancy is due to poor modelling of the Near Detector magnetic field, which leads to different levels of focussing in data and MC. Further discussion on this topic can be found in [100].

Figure 8.5(a) shows the number of track strips in data and MC events, revealing that there is a shift in the data towards fewer track strips. However, Figure 8.5(b) shows that the overall track length agrees closely between data and MC. These observations imply that the tracks in data are more likely to contain gap planes; planes between the track vertex and end plane that do not contain a reconstructed track strip. The track reconstruction software attempts to fill such gaps and only fails to do so when there are no plausible strips in the slice. The interpretation is therefore that track strips in data are more likely to be incorrectly divided into different slices than those in MC.

The track length distributions in Figure 8.5(b) show sharp rises at 197 planes, due to tracks with vertex at the very last plane in the fiducial volume that proceed to pass through the entire spectrometer region. There is a sudden cut-off in the distributions at 264 planes, which is associated with tracks that travel the maximum possible distance; from the first plane in the fiducial volume, through the spectrometer region and out of

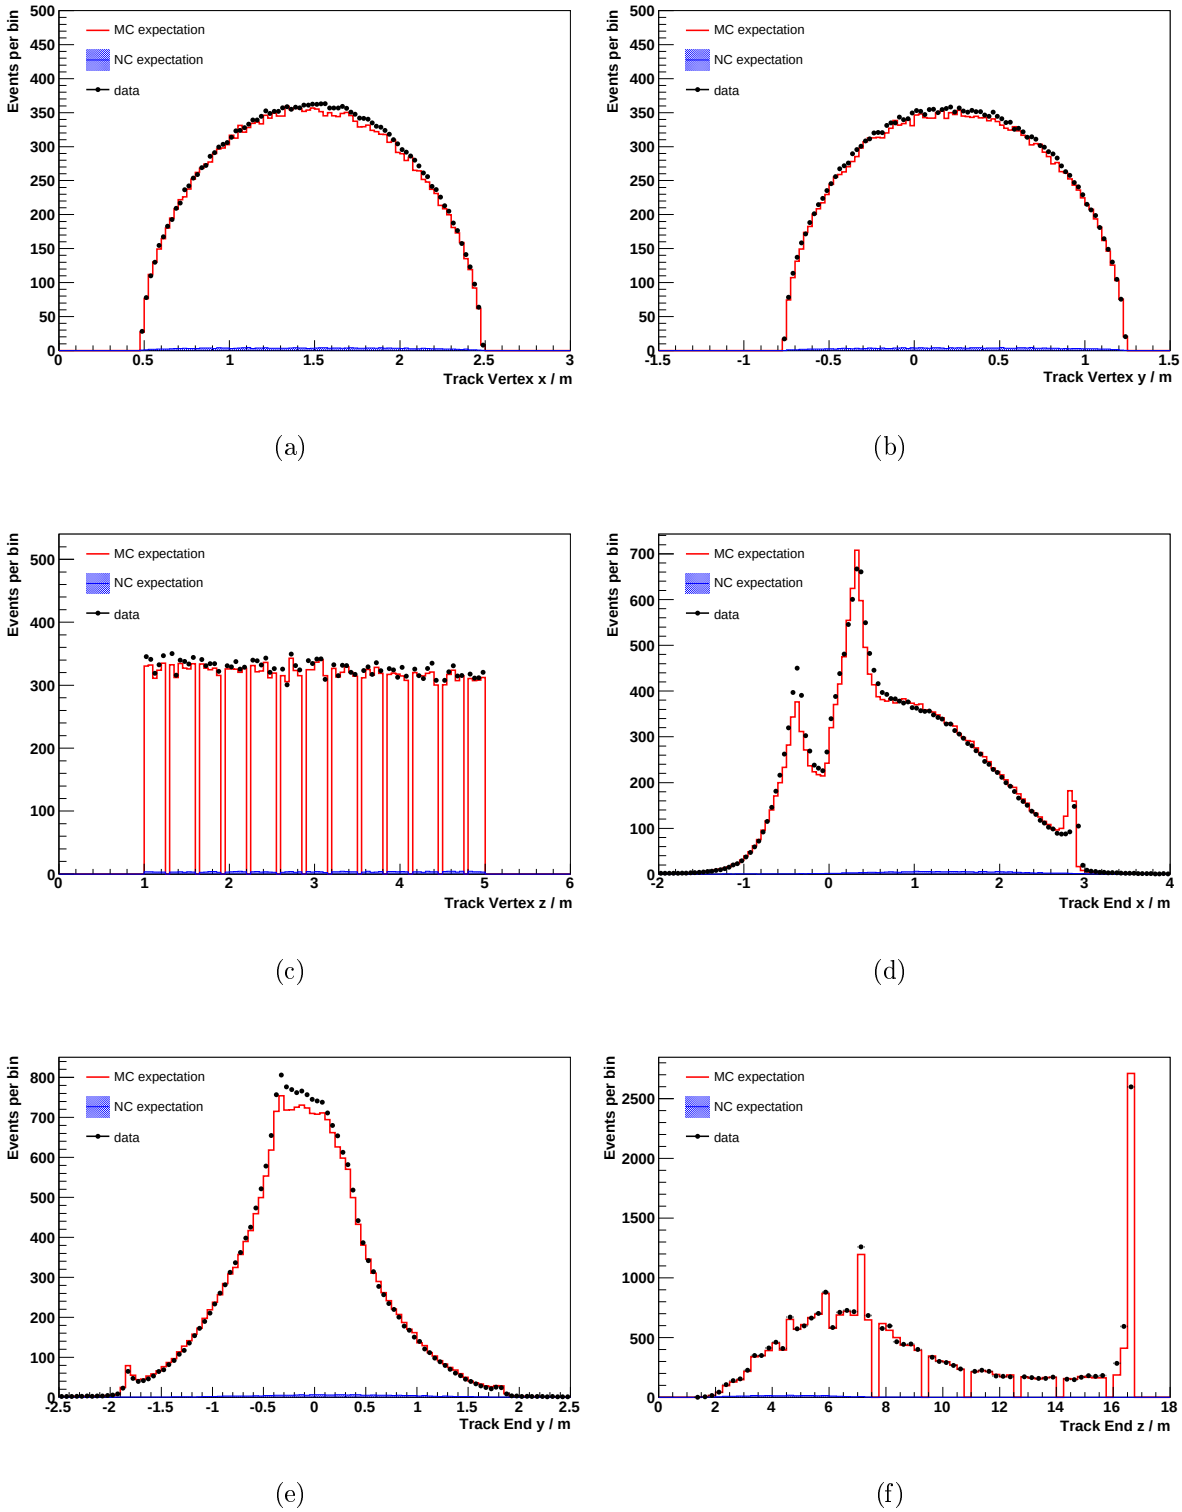

Figure 8.4: The distributions of track vertex and end positions for selected Near Detector data and MC events: (a) vertex x coordinate, (b) vertex y coordinate, (c) vertex z coordinate, (d) end x coordinate, (e) end y coordinate, and (f) end z coordinate.

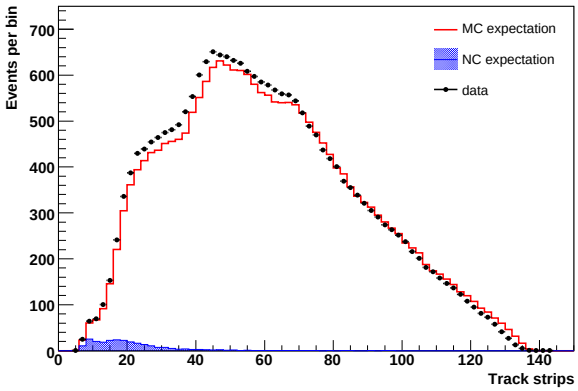

(a)

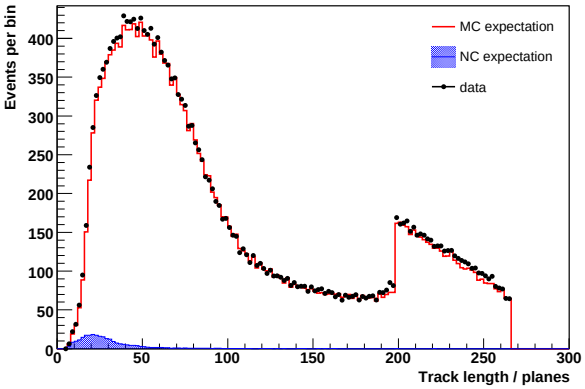

(b)

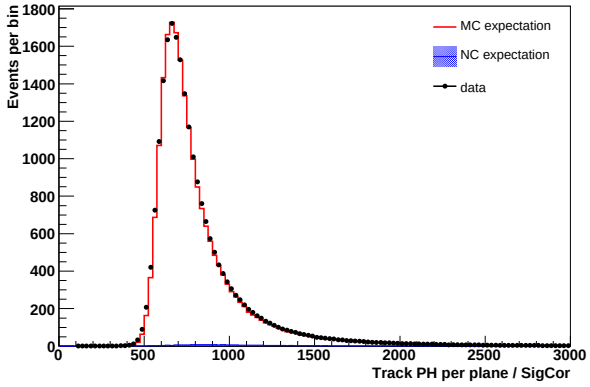

(c)

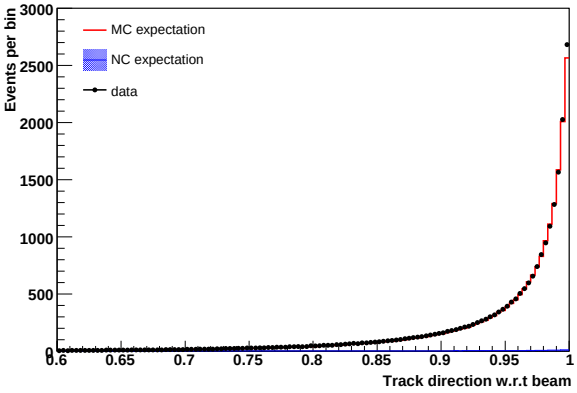

(d)

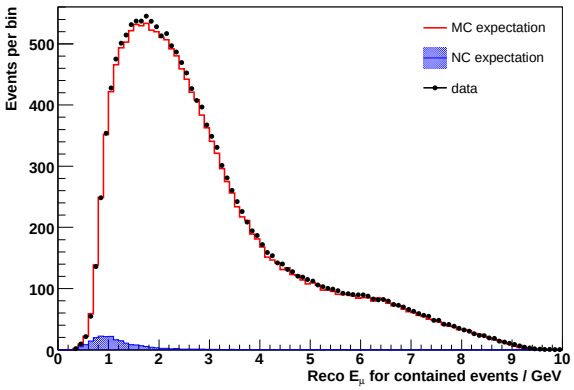

(e)

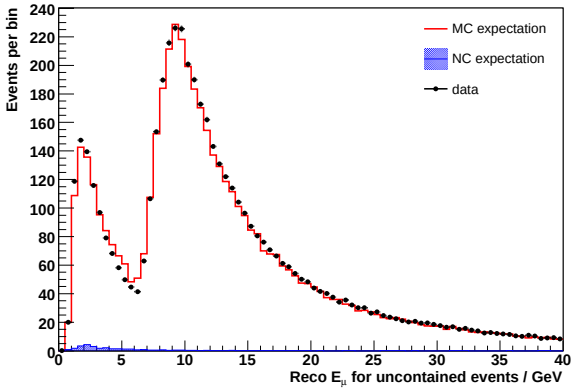

(f)

Figure 8.5: Comparison between Near Detector data and MC: (a) number of track strips, (b) track length, (c) track pulse height per plane, (d) track direction cosine with respect to the beam, (e) reconstructed muon energy for contained tracks and, (f) reconstructed muon energy for uncontained tracks.

the detector. A large fraction of the tracks that enter the spectrometer reach the end of the detector.

The remaining distributions in Figure 8.5 show the track pulse height per plane, the track direction with respect to the beam and the reconstructed muon energy. Figure 8.5(e) shows the reconstructed muon energy for contained tracks, for which the energy is calculated from track range, whilst Figure 8.5(f) shows the distribution for uncontained tracks, for which the energy is determined from track curvature in the magnetic field. Figure 8.5(f) reveals that there are two distinct classes of track for which the curvature measurement is used. The first class consists of low energy muons that pass out of the side of the detector, whilst the second class consists of high energy muons that pass all the way through the spectrometer and out of end of the detector. For both the range and curvature energy measurements, close agreement is observed between the data and MC. This is unsurprising, as the extrapolation weights fix the E_ν and y distributions to ensure correct estimation of the neutrino energy and correct division of the energy between the muon track and hadronic shower.

8.1.3 Shower related quantities

In Section 8.1.1, a comparison of the PID input variables for data and MC suggested an excess of low energy showers in the Near Detector data, relative to the MC expectation. This excess is confirmed by the distributions in Figure 8.6. Figure 8.6(a) shows that there is a 6% excess of single shower events in the data, whilst Figure 8.6(b) shows that the excess showers are very small; typically containing fewer than 10 strips. Figure 8.6(c) then shows the distributions of shower pulse height per strip and suggests that the strips in the excess showers have low pulse heights and represent small energy deposits. Finally, Figure 8.6(d) reveals a 10% excess of data showers with energies reconstructed in the range $0 - 125\,\mathrm{MeV}$.

These observations provide evidence for poor simulation of low energy showers. The low pulse heights and low strip multiplicities of the excess showers suggests that the problem is related to modelling of the thresholds for strip reconstruction and shower strip identification. However, the modelling of low pulse height strips is a particularly complex issue, which encompasses the uncertainties in modelling intra-nuclear rescattering and the response of the detector. The reconstruction of low pulse height showers is also strongly affected by calibration uncertainties, which affect the scintillator response, the light attenuation in the wavelength-shifting fibres and the PMT gains. Finally, pathologies in the reconstruction software could contribute to the discrepancy.

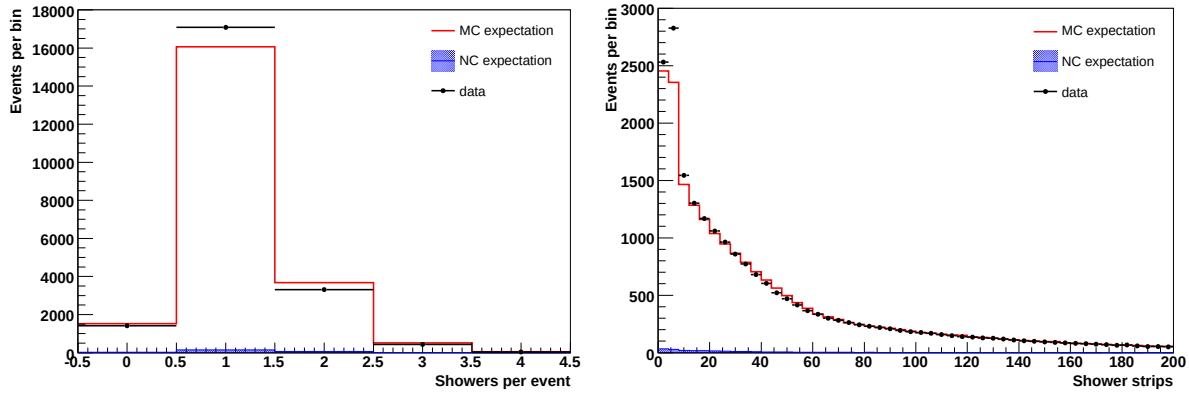

(a) (b)

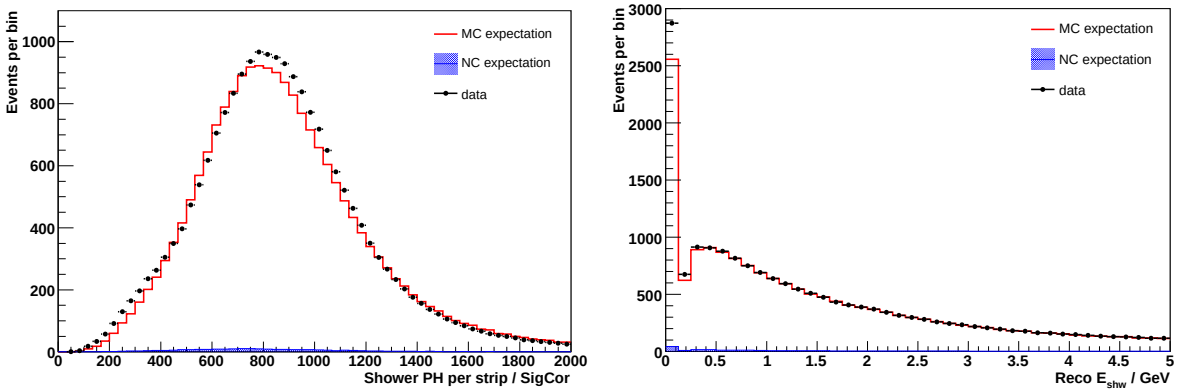

(c) (d)

Figure 8.6: Comparison between Near Detector data and MC: (*a*) number of showers per event, (*b*) number of shower strips, (*c*) shower pulse height per strip, and (*d*) reconstructed shower energy.

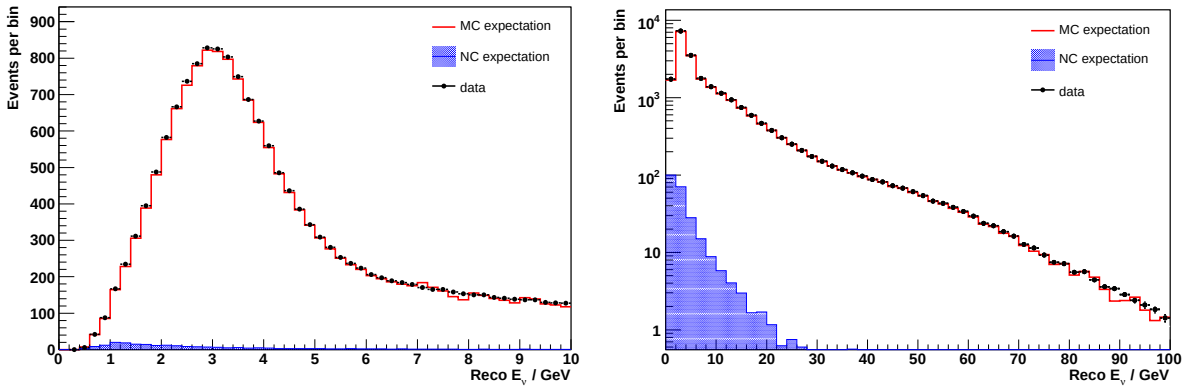

(a) (b)

Figure 8.7: Comparison of reconstructed neutrino energy distributions for Near Detector data and MC.

8.1.4 Neutrino energy

Figure 8.7 shows the reconstructed neutrino energy distributions for both data and MC events. As the extrapolation weights are designed to match the MC E_ν distribution to the data, the agreement between the data and MC is (by construction) almost perfect. This agreement holds both at low energies and at much higher energies, in the tail of the distribution.

8.2 Far Detector

The agreement between data and MC can also be assessed at the Far Detector. However, the volume of data collected at the Far Detector is quite small, so it is difficult to draw firm conclusions. This section therefore looks only to find general consistency between the data and the MC predictions, searching for areas where there is evidence of strong disagreement and poor simulation. To aid this examination, the agreement between the data and MC is quantified by calculating a χ^2/dof value for each distribution. This is defined as follows:

$$\chi^2 = \sum_{\text{bins}} \frac{(n_o - n_e)^2}{\sigma_e^2} \tag{8.1}$$

where n_o is the number of observed events in each bin of the distribution, n_e is the number of expected events in each bin and σ_e is the statistical uncertainty in the expected number

of events. No detailed statistical interpretation is drawn from the value of χ^2/dof, but values clustered around unity indicate that there are no major discrepancies between the data and MC.

For this study, the entire Far Detector Run I and Run IIa datasets were examined (1.27×10^{20} PoT for Run I and 1.23×10^{20} PoT for Run IIa) and compared with a MC sample equivalent to 3.41×10^{23} PoT. Each MC event was weighted by the PoT weighted mean of its Run I and Run IIa extrapolation weights. Only selected CC ν_μ events were considered and the data distributions were compared with both the MC predictions assuming no oscillations and the MC predictions assuming typical oscillation parameters of $\Delta m^2_{32} = 2.4 \times 10^{-3}\text{eV}^2$ and $\sin^2(2\theta_{23}) = 1.0$.

8.2.1 PID related quantities

Figure 8.8 shows the distributions of the PID input variables and the resulting PID parameter for Far Detector data and MC. As expected, the predictions in the absence of oscillations disagree with the data in both shape and normalisation. However, a comparison between the data and the oscillated predictions produces χ^2/dof values that indicate consistency between the data and the expectation. There is no indication of any large disagreements that would need to be addressed in the oscillation analysis.

8.2.2 Track related quantities

Figure 8.9 shows the number of tracks per selected event in Far Detector data and MC. There is close agreement between the data and the prediction assuming typical oscillation parameters. Figure 8.10 compares the distributions of track vertex and end positions in data and MC. The data distributions all prove to be consistent with the oscillated MC expectation.

Figure 8.11 shows the distributions of track vertex positions and track end positions for just the selected data events. It can be seen that the track vertex positions are spread quite uniformly across the x and y coordinates in the Far Detector. By contrast, the track end positions show the result of focussing in the detector magnetic field, with the events clustering around the coil hole. Some track end positions are also located along the edge of the detector. Examination of these events suggests that they are due to μ^- that pass out of the detector despite the magnetic field, although it is plausible that some are due to the anti-focussing of incorrectly selected μ^+ tracks.

