Three Flavour Neutrino Oscillations in MINOS+

Joseph O'Connor University College London

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

March 17, 2016

Declaration

I, Joseph O'Connor confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Joseph O'Connor

Abstract

MINOS+ is an extension of the MINOS long-baseline neutrino experiment. It makes measurements of the upgraded NuMI neutrino beam at two detectors: the Near Detector, 1 km from the beam at Fermilab, Illinois and the Far Detector, 735 km from the beam in the Soudan Mine, Minnesota. A comparison of the neutrino energy spectra in these two detectors is made in order to investigate the effect of neutrino oscillations.

This thesis documents the first MINOS+ beam ν_{μ} disappearance analysis. The data are combined with those from the MINOS beam disappearance, atmospheric ν_{μ} disappearance and beam ν_e appearance analyses, as well as additional atmospheric data taken after the MINOS shutdown, in order to estimate the parameters of the threeflavour neutrino oscillation model. The 68% confidence intervals on the atmospheric mixing angle obtained from this combined analysis are $\Delta m_{32}^2 = 2.45^{+0.08}_{-0.11} \times 10^{-3} \text{eV}^2$ and $\Delta m_{32}^2 = 2.42^{+0.09}_{-0.11} \times 10^{-3} \text{eV}^2$, assuming normal and inverted neutrino mass hierarchy respectively.

In addition, a measurement of the time of flight of the neutrino beam between the Near and Far Detector is described. The analysis is performed using data from the two detectors, along with a beam monitor at NuMI, various timers and long-distance synchronisation systems. This results in the most precise measurement of the flighttime of a neutrino beam ever made. The speed of the beam is calculated to be $(v - c)/c = (1.0 \pm 1.1) \times 10^{-6}$, consistent with v = c.

Acknowledgements

I would like to thank a great number of people who have supported me in one way or another. I have needed a lot of help!

Firstly, to my supervisor, Jenny Thomas, who has been constantly supportive, helpful and patient. Thank you for everything.

Thanks to Ryan Nichol, Leigh Whitehead and Alexander Radovic, who have all been incredibly helpful over the years. I have shared some good company at UCL – thanks especially to my distinguished elders, Adam, D'Arcy, Jon and Motty, who made me feel welcome.

Thanks also to my colleagues on the MINOS experiment, Rob Plunkett, Art Kreymer, Tricia Vahle, Alex Sousa, Phil Adamson, Justin Evans, Benton Pahlka, Adam Schreckenberger, Son Cao, Junting Huang, Ashley Timmons and Andy Perch, each of whom has helped with my work in some capacity and each of whom it has been a pleasure to meet. Special thanks goes to Michelle Medeiros, alongside whom much of the work in this thesis was carried out, and to Navaneeth Poonthottathil, whose sofa, cooking and bakwas I will miss.

I could not have finished this work without my family and friends. My brother Jack, who is a beautiful, kind man. My wonderful mum, who has done more for me than I can ever hope to return. My dad, who is a generous, remarkable person, and my sister Georgia, whom I love very much. Adam, Ed, I love you both too. And Mary, thank you for keeping me going; you are perfect.

Finally, to Dia, who really wanted to be on the acknowledgements page.

Preface

This thesis presents two analyses to which I made significant contributions. Chapter 4 presents the first ν_{μ} disappearance analysis of MINOS+ beam data, along with a combined analysis of the full MINOS and MINOS+ dataset. I updated the MINOS fitting software for this analysis, processed the data and Monte Carlo for the fit along with Michelle Medeiros and made the contours shown in this chapter, along with various sensitivity contours before the final analysis. Various other work was performed by members of the MINOS+ Standard Oscillations Group. I performed similar tasks for an updated MINOS analysis presented at the Neutrino 2014 conference in Boston, MA [1].

Two analyses measuring the speed of propagation of neutrinos from the NuMI beam is presented in Chapter 5. I made small contributions to the first analysis (labelled RTOF), including analysis of detector delays using the Auxiliary Detectors. I performed the majority of the data analysis for the second (HRTOF) analysis – published in [2] – with guidance from Phil Adamson.

Aside from these analyses, my work on the experiment has included smaller roles working on the 2012 MINOS sterile neutrino analysis and, during the MINOS-MINOS+ shutdown, work on upgrading and vaildating the Near Detector reconstruction algorithms to deal with the increased neutrino energies and rates in MINOS+, including developing the algorithm shown in section 2.8.1 along with Michelle Medeiros.