Figure 8.12 shows the distributions of track strips, track length, track pulse height per

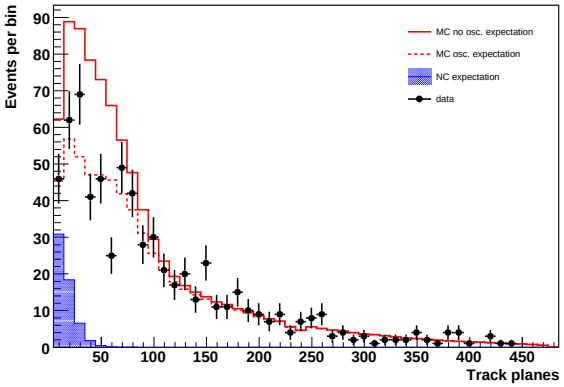

(a) $\chi^2/dof = 55.1/48$

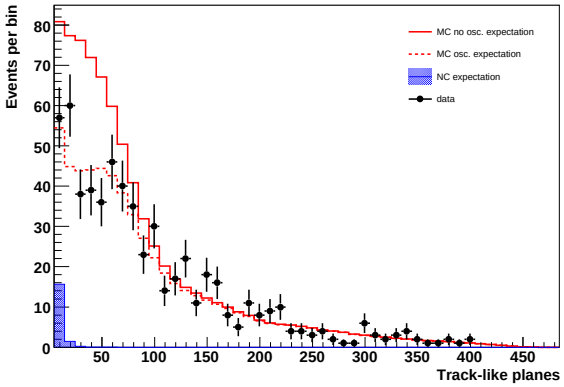

(b) $\chi^2/dof = 47.7/47$

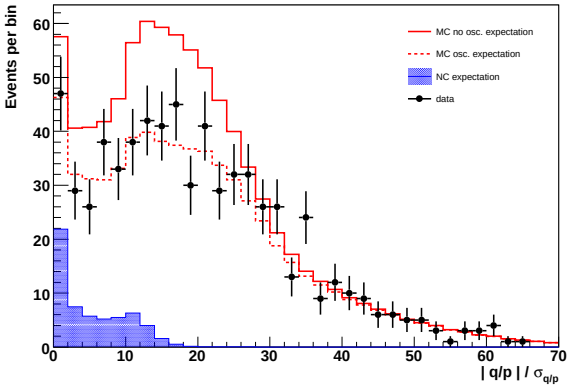

(c) $\chi^2/dof = 28.2/35$

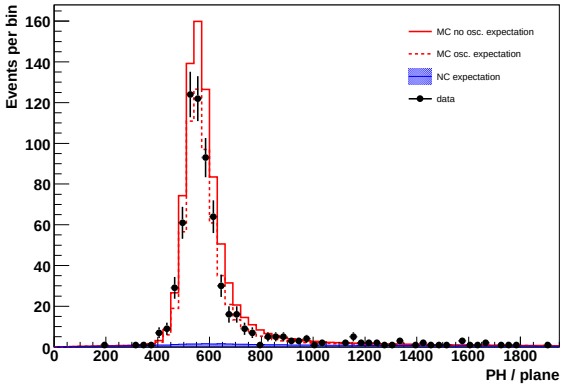

(d) $\chi^2/dof = 61.1/64$

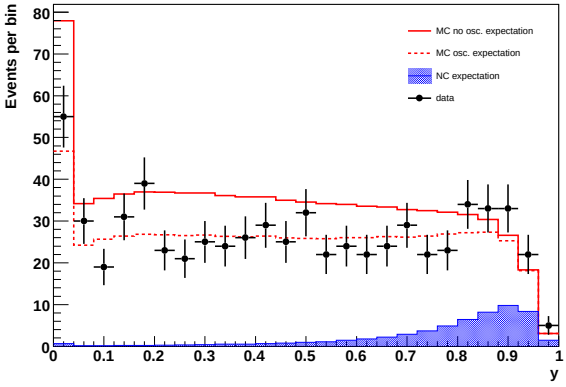

(e) $\chi^2/dof = 25.0/25$

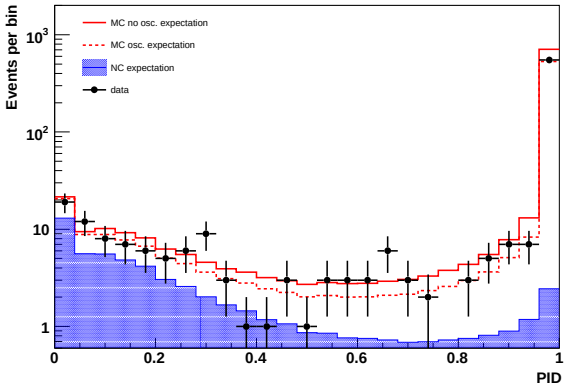

(f) $\chi^2/dof = 26.5/25$

Figure 8.8: Comparison between Far Detector data and MC for the PID parameter and its input variables: (a) track planes, (b) track-like planes, (c) goodness of muon track fit, (d) mean pulse height on track-like planes, (e) reconstructed y, and (f) PID parameter. The χ^2/dof values compare the agreement between the oscillated predictions and the data.

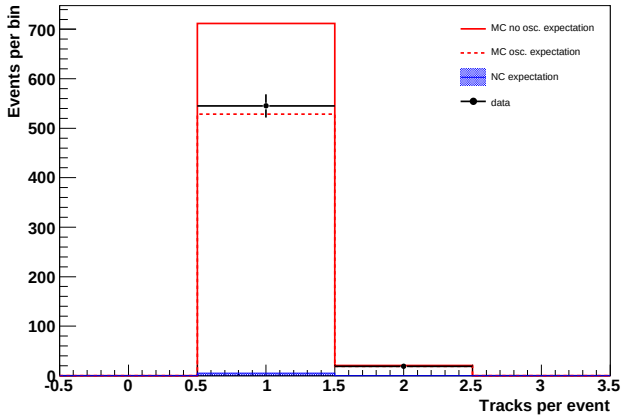

Figure 8.9: The number of reconstructed tracks observed in selected Far Detector data and MC events.

plane, track direction with respect to the beam axis and the reconstructed muon energy. Separate muon energy distributions are shown for contained events, for which the energy is reconstructed from track range, and for uncontained events, for which the energy is reconstructed from track curvature in the magnetic field. For all the distributions, the χ^2/dof values indicate that the data is consistent with the oscillated MC expectation.

8.2.3 Shower related quantities

Figure 8.13(a) shows the number of reconstructed showers in selected Far Detector data and MC events. The data are consistent with the oscillated MC expectation and there is no evidence of the excess of single shower events observed in the Near Detector data. Comparing the number of reconstructed showers at the Near and Far Detectors (Figures 8.6(a) and 8.13(a)) shows that more showers are typically reconstructed at the Far Detector. This feature is well simulated and can be attributed to the reduced instrumentation in the Near Detector; for instance, no showers are identified in the Near Detector spectrometer region.

The remaining distributions in Figure 8.13 show the number of shower strips, the shower pulse height per strip and the reconstructed shower energy. The χ^2/dof values indicate that the data distributions are consistent with the oscillated MC expectation. There is no indication of the discrepancies observed at the Near Detector, which involved an excess of low pulse height showers in the data.

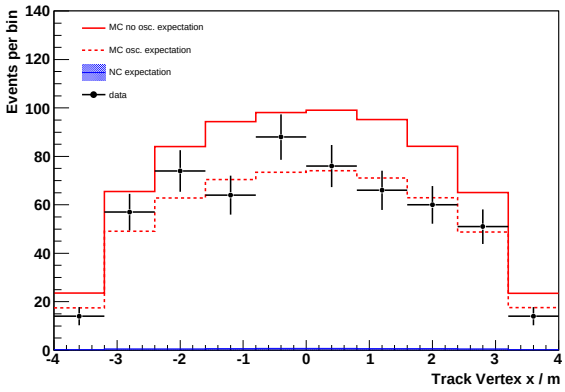

(a) $\chi^2/dof = 8.8/10$

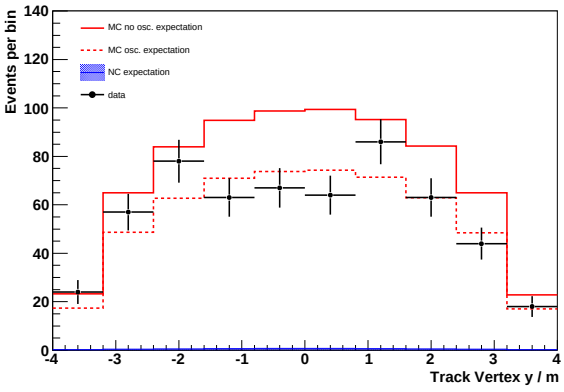

(b) $\chi^2/dof = 14.1/10$

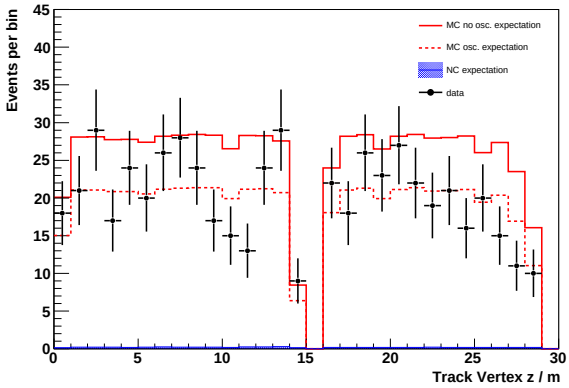

(c) $\chi^2/dof = 27.9/28$

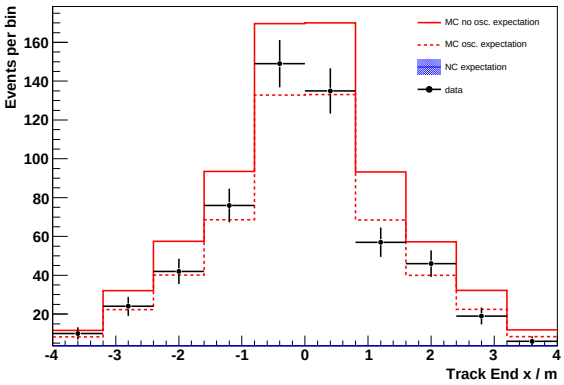

(d) $\chi^2/dof = 7.4/10$

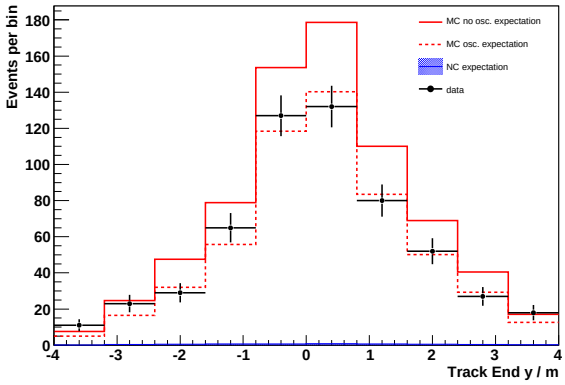

(e) $\chi^2/dof = 15.2/10$

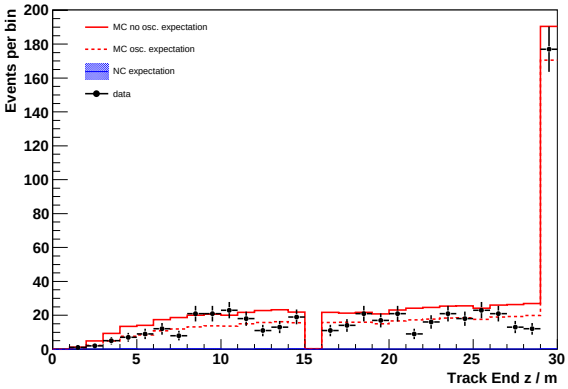

(f) $\chi^2/dof = 37.2/30$

Figure 8.10: The distributions of track vertex and end positions for selected Far Detector data and MC events: (a) vertex x coordinate, (b) vertex y coordinate, (c) vertex z coordinate, (d) end x coordinate, (e) end y coordinate, and (f) end z coordinate. The χ^2/dof values compare the agreement between the oscillated predictions and the data.

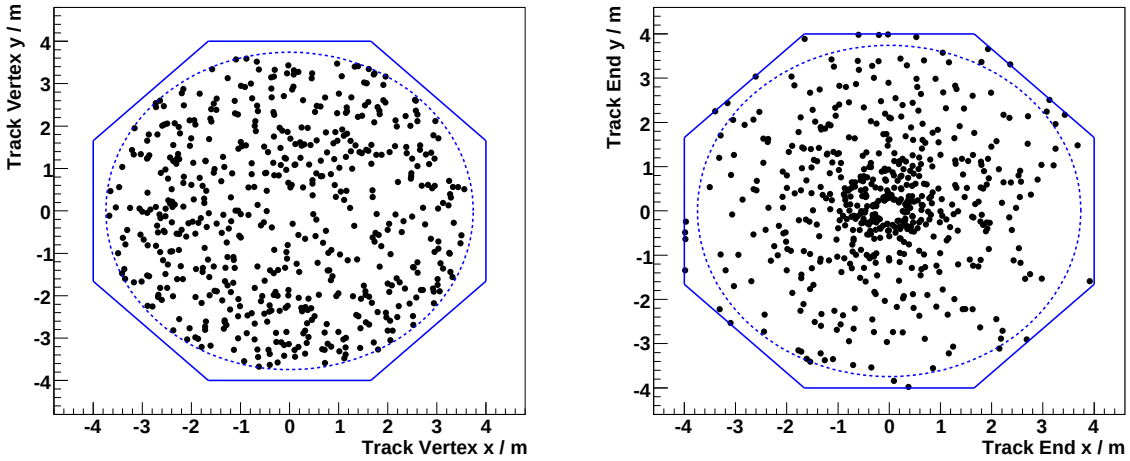

(a) (b)

Figure 8.11: The distributions for selected data events of (*a*) track vertex positions, and (*b*) track end positions across the face of the Far Detector. The dotted lines indicate the edge of the fiducial volume.

8.3 Conclusions

A comparison of the Near Detector data and MC reveals remarkable agreement, given that MINOS is a neutrino experiment. A small excess of single track events was observed in the Near Detector data, relative to the MC expectation, and some minor discrepancies between data and MC were observed in the PID input distributions. The largest discrepancy observed was an excess of low energy showers in the data. These excess showers contained small numbers of low pulse height strips. The reason for poor simulation of these showers is unclear, but is plausibly related to different thresholds for shower strip reconstruction in data and MC. This complex issue involves uncertainties in the modelling of intra-nuclear rescattering and the detector simulation, in addition to pathologies in the calibration and reconstruction software. In the Far Detector, the small data sample prevented a very detailed comparison of data and MC. However, for all the distributions examined, the data were found to be consistent with the MC expectation for typical oscillation parameters.

This study was intended to identify any major problems with the simulation, which would not be addressed by the treatment of systematic uncertainties developed in Chapter 7 and which could affect the results of an oscillation analysis. No such problems have

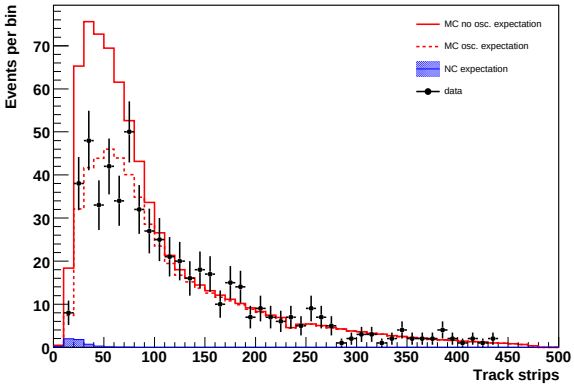

(a) $\chi^2/dof = 36.8/49$

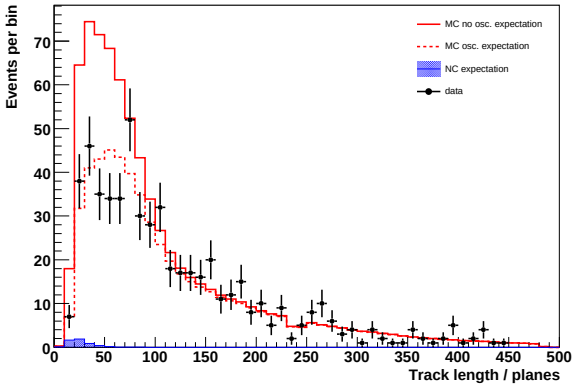

(b) $\chi^2/dof = 54.7/49$

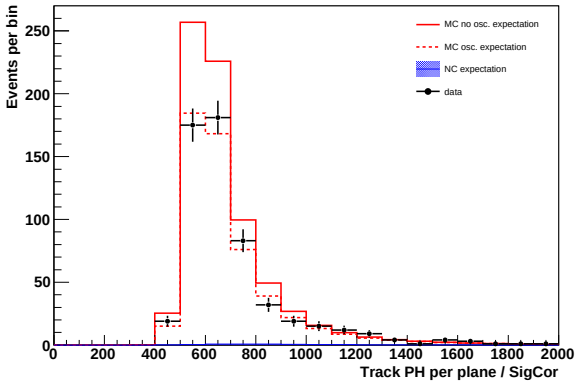

(c) $\chi^2/dof = 13.2/18$

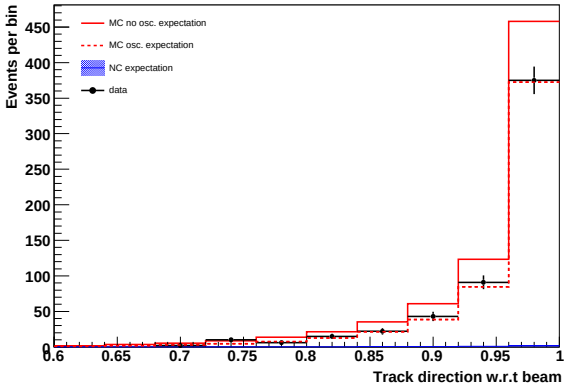

(d) $\chi^2/dof = 10.7/10$

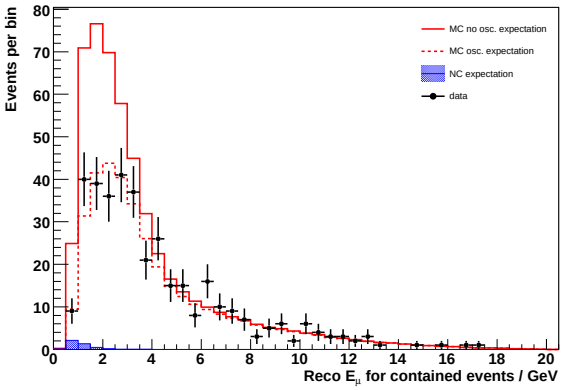

(e) $\chi^2/dof = 26.4/41$

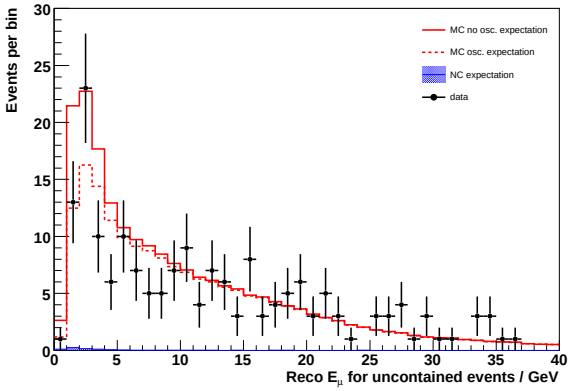

(f) $\chi^2/dof = 45.1/40$

Figure 8.12: Comparison between Far Detector data and MC: (a) number of track strips, (b) track length, (c) track pulse height per plane, (d) track direction cosine with respect to the beam, (e) reconstructed muon energy for contained tracks and, (f) reconstructed muon energy for uncontained tracks. The χ^2/dof values compare the agreement between the oscillated predictions and the data.

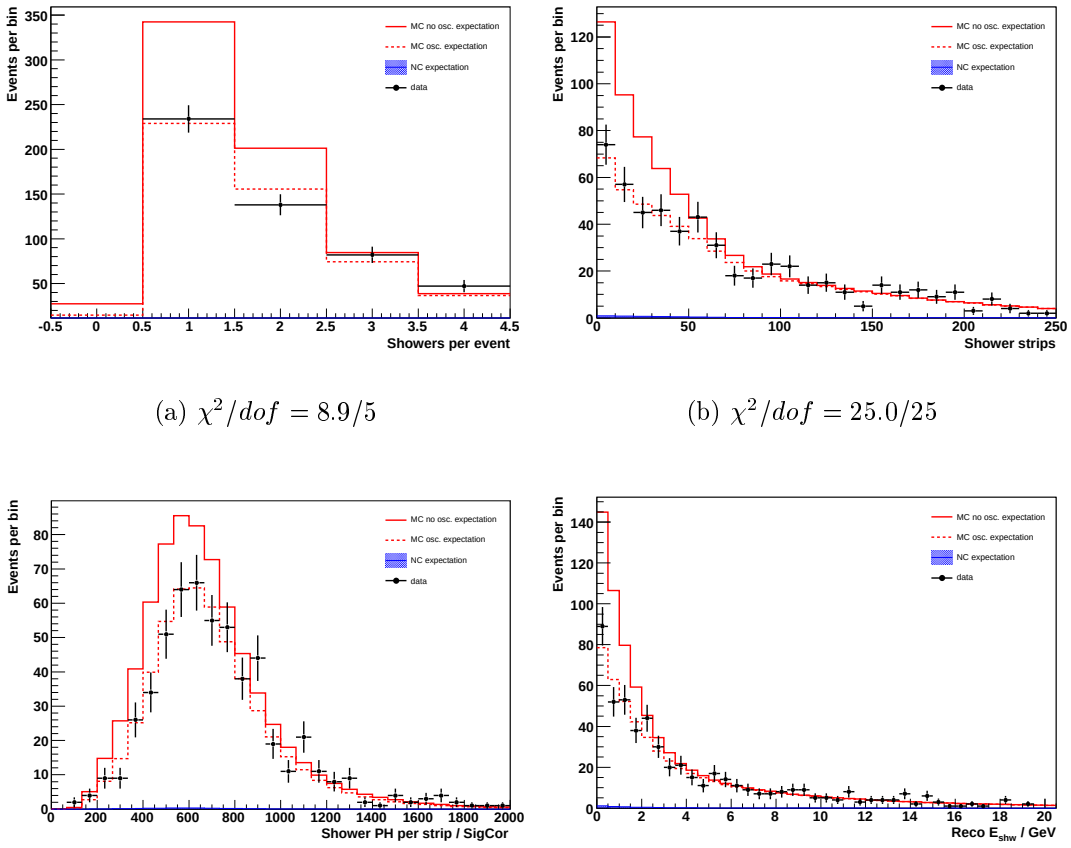

(a) $\chi^2/dof = 8.9/5$

(b) $\chi^2/dof = 25.0/25$

(c) $\chi^2/dof = 56.4/30$

(d) $\chi^2/dof = 60.6/60$

Figure 8.13: Comparison between Far Detector data and MC: (a) number of showers per event, (b) number of shower strips, (c) shower pulse height per strip, and (d) reconstructed shower energy. The χ^2/dof values compare the agreement between the oscillated predictions and the data.

been identified. The combined effect of the discrepancies in the Near Detector PID input distributions proved to have minimal effect on a MC test extrapolation and oscillation fit, whilst the excess of Near Detector data events will be addressed by the normalisation nuisance parameter ($\sigma = 4\%$) in the oscillation fit. The distortion of the Near Detector shower energy distribution, due to excess low energy showers, will be addressed by the shower energy scale nuisance parameter ($\sigma = 10\%$) and the relative shower energy scale nuisance parameter ($\sigma = 3\%$). The conclusion is that the few discrepancies observed between the data and MC will have negligible impact on the analysis results.

Chapter 9

Sensitivity Improvements

This Chapter discusses the techniques that can be used to extract the maximum amount of information about the oscillation parameters from a fixed data sample. These techniques can be broadly divided into the following categories:

- **Improving the neutrino energy reconstruction.** For optimal sensitivity to the oscillation parameters, a clear and sharp oscillation minimum (or 'dip') needs to be resolved. If the energy reconstruction smears the distribution too much, or reconstruction problems cause many events to be reconstructed from the energy peak down into dip, then the sensitivity will be reduced.

- **Identifying events with poor energy reconstruction.** This may simply involve vetoing classes of events that are likely to be poorly reconstructed. This is most applicable if serious reconstruction pathologies are identified. However, this category also includes more sophisticated approaches, in which events likely to be well reconstructed are assigned a larger weight in the oscillation fit than events that are likely to be poorly reconstructed.

- **Optimising the oscillation fit.** The oscillation fit needs to quantify how closely the Far Detector data spectrum matches the spectra predicted for different oscillation parameters. The spectra used in the fit must therefore be finely binned so that the distributions are not smeared and clear distinctions can be made between the predictions with different oscillation parameters.

Each technique aims to reduce the size of the sensitivity contours, which are the mean (68% or 90%) confidence level contours obtained from fitting a large number of independent data samples. These contours can be evaluated by using many MC 'fake data' samples at the specified PoT exposure. The predicted spectra are unchanged in each fit, but statistical fluctuations in the data samples mean that the best fit points and confidence level contours will display some variation.

A technique that reduces the size of the sensitivity contours allows the oscillation parameters to be extracted with greater precision from a typical data sample. If the technique does not introduce significant systematic uncertainties, it can be advocated for use in the oscillation analysis. In this Chapter a number of methods to improve the sensitivity are evaluated. These methods are then combined together and the expected increase in sensitivity is evaluated for the 2.5×10^{20} PoT MINOS dataset.

9.1 Binning in the Oscillation Fit

The first sensitivity improvement is simply to use fine binning for the reconstructed energy spectra used in the oscillation fit. The standard MINOS analysis[102] uses 1 GeV bin widths to cover the peak of the energy distribution and then successively wider bins to cover the tail. However, given that the energy resolution in the $1-2$ GeV region is typically 200 MeV, this can be considered as quite coarse binning. Coarse binning conceals information that could be used to distinguish between the predictions for two different sets of oscillation parameters. At the other extreme, it is unwise to use bin widths narrower than the typical energy resolution, and the number of bins also affects the memory requirements in the current implementation of the oscillation fit.

Sensitivity contours were obtained using the standard MINOS binning and two new binning configurations, which covered the spectrum up to 30 GeV. The region up to 30 GeV can be considered as that in which the spectral distortion caused by the effects of oscillations is non negligible. Above 30 GeV, an overflow bin was used to ensure that normalisation information was used correctly in the fit. The two new configurations used 0.50 GeV bin widths and 0.25 GeV bin widths. The standard MINOS binning closely corresponds to using 1.0 GeV bin widths and an overflow bin. The sensitivity contours obtained with typical input oscillation parameters of $\Delta m^2_{32} = 2.4 \times 10^{-3}\,\mathrm{eV}^2$ and $\sin^2(2\theta_{23}) = 1.0$ are shown in Figure 9.1.

An increase in sensitivity is seen when moving from the standard MINOS binning to 0.50 GeV bin widths and a further improvement is seen when the bin widths are reduced again to 0.25 GeV. This is as expected, suggesting that smaller bin widths allow

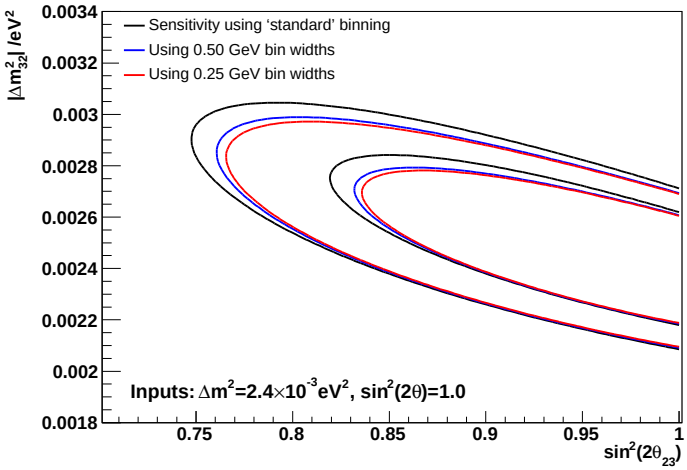

Figure 9.1: The mean 68% and 90% confidence level contours obtained using different binning for the energy spectra in the oscillation fit.

better resolution of the oscillation dip and better distinction between the predictions for different oscillation parameters. By moving from 1 GeV binning to 0.25 GeV binning the sensitivity to the oscillation parameters improves by about 11%.

For all oscillation fits performed in this thesis 0.25 GeV bin widths are used. This includes all the example fits in Chapter 7 and all the results marked as pre-sensitivity improvements in Chapter 10.

9.2 Deweighted Shower Energy

The next technique for improving the sensitivity is to optimise the hadronic shower energy reconstruction in the Far Detector. Shower strips are identified during the standard reconstruction and the detector response to a shower is typically characterised by the total charge deposited in these strips. This response can be calibrated to provide an estimate of the shower energy in GeV. However, this method of energy reconstruction is strongly affected by Landau fluctuations, which allow the energy deposited by a particle to vary considerably from the mean value. To address this problem, a new approach to hadronic calorimetry can be used[89, 103], which increases the importance of the number of shower strips over the charge that is actually deposited in the strips.

9.2.1 Detector Response to Hadronic Showers

The response of the detector to a hadronic shower is calculated in Minimally Ionising Particles (MIP), as defined in Section 3.3. For the purposes of neutrino energy recon-

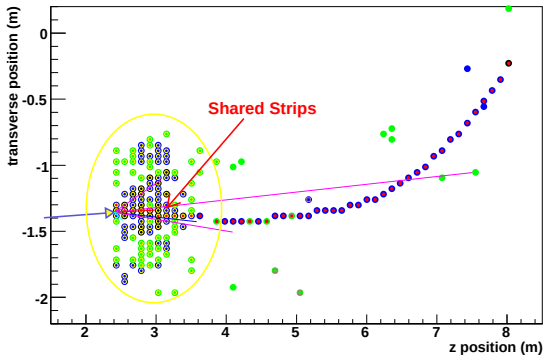

Figure 9.2: Event display for a typical CC ν_μ interaction, highlighting the strips that are shared between the reconstructed track and the vertex shower.

struction it is the response to the vertex hadronic shower that must be calculated. For a given event, the vertex shower can be identified by examining the proximity of different showers to the track vertex and by considering their charge deposits. Each strip in the vertex shower has an associated record of the pulse height observed at the strip ends, calibrated to MIP values. If a strip is not shared with the track, these MIP values can simply be added to a total MIP value that represents the detector response.

If a strip is shared by the vertex shower and the track, as illustrated in Figure 9.2, the muon and hadronic contributions must be carefully separated. The Bethe-Bloch equation [27] is used to estimate the charge deposited by the muon and this charge is subtracted from the strip MIP value. After track subtraction, the remaining MIP value is added to the total detector response. The final sum of (track-subtracted) MIP values should accurately represent the total detector response to the shower and can be calibrated to estimate the shower energy in GeV.

9.2.2 Linear Calibration

To convert the detector response into an estimate of the shower energy, a large sample of selected Far Detector MC events is required. For each event the detector response to the vertex shower is calculated. These responses are then divided into bins of true shower energy and, in each bin, a Gaussian fit is used to determine the mean and standard deviation. A 2σ cut restricts the fits to the centre of the distributions and removes outlying MIP values. Figure 9.3 plots the true shower energy value against the mean MIP value for each bin. The true shower energy is observed to vary almost linearly with the mean detector response. Best results are obtained using a linear fit above 100 MIP and a quadratic fit (constrained to pass through zero) below 100 MIP. This 'linear' calibration essentially reproduces the standard MINOS shower energy.

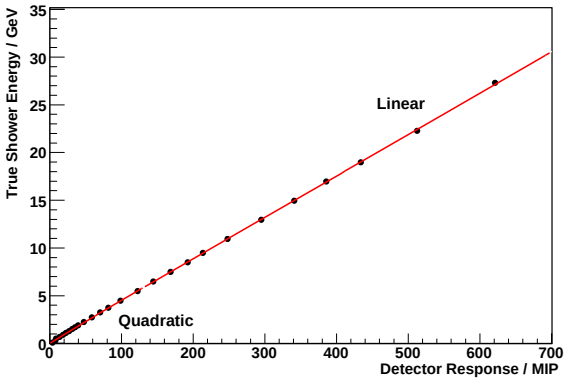

Figure 9.3: Mean detector response for different bins of true shower energy. The fitted polynomials provide a calibration of the detector response.

9.2.3 Deweighted Calibration

The linear calibration is sensitive to large tails in the distributions of shower pulse heights, which decrease the shower energy resolution at low energies. These tails are associated with Landau fluctuations and large energy deposits from stopping protons. One approach to this problem is to characterise the detector response by simply counting the number of shower strips, instead of using the total charge in the strips. This digital approach to hadronic calorimetry has a number of benefits[103], but is not optimal at high energies.

To obtain the optimal shower energy resolution at all energies, each shower MIP value can be deweighted (raised to a power between 0 and 1, as shown in Equation 9.1) before it is added to the total detector reponse. This deweighting has the effect of altering the relative importance of the number of shower strips over the charge deposited in the strips; using a deweighting power of zero is equivalent to counting the strips, whilst using a power of one simply adds the total strip charges. Importantly, the deweighting power can be varied as a function of energy and an optimal deweighting power can be parameterised as a function of the linear shower energy.

$$\text{Detector Response} = \sum_{\text{Strips}} Q^{\alpha(E)} \tag{9.1}$$

where Q is the track-subtracted MIP value for a shower strip and $\alpha(E)$ is the optimal deweighting power, which varies with the linear shower energy, E.

A deweighted shower energy calibration reduces the sensitivity to systematic errors in the MC modelling of low energy charge deposition. However, it increases the sensitivity to systematic differences between the strip reconstruction thresholds in data and MC.

If a deweight power of zero is used (i.e. counting strips), a strip just above threshold has equal importance to a strip representing a large amount of charge. A change in a quantity such as the PMT gain could therefore have a large systematic effect on the deweighted shower energy. For this reason, the minimum deweighting power considered in this thesis is 0.25.

To examine the variation of the optimal deweighting power with energy, fixed deweighting powers of 0.25, 0.50 and 0.75 were considered. A power of 1.0 (the linear calibration) was also considered and was expected to be optimal at high energies, where fluctuations are less important. For each deweighting power, the detector response was calculated and calibrated (in the same way as for the linear calibration). The shower energy resolution, σ_E/E, was then evaluated for each deweighting power by fitting a Gaussian to the $(E_{shw}^{Reco} - E_{shw}^{True})/E_{shw}^{True}$ distribution.

Figure 9.4(a) shows the shower energy resolution for each deweighting power as a function of the linear shower energy. There are substantial variations in the resolution achieved with different deweighting powers. As expected, low deweighting powers are optimal at low values of the linear shower energy. There is then a smooth increase in the optimal deweighting power until the best resolution values are obtained with no deweighting, at linear shower energies of greater than 14 GeV. This variation of the optimal deweighting power can be smoothly parameterised as a function of the linear shower energy using a sigmoid function, as shown in Figure 9.4(b).

The parameterisation of the optimal deweighting power enables optimal deweighted detector responses to be calculated for Far Detector CC ν_μ events. These responses can be calibrated by dividing the events into bins of true shower energy, calculating the truncated mean response in each bin and fitting polynomials to the observed variation. As shown in Figure 9.5, the true shower energy is found to initially vary as a fourth order polynomial of the optimal deweighted detector response. At higher optimal deweighted responses, the variation reverts to the same linear parameterisation observed in the linear shower calibration. This is as expected, as there is no deweighting at large values of the linear calibrated shower energy.

9.2.4 Results

Figure 9.6(a) shows that the deweighted calibration considerably improves the shower energy resolution at low energies; for 1 GeV showers the resolution improves by 22%. Figure 9.6(b) shows that the parameterisation of the optimal deweighting power with linear calibrated energy works effectively. At all energies, the deweighted calibration

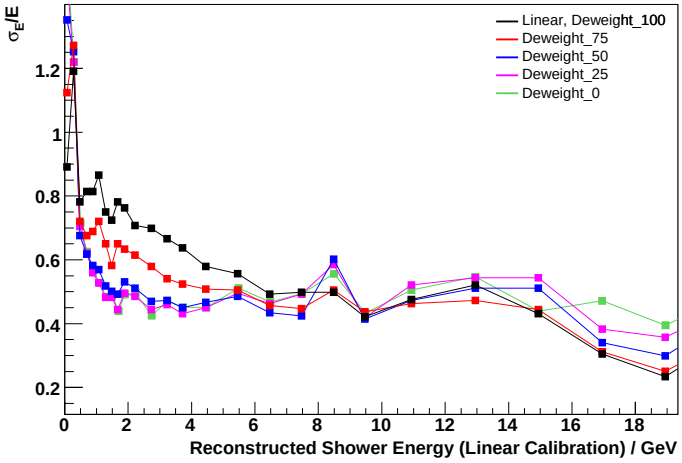

(a)

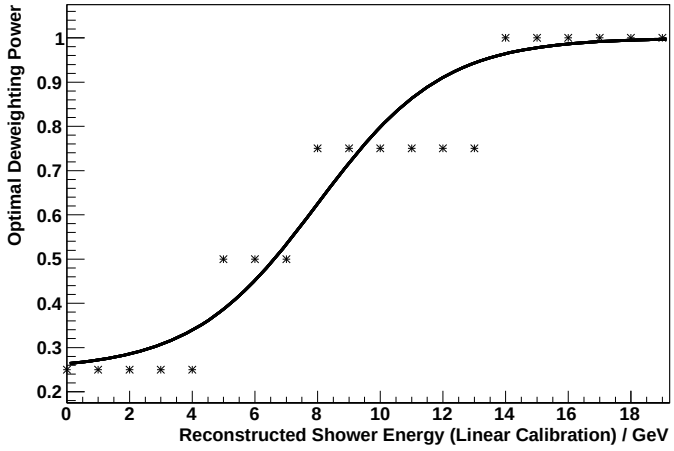

(b)

Figure 9.4: (a) Variation of the shower energy resolution with the linear reconstructed shower energy for a number of different deweight powers. (b) The variation of the optimal deweighting power with the linear reconstructed shower energy. This variation can be smoothly parameterised, using a sigmoid function.

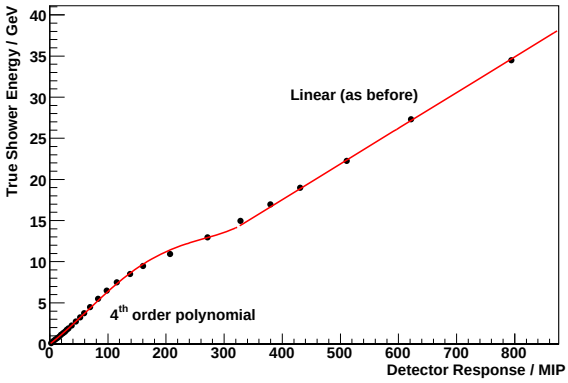

Figure 9.5: Mean deweighted detector response for different bins of true shower energy. The fitted polynomials provide a calibration of the optimal deweighted detector response.

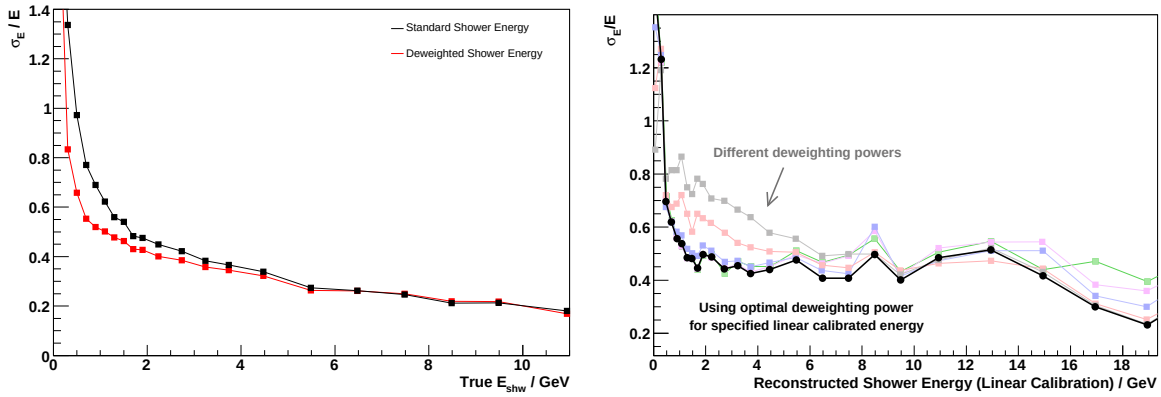

(a) (b)

Figure 9.6: (*a*) Shower energy resolution as a function of true shower energy. (*b*) Shower energy resolution as a function of linear reconstructed shower energy. The coloured lines show the resolutions obtained with different fixed deweighting powers. By varying the deweighting power, the deweighted shower energy (black line) achieves the optimal resolution at all energies.

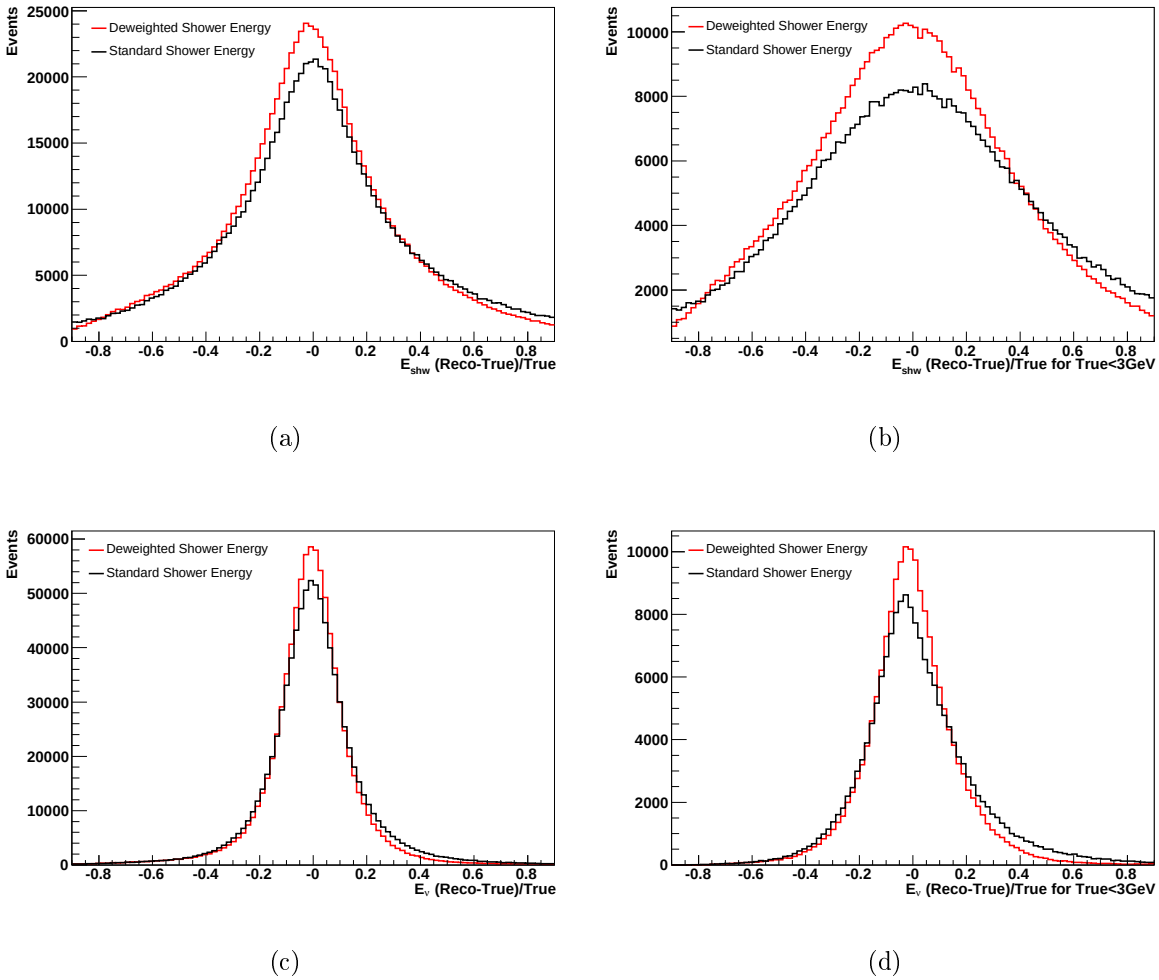

Figure 9.7: Performance comparison for the deweighted and linear shower energy calibrations, using selected Far Detector beam MC events. (*a*) and (*b*) show the fractional discrepancy in reconstructed shower energy for all events and low energy events respectively. (*c*) and (*d*) show the fractional discrepancy in reconstructed neutrino energy for all events and low energy events respectively.

matches the best possible resolution achieved with the different fixed deweighting powers. Figure 9.7 shows that using the final deweighted shower energy values improves the shower energy reconstruction and hence the neutrino energy reconstruction. The improvement is particularly large for low energy events, which are those that have the largest impact on the sensitivity.

Use of the standard shower energy reconstruction and the deweighted shower energy reconstruction results in the sensitivity contours in Figure 9.8. The deweighted shower energy improves the sensitivity to the oscillation parameters by about 5% and is therefore advocated for use in an optimised analysis of the MINOS data.

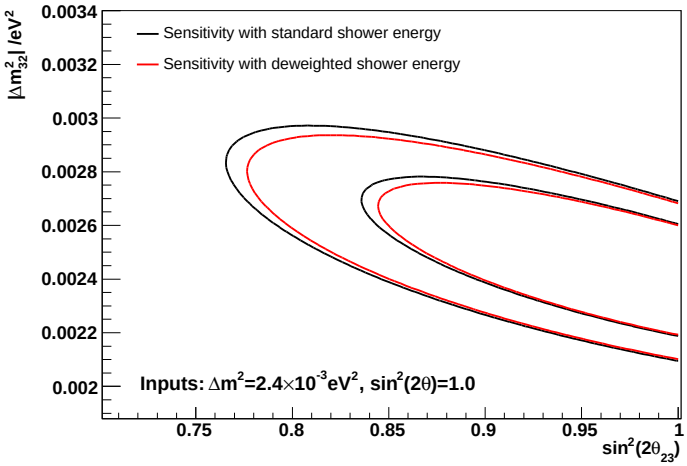

Figure 9.8: The mean 68% and 90% confidence level contours obtained using the deweighted and standard shower energy calibrations.

9.3 Use of Event Energy Resolution

The next technique to improve the sensitivity is to try to identify whether events have been reconstructed accurately and to incorporate this information into the oscillation fit. The important stage in this process is to obtain a useful energy resolution parameterisation, so that a resolution value can be assigned to each event, based on its reconstructed properties. The reconstructed neutrino energy is the sum of the reconstructed shower energy and the reconstructed muon energy. The latter may be reconstructed from track range or track curvature. The different nature of these contributions to the neutrino energy must be reflected by separate resolution parameterisations.

The deweighted shower energy resolution, σ_{Eshw}, can be parameterised as a function of the deweighted shower energy, whilst, for fully contained tracks, the muon energy resolution, $\sigma_{E\mu}$, can be parameterised as a function of the track range energy. For partially contained tracks $\sigma_{E\mu}$ can be parameterised as a function of $p^2\sigma_{q/p}$, where $\sigma_{q/p}$ is the Kalman Filter estimate of the error in q/p and the factor of p^2 then allows an estimate of the error in p. These parameters are reconstructed quantities that strongly scale with the expected error in the shower or muon energy measurements. Once values are obtained for σ_{Eshw} and $\sigma_{E\mu}$, the resolution for the neutrino energy reconstruction can be evaluated as $\sigma^2_{E\nu} = \sigma^2_{E\mu} + \sigma^2_{Eshw}$.

Energy resolution parameterisations were obtained by applying Gaussian fits to the $(E_{reco} - E_{true})$ distributions obtained for different bins of the relevant parameter. The resulting Gaussian widths were plotted as a function of the parameter and an appropriate function fitted using MINUIT. The decision to fit the $(E_{reco} - E_{true})$ distribution, rather

than the more typical $(E_{reco} - E_{true})/E_{true}$ distribution, was made in order to ensure a more stable parameterisation at low energies.

The parameterisation used for the deweighted shower energy resolution is shown in Figure 9.9(a). This parameterisation takes the functional form of $\sigma_{Eshw} = A\sqrt{E} + BE$, with A and B constants. The parameterisation for muon energy from range is shown in Figure 9.9(b). This Figure shows a linear dependence of the resolution with the muon range energy. The parameterisation reflects this linear dependence and is consistent with the 5% error expected for energy from range measurements. The final parameterisation is for muon energy from curvature. This parameterisation is shown in Figure 9.9(c) and takes the form of $\sigma_{E\mu} = C\sqrt{p_{err}} + Dp_{err}$, where C and D are constants and $p_{err} = p^2\sigma_{q/p}$

Using the resolution parameterisations, every event can be assigned a $\sigma_{E\nu}$ value. The quantity $\sigma_{E\nu}/E_\nu$ can then be used to divide events into bins of resolution. This division is performed by examining the $\sigma_{E\nu}/E_\nu$ distribution *for each bin of reconstructed neutrino energy* and identifying the quantiles of this $\sigma_{E\nu}/E_\nu$ distribution. These quantiles act as the divisions between the different resolution bins. For instance, with five bins of resolution, the events in the lowest quantile of $\sigma_{E\nu}/E_\nu$ can be considered as the 20% of events with the best energy resolution. The events with high $\sigma_{E\nu}/E_\nu$ values, in the highest quantile, can be considered as the 20% of events with the worst energy resolution.