Contents

Li	List of Figures 10				
\mathbf{Li}	st of	Tables	5	19	
1	Neu	itrino 1	Physics	21	
	1.1	Three	Flavours of Neutrino	21	
	1.2	Neutri	nos and the Standard Model	23	
	1.3	Neutri	no Mass	24	
	1.4	Three-	Flavour Neutrino Oscillations	25	
		1.4.1	Mass Hierarchy	27	
		1.4.2	Oscillations in Matter	28	
		1.4.3	Two-Flavour Approximation	30	
	1.5	Evider	nce for Neutrino Oscillations	31	
		1.5.1	θ_{12} : Solar Neutrino Oscillations	31	
		1.5.2	θ_{23} : Atmospheric Neutrino Oscillations	35	
		1.5.3	θ_{13} : Reactor ν_e Disappearance, Beam ν_e Appearance \ldots	36	
		1.5.4	Mass Hierarchy	39	
		1.5.5	δ_{CP}	39	
		1.5.6	θ_{23} Octant	40	
	1.6	Summ	ary of Three-Flavour Oscillation Parameters	41	
0	T I	NTNA	I D MINIOS D. tt	49	
2	I ne		Beam, MINOS Detectors and MINOS+	43	
	2.1	Overv		43	
	2.2	MINO	S+ Physics Goals	44	
	2.3	The N	uMI Beam	45	
		2.3.1	FINAL Accelerator Complex	46	
		2.3.2	Neutrino Production	48	
		2.3.3	Summary of NuMI Upgrades	51	

		2.3.4	Properties of the NuMI Beam	51
	2.4	The N	IINOS Detectors	53
		2.4.1	Overview	53
		2.4.2	Near Detector	55
		2.4.3	Far Detector	57
		2.4.4	Scintillator and Readout	58
	2.5	Calibr	ation	60
	2.6	Monte	e Carlo Simulation	62
		2.6.1	Beam	62
		2.6.2	Detectors	64
	2.7	Intera	ction Topologies	65
	2.8	Recon	struction	67
		2.8.1	Slicing	67
		2.8.2	Tracking	70
		2.8.3	Showering	70
		2.8.4	Near Detector Reconstruction Performance in MINOS+	72
0	T L.		Negative Operation of MINIOC	F 4
3	1 nr	ee-ria	Nour Neutrino Oscillations at MINOS	74
	3.1	Beam	ν_{μ} Disappearance	(4 70
		3.1.1 2.1.0	Data	76 76
		3.1.2		70 20
		3.1.3	Far Detector Prediction	80
	2.0	3.1.4	Systematics	85
	3.2	Atmos	spheric ν_{μ} Disappearance	88
		3.2.1		88
		3.2.2	Far Detector Prediction	89
	0.0	3.2.3 D	Systematics	90
	პ.პ	Beam	Neutrino Appearance	90
		3.3.1 2.2.0		91
		332	Selection	91
		0.0.2		0.0
		3.3.3	Far Detector Prediction	93
	0.4	3.3.3 3.3.4	Far Detector Prediction	93 94
	3.4	3.3.3 3.3.4 Fitting	Far Detector Prediction	93 94 95
	3.4	3.3.3 3.3.4 Fitting 3.4.1	Far Detector Prediction	93 94 95 95

4	\mathbf{MII}	NOS+	Beam Disappearance Analysis	100
	4.1	Data		. 100
	4.2	Select	ion	. 101
		4.2.1	roID retuning	. 102
		4.2.2	ND Coil Hole Cut	. 102
	4.3	Fittin	g	. 104
		4.3.1	Binning	. 104
		4.3.2	Systematics	. 105
	4.4	Result	${ m ts}$. 106
	4.5	Comb	ined MINOS+ and MINOS Analysis	. 110
		4.5.1	Data	. 110
		4.5.2	Systematic Parameter Correlation	. 114
		4.5.3	Results	. 115
5	Mea	asuring	g Neutrino Speed at MINOS	119
	5.1	Introd	luction	. 119
	5.2	MINC	OS Time of Flight Analyses	. 120
		5.2.1	Timing in the NuMI Beam	. 121
		5.2.2	Timing in the MINOS Detectors	. 122
	5.3	Retro	spective Time of Flight Analysis	. 123
		5.3.1	Antenna Cable Delays	. 123
		5.3.2	Detector Latencies	. 124
		5.3.3	Selection	. 125
		5.3.4	Fitting	. 125
		5.3.5	GPS Reset Instability	. 127
		5.3.6	Result	. 127
	5.4	High-	Resolution TOF Analysis	. 127
		5.4.1	Resistive Wall Current Monitor	. 128
		5.4.2	Synchronisation	. 129
		5.4.3	Validation of the Timing System	. 132
		5.4.4	ND-to-FD Distance Survey	. 133
		5.4.5	Systematic Uncertainties	. 134
		5.4.6	Selection	. 134
		5.4.7	Fitting	. 135
		5.4.8	Results	. 137