There is a clear difference between the oscillation dips observed for events identified as the 20% with best energy resolution and the 20% with worst energy resolution, as is shown in Figure 9.10. The oscillation dip observed for the best resolution events is much deeper and sharper than that for the worst resolution events. There is also a small shift in the minimum (away from the expected position) for the worst resolution events, which is associated with smearing of the steeply falling distribution. Incorporating resolution information into the oscillation fit will therefore bring about an improvement in sensitivity.

Resolution information can be included in the oscillation fit by carrying out separate shape fits for each resolution bin and an overall normalisation fit. This requires that predictions of the Far Detector data spectrum are made for each resolution bin and that the actual Far Detector data sample is divided into the resolution bins. The likelihood function used in the fit is:

$$-\ln\mathcal{L} = N^{MC} - N^{Data}\ln N^{MC} + \sum_{i}^{n_{Res}}\sum_{j}^{n_{Ebins}}(\mu_j^i - N_j^i\ln\mu_j^i) + \sum_{k}^{n_{Syst}}\frac{a_k^2}{2\sigma_k^2} \qquad (9.2)$$

where N^{Data} is the total number of events in the data sample and N^{MC} is the total

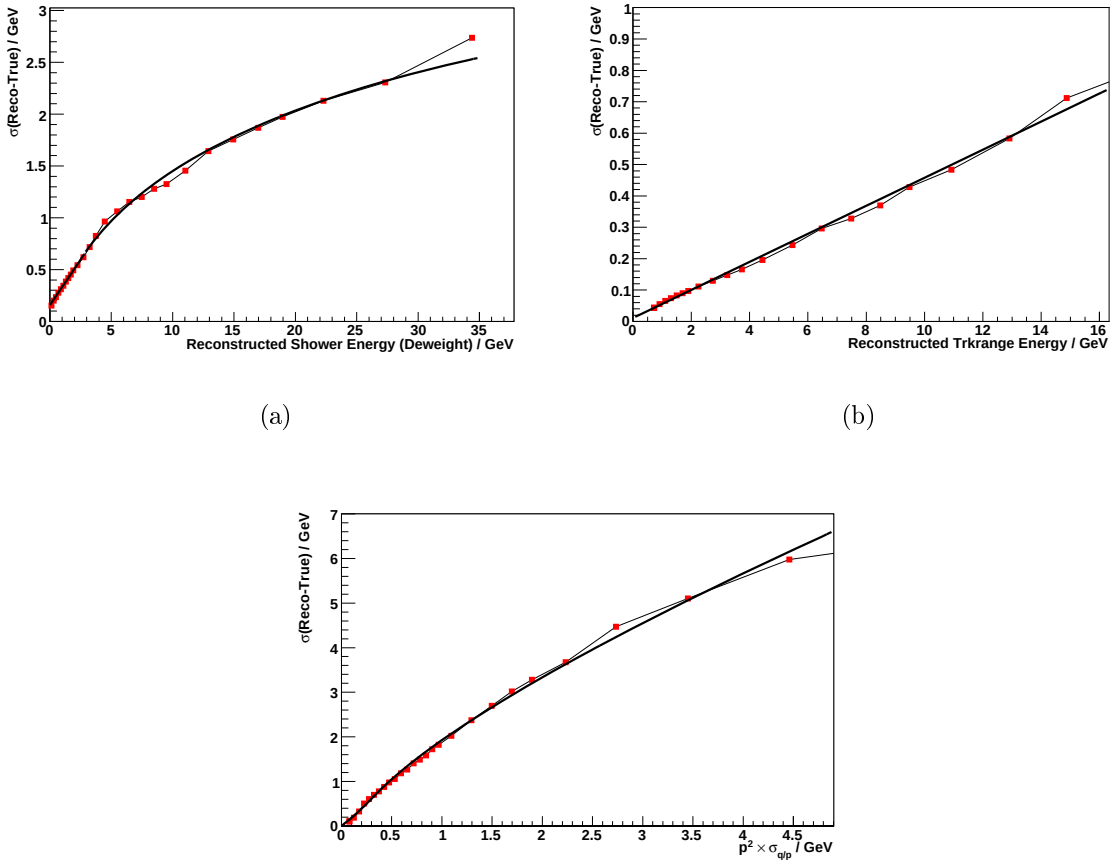

(a) (b)

(c)

Figure 9.9: Energy resolution parameterisations: (a) deweighted shower energy resolution as a function of the deweighted shower energy, (b) muon energy from range resolution (for contained tracks) as a function of track range energy, and (c) muon energy from curvature resolution (for uncontained tracks) as a function of $p^2\sigma_{q/p}$.

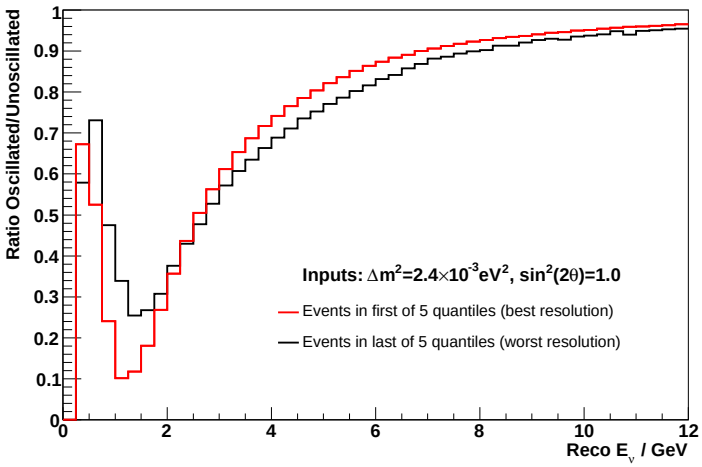

Figure 9.10: The ratios of oscillated to unoscillated neutrino energy spectra for events in the best and worst of the five energy resolution quantiles.

number of events in the MC prediction. N_j^i is the number of events in the j^{th} energy bin of the data spectrum for the i^{th} resolution bin. μ_j^i is the value of the j^{th} energy bin of the MC prediction PDF for the i^{th} resolution bin. a_k is the magnitude of the shift of nuisance parameter k from its nominal MC value, whilst σ_k is the typical uncertainty in this parameter.

Figure 9.11(a) shows that dividing events into five resolution bins and using this new likelihood function improves the sensitivity to the oscillation parameters by about 3%. Simply by making better use of the available event information, the oscillation parameters can be determined more accurately. The final decision is how many bins of resolution to use. The expectation is that more resolution bins will lead to larger sensitivity improvements. However, Figure 9.11(b) shows that with more than five resolution bins, the successive improvements observed when moving to higher numbers of bins become very small. In addition, it should be remembered that the number of resolution bins is a direct multiple of the memory and CPU time required in the oscillation fit. For this reason, a total of five resolution bins is advocated for use in an optimised analysis of the MINOS data.

9.4 Removal of Events in Coil Hole

The resolution parameterisations described above can also be useful for identifying classes of poorly reconstructed events. For instance, a sample of events with reconstructed shower energy more than $3\sigma_{Eshw}$ away from the true shower energy can be selected. These events can then be examined in order to identify any common features

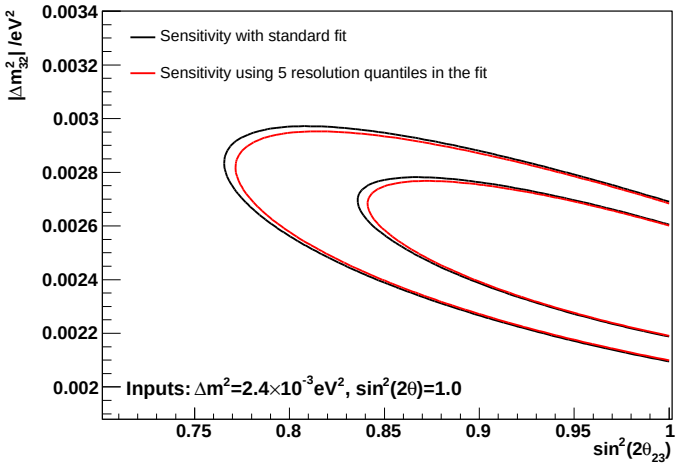

(a)

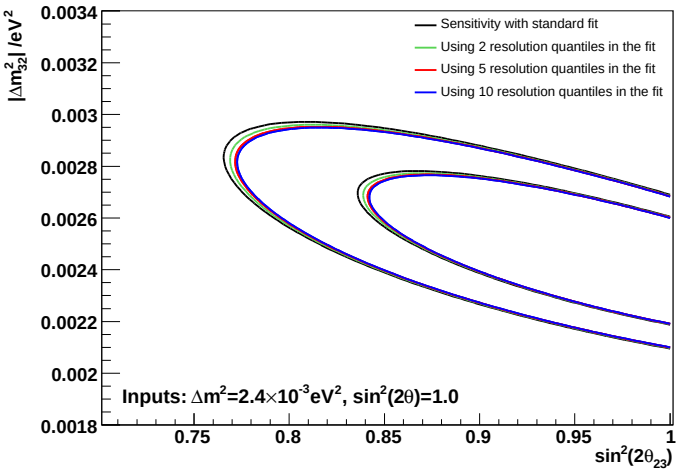

(b)

Figure 9.11: (a) The mean 68% and 90% confidence level contours obtained using a standard oscillation fit and using five resolution bins in the fit. (b) The changes in sensitivity observed with use of different numbers of resolution bins in the fit.

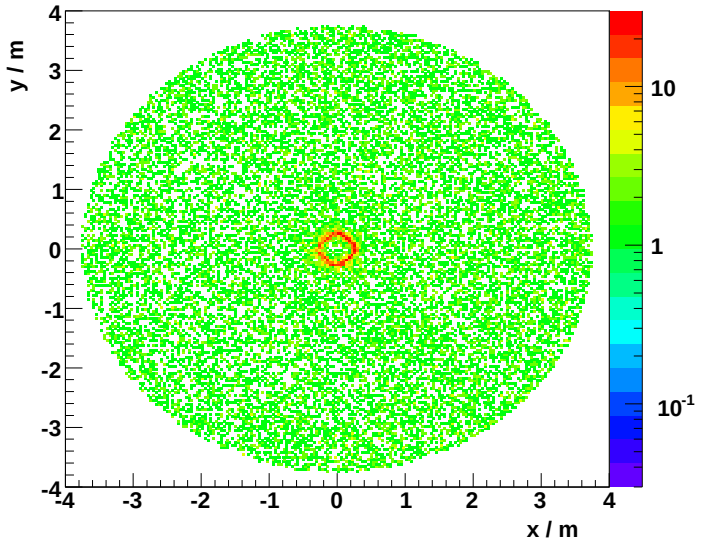

Figure 9.12: The distribution of track vertex x and y positions for events with poor shower energy reconstruction.

that could explain the poor shower energy reconstruction. Of particular interest is the position of these events in the detector, as a common reason for reconstruction failure is proximity to the supermodule gap, the outer edge of the detector or the coil hole.

A plot of track vertex position for the events with poor shower energy reconstruction is shown in Figure 9.12. This shows a large clustering of poorly reconstructed events around the detector coil hole, superimposed on an otherwise uniform distribution across the detector. The uniformly distributed events represent events from the tails of the Gaussian $(E_{reco} - E_{true})$ distribution, whilst the large number of events around the coil hole represents a genuine reconstruction pathology.

When a neutrino interaction occurs in (or nearby) the coil hole, a large amount of the shower pulse height can be lost in an uninstrumented region of the detector, and this will clearly harm the reconstruction of the shower energy. An example MC event display is shown in Figure 9.13. Here, the true neutrino interaction vertex is inside the uninstrumented coil hole region. Despite there being a 2 GeV shower in the generated event, the vertex shower is entirely contained in the coil hole and so no shower energy is reconstructed. This kind of reconstruction problem leads to an underestimate of the shower energy and so can incorrectly reconstruct events in the neutrino energy peak down into the oscillation dip.

The problem of poor reconstruction near the coil hole indicates an omission from the standard MINOS detector fiducial volume. In order to improve the sensitivity to the oscillation parameters, the best procedure is simply to remove events near the coil

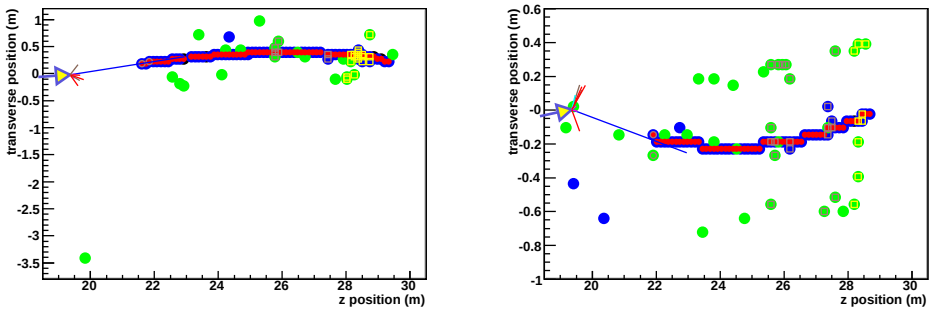

Figure 9.13: Event display showing a selected event with true interaction vertex (indicated by the blue arrow) inside the coil hole, leading to incorrect shower energy reconstruction. The red markers indicate strips in a reconstructed track.

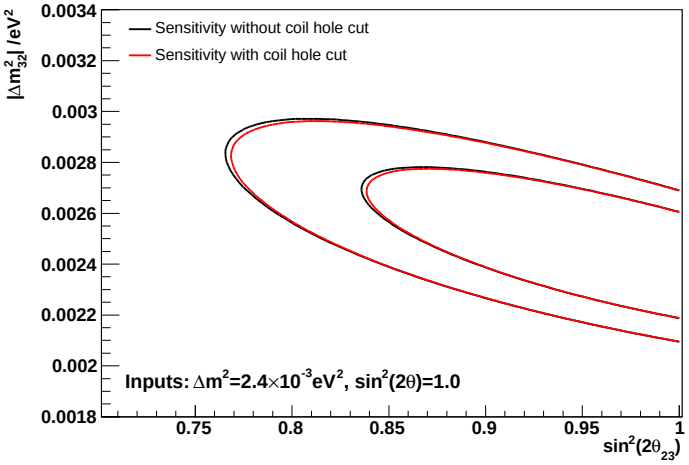

Figure 9.14: The mean 68% and 90% confidence level contours obtained with and without a cut on the proximity of the track vertex to the coil hole.

hole from the data and MC samples. An optimal coil hole cut is specified in [84] and demands that the track vertex is greater than 0.4 m from the centre of the detector.

Implementing the coil hole cut and evaluating the sensitivity produces the contours in Figure 9.14. Removing events near the coil hole produces a modest sensitivity improvement of about 2%. This improvement is consistent with the fraction of poorly reconstructed events that are removed by the cut, as tested by using random numbers to discard the same fraction of events with reconstructed shower energy more than $3\sigma_{Eshw}$ away from the true shower energy and evaluating the sensitivity. Use of a coil hole cut is advocated for use in an optimised analysis of the MINOS data.

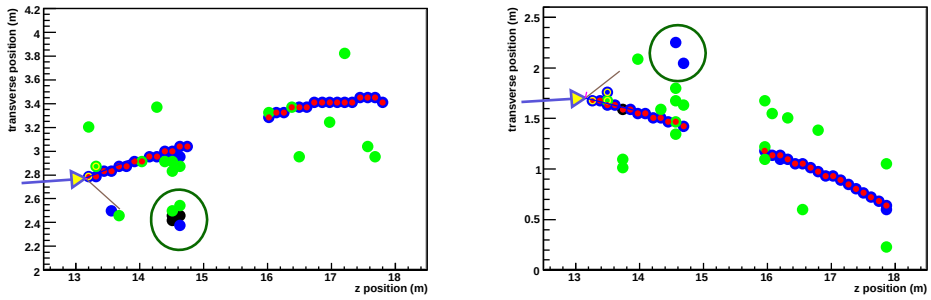

Figure 9.15: Event display showing a selected event in which several large pulse height strips (highlighted) are neglected by the shower energy reconstruction. The red markers indicate strips in a reconstructed track.

9.5 Pulse Height Omitted by Shower Finder

The final sensitivity improvement again addresses the issue of shower energy reconstruction. Examining the displays of events with poorly reconstructed showers, it was observed that poor reconstruction is often due to the incomplete identification of the shower strips by the shower reconstruction software. In particular, large pulse height strips, that should be part of the vertex shower, are often omitted. An example event display is shown in Figure 9.15. In this example, groups of large pulse height strips from a neutron interaction are not included in a shower. The strips are not included in either a track or a shower, so the energy deposited in these strips is completely neglected by the standard energy reconstruction. This clearly leads to an underestimate of the neutrino energy, reconstructing events down into the oscillation dip and decreasing the sensitivity.

In order to account for the energy deposited in the unassigned strips, the strips not assigned to either a track or a shower were identified in each event. Strips due purely to noise or cross-talk were rejected by using a pulse height threshold cut of approximately three photoelectrons. The pulse height from the remaining unassigned strips could then be added to a total unassigned pulse height value for the event. This value is measured in SigCor, so it can be considered as proportional to the deposited energy in GeV.

A plot of the total unassigned pulse height against the discrepancy between reconstructed shower energy and true shower energy is shown in Figure 9.16. The plot shows that there is a strong linear correlation between the discrepancy and the unassigned pulse height. This allows a correction to be made to the reconstructed shower energy. The correction is linear with unassigned pulse height and must always be above zero.

Applying an unassigned pulse height correction to the reconstructed shower energy

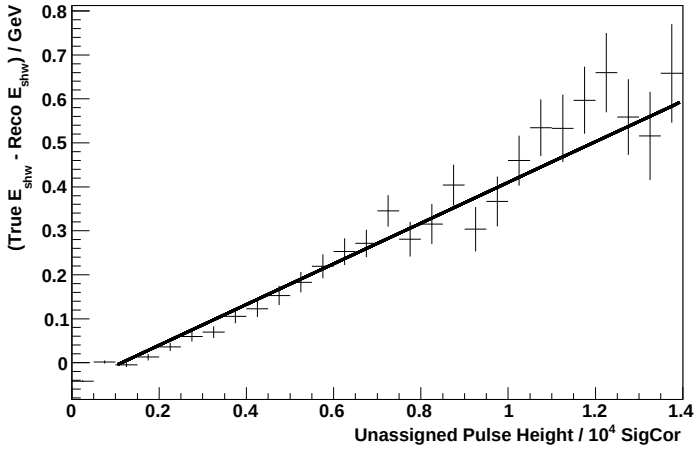

Figure 9.16: Mean discrepancy between the reconstructed and true shower energy *vs.* total unassigned strip pulse height above a threshold of three photo-electrons.

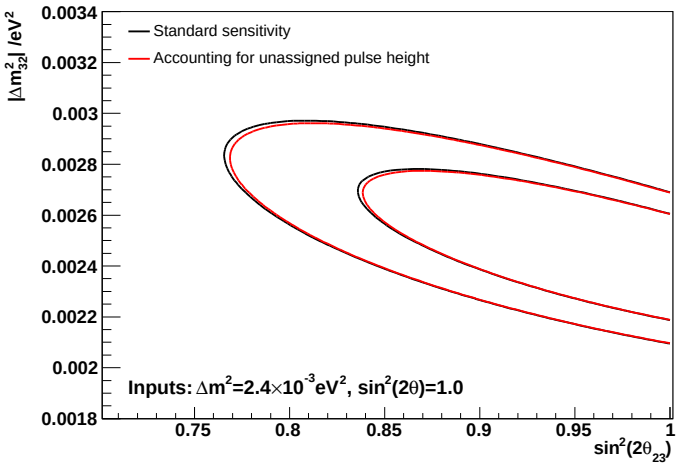

Figure 9.17: The mean 68% and 90% confidence level contours obtained before and after correction of the reconstructed shower energy to account for unassigned pulse height.

produces the contours shown in Figure 9.17. This figure shows that the correction improves sensitivity to the oscillation parameters by about 2%. Use of the unassigned pulse height correction is therefore recommended for use in an optimised analysis of the MINOS data.

9.6 Collective Improvement

The sensitivity improvements detailed above can be combined together to produce the contours shown in Figure 9.18. This Figure shows that, step by step, a very large and impressive increase in sensitivity is obtained. By making all five of the suggested improvements, the sensitivity to Δm^2_{32} improves by approximately 21%, whilst the sensitivity to $\sin^2(2\theta_{23})$ improves by approximately 23%. This can be considered as equivalent to increasing the Far Detector dataset by approximately 50%. Full details of the improvements in sensitivity are summarised in Table 9.1.

Configuration	$\mathbf{\Delta m^2_{32}}$ **1σ limits** $\mathbf{/10^{-3}eV^2}$	$\mathbf{\Delta m^2_{32}}$ **$\pm$1σ range** $\mathbf{/10^{-3}eV^2}$	$\mathbf{\sin^2(2\theta_{23})}$ **1σ limits**
Standard	2.255−2.688	0.433	≥0.8784
+ Fit in 0.25 GeV bins	2.261−2.643	0.382	≥0.8921
+ Deweighted shower energy	2.264−2.628	0.364	≥0.8984
+ 5 quantiles of resolution	2.266−2.620	0.354	≥0.9023
+ Coil hole cut	2.266−2.616	0.350	≥0.9045
+ Unassigned pulse height	2.267−2.611	0.344	≥0.9065

Table 9.1: Summary of the sensitivities observed with the successive introduction of the techniques described in this Chapter. Input values for MC fake data were $\Delta m^2_{32} = 2.4 \times 10^{-3}\,\mathrm{eV}^2$, $\sin^2(2\theta_{23}) = 1.0$

9.7 Associated Systematics

Before the sensitivity improvements can be used in an analysis of the MINOS data, it is important to check that they do not introduce any significant systematic errors.

Of particular concern is the potential for systematic errors associated with use of the deweighted shower energy. The idea of the deweighted shower energy calibration is to increase the importance of the number of shower strips and to reduce the importance of the charge deposited in the shower strips. This is very useful at low energies, as it reduces the impact of charge fluctuations and it reduces the systematic errors associated with poor MC modelling of charge deposition. However, the dependence on the number of shower strips means that the deweighted shower energy is strongly affected by changes in the threshold for strip reconstruction.

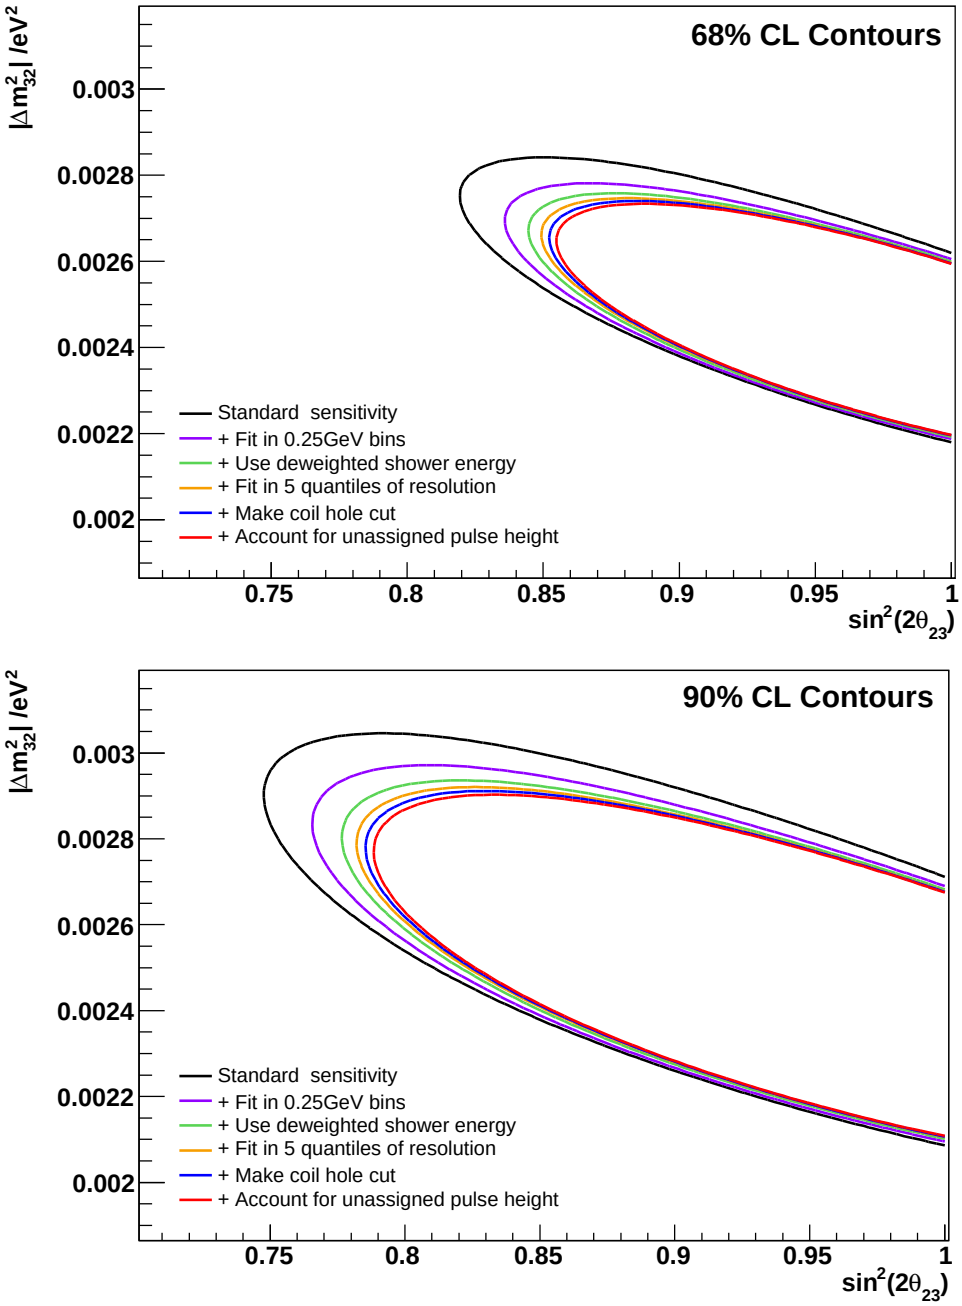

Figure 9.18: The improvements in sensitivity obtained with the successive introduction of the techniques described in this Chapter. Input values for MC fake data were $\Delta m^2_{32} = 2.4 \times 10^{-3}\,\mathrm{eV}^2$, $\sin^2(2\theta_{23}) = 1.0$. The contours account for statistical errors only.

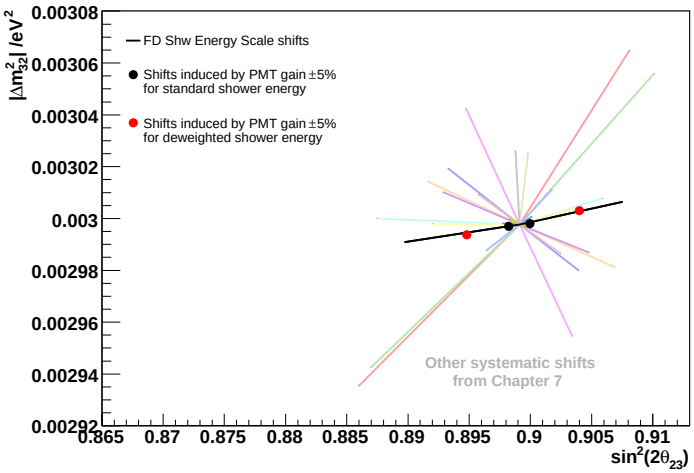

Figure 9.19: The changes in best fit parameters induced by ±5% systematic shifts in the PMT gains. Results are shown for oscillation fits using both standard and deweighted shower energy calibrations. The black line indicates the changes induced by ±3% shifts in the relative Far Detector shower energy scale.

Problems can occur, for instance, if the PMT gains are different in data and MC. Changes in the PMT gains affect the number of low pulse height strips in an event, altering the number of strips whilst leaving the total strip charge almost unchanged. Such changes therefore have a larger impact on the deweighted shower energy than on the standard shower energy.

To reduce the impact of different data/MC strip reconstruction thresholds, a minimum deweight power of 0.25 was used in the calibration. The remaining impact was then assessed by generating Far Detector fake data samples with ±5% shifts in the simulated PMT gains, as suggested in [104]. Oscillation fits were performed using both the standard and deweighted shower energies and the shifts in best fit parameters were examined. The results are shown in Table 9.2 and Figure 9.19. As expected, the shifts in PMT gains have a larger effect on the deweighted shower energy than the standard shower energy. However, the shifts in best fit parameters actually prove to be rather small and unimportant for both shower energy measurements. The shifts are much smaller than those observed for any of the six most important systematic uncertainties identified in Chapter 7.

Figure 9.19 also suggests that the effects of changing the PMT gains are very similar to the effects of changing the Far Detector shower energy scale. This is not too surprising, as the two quantities are closely related. The Far Detector shower energy scale is included as a nuisance parameter in the oscillation fit and so should also address any differences between strip reconstruction thresholds in the data and MC. These results

suggest that the deweighted shower energy can be safely used in the analysis; the effects of strip reconstruction thresholds actually prove to be quite small and should be directly addressed in the oscillation fit.

Configuration	**PMT Gain**	$\Delta(\Delta m^2_{32})\,/10^{-4}\,\mathrm{eV}^2$	$\Delta(\sin^2(2\theta_{23}))$
Standard shower energy	+5%	−0.011	−0.0008
	−5%	+0.003	+0.0008
Deweighted shower energy	+5%	−0.041	−0.0043
	−5%	+0.053	+0.0049

Table 9.2: The effects of systematic changes in the PMT gains, as observed in oscillation fits using the standard shower energy and the deweighted shower energy calibrations.

The next issue requiring examination is the division of the Far Detector data into a large number of energy bins and resolution bins during the oscillation fit. Whilst the use of resolution information and high frequency sampling of the energy distribution can increase the sensitivity, the risk is that using such detailed features could increase the impact of systematic differences between the data and MC. As a simple example, consider a normalisation fit. This has limited sensitivity, but is unaffected by distortions in the shower energy. Using information from the shape of the energy distribution greatly increases the sensitivity, but allows the shower energy distortion to affect the best fit parameters. There is a balance between sensitivity and potential systematic errors which must also be considered with use of resolution information and with the sampling of the energy distributions.

To assess whether there are problems with the binning configurations suggested for this analysis, a 30% smearing was introduced to the shower energy. This smearing changed the shower energy resolution and, as shown in Figure 9.20, significantly degraded the Near Detector data/MC agreement for the neutrino energy. The effects of this smearing (when implemented first in the Far Detector and then in both detectors) were investigated by performing oscillation fits with different binning configurations. The resulting shifts in best fit parameters are shown in Table 9.3.

The results in Table 9.3 show that reducing the width of the energy bins from 1.0 GeV to 0.25 GeV did not increase the shift in best fit oscillation parameters. Similarly, the best fit shifts were not increased by using five energy resolution bins, despite the disagreement between the resolution parameterisation and the actual resolution of the events

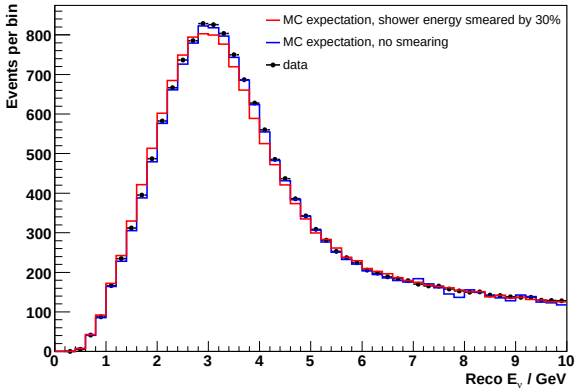

Figure 9.20: A 30% smearing of the shower energy in a Near Detector MC sample produces a noticeable disgreement between data and MC. The impact of this disagreement can be examined for different oscillation fit configurations.

in the smeared fake data sample. The division of events into bins of energy resolution is not expected to introduce any problems; the resolution parameterisation treats data and MC events in the same way, and the extrapolation method ensures similar y distributions for data and MC. The relative normalisation of the resolution bins is not used in the oscillation fit; only the shape information is used. The findings of this study suggest that the configuration of 0.25 GeV bin widths (plus overflow bin) and five resolution bins can be used safely in the analysis.

Configuration	**Energy bin width / GeV**	**Resolution bins**	$\boldsymbol{\Delta(\Delta m^2_{32})}$ $\boldsymbol{/10^{-4}\,\mathrm{eV}^2}$	$\boldsymbol{\Delta(\sin^2(2\theta_{23}))}$
Smear FD Shower Energy	1.0	1	$+0.841$	-0.0414
	0.25	1	$+0.727$	-0.0413
	0.25	5	$+0.706$	-0.0388
Smear ND and FD Shower Energy	1.0	1	$+0.341$	-0.0301
	0.25	1	$+0.222$	-0.0297
	0.25	5	$+0.231$	-0.0276

Table 9.3: Impact of a 30% smearing of the shower energy, as observed using oscillation fits with different binning configurations

Of the remaining sensitivity improvements, there are no concerns associated with the use of a coil hole cut or with the identification of unassigned shower strips. The coil hole cut simply represents a new definition of the Far Detector fiducial volume, whilst the

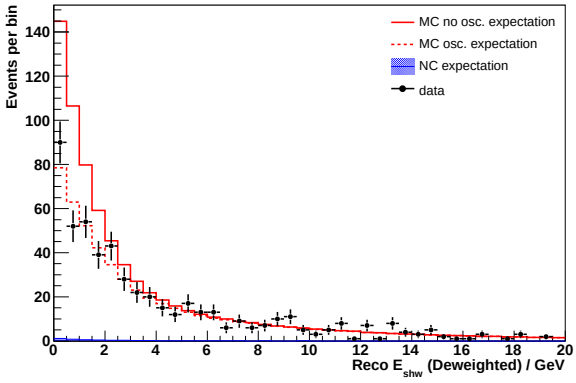

Figure 9.21: The agreement between Far Detector data and MC for the deweighted shower energy. The oscillated MC prediction assumes typical oscillation parameters of $\Delta m^2_{32} = 2.4 \times 10^{-3}\text{eV}^2$ and $\sin^2(2\theta_{23}) = 1.0$.

use of unassigned strips can be considered as a correction to the shower reconstruction software and should not change the systematic errors already assigned for the shower energy.

Finally, to complete this investigation, the agreement between Far Detector data and MC is examined for the deweighted shower energy in Figure 9.21. Although the data sample is small, good agreement is observed between the data distribution and the oscillated MC prediction.

9.8 Implementation

The success of the sensitivity improvements means that they will be used in Chapter 10 in an optimised fit to the MINOS Run I and Run IIa 2.5×10^{20} PoT dataset. The improvements therefore need to be included in the extrapolation and the oscillation fit that were described in Chapter 7.

The decision was made to leave the mechanics of the extrapolation unchanged, as the method had already been tested extensively and clearly worked effectively. This meant that the extrapolation would continue to use the standard detector fiducial volumes and also the standard reconstructed shower energy in the Near Detector. It should be emphasised that the Far Detector reconstructed energy never plays any part in determining the final event weights in the extrapolation. The more complicated instrumentation of the Near Detector also makes it very difficult to calculate a Near Detector deweighted shower energy.

Instead of making a large change to the extrapolation mechanics, the sensitivity

changes were all implemented at the point just after the final extrapolation weight has been calculated for a Far Detector MC event. After determination of the final weight, the event in question is examined to see if it passes the coil hole cut. If the event passes this cut, its muon energy and final shower energy are identified. The final shower energy is the deweighted calibration value, plus the contribution from any unassigned pulse height. After this identification of the event energy, the resolution of the event is calculated and the correct resolution bin is determined. The event is then ready to contribute to the 'prediction source histogram' for the relevant resolution bin.

These changes to the extrapolation are trivial to implement, as they occur entirely after the final extrapolation weight has been calculated. The output source histograms can then be used, as usual, as an input to the oscillation fit. In this way, the oscillation fit instantly adapts to all of the sensitivity improvements, apart from the use of resolution bins. A few small changes are needed to ensure that that oscillation fit can produce separate predicted energy spectra for each resolution bin. The final change is the adaptation of the likelihood function used in the fit, so that it is as described in Equation 9.2.

The changes to both the extrapolation and the oscillation fit are therefore quite minor. An example oscillation fit was performed for a MC sample representing 2.5×10^{20} PoT. This test produced the results shown in Figure 9.22, which also displays the results obtained without the inclusion of any sensitivity improvements. Both sets of results represent a full treatment of the six nuisance parameters detailed in Chapter 7. The confidence level contours obtained after implementation of the sensitivity improvements clearly cover a much reduced region of oscillation parameter space. This indicates that the improvements work effectively and that they allow the oscillation parameters to be determined with greater precision.

With input parameters of $\Delta m^2_{32} = 3.0 \times 10^{-3}\,\mathrm{eV}^2$ and $\sin^2(2\theta_{23}) = 0.9$, the MC test without sensitivity improvements identified the oscillation parameters as:

$$\Delta m^2_{32} = (3.00^{+0.26}_{-0.24}) \times 10^{-3}\,\mathrm{eV}^2,\ \sin^2(2\theta_{23}) = 0.899^{+0.075}_{-0.080}.$$

However, with the sensitivity improvements, the oscillation parameters could be specified more accurately as:

$$\Delta m^2_{32} = (3.00^{+0.24}_{-0.22}) \times 10^{-3}\,\mathrm{eV}^2,\ \sin^2(2\theta_{23}) = 0.902^{+0.064}_{-0.071}.$$

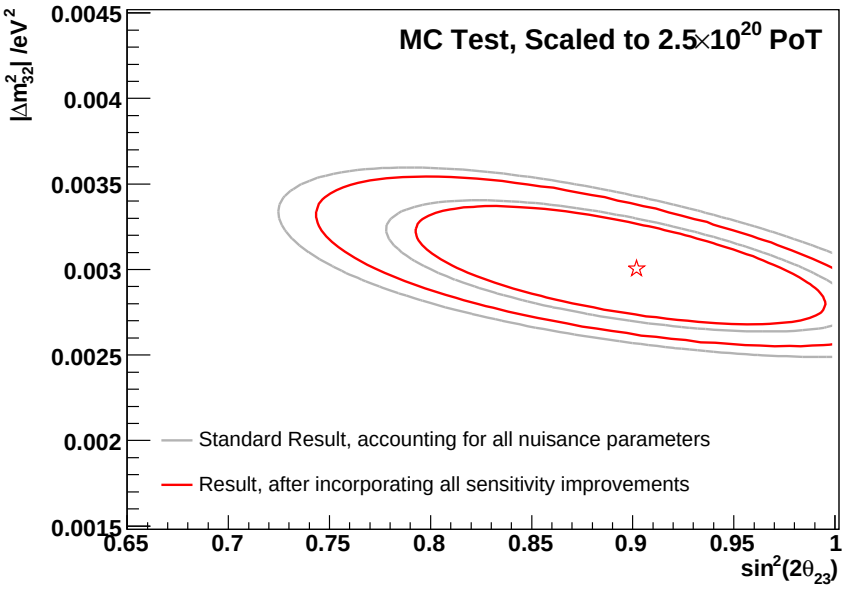

Figure 9.22: The results of oscillation fits to a MC fake data sample representing 2.5×10^{20} PoT. 68% and 90% confidence level contours are shown for fits with and without use of the sensitivity improvements. The results account for the six most important nuisance parameters, as identified in Chapter 7.

Chapter 10

Results

This Chapter brings together the ideas and techniques that have been developed throughout this thesis in order to extract a measurement of the neutrino oscillation parameters from the MINOS data. Development of these analysis tools and techniques has used only MC fake data for testing purposes, with no exposure to the real data. This means that the analysis can be considered as a 'blind analysis', which reduces conscious and unconscious bias in the results by using the data only as a final input to well understood algorithms.

The analysis uses data collected during the first two periods of operation of the MINOS experiment, as described in Chapter 6. Run I commenced in May 2005 and ended in February 2006, whilst Run IIa commenced in September 2006 and ended in March 2007. After the application of data quality cuts, the Far Detector data samples represent 1.27×10^{20} PoT from Run I and 1.23×10^{20} PoT from Run IIa. The combined Run I and Run IIa Far Detector data sample therefore represents a total of 2.5×10^{20} PoT.

At the start of Run IIa, the NuMI target was replaced and this resulted in a noticeable change in the neutrino energy spectrum, as observed in the Near Detector. For this reason, and to account for any other changes in the beam conditions, the extrapolation stage of the analysis is performed separately for each dataset. This means that separate Run I and Run IIa predictions are made for the Far Detector energy spectrum. The separate predictions can be used to make separate measurements of the oscillation parameters for each dataset, or they can be combined in order to extract a measurement from the full dataset.

Results are presented in this Chapter both with and without the inclusion of the sensitivity improvements detailed in Chapter 9.

10.1 Selected Events

The effects of the selection cuts on the Far Detector data sample are shown in Table 10.1. The final analysis sample consists of 275 selected CC candidate events from Run I and 289 events from Run IIa, representing a total sample of 564 events. In Table 10.2, this number of observed events is compared with the number of events expected in the absence of oscillations.

Selection Cut	**Number of Events**		
	Run I	**Run IIa**	**Combined**
Track in fiducial volume	431	416	847
Data quality cuts	416	414	830
Timing cuts	414	414	828
Beam quality cuts	402	410	812
Track passed by Kalman filter	401	410	811
Track charge ≤ 0	333	339	672
PID parameter > 0.85	**275**	**289**	**564**

Table 10.1: The number of Far Detector data events passing successive cuts.

Data sample	**Observed Events**	**Expected Events (No Oscillations)**	**Observed / Expected**
Run I	275	381±15 (syst.)	0.72 (4.3σ)
Run IIa	289	359±14 (syst.)	0.81 (2.9σ)
Combined	564	740±30 (syst.)	0.76 (4.4σ)
Combined <10 GeV	310	502±20 (syst.)	0.62 (6.4σ)
Combined <5 GeV	198	357±14 (syst.)	0.55 (6.7σ)

Table 10.2: Comparison of the number of observed events with the expectation in the absence of oscillations. The statistical significances shown account for both statistical and systematic uncertainties in the expectation.

Table 10.2 shows that there is a strong deficit in the number of observed events compared with the MC expectations in the absence of oscillations. The Table also shows that the deficit is energy dependent. For the combined dataset, across the entire energy spectrum, a deficit of 4.4σ relative to the no oscillation prediction is observed. Below 10 GeV, this deficit increases to 6.4σ relative to the no oscillation prediction. In calculating these statistical significances, both the statistical and systematic errors in the prediction have been considered. The dominant contribution to the systematic error is the 4% normalisation uncertainty. Uncertainties in the NC background also make a small contribution to the event rate uncertainty, but this contribution is smaller than 0.5%. Based purely on the event rate measurements, the no oscillation hypothesis is strongly disfavoured.

10.2 Results without Sensitivity Improvements

The first set of results presented do not make use of the sensitivity improvements described in Chapter 9. Instead, the standard energy reconstruction is used and standard selection cuts are used. The energy spectra used in the oscillation fits use a fine 0.25 GeV bin width up to 30 GeV and a single overflow bin for higher reconstructed neutrino energies.

As the volume of data collected at the Near Detector during Run I and Run IIa is so large, the extrapolation did not need to make use of the full Near Detector datasets. As described in Chapter 6, the beam is sampled accurately, without using all the data, by selecting approximately one complete Near Detector run per week of operation. This resulted in a Run I Near Detector data sample representing 2.37×10^{19} PoT and a Run IIa sample representing 2.21×10^{19} PoT.

The separate Run I and Run IIa predictions were used together in a simultaneous oscillation fit to the combined dataset. In this fit, Run I predictions were compared only with Run I data and Run IIa predictions were compared only with Run IIa data. Each comparison produced a separate log likelihood contribution and the overall log likelihood minimised in the fit was the sum of the Run I and Run IIa contributions. At each point in oscillation parameter space, the nuisance parameters (which are correlated for Run I and Run IIa) were varied to minimise the overall log likelihood.

The results obtained for the Run I, Run IIa and combined data samples are summarised in Table 10.3. The 68% and 90% confidence level contours are shown in Figure 10.1. Also shown in Figure 10.1 are the energy distributions for the data samples and the best fit predicted spectra.

Parameter	**Run I**	**Run IIa**	**Combined**
Δm^2_{32} / $10^{-3}\,\mathrm{eV}^2$	$2.63^{+0.25}_{-0.22}$	$2.20^{+0.21}_{-0.19}$	$2.41^{+0.16}_{-0.15}$
$\sin^2(2\theta_{23})$ (68% C.L.)	> 0.916	> 0.910	> 0.950
E_{shw} Scale	-0.081 (-0.8σ)	-0.006 (-0.06σ)	-0.070 (-0.7σ)
Relative E_{shw} Scale	-0.007 (-0.2σ)	$+0.009$ $(+0.3\sigma)$	$+0.003$ $(+0.1\sigma)$
E_μ Range Scale	-0.004 (-0.2σ)	$+0.0002$ $(+0.01\sigma)$	-0.003 (-0.2σ)
E_μ Curvature Scale	-0.016 (-0.2σ)	$+0.011$ $(+0.1\sigma)$	$+0.020$ $(+0.2\sigma)$
NC Background	-0.030 (-0.06σ)	-0.073 (-0.1σ)	-0.098 (-0.2σ)
Normalisation	-0.002 (-0.05σ)	$+0.010$ $(+0.3\sigma)$	$+0.003$ $(+0.08\sigma)$

Table 10.3: Summary of the oscillation fit results for the Run I, Run IIa and combined data samples. The pulls on the nuisance parameters at the best fit point are also indicated.

There are a few points concerning the nuisance parameters that are worth highlighting. Firstly, it can be seen that the NC background shifts at the best fit point are consistent with the findings of Section 6, which suggested that the NC level in the MC be adjusted by $(-6.3 \pm 1.3)\%$. Secondly, in the Run I and combined sample, a shift approaching -1σ in the E_{shw} scale significantly broadens the width of the confidence level contours in Δm^2_{32} and shifts the best fit point to a higher Δm^2_{32} value. The other nuisance parameters do little other than to induce some small and symmetric broadening of the contours.

Figure 10.2 compares the data spectrum and best fit oscillation prediction with the spectrum expected in the absence of oscillations. The Figure includes a ratio plot, which shows a clear oscillation dip with minimum in the region $1-2\,\mathrm{GeV}$ and provides compelling evidence for neutrino oscillations. At the minimum, the ratio of the observed energy distribution to the no oscillation expectation is less than 0.2.