Contents

6	Summary and Conclusion	143
Bi	bliography	144

List of Figures

1.1	The cross-section for Z production as a function of the LEP centre- of-mass energy. Predictions are shown for two, three and four light, active neutrinos [3]	23
1.2	Feynman diagrams for the two types of neutrino interaction. Charged- current interactions (left) involve the exchange of a charged W boson, while neutral-current interactions (right) have a neutral Z boson as propagator.	24
1.3	Illustration of the two neutrino mass hierarchy scenarios allowed by current data. The value Δm_{21}^2 is known to be positive, but the sign of Δm_{32}^2 is unknown. Taken from [4]	28
1.4	Comparison of the ν_{μ} disappearance probability in the two- and three- flavour oscillation models. The mixing angles used (in radians) are $\theta_{12} = 1.17$, $\theta_{23} = 0.68$ and $\theta_{13} = 0.16$ and the mass-splittings (in eV ²) are $\Delta m_{21}^2 = 7.58 \times 10^{-5}$ and $\Delta m_{32}^2 = 2.41 \times 10^{-3}$.	31
1.5	Components of the solar neutrino spectrum as predicted by the BSO5(OP) SSM [5]	32
1.6	Best fit along with 68%, 95% and 99.73% C.L. contours for the parameters Δm_{21}^2 and $\tan^2 \theta_{12}$ from a combined analysis of KamLAND and SNO data [6]	34
1.7	Allowed regions for $\sin^2(\theta_{23})$ and Δm_{32}^2 at the 68% (dashed lines) and 90% (solid lines) in MINOS and T2K. Two years of MINOS+ atmospheric data are included in the MINOS contour. The best fits are shown by the red and blue dots for MINOS and T2K respectively.	37

1.8	Allowed regions for $\sin^2 2\theta_{13}$ and $ \Delta m_{ee}^2 $ at the 68.3%, 95.5%, and 99.7% confidence level from the Daya Bay experiment. The best estimate of the oscillation parameters is given by the black dot. The adjoining panels show the profile likelihoods for each of the parameters. Taken from [7].	38
1.9	Sensitivity to the neutrino mass hierarchy for 6 years of NO ν A running in combination with T2K $\nu_{\mu} \rightarrow \nu_{e}$ data [8]	40
1.10	Sensitivity to δ_{CP} for 6 years of NO ν A running in combination with T2K $\nu_{\mu} \rightarrow \nu_{e}$ data [8].	41
2.1	Geographical layout of the MINOS+ experiment. The NuMI beam and Near Detector are located at Fermilab in north-east Illinois, the Far Detector at the Soudan Mine in northern Minnesota [9]	44
2.2	The $\nu_{\mu} \rightarrow \nu_{\mu}$ and $\nu_{\mu} \rightarrow \nu_{e}$ oscillation probabilities are shown as a function of L/E. The shaded bands show the energy range within $\pm 1 \sigma$ of the mean beam neutrino energy in MINOS+, MINOS and NO ν A at their respective Far Detectors. For MINOS+, the mean is taken to be 6.6 \pm 2.5 GeV, for MINOS 3.5 \pm 1.2 GeV and for NO ν A 2.1 \pm 0.5 GeV.	45
2.3	Schematic of the Tevatron accelerator complex. The Tevatron stopped operation in September 2011. Protons are extracted from the Main Injector to the NuMI target [10].	46
2.4	Time profile of the the neutrino beam as seen at the MINOS+ Near Detector, April 2015. The first two batches are doubled in intensity by slip stacking.	48
2.5	Schematic of the NuMI beamline. Protons from the Main Injector hit a graphite target, producing charged hadrons that are focused by a pair of magnetic horns. These hadrons traverse a decay pipe in which they may decay to produce neutrinos, before absorbers remove contaminants such as muons. The horn positions seen here correspond to the low-energy running seen for most of the lifetime of MINOS [11].	49

2.6	Drawing of the new NuMI target. Accelerated protons interact with thin graphite plates, which are clamped between two aluminium sheets, cooled by water pipes [12]	50
2.7	Unoscillated Far Detector ν_{μ} energy spectra are shown for 6×10^{20} POT for typical MINOS, MINOS+ and NO ν A running. NO ν A sharply- peaked beam at around 2 GeV due to its position, 14 mrad from the beam axis, while MINOS+, which is on-axis, sees a much higher flux and energy peak [11].	52
2.8	General structure of the MINOS detectors. Scintillator strips are mounted onto steel planes ad read out to PMTs. Adjacent planes are aligned orthogonally to one another to enable 3D reconstruction of interactions in the detector [13]	54
2.9	Orientation of scintillator strips on consecutive Far Detector planes. The Near Detector follows the same pattern, with orthogonally aligned strips on consecutive planes [14]	55
2.10	Schematic and photograph of the MINOS Near Detector. A: The region of the detector in the centre of the beam. B: Entrance point for the coil used to magnetise the detector. C: Electronics racks used to read out the scintillator strips [14].	55
2.11	ND instrumentation layout. Top: Half-instrumented u and v planes in the calorimeter region. These make up 80% of the calorimeter planes. Bottom: Fully instrumented u and v Near Detector planes. These account for every 5th calorimeter plane and every 6th spectrometer plane. The remainder of the spectrometer planes have no scintilla- tor coverage. The colours distinguish different shapes of scintillator	
	module [14]. \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots	56

7
9
C
3
5
6
8

2.19	A Near Detector time slice in u and v . Three distinct events can be	
	picked out by eye, however they cannot be separated using only timing	
	information	<u> 5</u> 9