10.3 Results with Sensitivity Improvements

The final results presented are those incorporating the sensitivity improvements introduced in Chapter 9; these represent the central results of this thesis. The deweighted shower energy calibration is used, and the shower energy measurement also accounts for high pulse height strips neglected by the shower finding software. A coil hole cut is made, in order to remove a large source of poorly reconstructed events, and the oscillation fit is performed using five bins of energy resolution. The energy spectra used in

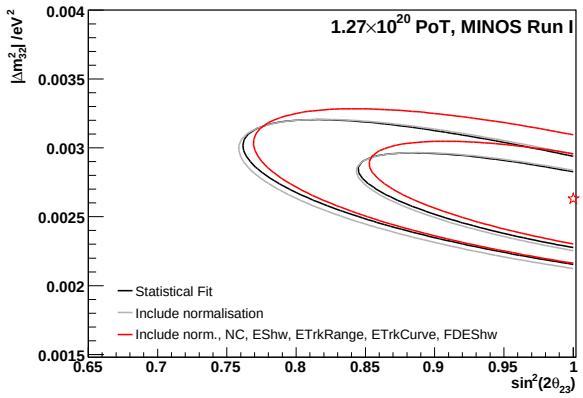

(a)

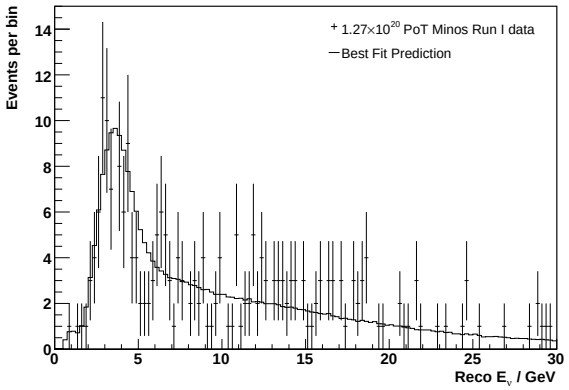

(b)

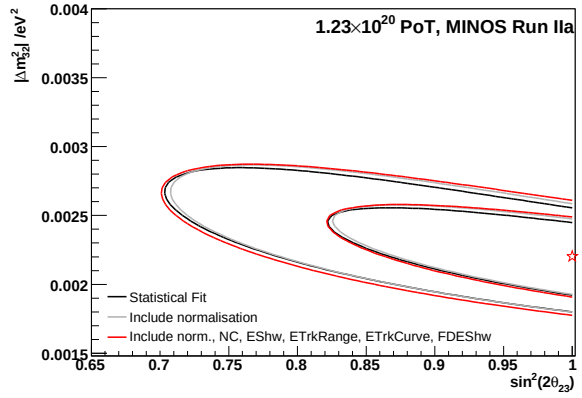

(c)

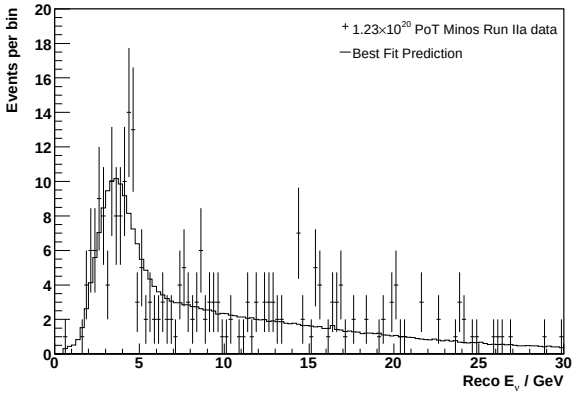

(d)

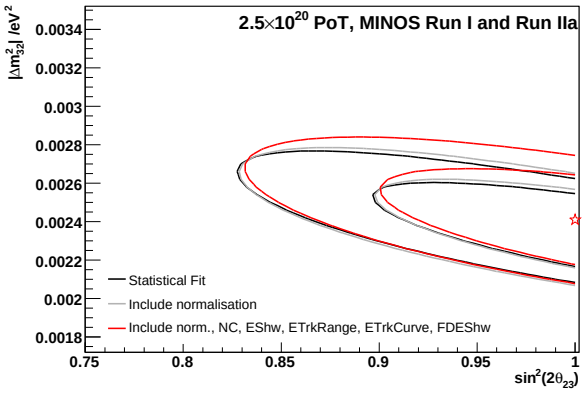

(e)

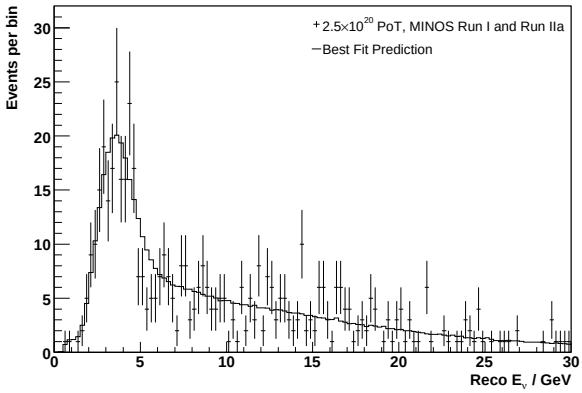

(f)

Figure 10.1: The results of oscillation fits to the Run I, Run IIa and combined data samples. The 68% and 90% confidence level contours are shown for a statistical fit, a fit including normalisation as a nuisance parameter and a full fit including all nuisance parameters. Also shown are the data distributions and the best fit predictions from the full fit.

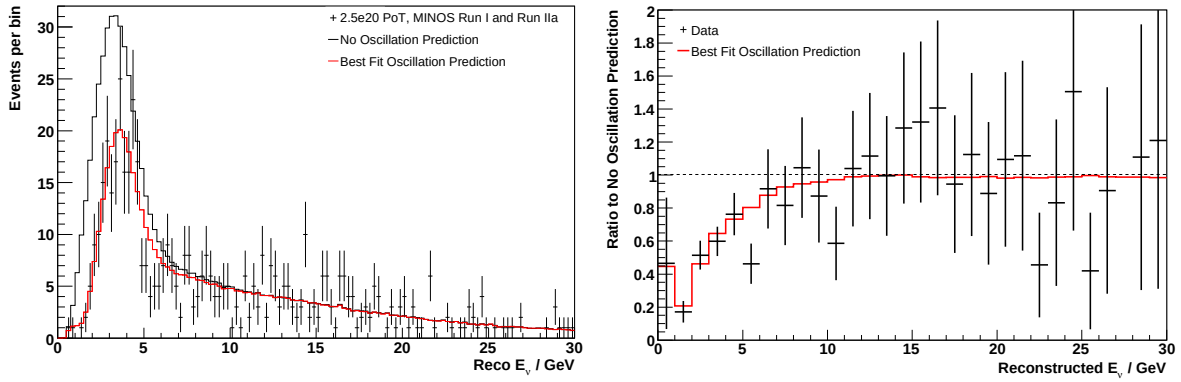

(a) (b)

Figure 10.2: Comparison of the combined data spectrum and best fit prediction with the energy distribution expected in the absence of oscillations.

the oscillation fit use a fine 0.25 GeV bin width up to 30 GeV and a single high energy overflow bin.

The results of the NC background study in Chapter 6 are also used, reducing the uncertainty in the NC background to 25% and implementing a default NC nuisance parameter shift of −6.3%. The NC background nuisance parameter term in the likelihood expression then penalises large deviations from this new default value.

As discussed in Chapter 9, the inclusion of the sensitivity improvements leaves the mechanics of the extrapolation unchanged. The Near Detector data samples therefore remained unchanged. However, the different shower energy reconstruction and the inclusion of a coil hole cut does change the shape and normalisation of the Far Detector predicted energy spectra and the actual data spectrum. Table 10.4 shows the new comparison between the number of observed events and the expectation in the absence of oscillations. The changes in this display of event rate information are only small.

The new results obtained by performing oscillation fits to the Run I, Run IIa and combined data samples are summarised in Table 10.5. The 68% and 90% confidence level contours for the Run I and Run IIa datasets are shown in Figure 10.3, together with the energy spectra for these data and the best fit oscillation predictions. For reference, the Figures also show the confidence level contours obtained without including the sensitivity improvements. The outcome of the improved fit to the combined data sample can be considered as the final result for this thesis, and the final contours and energy spectra are shown in Figure 10.4.

Data sample	Observed Events	Expected Events (No Oscillations)	Observed / Expected
Run I	274	378±15 (syst.)	0.72 (4.2σ)
Run IIa	289	355±14 (syst.)	0.81 (2.8σ)
Combined	563	733±29 (syst.)	0.77 (4.3σ)
Combined <10 GeV	311	498±20 (syst.)	0.62 (6.3σ)
Combined <5 GeV	197	358±14 (syst.)	0.55 (6.8σ)

Table 10.4: Comparison (after the inclusion of sensitivity improvements) of the number of observed events with the expectation in the absence of oscillations. The statistical significances shown account for both statistical and systematic uncertainties in the expectation.

Parameter	Run I	Run IIa	Combined
Δm^2_{32} / $10^{-3}\,\mathrm{eV}^2$	$2.47^{+0.24}_{-0.19}$	$2.15^{+0.20}_{-0.18}$	$2.29^{+0.14}_{-0.14}$
$\sin^2(2\theta_{23})$ (68% C.L.)	> 0.920	> 0.904	> 0.953
E_{shw} Scale	-0.099 (-1.0σ)	$+0.074$ ($+0.7\sigma$)	-0.027 (-0.3σ)
Relative E_{shw} Scale	-0.034 (-1.1σ)	$+0.023$ ($+0.8\sigma$)	$+0.023$ ($+0.8\sigma$)
E_μ Range Scale	-0.004 (-0.2σ)	$+0.003$ ($+0.2\sigma$)	-0.001 (-0.05σ)
E_μ Curvature Scale	-0.020 (-0.2σ)	-0.064 (-0.6σ)	-0.069 (-0.7σ)
NC Background	-0.007 (-0.03σ)	-0.017 (-0.07σ)	-0.029 (-0.1σ)
Normalisation	-0.010 (-0.3σ)	$+0.014$ ($+0.4\sigma$)	$+0.001$ ($+0.03\sigma$)

Table 10.5: Summary of the oscillation fit results for the Run I, Run IIa and combined data samples, after inclusion of the sensitivity improvements. The pulls on the nuisance parameters at the best fit point are also indicated.

The results indicate oscillation parameters in regions of parameter space compatible with the findings of the K2K and Super-Kamiokande experiments. The best fit parameters for the combined sample are $\Delta m^2_{32} = (2.29^{+0.14}_{-0.14}) \times 10^{-3}\,\mathrm{eV}^2$ and $\sin^2(2\theta_{23}) > 0.953$ (68% confidence level). Comparing the results obtained with and without sensitivity improvements shows that the improvements have decreased the best fit Δm^2_{32} values. Such changes are to be expected, as the different shower energy measurement alters the shape of the neutrino energy distributions, particularly at low energies. However, the new best fit Δm^2_{32} values are all within 1σ of those obtained without sensitivity improvements.

A second observation is that there is a clear improvement in the sensitivity to Δm^2_{32}, as can be seen from the width of the confidence level contours, or from the errors quoted

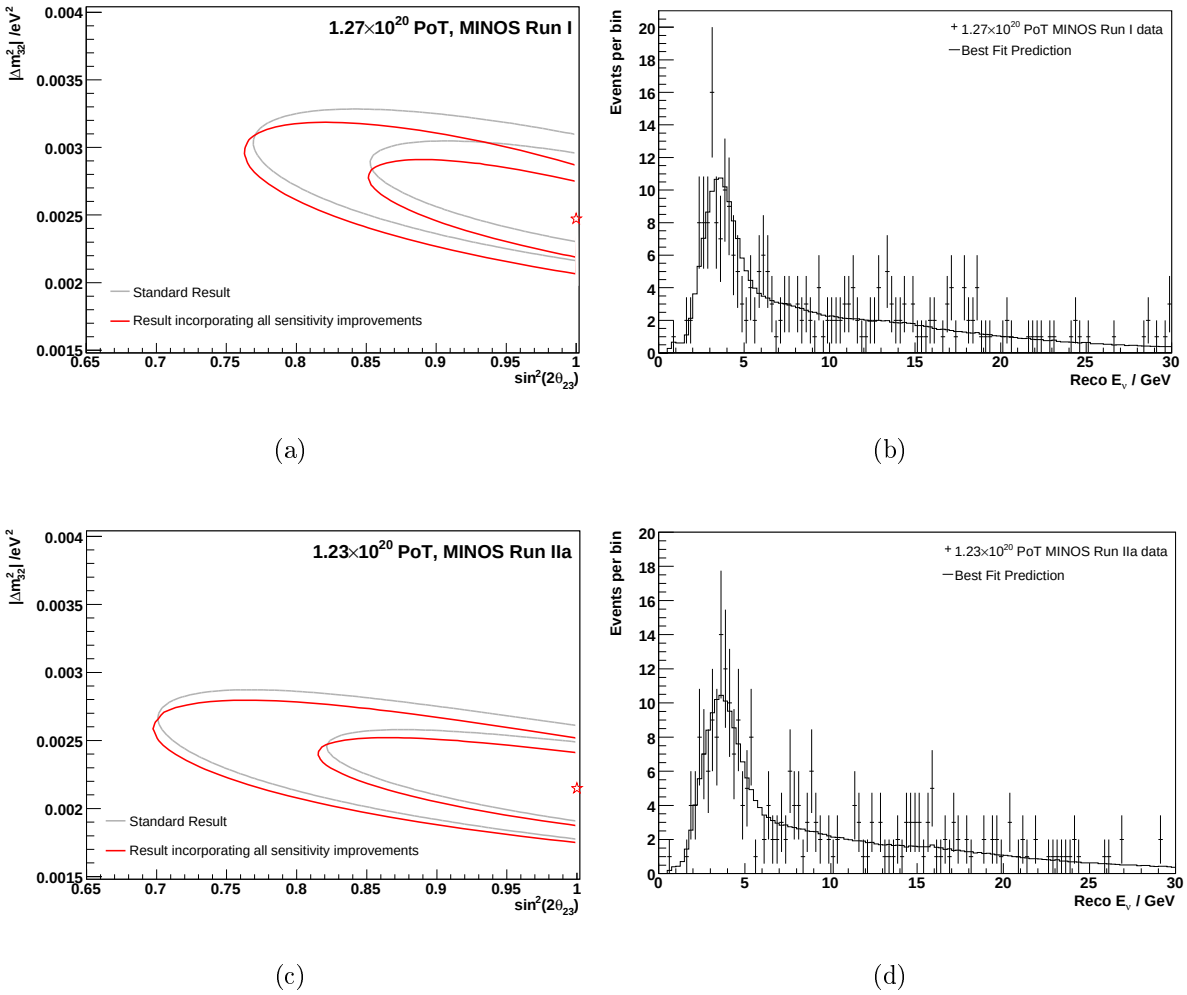

Figure 10.3: The results of oscillations fits to the Run I and Run IIa data samples, after inclusion of the sensitivity improvements. The 68% and 90% confidence level contours are shown, together with the data distributions and best fit predictions. For comparison, the confidence level contours produced without sensitivity improvements are also shown.

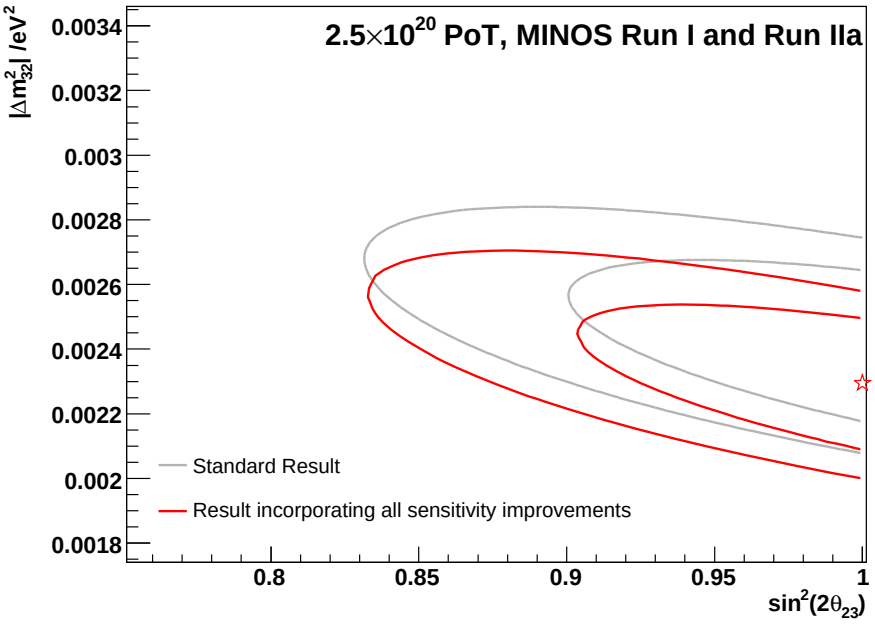

(a)

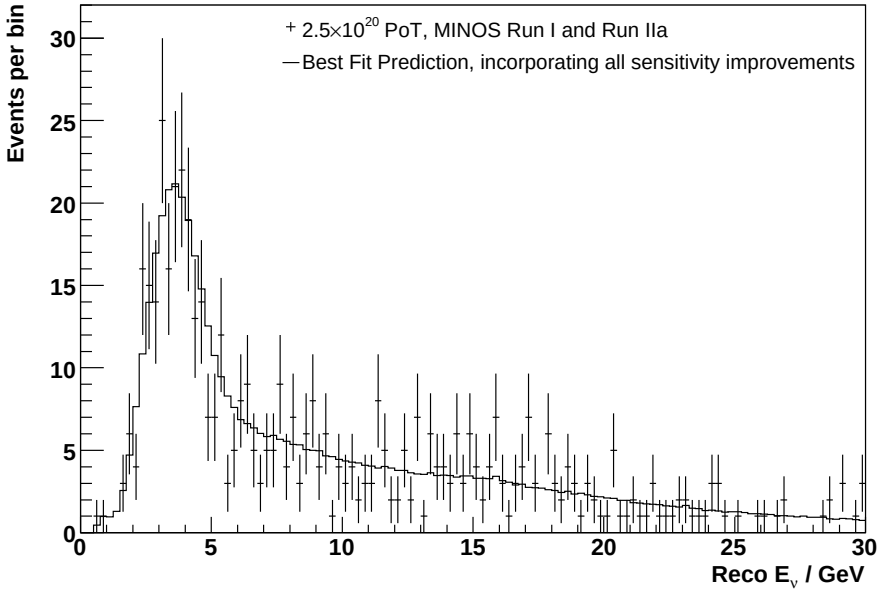

(b)

Figure 10.4: The results of oscillations fits to the combined data sample, after inclusion of the sensitivity improvements. The 68% and 90% confidence level contours are shown, together with the data distribution and best fit prediction. For comparison, the confidence level contours produced without sensitivity improvements are also shown.

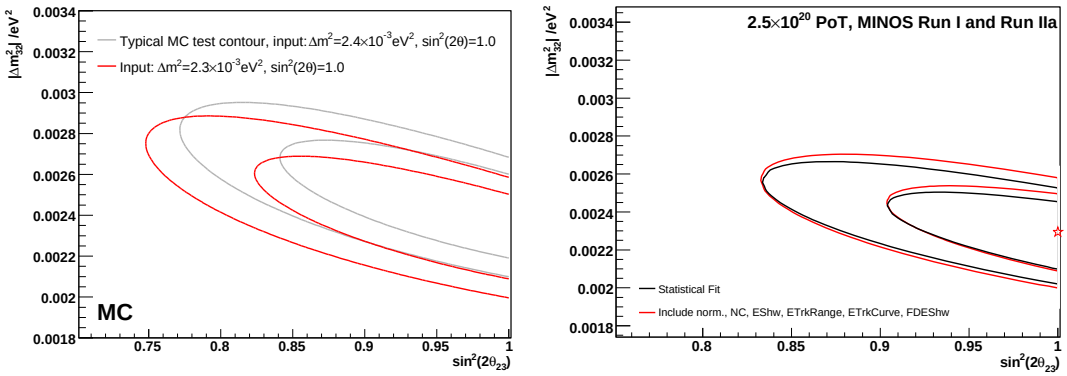

(a) (b)

Figure 10.5: (*a*) The typical broadening of confidence level contours associated with a decrease in best fit Δm^2_{32}. (*b*) The effects of accounting for systematic uncertainties in the final fit to the combined dataset.

on the best fit values. However, there is little change in the sensitivity to $\sin^2(2\theta_{23})$. This can be explained by the fact that the $\sin^2(2\theta_{23})$ sensitivity decreases noticeably at lower values of Δm^2_{32}. Lower Δm^2_{32} values push the oscillation dip to lower neutrino energies, where there are fewer events and where accurate energy reconstruction becomes more difficult. This leads to poor dip resolution and so to reduced sensitivity. In particular, the filling of the oscillation dip with poorly reconstructed events greatly decreases the sensitivity to $\sin^2(2\theta_{23})$.

Figure 10.5(a) shows typical confidence level contours, obtained with a MC test sample, for two different Δm^2_{32} input values. These input values are close to the best fit points for the combined data sample with and without sensitivity improvements, and so the contours illustrate the broadening expected with the observed shift in Δm^2_{32}. The effect is rather large, and it is purely due to the sensitivity improvement techniques that the sensitivity to $\sin^2(2\theta_{23})$ remains high in the new results, despite the decrease in best fit Δm^2_{32}.

Figure 10.5(b) indicates the effects of accounting for systematic uncertainties in the final fit to the combined dataset. It compares the 68% and 90% confidence level contours obtained from a purely statistical fit to those obtained with a full treatment of the nuisance parameters. The inclusion of the nuisance parameters results in a modest broadening of the contours in Δm^2_{32}, but very little change in $\sin^2(2\theta_{23})$. Unlike the results without sensitivity improvements, there is very little change to the best fit parameters.

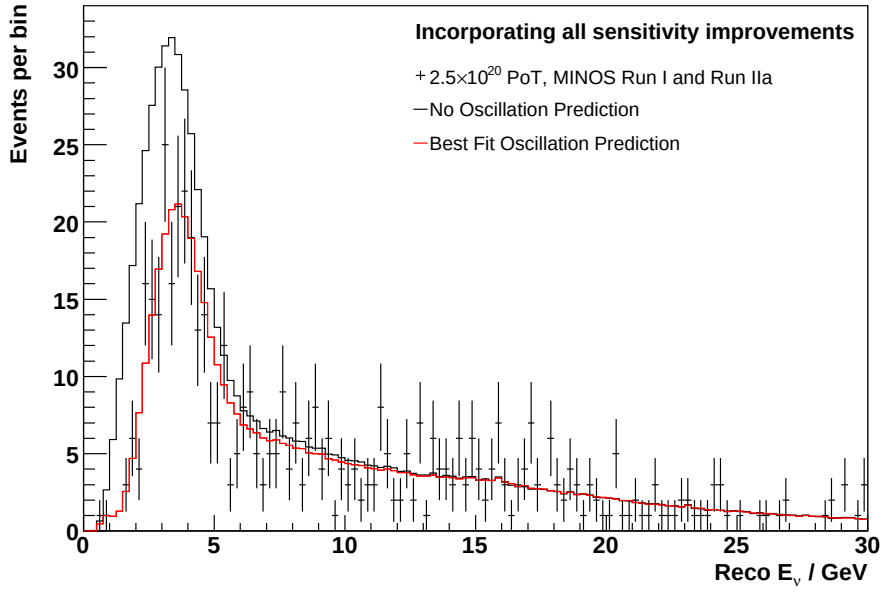

(a)

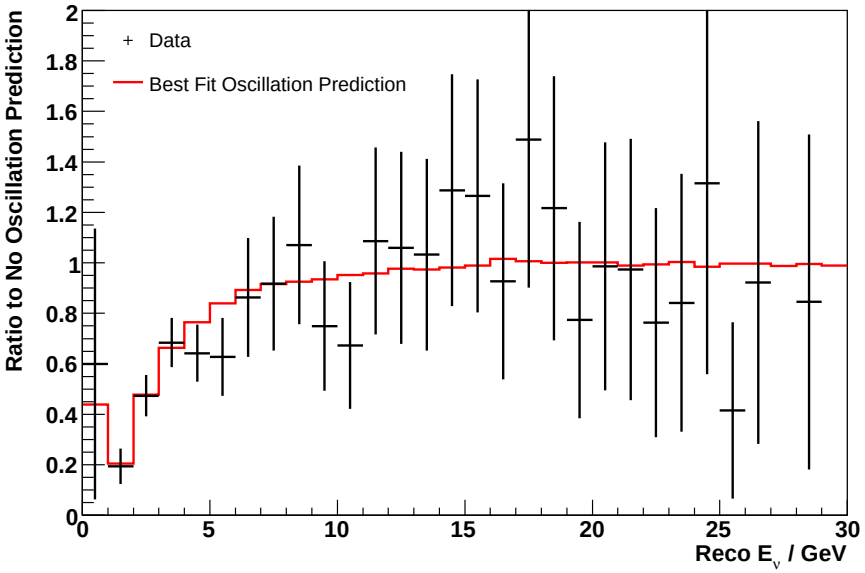

(b)

Figure 10.6: Comparison (after the inclusion of sensitivity improvements) of the combined data spectrum and best fit prediction with the energy distribution expected in the absence of oscillations.

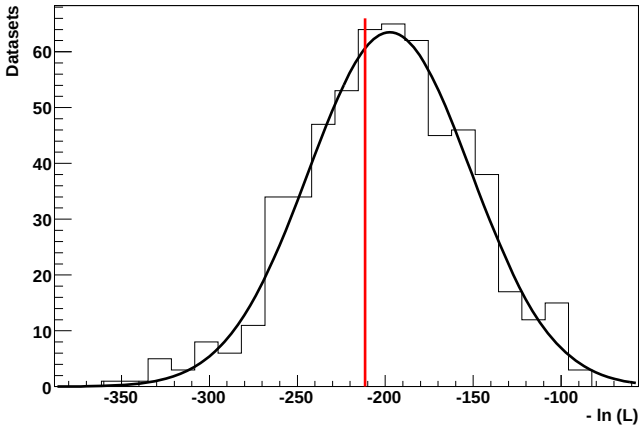

Figure 10.7: The range of minimum negative log likelihood values expected from natural fluctuations of a 2.5×10^{20} PoT Far Detector dataset. The value obtained in the fit to the real dataset is indicated.

An interesting observation is that the relative E_{shw} shifts at the best fit points are noticeably larger in the new results than in the results without sensitivity improvements. This is to be expected with different shower energy calibrations in the Near and Far Detectors. Finally, the inclusion of the results from the NC background study has, as expected, reduced the NC background shifts at the best fit point to very small values.

Figure 10.6 compares the data spectrum and the final best fit prediction with the Far Detector energy distribution expected in the absence of oscillations. A clear oscillation dip can be observed in the ratio of the data spectrum to the no oscillation prediction.

In order to assess the goodness of fit, the minimum negative log likelihood obtained in the fit to the data was compared with the range of values expected from natural fluctuations of the Far Detector dataset. Many representative 2.5×10^{20} PoT Far Detector energy distributions were produced by using the standard extrapolation method, together with the best fit oscillation parameters and different values for the nuisance parameters, drawn from Gaussian distributions. Each bin in these predicted energy distributions was fluctuated using Poisson statistics. The range of different datasets was then passed to the standard oscillation fit and the minimum negative log likelihood values were recorded. The distribution of these values was found to be approximately Gaussian and is shown in Figure 10.7. The red line in the Figure indicates the value obtained for the real data sample; this value is located just below the peak of the distribution, indicating a good fit.

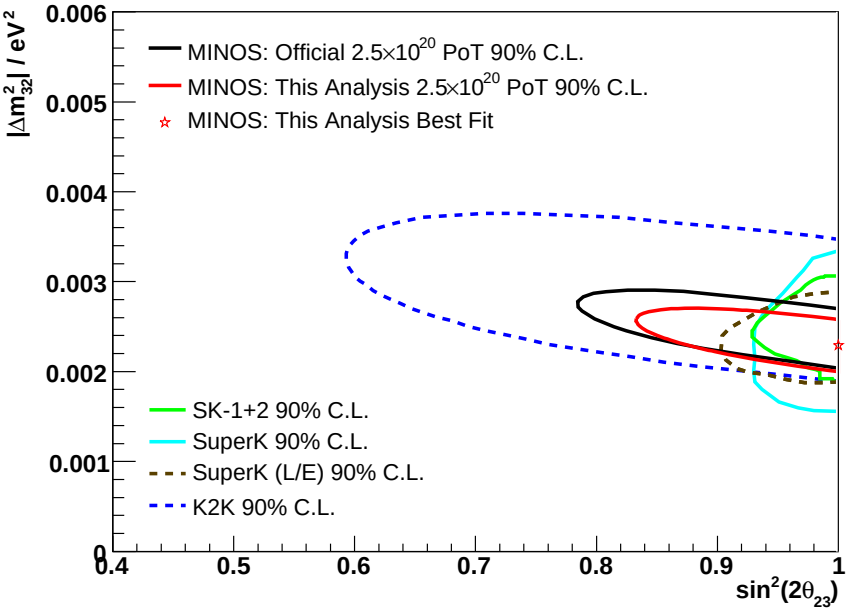

Figure 10.8: Comparison of the final 90% confidence level contour with the results from K2K, Super-Kamiokande and the official MINOS analysis.

10.4 Comparison with Other Results

Figure 10.8 compares the 90% confidence level contour from this analysis with the results from the K2K [24] and Super-Kamiokande experiments [21, 22, 105]. The results are also compared to the official MINOS 2.5×10^{20} PoT results [106]. It can be seen that this analysis represents the most accurate measurement of Δm^2_{32}, and that the sensitivity to $\sin^2(2\theta_{23})$ is much greater than that of the official MINOS result.

The differences between this analysis and the official MINOS analysis are due to the use of completely independent extrapolation methods and oscillation fitting techniques. This analysis also accounts for the six most important nuisance parameters, whilst the official analysis only considers the E_{shw} scale, normalisation and NC background. Importantly, this method makes use of the sensitivity improvement techniques described in Chapter 9 and also the result of the NC background study in Chapter 6.

10.5 Decoherence and Decay Models

The oscillation survival probability used in the fit can be replaced with the relevant survival probabilities for the decoherence and decay models of neutrino propagation (Equations 2.16 and 2.18). This allows fits to the data to be performed for these exotic

models. The results of these fits can be compared to the results of the oscillation fit, and the model that provides the best description of the data can be identified.

For both the decay and decoherence models, it should be noted that the survival probabilities are exponential functions of energy. This means that there is an important difference between the Far Detector predictions for these models and the predictions for the oscillations model: the decay and decoherence predictions cannot produce a dip in the ratio to the simple propagation (i.e. no oscillation) prediction.

Figure 10.9 shows the ratio to the simple propagation prediction for the data distribution and for the best fit prediction produced by each model. The dip in the data at low energies means that the most accurate description of the data is clearly provided by the oscillations model. The oscillations model also provides the best agreement at high energies. In the decoherence model the ratio cannot pass below 0.5, so there are large discrepancies between the data and the best fit prediction below 3 GeV. The low energy agreement is better for the decay model, but the model still cannot reproduce the dip in the data, so there is clear disagreement in the lowest energy bin.

Table 10.6 shows the negative log likelihood values obtained at the best fit point for each model, and also the minimum value achieved with the simple propagation model (after variation of the nuisance parameters). The best description of the data is provided by neutrino oscillations, so this provides the overall minimum negative log likelihood. Comparing this value with those obtained for the decay and decoherence models indicates that the decay model is disfavoured at the 3.2σ level and the decoherence model is disfavoured at the 5.0σ level. The no oscillation model is excluded at the 9.4σ level.

Propagation Model	$-\ln\mathcal{L}$ **at best fit**	$\Delta\ln\mathcal{L}$	**Comments**
Oscillations	−211.6	0	Best fit model for data
Decay	−206.6	5.0	Disfavoured at 3.2σ
Decoherence	−199.3	12.3	Disfavoured at 5.0σ
Simple/No Oscillation	−167.4	44.2	Disfavoured at 9.4σ

Table 10.6: The minimum negative log likelihood values obtained in fits to the oscillation, decoherence, decay and simple models of neutrino propagation. The standard deviations, σ, are related to the difference between negative log likelihood values: $\Delta\ln\mathcal{L} = \sigma^2/2$.

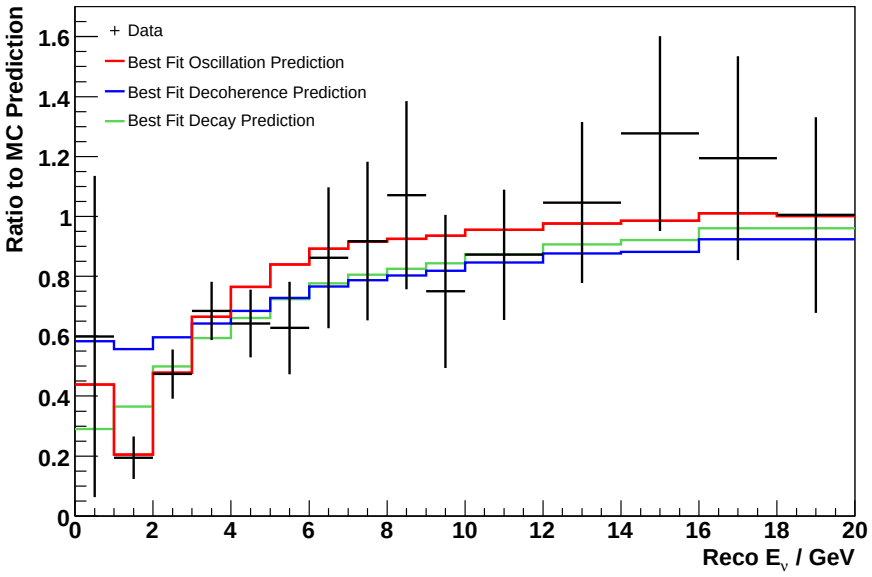

Figure 10.9: The ratios of the best fit oscillation, decoherence and decay predictions to the simple/no oscillation prediction.

Chapter 11

Summary and Outlook

An analysis of ν_μ disappearance has been performed using the MINOS detectors to observe the interactions of neutrinos in the NuMI beam. The two detectors measure the rates and energy spectra of CC ν_μ interactions at distances of 1 km and 735 km along the NuMI beamline. The Near Detector provides information about the initial composition and energy of the beam, whilst the Far Detector is sensitive to the effects of neutrino oscillations. The analysis uses data collected with the NuMI beam in its low energy configuration, between May 2005 and March 2007. The beam and detectors were continuously monitored to ensure the quality of the data, and, after data quality cuts, the total analysis sample corresponds to 2.5×10^{20} protons on target.

GPS timing signals were used to ensure data was recorded whenever beam neutrinos passed through the detectors. The signature of a CC ν_μ interaction is a muon track, and a series of algorithms has been developed to efficiently identify these tracks. In the first stage of track reconstruction, pattern recognition techniques are used to locate seed tracks, which specify the basic features of a muon trajectory. The seed tracks are then passed to a Kalman filter track fitter, which uses understanding of muon propagation to improve the track strip identification and reconstruct the muon kinematics. The algorithms identify tracks with high purity and completeness and accurately reconstruct the muon momentum and charge sign.

To isolate a sample of CC beam ν_μ interactions, events were required to contain a track with negative reconstructed charge. The track vertex was required to be within a specified detector fiducial volume in order to ensure accurate reconstruction and to

veto cosmic muon and rock muon interactions. A further cut on the track direction reduced the cosmic muon background to an insignificant level. The selection process was completed by using a multivariate technique to remove NC background events. This technique examined the properties of reconstructed tracks to assess whether they were really produced by a muon. The different selection efficiencies for NC background events in data and MC were then examined by removing muon tracks from samples of CC events, to produce fake NC-like events. This study suggested a normalisation correction for NC events in the nominal MC and provided information about the systematic uncertainty in the NC background.

Sensitivity to neutrino oscillations was obtained by comparing the reconstructed neutrino energy spectrum for selected Far Detector data events with predictions made for different oscillation parameters. Absolute predictions of the Far Detector neutrino spectrum are very difficult, with poor understanding of the neutrino beam representing a major source of systematic error. For this reason, an extrapolation method was implemented. This method uses Near Detector data to establish the characteristics of the neutrino beam and then exploits this understanding of the beam in order to produce an accurate prediction of the Far Detector spectrum. The extrapolation method was tested extensively and confirmed to be robust against even very poor modelling in the MC simulation.

The most important systematic uncertainties in the analysis were identified by applying systematic shifts to fake data samples. The fake Near Detector data was extrapolated to produce Far Detector predicted spectra. These were then used to identify the oscillation parameters introduced to the fake Far Detector data sample. Applying a systematic shift to the data would produce a change in the identified best fit parameters. The magnitude of this change characterised the importance of the systematic uncertainty. The most important uncertainties were identified as the shower energy scale, the relative normalisation, the muon range and curvature energy scales, the relative Far Detector shower energy scale and the NC background. An oscillation fit that fully accounted for these systematic uncertainties was implemented.

The extrapolation method calculates weights for MC events in order to improve agreement with the data. A comparison of the weighted MC events with the data was performed and no significant problems were identified. Some minor discrepancies were noted, including evidence of incomplete shower simulation and inaccurate modelling of the Near Detector spectrometer region, but these issues are fully addressed by the treatment of systematic uncertainties in the analysis.

A number of techniques to improve the sensitivity of an oscillation measurement

have also been developed. These improvements include optimising the oscillation fit and improving the shower energy reconstruction. Improvements were obtained by using a deweighted shower energy calibration, which suppresses fluctuations in the energy deposited in individual shower strips and increases the importance of the overall number of shower strips. Increases in sensitivity were also obtained by refining the Far Detector fiducial volume and by incorporating energy resolution information into the oscillation fit. For an exposure of 2.5×10^{20} PoT, these techniques collectively improved the sensitivity to the squared mass splitting Δm^2_{32} by approximately 21% and increased sensitivity to the mixing angle $\sin^2 2\theta_{23}$ by about 23%.

Extrapolating from the Near Detector data, 733 ± 29 events were expected in the Far Detector in the absence of neutrino oscillations, but only 563 events were observed. This deficit has a significance of 4.3σ; based purely on the event rate information, the no oscillation hypothesis is strongly disfavoured. The deficit was observed to be energy dependent and significant distortion of the Far Detector neutrino energy spectrum was detected. Using the optimised oscillation fit, assuming two flavour $\nu_\mu \rightarrow \nu_\tau$ oscillations, the allowed regions were determined for the oscillation parameters. The best fit values were identified as $\Delta m^2_{32} = 2.29^{+0.14}_{-0.14} \times 10^{-3}\text{eV}^2$ and $\sin^2 2\theta_{23} > 0.953$ (68% C.L.). The models of neutrino decoherence and decay are disfavoured at the 5.0σ and 3.2σ levels respectively, whilst the no oscillation model is excluded at the 9.4σ level.

This analysis currently provides the leading measurement of the parameter Δm^2_{32}, and represents a significant improvement in $\sin^2 2\theta_{23}$ sensitivity compared to the official MINOS analysis. However, data continues to be collected by the MINOS detectors and an analysis with increased statistics (3.25×10^{20} PoT) will soon be completed. This analysis will adopt many of the described sensitivity improvements, and so should significantly constrain the oscillation parameters. The thorough treatment of systematic uncertainties developed in this analysis will prove to be crucial in all future analyses, as the statistical errors shrink and systematic errors begin to dominate. Aside from increasing statistics, there are also numerous exciting extensions to the analysis.

Additional sensitivity to the oscillation parameters could be provided by including information from rock muons. These are muons produced by the interactions of beam neutrinos in the rock surrounding the detector. They can be identified using tight cuts around the predicted neutrino arrival time and on the direction of the reconstructed track. Including rock muon interactions would greatly increase the available event rate information. Other datasets that could be included are the samples collected when the NuMI beam was running in higher energy configurations. These small data samples would need to be treated separately in the extrapolation stage of the analysis, but

the high energy data would help to discriminate between the models of oscillations, decoherence and decay.

A natural extension to the analysis is to investigate the oscillations of anti-neutrinos. Such an analysis could investigate disappearance of the 6% of the beam that is naturally $\bar{\nu}_\mu$, or could reverse the direction of the current in the magnetic focussing horns in order to produce a $\bar{\nu}_\mu$ beam. The analysis would then search for differences between the neutrino oscillation parameters and the anti-neutrino oscillation parameters $\Delta\bar{m}^2_{32}$ and $\sin^2 2\bar{\theta}_{23}$. MINOS is also collecting atmospheric neutrino and anti-neutrino data, which could be used to complement this analysis.

Moving away from the standard oscillation analysis, MINOS is also searching for evidence of the existence of sterile neutrinos. Sterile neutrinos do not couple to the Z boson, and so are not directly observable in experiments. However, their presence could be detected by measuring a deficit of NC events in the Far Detector, relative to the prediction extrapolated from the Near Detector. Detection of sterile neutrinos would be very interesting, as they represent a way to introduce masses for the active neutrinos into the Standard Model, via a seesaw mechanism[107].

Finally, MINOS is also performing a search for $\nu_\mu \to \nu_e$ oscillations, by looking for ν_e appearance in the Far Detector. These oscillations are only possible with a non-zero value of θ_{13}, so this analysis should help to constrain the value of this important parameter. CC ν_e events consist of a compact shower, with a characteristic electromagnetic profile. Unfortunately the coarse MINOS detectors make these interactions difficult to identify. It is very difficult to distinguish signal events from the interactions of intrinsic beam ν_e, NC events or high-y CC ν_μ events. In the near future, MINOS is expected to achieve a reach comparable[108] to that of the CHOOZ experiment.

Future neutrino oscillation experiments are also focussing on obtaining a measurement of θ_{13}. The Double CHOOZ experiment[109] proposes to use two identical detectors, at distances of 300 m and 1.05 km from the CHOOZ nuclear reactors. With two detectors, the experiment should improve on the original CHOOZ result by increasing the statistics, reducing the systematic errors and by improving understanding of the backgrounds.

The T2K[110] and NOνA[111] experiments both propose to use off-axis detectors to measure θ_{13} and to improve measurements of the atmospheric neutrino oscillation parameters. By positioning the detectors a few degrees off-axis, the flux of neutrinos at the detector is reduced, but the neutrino energy spectrum is more suitable for studying oscillations. The off-axis neutrino energy spectrum is narrower than the on-axis spectrum and is peaked at a lower energy.

Bibliography

[1] W. Pauli, a letter to the Tubingen Conference on radioactivity, December, 1930.

[2] J. Chadwick, Distribution in intensity in the magnetic spectrum of the β-rays of radium, *Ver. Dtsch. Physik. Ges.* **16**, 383–391 (1914).

[3] J. Chadwick, Possible Existence of a Neutron, *Nature* **129**, 312 (1932).

[4] E. Fermi, An attempt of a theory of beta radiation. 1, *Z. Phys.* **88**, 161–177 (1934).

[5] F. Reines and C. L. Cowan, Detection of the free neutrino, *Phys. Rev.* **92**, 830–831 (1953).

[6] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, Detection of the free neutrino: A Confirmation, *Science* **124**, 103–104 (1956).

[7] M. Goldhaber, L. Grodzins, and A. W. Sunyar, Helicity of Neutrinos, *Phys. Rev.* **109**, 1015–1017 (1958).

[8] R. Davis and D. S. Harmer, Attempt to observe the $Cl^{37}(\bar{\nu}, e^-)Ar^{37}$ reaction induced by reactor antineutrinos, *Bull. Am. Phys. Soc.* **4**, 217 (1959).

[9] G. Danby *et al.*, Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos, *Phys. Rev. Lett.* **9**, 36–44 (1962).

[10] J. J. Aubert *et al.*, Experimental Observation of a Heavy Particle J, *Phys. Rev. Lett.* **33**, 1404–1406 (1974).

[11] J. E. Augustin *et al.*, Discovery of a Narrow Resonance in e^+ e^- Annihilation, *Phys. Rev. Lett.* **33**, 1406–1408 (1974).

[12] M. L. Perl *et al.*, Evidence for anomalous lepton production in e^+ e^- annihilation, *Phys. Rev. Lett.* **35**, 1489–1492 (1975).

[13] First direct observation of the tau-neutrino, *Europhys. News* **32**, 56–57 (2001).

[14] The LEP Collaborations, A combination of preliminary electroweak measurements and constraints on the standard model, (2005), hep-ex/0511027.

[15] Ch. Kraus *et al.*, Final results from phase II of the Mainz neutrino mass search in tritium beta decay, *Eur. Phys. J.* **C40**, 447–468 (2005).

[16] K. Assamagan *et al.*, Upper limit of the muon-neutrino mass and charged pion mass from momentum analysis of a surface muon beam, *Phys. Rev.* **D53**, 6065–6077 (1996).

[17] J. M. Roney, Review of the tau neutrino mass, *Nucl. Phys. Proc. Suppl.* **91**, 287–292 (2001).

[18] D. N. Spergel *et al.*, Wilkinson Microwave Anisotropy Probe (WMAP) three year results: Implications for cosmology, *Astrophys. J. Suppl.* **170**, 377 (2007).

[19] B. Pontecorvo, Inverse beta processes and non conservation of lepton charge, *Sov. Phys. JETP* **7**, 172–173 (1958).

[20] B. Aharmim *et al.*, Electron energy spectra, fluxes, and day-night asymmetries of ^{8}B solar neutrinos from the 391-day salt phase SNO data set, *Phys. Rev.* **C72**, 055502 (2005).

[21] Y. Ashie *et al.*, A measurement of atmospheric neutrino oscillation parameters by Super-Kamiokande I, *Phys. Rev.* **D71**, 112005 (2005).

[22] Y. Ashie *et al.*, Evidence for an oscillatory signature in atmospheric neutrino oscillation, *Phys. Rev. Lett.* **93**, 101801 (2004).

[23] S. Abe *et al.*, Precision Measurement of Neutrino Oscillation Parameters with KamLAND, (2008), arXiv:0801.4589 [hep-ex].

[24] M. H. Ahn *et al.*, Measurement of neutrino oscillation by the K2K experiment, *Phys. Rev.* **D74**, 072003 (2006).

[25] Z. Maki, M. Nakagawa, and S. Sakata, Remarks on the unified model of elementary particles, *Prog. Theor. Phys.* **28**, 870 (1962).

[26] B. Kayser, On the Quantum Mechanics of Neutrino Oscillation, *Phys. Rev.* **D24**, 110 (1981).

[27] W.-M. Yao *et al.*, Review of Particle Physics, *J. Phys. G* **33** (2006).

[28] A. S. T. Blake, *A study of atmospheric neutrino oscillations in the MINOS far detector*, PhD thesis, University of Cambridge, FERMILAB-THESIS-2005-77.

[29] E. Lisi, A. Marrone, and D. Montanino, Probing possible decoherence effects in atmospheric neutrino oscillations, *Phys. Rev. Lett.* **85**, 1166–1169 (2000).

[30] G. L. Fogli, E. Lisi, A. Marrone, and D. Montanino, Status of atmospheric $\nu_\mu \to \nu_\tau$ oscillations and decoherence after the first K2K spectral data, *Phys. Rev.* **D67**, 093006 (2003).