A Near Detector slice in u and v using the new algorithm. A combina- tion of timing and position information has been used to separate the two tracks from the NC-like event. The algorithm is conservative: two CC-like have significant overlap and are not separated. The z-density of hits in these tracks could potentially be used to separate them but this is not attempted here	69
Reconstruction efficiency vs. energy for MINOS+ simulated data using the final version of the MINOS reconstruction ("Dogwood") and the first version of the MINOS+ reconstruction ("Elm"). A modest improvement in efficiency is seen above 1 GeV [17]	72
Rate of protons delivered to NuMI per week over the lifetime of the MINOS experiment. RHC (orange) and high-energy (red) running are shown separately to standard low-energy FHC running (green). The integrated POT count over time is shown by the blue line [18]	76
Front-on view of fiducial volumes in the Near (left) and Far Detector (right). Preselection eliminates events with vertices outside of these volumes [13]	78
Flow chart illustrating the beam extrapolation procedure. Taken from [19].	81
Diagram demonstrating the difference in angular acceptance between the Near and Far Detectors (not to scale) [20]	83
Example ν_{μ} energy spectra in the Near and Far Detector. Hadrons that produce events of one energy at the Near Detector will produce events of different energies at the Far Detector due to acceptance effects, as illustrated by the hatched bands [21]	84
Figure illustrating the ν_e analysis preselection cuts on reconstructed energy. Below 1 GeV, the sample is poorly reconstructed and domi- nated by NC events, while above 8 GeV, the signal drops to almost zero [22]	93
	A Near Detector slice in u and v using the new algorithm. A combination of timing and position information has been used to separate the two tracks from the NC-like event. The algorithm is conservative: two CC-like have significant overlap and are not separated. The z-density of hits in these tracks could potentially be used to separate them but this is not attempted here

4.1	Rate of protons delivered per week to MINOS and MINOS+ from start of running to 4 May 2015. Protons from the upgraded NuMI beam used in MINOS+ are shown in pink, which has only been operated in FHC mode. FHC MINOS data is shown in green, with RHC data shown in orange. Special, short, high-energy MINOS runs are marked in red [23]
4.2	Rate of selected events per 10^{16} POT in the Near Detector during Run XI [24]
4.3	Rate of selected events per 10^{18} POT in the Far Detector during Run XI. The first bin is approximately 2.4σ from the mean, but is consistent with the statistical expectation for a 70-bin histogram [25]. 103
4.4	Fake data sensitivities to the parameters of the 2-flavour approximate disappearance model in MINOS+ for the old (blue) and new (red) kNN. A third line (black) shows the sensitivity when using perfect PID information. Taken from [26]
4.5	Shift in the MINOS+ two-flavour best fit point caused by $\pm 1\sigma$ shifts to five systematic parameters
4.6	Selected MINOS+ Near Detector data binned in reconstructed energy, along with ND simulated data. The difference between data and simulation is covered by the 1σ systematic error on the simulation [25].107
4.7	Selected MINOS+ Far Detector data binned in reconstructed energy. Predictions at the MINOS+ best fit point and a prediction assuming the best-fit oscillation parameters from the MINOS analysis are also shown
4.8	The number of selected MINOS+ events per reconstructed energy bin over the number expected given no neutrino oscillations is shown. The ratio of the MINOS+ best-fit prediction and MINOS prediction with respect to the no-oscillation prediction are also shown
4.9	Likelihood contour in Δm_{32}^2 - $\sin^2(\theta_{23})$ for MINOS+

4.10	D Likelihood profiles over Δm_{32}^2 (left) and $\sin^2(\theta_{23})$ (right) are shown for the MINOS+ data. Separate profiles are presented assuming normal and inverted mass hierarchy
4.1	1 Selected MINOS+ and MINOS data binned in reconstructed energy. Predictions at the MINOS+ best fit point and a prediction assuming the best-fit oscillation parameters from the MINOS analysis are also shown
4.12	2 Number of events selected for atmospheric analyses over time in MINOS and MINOS+. The blue points show the data analysed for the last published MINOS analysis [27]. An analysis including the MINOS and MINOS+ data has since been performed [1]
4.1	³ The reconstructed energy distributions for three α_{LEM} ranges. The events with energy greater than 5 GeV are combined into a single bin for the fits. The vertical bars through the data points denote statistical uncertainties. The signal predictions assume $\sin^2(2\theta_{13}) =$ $0.051, \Delta m_{32}^2 > 0, \delta = 0, \text{ and } \theta_{23} = \pi/4$. The plots in the left column correspond to data collected in the ν beam mode. The plots in the right column correspond to data collected in the $\bar{\nu}$ beam mode [28]. 113
4.14	4 A two-dimensional $(\Delta m_{32}^2 - \sin^2(\theta_{23}))$ profile of the log-likelihood surface from the MINOS ν_e appearance analysis used in the combined MINOS and MINOS+ result. Note that the surface is asymmetric in the octant of θ_{23} and the sign of Δm_{32}^2 . This surface contributes the majority of the small sensitivity to the octant and hierarchy in the MINOS and MINOS+ combined analyses
4.1	5 68% and 90% confidence limits in $\Delta m_{32}^2 - \sin^2(\theta_{23})$ for the MINOS, MINOS+ combined analysis. The best fit point – marked with the star – is found at $\Delta m_{32}^2 = -2.46 \times 10 - 3 \text{ eV}^2$, $\sin^2(\theta_{32}) = 0.404$ 117
4.10	5 68% and 90% confidence limits in $\Delta m_{32}^2 - \sin^2(\theta_{23})$ for the MINOS, MINOS+ combined analysis are shown along with 90% confidence limits for the MINOS-only analysis and the 2014 T2K ν_{μ} disappearance analysis [29]