[31] V. D. Barger, J. G. Learned, S. Pakvasa, and T. J. Weiler, Neutrino decay as an explanation of atmospheric neutrino observations, *Phys. Rev. Lett.* **82**, 2640–2643 (1999).

[32] V. D. Barger *et al.*, Neutrino decay and atmospheric neutrinos, *Phys. Lett.* **B462**, 109–114 (1999).

[33] J. N. Bahcall, A. M. Serenelli, and S. Basu, New solar opacities, abundances, helioseismology, and neutrino fluxes, *Astrophys. J.* **621**, L85–L88 (2005).

[34] B. T. Cleveland *et al.*, Measurement of the solar electron neutrino flux with the Homestake chlorine detector, *Astrophys. J.* **496**, 505–526 (1998).

[35] J. N. Bahcall and M. H. Pinsonneault, What do we (not) know theoretically about solar neutrino fluxes?, *Phys. Rev. Lett.* **92**, 121301 (2004).

[36] J. N. Abdurashitov *et al.*, Measurement of the solar neutrino capture rate by the Russian-American gallium solar neutrino experiment during one half of the 22-year cycle of solar activity, *J. Exp. Theor. Phys.* **95**, 181–193 (2002).

[37] W. Hampel *et al.*, GALLEX solar neutrino observations: Results for GALLEX IV, *Phys. Lett.* **B447**, 127–133 (1999).

[38] M. Altmann *et al.*, Complete results for five years of GNO solar neutrino observations, *Phys. Lett.* **B616**, 174–190 (2005).

[39] J. Boger *et al.*, The Sudbury Neutrino Observatory, *Nucl. Instrum. Meth.* **A449**, 172–207 (2000).

[40] S. Fukuda *et al.*, Determination of solar neutrino oscillation parameters using 1496 days of Super-Kamiokande-I data, *Phys. Lett.* **B539**, 179–187 (2002).

[41] L. Wolfenstein, Neutrino oscillations in matter, *Phys. Rev.* **D17**, 2369 (1978).

[42] S. P. Mikheev and A. Yu. Smirnov, Resonance enhancement of oscillations in matter and solar neutrino spectroscopy, *Sov. J. Nucl. Phys.* **42**, 913–917 (1985).

[43] J. Shirai, Start of KamLAND, *Nucl. Phys. Proc. Suppl.* **118**, 15–22 (2003).

[44] T. Kajita and Y. Totsuka, Observation of atmospheric neutrinos, *Rev. Mod. Phys.* **73**, 85–118 (2001).

[45] T. K. Gaisser and M. Honda, Flux of atmospheric neutrinos, *Ann. Rev. Nucl. Part. Sci.* **52**, 153–199 (2002).

[46] C. L. F. Howcroft, *Atmospheric Neutrinos in the MINOS Far Detector*, PhD thesis, University of Cambridge, FERMILAB-THESIS-2004-71.

[47] R. Becker-Szendy *et al.*, A Search for muon-neutrino oscillations with the IMB detector, *Phys. Rev. Lett.* **69**, 1010–1013 (1992).

[48] K. S. Hirata *et al.*, Observation of a small atmospheric ν_μ/ν_e ratio in Kamiokande, *Phys. Lett.* **B280**, 146–152 (1992).

[49] W. A. Mann, Atmospheric neutrinos and the oscillations bonanza, (1999), hep-ex/9912007.

[50] Y. Fukuda *et al.*, The Super-Kamiokande detector, *Nucl. Instrum. Meth.* **A501**, 418–462 (2003).

[51] S. H. Ahn *et al.*, Detection of accelerator produced neutrinos at a distance of 250-km, *Phys. Lett.* **B511**, 178–184 (2001).

[52] M. Apollonio *et al.*, Search for neutrino oscillations on a long base-line at the CHOOZ nuclear power station, *Eur. Phys. J.* **C27**, 331–374 (2003).

[53] P. Adamson *et al.*, Study of muon neutrino disappearance using the fermilab main injector neutrino beam, *Phys. Rev.* **D77**(7), 072002 (2008).

[54] V. Garkusha *et al.*, Design Study of the NuMI Low Energy Target for Higher Power beam, *MINOS-doc-2778* (2007).

[55] S. Kopp *et al.*, Blessed Beam Plots, *MINOS-doc-1485* (2006).

[56] Z. Pavlovic, MINOS Results from the NuMI Beam, *MINOS-doc-3418* (2007).

[57] D. G. Michael *et al.*, The magnetized steel and scintillator calorimeters of the minos experiment, to be submitted to Nucl. Instrum. Meth. A.

[58] A. Belias *et al.*, The MINOS data acquisition system, *IEEE Trans. Nucl. Sci.* **51**, 451–455 (2004).

[59] J. Oliver, N. Felt, G. Feldman, A. Lebedev, and R. Lee, Design and performance of the readout system of the MINOS Far Detector, *IEEE Trans. Nucl. Sci.* **51**, 2193–2195 (2004).

[60] T. Cundiff *et al.*, The MINOS near detector front end electronics, *IEEE Trans. Nucl. Sci.* **53**, 1347–1355 (2006).

[61] J. J. Hartnell, *Measurement of the calorimetric energy scale in MINOS*, DPhil thesis, University of Oxford, FERMILAB-THESIS-2005-51.

[62] MINOS Calibration Group, Calibration Position Paper, *MINOS-doc-3137* (2007).

[63] C. P. Ward, Far Detector Linearity Calibration, *MINOS-doc-1886* (2006).

[64] N. Tagg, Creating Near Detector Linearity Constants, *MINOS-doc-2176* (2006).

[65] P. Adamson *et al.*, The MINOS calibration detector, *Nucl. Instrum. Meth.* **A556**, 119–133 (2006).

[66] M. A. Kordosky, *Hadronic interactions in the MINOS detectors*, PhD thesis, University of Texas at Austin, FERMILAB-THESIS-2004-34.

[67] P. L. Vahle, *Electromagnetic interactions in the MINOS detectors*, PhD thesis, University of Texas at Austin, FERMILAB-THESIS-2004-35.

[68] A. Ferrari, P. R. Sala, A. Fasso, and J. Ranft, FLUKA: A multi-particle transport code (Program version 2005), CERN-2005-010.

[69] R. Brun *et al.*, Geant detector description and simulation tool, (1994), CERN Program Library Long Writeup W5013.

[70] H. Gallagher, The NEUGEN neutrino event generator, *Nucl. Phys. Proc. Suppl.* **112**, 188–194 (2002).

[71] C. Zeitnitz and T. Gabriel, The GEANT-CALOR interface and benchmark calculations of ZEUS test calorimeters, *Nucl. Instrum. Meth.* **A349**, 106–111 (1994).

[72] M. A. Thomson, Alternative DeMuxer, *MINOS-doc-2535* (2006).

[73] G. Welch and G. Bishop, An Introduction to the Kalman Filter, *Technical Report TR 95-041* (2006).

[74] R. Lee, Event Reconstruction in the Near Detector, *NuMI-917* (2000).

[75] S. Seun, Muon swimmer package, based on 'Numerical Recipes' section 16.2: Adaptive Stepsize Control for Runge-Kutta.

[76] E. J. Wolin and L. L. Ho, Covariance Matrices for Track Fitting with the Kalman Filter, *Nucl. Instrum. Meth* (1993).

[77] J. Musser, Swimmer Energy Loss, *MINOS-doc-2081* (2006).

[78] J. Musser, Spectrometer Track Finding using the Kalman Filter, *MINOS-doc-1676* (2006).

[79] J. S. Marshall, Near Detector Spectrometer Demultiplexing Improvements, *MINOS-doc-2058* (2006).

[80] A. S. T. Blake, Far Detector Data Quality Software, *MINOS-doc-3486* (2007).

[81] D. A. Petyt, Near detector data quality checks, *MINOS-doc-3883* (2007).

[82] S. Kopp and L. Loiacono, Beam Quality Cuts, *MINOS-doc-2747* (2007).

[83] D. A. Petyt, Recent results from the MINOS experiment, *MINOS-doc-3778* (2007).

[84] A. S. T. Blake, Far Detector Fiducial Volume and Background Studies, *MINOS-doc-2728* (2007).

[85] A. S. T. Blake, Cosmic Muon and Rock Muon Backgrounds in the CC Event Selection, *MINOS-doc-3386* (2007).

[86] A. S. T. Blake, A Likelihood Method for CC/NC Separation, *MINOS-doc-2720* (2007).

[87] D. A. Petyt, PID Sensitivity Study, *MINOS-doc-2743* (2007).

[88] G. P. Zeller, Low energy neutrino cross sections: Comparison of various Monte Carlo predictions to experimental data, (2003), hep-ex/0312061.

[89] A. J. Culling, *An optimised oscillation analysis of MINOS beam data*, PhD thesis, University of Cambridge, FERMILAB-THESIS-2007-27.

[90] A. J. Culling and M. A. Thomson, NC Background Studies, *MINOS-doc-2758* (2007).

[91] C. L. F. Howcroft, Muon Removal, *MINOS-doc-1333* (2005).

[92] S. Kopp *et al.*, Constraining the Beam Neutrino MC Flux Using the MINOS Near Detector Data, *MINOS-doc-2965* (2007).

[93] M. Szleper and A. Para, Neutrino spectrum at the far detector - systematic errors, *NuMI-781* (2001).

[94] C. Andreopoulos *et al.*, Cross section model uncertainties for the charged current analysis, *MINOS-doc-2989* (2007).

[95] F. James, MINUIT: Function Minimization and Error Analysis, (1994), CERN Program Library Long Writeup.

[96] A. D. Marino, Mock Data Challenge III Revealed, *MINOS-doc-1820* (2006).

[97] S. Kopp *et al.*, Fitting the Beam MC to the ND Data, *MINOS-doc-1548* (2006).

[98] N. Saoulidou, Mock Data Challenge III Results and Errata, *MINOS-doc-1824* (2006).

[99] N. Saoulidou, Beam Matrix Method Studies, *MINOS-doc-3086* (2007).

[100] R. Hatcher *et al.*, Range-Curvature Task Force Position Paper, *MINOS-doc-3134* (2007).

[101] S. Dytman *et al.*, Estimated Uncertainty on the Shower Energy Scale in Muon Neutrino CC Events, *MINOS-doc-3362* (2007).

[102] D. G. Michael *et al.*, Observation of muon neutrino disappearance with the MINOS detectors and the NuMI neutrino beam, *Phys. Rev. Lett.* **97**, 191801 (2006).

[103] U. Amaldi, Fluctuations in Calorimetry Measurements, *Phys. Scripta* **23**, 409 (1981).

[104] P. A. Rodrigues, Shower Energy Estimators and Gain Calibration, *MINOS-doc-3566* (2007).

[105] K. Nakamura, Atmospheric and long-baseline neutrinos, presented at TAUP 2007.

[106] Preliminary Results from MINOS on Muon Neutrino Disappearance Based on an Exposure of 2.5×10^{20} 120 GeV Protons on the NuMI Target, 2007, arXiv:0708.1495 [hep-ex].

[107] A. Kusenko, Sterile neutrinos, *AIP Conf. Proc.* **917**, 58–68 (2007).

[108] M. Sanchez, ν_e appearance in MINOS, *MINOS-doc-3347* (2007).

[109] F. Ardellier *et al.*, Double Chooz: A search for the neutrino mixing angle θ_{13}, (2006), hep-ex/0606025.

[110] Y. Hayato, T2K at J-PARC, *Nucl. Phys. Proc. Suppl.* **143**, 269–276 (2005).

[111] D. S. Ayres *et al.*, NOvA proposal to build a 30-kiloton off-axis detector to study neutrino oscillations in the Fermilab NuMI beamline, (2004), hep-ex/0503053.

List of Figures

2.1 The pp chain of solar fusion reactions 11
2.2 The solar neutrino energy spectrum, predicted by the SSM 11
2.3 Measurements of the solar neutrino flux at SNO 14
2.4 Confidence limits obtained in the SNO oscillation analysis 15
2.5 Prompt energy spectrum of selected $\bar{\nu}_e$ events at KamLAND 16
2.6 Confidence limits obtained in the KamLAND oscillation analysis 17
2.7 KamLAND L_0/E analysis 17
2.8 Production of atmospheric neutrinos 18
2.9 Summary of measurements of the double ratio, R' 19
2.10 Super-Kamiokande atmospheric neutrino zenith angle distributions . . . 21
2.11 Confidence limits obtained in Super-Kamiokande zenith angle analysis . . 22
2.12 Super-Kamiokande L/E analysis 23
2.13 Confidence limits obtained in Super-Kamiokande analyses 23
2.14 Energy distribution for beam neutrino events at K2K 25
2.15 Confidence limits obtained in K2K oscillation analysis 25
2.16 Prompt e^+ spectrum measured at CHOOZ 27
2.17 Confidence limits obtained in CHOOZ analysis 27

3.1 The MINOS neutrino beamline 29
3.2 The NuMI beam facility 30
3.3 The NuMI target and its cooling mechanisms 31
3.4 The magnetic focussing horns 32
3.5 Near Detector CC ν_μ spectra for different target/horn configurations . . 32
3.6 Construction of the detectors from steel/scintillator planes 34
3.7 Cross section through a scintillator strip 35
3.8 The optical readout system 35
3.9 The MINOS Far Detector 36
3.10 The MINOS Near Detector 37

3.11 Instrumentation of the Near Detector steel planes 38
3.12 Arrangement of scintillator planes in the Near Detector 39
3.13 Summary of the detector calibration 40

4.1 Typical topology of a CC beam ν_μ interaction 47
4.2 Summary of the track finding algorithm 48
4.3 Possible Cluster configurations for Triplet formation 50
4.4 Identification of associations between Triplets 51
4.5 Identification of preferred associations between Triplets 52
4.6 Formation of matched associations and/or chains of TrackSegments . . . 53
4.7 Identification of the longest paths of matched associations 54
4.8 Selection of the best 2D Tracks 55
4.9 Evaluation of the overlap between 2D Tracks in the u and v views 56
4.10 Event display for typical 2D Tracks, which consist of Clusters 56
4.11 Selection of Hits from track Clusters 57
4.12 Near Detector modifications for the track finding algorithm 60
4.13 Identification of isolated Clusters for formation of ± 10 plane Triplets . . 61
4.14 Tracking efficiency for Far Detector beam MC events 62
4.15 Track purity for Far Detector beam MC events 63
4.16 Track completeness for Far Detector beam MC events 64
4.17 Tracking efficiency for Near Detector beam MC events 65
4.18 Track purity for Near Detector beam MC events 66
4.19 Track completeness for Near Detector beam MC events 66

5.1 Information required as an input to the Kalman filter 73
5.2 Identification of strips matching the Kalman state vectors 74
5.3 Division of a CC beam ν_μ interaction into 'clean' track and vertex shower 75
5.4 Example of improved track strip selection in a large vertex shower 76
5.5 Summary of the track fitting process 77
5.6 Performance of q/p reconstruction for Far Detector beam MC events . . . 79
5.7 Performance of p_{range} reconstruction for Far Detector beam MC events . 80
5.8 Offset and resolution for Far Detector p_{range} reconstruction 81
5.9 Offset and resolution for Far Detector p_{curve} reconstruction 81
5.10 Fitted track purity for Far Detector beam MC events 82
5.11 Fitted track completeness for Far Detector beam MC events 82
5.12 Performance of q/p and p_{range} reconstruction for new and old software . . 84
5.13 Fitted track purity and completeness for new and old software 84

5.14 Track fitting failures and required CPU time for new and old software . . 85
5.15 Reconstruction performance for Far Detector cosmic muon MC events . . 86
5.16 Event display for typical multiple muon event in the Far Detector 87
5.17 Event display for complex cosmic muon event in the Far Detector 87
5.18 Summary of Near Detector spectrometer demultiplexing 89
5.19 Event displays for typical events that enter the spectrometer 90
5.20 Problems posed by tracks entirely contained in the spectrometer 91
5.21 Performance of q/p reconstruction for Near Detector beam MC events . . 92
5.22 Performance of p_{range} reconstruction for Near Detector beam MC events . 93
5.23 Offset and resolution for Near Detector p_{range} reconstruction 94
5.24 Offset and resolution for Near Detector p_{curve} reconstruction 94
5.25 Fitted track purity for Near Detector beam MC events 95
5.26 Fitted track completeness for Near Detector beam MC events 95

6.1 Summary of protons delivered to the NuMI target 99
6.2 Difference between Far Detector snarl time and SpillServer prediction . . 102
6.3 Distribution of reconstructed track vertices in the Near Detector 104
6.4 Distribution of reconstructed track vertices in the Far Detector 106
6.5 CC/NC PDFs of reconstructed track length 110
6.6 CC/NC PDFs of pulse height per track-like plane *vs.* track length 111
6.7 CC/NC PDFs of number of track-like planes *vs.* track length 111
6.8 CC/NC PDFs of goodness of muon track fit *vs.* track length 112
6.9 CC/NC PDFs of reconstructed y *vs.* reconstructed neutrino energy . . . 112
6.10 Relative CC/NC normalisations and PDFs of reconstructed track charge 113
6.11 Near Detector PID distribution and purity-efficiency curve 114
6.12 Far Detector PID distribution and purity-efficiency curve 114
6.13 Selected Near Detector events per spill *vs.* PoT in spill 115
6.14 Number of selected Near Detector data events each month 116
6.15 Energy spectrum of selected Near Detector events each month 116
6.16 Mean Near Detector energy spectra for Run I and Run IIa 117
6.17 Creation of NC-like events through muon removal in CC events 118
6.18 Removal of the reconstructed muon track from a typical CC event 120
6.19 Near Detector data/MC PID distributions 122
6.20 Comparison of MC muon removed CC events and MC NC events 123
6.21 Comparison of muon removed events in MC and data 124
6.22 Correction for different data/MC NC selection efficiencies 126

6.23 Data/MC PID distributions after MC selection efficiency corrections . . . 127
6.24 Data/MC PID distributions after efficiency and normalisation corrections 129

7.1 Decay angles required for neutrinos to interact at Near or Far Detectors . 133
7.2 Summary flow diagram for the extrapolation method 134
7.3 The steps required to obtain the beam particle weights 136
7.4 Motivation for 'flattening' the MC 137
7.5 MC flattening weights as a function of the matching properties. 138
7.6 Propagation of beam particle weights to the Far Detector 139
7.7 Selection efficiencies of true CC events in the Near and Far Detectors . . 141
7.8 Typical contribution of CC ν_τ events to a Far Detector prediction 143
7.9 Ratio of generated MC events to reconstructed MC events 143
7.10 Test of extrapolation using truth quantities and selections 146
7.11 Test of extrapolation using reconstructed quantities and selections 147
7.12 Tests of high energy prediction and NC background prediction 148
7.13 Results of a basic oscillation fit to MC fake data 148
7.14 Results for 100 Mock Data Challenge II 'experiments' 151
7.15 Overall results for Mock Data Challenge II 151
7.16 Results for 100 Mock Data Challenge III 'experiments' 152
7.17 Overall results for Mock Data Challenge III 152
7.18 Impact of different systematic shifts in a fake data sample 156
7.19 Far Detector predictions with removal of different interaction types . . . 158
7.20 Results of oscillation fits with removal of different interaction types . . . 159
7.21 Far Detector predictions with enhanced NC fraction 160
7.22 Far Detector NC background predictions with enhanced NC fraction . . . 161
7.23 Results of an oscillation fit with enhanced NC fraction 161
7.24 Results of an extrapolation without Near Detector PID cut 162
7.25 Demonstration that systematics are uncorrelated in the extrapolation . . 164
7.26 Effects of including systematic uncertainties in the fit 166
7.27 Near Detector data predictions for Run I and Run IIa 167
7.28 Far Detector no oscillation data predictions for Run I and Run IIa 167

8.1 Near Detector data/MC agreement: PID input variables 170
8.2 Near Detector data/MC agreement: PID parameter 172
8.3 Near Detector data/MC agreement: number of reconstructed tracks . . . 173
8.4 Near Detector data/MC agreement: track vertex and end positions . . . 174
8.5 Near Detector data/MC agreement: reconstructed track properties 175

8.6 Near Detector data/MC agreement: reconstructed shower properties . . . 177
8.7 Near Detector data/MC agreement: reconstructed neutrino energy 178
8.8 Far Detector data/MC agreement: PID and its input variables 180
8.9 Far Detector data/MC agreement: number of reconstructed tracks 181
8.10 Far Detector data/MC agreement: track vertex and end positions 182
8.11 Far Detector data: track vertex and end positions 183
8.12 Far Detector data/MC agreement: reconstructed track properties 184
8.13 Far Detector data/MC agreement: reconstructed shower properties . . . 185

9.1 Impact on sensitivity of binning used in the oscillation fit 188
9.2 Typical example of strips shared between the track and vertex shower . . 189
9.3 Calibration of MIP values to obtain linear shower energy estimation . . . 190
9.4 Selection of optimal deweighting power 192
9.5 Calibration of deweighted MIP values 193
9.6 Deweighted shower energy resolution 193
9.7 Performance comparison for deweighted and linear shower energies 194
9.8 Impact on sensitivity of using deweighted shower energy 195
9.9 Energy resolution parameterisations 197
9.10 Oscillation dips for events with best and worst resolution 198
9.11 Impact on sensitivity of using resolution information in oscillation fit . . 199
9.12 Positions of events with poor shower energy reconstruction 200
9.13 Typical event in which shower pulse height is lost in coil hole 201
9.14 Impact of coil hole cut on sensitivity 201
9.15 Typical omission of large pulse height strips by shower reconstruction . . 202
9.16 Offset in reconstructed shower energy *vs.* unassigned strip pulse height . 203
9.17 Impact on sensitivity of accounting for unassigned pulse height 203
9.18 Overall improvements in sensitivity 205
9.19 The effects of $\pm 5\%$ systematic shifts in the PMT gains 206
9.20 The effects of smearing the Near Detector shower energy 208
9.21 Far Detector data/MC agreement for deweighted shower energy 209
9.22 Impact of sensitivity improvements on a typical oscillation fit 211

10.1 Results for the Run I, Run IIa and combined data samples 216
10.2 Comparison of data and best fit prediction with no oscillation spectrum . 217
10.3 Results for Run I and Run IIa samples after sensitivity improvements . . 219
10.4 Results for combined data sample after sensitivity improvements 220
10.5 Broadening of C.L. contours associated with systematic uncertainties . . 221

10.6 Comparison of final prediction and data with no oscillation spectrum . . 222
10.7 Expected spread of best fit negative log likelihood values 223
10.8 Comparison of results with those from other experiments 224
10.9 Comparison of data with predictions for oscillations, decoherence & decay 226

List of Tables

6.1 Summary of the Far Detector analysis dataset 100
6.2 Summary of the Near Detector analysis dataset 101
6.3 Near Detector data and MC used in the muon removal study 121
6.4 Selection efficiency corrections for MC NC events 125
6.5 Combined normalisation and efficiency corrections for MC NC events . . 128

7.1 Summary of the MC used in the extrapolation 145
7.2 Impact on extrapolation of shifts in different systematic uncertainties . . 155

9.1 Summary of the observed sensitivity improvements 204
9.2 Effects of systematic changes in PMT gains. 207
9.3 Response of different fit configurations to energy smearing 208

10.1 The number of Far Detector data events passing successive cuts 213
10.2 The number of observed and expected Far Detector events 213
10.3 Summary of results for the Run I, Run IIa and combined data samples . 215
10.4 Number of observed and expected events after sensitivity improvements . 218
10.5 Summary of results after inclusion of sensitivity improvements 218
10.6 Results of fits to oscillation, decoherence, decay and simple models 225