4.17	Likelihood profiles over Δm_{32}^2 (left) and $\sin^2(\theta_{23})$ (right) are shown for the combined MINOS and MINOS+ analysis. Separate profiles are presented assuming normal and inverted mass hierarchy
5.1	Illustration of the main time structure of a NuMI beam spill, showing Booster batches and RF bunches. Taken from [2]
5.2	Comparison of the "wrapped batch" and "full spill" fit results for each of 27 data subsets. The results vary over time by more than expected from the statistical error on the fit, due to a problem with the ND-FD GPS connection
5.3	Distribution of the time of the first proton bunch relative to the kicker fire signal, measured by the RWCM
5.4	Diagram of the main components of the RWCM-ND-FD timing sys- tem. Original timing equipment is shown below the detectors, with equipment installed for this analysis shown above. Connection points (marked "C"), show points where equipment is connected temporarily, e.g. travelling GPS units and the TWSTT equipment. Taken from [2].129
5.5	Differences between four measurements of the ND-FD synchronisation are plotted as a function of the Modified Julian Date over the duration the USNO TWSTT test (April 18-19, 2012). The four measurements are direct ND-FD TWSTT (TWdirect); indirect ND-USNO-FD TW- STT (TW); and synchronisation using each pair of ND and FD GPS receivers (GPS_{2-4} and GPS_{8-3}). Taken from [2]
5.6	A single spill (top-left) as seen by the Resistive Wall Current Monitor, with a close-up of the first few bunches (bottom-left). The final PDF generated from this spill is also shown (top-right and bottom- right). Note that the last three batches of this spill are slip-stacked: the unsmeared measurement resolves two peaks for each slip-stacked bunch, while the smeared PDF sees one peak of double intensity 136
5.7	Distribution of the difference between reconstructed and true RAF neutrino interaction times [30]

5.8	Distribution of reconstructed Near Detector event time with respect to the start of a beam spill, modulo the bunch length (18.83 ns). The ND event time resolution is demonstrated by the clear emergence of this peak
5.9	Best fit time of flight, along with 1σ uncertainty, for one-day subsets of the HRTOF dataset. The fit is performed both using two-way time-transfer and GPS for ND-FD synchronisation
5.10	Distribution of RWCM-FD event time modulo 18.83 ns. The RAF distribution is shown uncorrected, and hence has a slight delay and smearing with respect to the CC distribution. The clear appearance of this structure demonstrates the resolution achieved on the event time and RWCM-FD synchronisation, which was performed using TWSTT for this plot
5.11	Likelihood profile in the RWCM-FD time of flight over a range of $\pm 3 \mu s$ from the $v = c$ prediction
5.12	Likelihood profile in the RWCM-FD time of flight, showing the best fit point at $2453935.0\pm0.1\mathrm{ns}$ and surrounding local minima. The early and late minima are excluded at 8.15σ and 4.5σ respectively 141

List of Tables

1.1	Table of the best-known values of the three-flavour neutrino oscillation parameters. Δm^2 is defined such that the vacuum ν_{μ} disappearance probability is symmetric about $\Delta m^2 = 0$, $\Delta m^2 = m_3^2 - (m_2^2 + m_1^2)/2$. Where the best fit depends on the sign of Δm^2 , the value assuming normal hierarchy is given followed in brackets by the value assuming
	inverted hierarchy $[31-33]$
2.1	Various beam parameters before and after ANU upgrades. Taken from [12]
2.2	Percentage compositions of the beam in each running mode in MINOS and MINOS+. Negatively charged mesons are more strongly defocused by the magnetic horns in the new NuMI configuration, leading to a purer ν_{μ} beam. These compositions are estimated from simulated data. 53
2.3	Table showing the distribution of slice completeness and purity for MINOS+ simulated data using the old and new slicing algorithms 73
4.1	Best-fit parameters from fit to MINOS+ data. The negative value of Δm^2_{32} indicates a preference for the inverted hierarchy, however MI- NOS+ has no real sensitivity to the mass ordering. The best fit values of all nuisance parameters are also given. Systematic parameter values are given as shifts from the nominally estimated value in standard deviations
4.2	$-2\Delta \ln L$ for the four degenerate minima in the MINOS+ Δm_{32}^2 - $\sin^2(\theta_{23})$ likelihood surface with respect to the global minimum 109

4.3	Summary of MINOS and MINOS+ runs used in this analysis. A total of 10.71×10^{20} POT of FHC data and 3.36×10^{20} POT of FHC data was taken during MINOS running
4.4	Best-fit parameters from a combined fit to MINOS and MINOS+ data. The MINOS and MINOS+ NC background uncertainties are treated as uncorrelated
4.5	$-2\Delta \ln L$ for the four degenerate minima in the combined MINOS+ and MINOS $\Delta m_{32}^2 - \sin^2(\theta_{23})$ likelihood surface with respect to the global minimum
5.1	Table of systematic uncertainties for 2007 time of flight analysis 123
5.2	Dominant systematic uncertainties in the 2012 high-resolution TOF analysis

Chapter 1

Neutrino Physics

1.1 Three Flavours of Neutrino

The existence of the neutrino was first suggested by Wolfgang Pauli in 1930 in the study of beta decay [34]. The decay was believed at the time to be simply the ejection of an electron from the nucleus,

$$N \to N' + e^-. \tag{1.1}$$

However, the observed energy spectrum of the emitted electrons appeared to follow that of a three-body decay, rather than the sharply-peaked two-body decay spectrum expected. Pauli's solution to the problem was to postulate the existence of a new particle, the 'neutron', which was a very light, very weakly-interacting, neutral fermion. James Chadwick's discovery of a distinct neutral particle in the nucleus in 1932 [35], now known as the neutron, led Pauli's particle to be renamed the *neutrino*.

Soon after, Enrico Fermi developed a model of the beta decay reaction [36] – an early description of the weak force. This model – a four-point interaction – would later be developed by Glashow, Salam and Weinberg into the electroweak interaction used in the Standard Model today. In 1937, Bethe and Peierls used Fermi's interaction to calculate the cross section for inverse beta decay [37], concluding that the observation of this reaction – and hence the direct detection of a neutrino – was "impossible".

This prediction was proved wrong nearly 20 years later, in 1956, when Clyde Cowan and Frederick Reines observed beta-decay antineutrinos generated by a nearby nuclear reactor. A large tank of water mixed with cadmium chloride was placed in the close vicinity of the reactor, sandwiched between two tanks of liquid scintillator. Neutrinos from the reactor interact with protons in the water through inverse beta decay, $\bar{\nu}_e + p^+ \rightarrow n^0 + e^+$. A prompt signal of two photons would come from the annihilation of the positron, followed by a further photon from neutron capture by the cadmium $5\,\mu$ s later. The pair observed around three signal events per hour with a predicted background of less than one per hour [38].

That this neutrino was uniquely partnered with the electron was not known until the 1960s. The muon was discovered in 1936 by Carl Anderson and Seth Neddermeyer [39] in the study of cosmic rays. The question of whether this implied the existence of two distinct types of neutrino was solved in the world's first accelerator neutrino experiment in Brookhaven in 1962 [40]. A pion-rich beam was generated by colliding accelerated protons with a beryllium target. The pions were known to decay by the reaction $\pi^{\pm} \rightarrow \mu^{\pm} + (\nu_{\mu}/\bar{\nu}_{\mu})$. Shielding was used to stop the muons and the interactions of the neutrinos were studied in a nearby spark chamber. The generated neutrinos were seen to exclusively produce muons and not electrons.

The third charged lepton, the tau, was hinted at in 1975 by measurements of missing energy in electron-positron collisions at SLAC in 1975 [41]. It would not be until 2001 that the tau neutrino was directly detected, however its existence could be inferred much sooner using data from the Large Electron-Positron (LEP) collider at CERN. The width of the Z resonance peak is sensitive to the number of weakly-interacting fermions with mass smaller than half that of the Z. Fits to the Z peak at LEP with respect to the number of neutrino generations, as shown in Figure 1.1 resulted in a measurement of $N_{\nu} = 2.9840 \pm 0.0082$.

Direct detection of the tau neutrino was achieved by the DONUT experiment [41] at Fermi National Accelerator Laboratory (FNAL/Fermilab) in Illinois, USA. A detector comprising alternating layers of steel and photographic emulsion was placed downstream of a neutrino beam created from 800 GeV protons. The tau neutrino component of the beam was maximised by designing a thick tungsten target that would absorb pions and kaons – usually used to create a ν_{μ} beam – and leave D_s mesons, which decay to taus 6% of the time. Tau neutrino interactions were selected by their unique topology in the detector. The produced tau is very short-lived and decays almost immediately into pions (and a neutrino), leaving a track with a sharp kink very close to the event vertex.



Figure 1.1: The cross-section for Z production as a function of the LEP centre-of-mass energy. Predictions are shown for two, three and four light, active neutrinos [3].

1.2 Neutrinos and the Standard Model

The Standard Model is a theory describing all 17 known fundamental particles – six quarks, three charged leptons, three neutrinos, four exchange bosons and a Higgs boson. The model combines the theories of the strong force – quantum chromodynamics – with a unified theory of the weak and electromagnetic forces, the electroweak interaction.

Neutrinos interact exclusively through the weak force, via the massive exchange bosons W^{\pm} and Z^{0} . Feynman diagrams for these interactions are shown in Figure 1.2. Charged-current (CC) interactions, mediated by the W boson, may convert between quarks differing by one unit of charge and between leptons and their corresponding neutrinos. Neutral-current (NC) interactions do not carry charge or change flavour, working analogously to photon exchange in the electromagnetic interaction.

The weak interaction maximally violates parity [42, 43], coupling only to lefthand-chiral particles, hence the Standard Model contains only left-handed neutrinos



Figure 1.2: Feynman diagrams for the two types of neutrino interaction. Charged-current interactions (left) involve the exchange of a charged W boson, while neutral-current interactions (right) have a neutral Z boson as propagator.

(and right-handed antineutrinos). Masses in the Standard Model are generated by the Higgs mechanism, which couples left- and right-handed chiral fields. Generating neutrino masses means either including right-handed neutrino fields or including an alternative mass term, e.g. Majorana mass, or both [44].

1.3 Neutrino Mass

Non-zero neutrino masses remain, at time of writing, not directly detected. However they can be inferred by the observation of *neutrino oscillations*, discussed later in this chapter.

Neutrino oscillations necessitate the refining of the concept of neutrino mass. The neutrino states that interact with matter – the weak eigenstates – do not have a well-defined mass, but can be described as a superposition of mass eigenstates. Current knowledge fits the model that three neutrino mass eigenstates exist, with no more than one having zero mass. Two independent mass-squared differences are measured and provide a lower bound to the heaviest two eigenstates.

The best constraints on the sum of the neutrino masses currently come from cosmological observations, though the limits can be highly model-dependent. The cosmological upper limit is given by the latest results from the Planck collaboration, based on observations of temperature and polarisation anisotropies in the cosmic microwave background combined with other astrophysical data [45]:

$$\sum_{j} m_j < 0.23 \,\text{eV}, \quad (95\% \ C.L.). \tag{1.2}$$

Attempts at direct measurements are usually made by precise observations of the beta decay spectrum. The maximum energy of the beta electron is reduced by a non-zero $\bar{\nu}_e$ mass. Precise measurements of the endpoints of this spectrum should be sensitive to this effect. The Troitzk [46] experiment set the current best direct-measurement limit on the $\bar{\nu}_e$ mass of < 2.05 eV at 95% CL. The KATRIN experiment [47], which is due to begin running in 2016, is expected to push this limit down to < 0.2 eV.

A further set of experiments may be able to measure an effective neutrino mass using neutrinoless double beta decay $(0\nu\beta\beta)$. This is a process only allowed if the neutrino is a *Majorana* particle – it is identical to its own antiparticle [44]. These experiments are sensitive to the effective Majorana neutrino mass $m_{\beta}\beta = \sum_{i=1}^{3} U_{ei}^2 m_{\nu_i}$, the weighted sum of the neutrino mass eigenstates with the ν_e elements of the PMNS matrix, defined later in the chapter. Current combined limits on this mass range from 130 to 310 meV depending on the choice of Nuclear Matrix Element used to ralate the half-life limit to a mass limit. [48].

1.4 Three-Flavour Neutrino Oscillations

Massive neutrinos enable the possibility of neutrino mass mixing, analogous to that seen in the quark sector. This is the idea that the distinct flavours of neutrino we create and detect do not each have a uniquely defined mass. Rather, there are three mass eigenstates and the creation of a neutrino forces it into one of these states (i = 1, 2, 3) with a certain probability, $U_{\alpha i}$, determined uniquely by the flavour, α . More precisely, given only small differences between these masses, a coherent superposition of mass eigenstates is created. A direct consequence of this mass mixing is the phenomenon of neutrino flavour oscillation – a neutrino created with flavour α may be later detected as having a flavour β with a probability dependent on the mass mixing coefficient, the neutrino energy and the propagation distance. That this phenomenon follows from neutrino mass can be shown with standard (non-relativistic) quantum mechanics.

A neutrino flavour eigenstate, $|\nu_{\alpha}\rangle$ may be considered a superposition of mass eigenstates $|\nu_i\rangle$, with mixing coefficients $U_{\alpha i}$.

$$\left|\nu_{\alpha}\right\rangle = \sum_{i=1}^{3} U_{\alpha i} \left|\nu_{i}\right\rangle, \qquad (1.3)$$

and vice versa,

$$|\nu_i\rangle = \sum_{i=1}^3 U^*_{\alpha i} |\nu_\alpha\rangle.$$
(1.4)

The time-dependent Schrödinger equation,

$$i\hbar\frac{\partial}{\partial t}\Psi = \hat{H}\Psi \tag{1.5}$$

r

has plane-wave solutions for the mass eigenstates of the form

$$\left|\nu_{i}(t)\right\rangle = e^{-i(E_{i}t - \vec{p_{i}} \cdot \vec{x})} \left|\nu_{i}(0)\right\rangle.$$
(1.6)

Making relativistic approximations,

$$E_i = \sqrt{p_i^2 + m_i^2} \simeq p_i + \frac{m_i^2}{2p_i} \approx E + \frac{m_i^2}{2E}, \quad p \approx E \quad \text{and} \quad t \approx L, \tag{1.7}$$

the propagation of the flavour eigenstates can be described as

$$|\nu_{\alpha}(t)\rangle = \sum_{i=1}^{3} U_{\alpha i} e^{-im_{i}^{2}L/2E} |\nu_{i}(0)\rangle = \sum_{i} \sum_{\beta} U_{\alpha i} U_{\beta i}^{*} e^{-im_{i}^{2}L/2E} |\nu_{\beta}(0)\rangle$$
(1.8)

Here it can already be seen that the amplitude of the flavour state oscillates as a function of propagation distance L. Calculating the probability of measuring this superposition to be in flavour state β gives

$$P_{\alpha \to \beta}(L, E) = |\langle \nu_{\beta}(t) | \nu_{\alpha}(0) \rangle|^{2} = \sum_{i} \sum_{j} U_{\beta i}^{*} U_{\beta j} U_{\alpha j}^{*} U_{\alpha i} \exp(i\frac{L}{2E}\Delta m_{ij}^{2})$$
(1.9)

where $\Delta m_{ij}^2 = m_i^2 - m_j^2$.

The neutrino transition probability is a function of neutrino energy, E, and propagation distance L. The behaviour of the equation is governed by a set of physical constants, Δm_{ij}^2 and $U_{\alpha i}$. The determination of these constants is the primary purpose of experimental neutrino oscillation physics. The coefficients $U_{\alpha i}$ make up U, the PMNS matrix, after Pontecorvo, Maki, Nakagawa and Sakata [49]. By convention, the matrix is parametrised into three *mixing angles* and a phase, δ , as follows.

$$U = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix}$$
(1.10)
$$= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}$$
(1.11)
$$= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix}$$
(1.12)

The three component matrices in this factorised form have historically been separately probed by different classes of experiment. The θ_{23} sector is identified with atmospheric neutrino oscillations, the θ_{12} sector with solar neutrino oscillations, and the θ_{13} sector with electron-neutrino appearance. A non-zero δ generates CPviolation in the oscillation equations and is subsequently referred to as δ_{CP} .

1.4.1 Mass Hierarchy

Neutrino oscillations allow us to infer the differences between the squares of the three neutrino mass eigenvalues, $\Delta m_{21}^2 = m_2^2 - m_1^2$, $\Delta m_{32}^2 = m_3^2 - m_2^2$ and $\Delta m_{31}^2 = \Delta m_{32}^2 - \Delta m_{21}^2$. The two-flavour oscillation probabilities in vacuum are functions of the squares of these squared mass differences, and therefore independent of the sign of the mass differences or the mass hierarchy. The probabilities including matter effects are sensitive to the sign of the Δm^2 parameters and hence experiments



Figure 1.3: Illustration of the two neutrino mass hierarchy scenarios allowed by current data. The value Δm_{21}^2 is known to be positive, but the sign of Δm_{32}^2 is unknown. Taken from [4].

probing neutrino oscillations in matter may be used to determine the mass hierarchy. The value of Δm_{21}^2 has been determined by solar neutrino experiments, in which significant matter effects also allow determination of its sign. Δm_{21}^2 is positive given that by convention ν_1 is defined as the dominant component of the electron neutrino. The lack of such significant matter effects in atmospheric and beam experiments means that the sign of Δm_{32}^2 remains uncertain. This leads to two possible orderings of the neutrino mass spectrum: $m_1 < m_2 < m_3$, named the normal hierarchy, and $m_3 < m_1 < m_2$, the inverted hierarchy, illustrated in Figure 1.3.

1.4.2 Oscillations in Matter

The interaction of neutrinos with matter, though weak, has a measurable effect on neutrino oscillation probabilities that must be taken into account when describing neutrinos that travel a significant distance in a dense medium – for example the Earth or Sun.

Coherent forward scattering of neutrinos with matter leads to the inclusion of interaction potential terms on the neutrino's energy, separated into an NC and a CC term, $V = V_{NC} + V_{CC}$:

$$V_{NC} = \frac{\pm 1}{\sqrt{2}} G_F n_n \tag{1.13}$$

$$V_{CC} = \pm \sqrt{2}G_F n_e \tag{1.14}$$

where G_F is Fermi's Constant and n_n and n_e are the number density of neutrons and electrons in the medium respectively. The sign on the potential is opposite for neutrinos and antineutrinos. Importantly, the NC potential affects each flavour eigenstate of neutrino identically, whereas the CC potential affects only the ν_e eigenstate. This leads to a change of phase between the mass eigenstates.

Using a two-flavour approximation for simplicity, the neutrino flavour eigenstates in vacuum have the following Hamiltonian:

$$H_V = \frac{\Delta m^2}{4E} \begin{pmatrix} -\cos 2\theta & \sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix}$$
(1.15)

where multiples of the identity matrix have been neglected since they will have no effect on oscillations. Adding the CC scattering potential gives

$$H_M = H_V + \begin{pmatrix} V_{CC} & 0\\ 0 & 0 \end{pmatrix} \tag{1.16}$$

We would like to recast this Hamiltonian into the form of Equation 1.15, with an effective squared-mass difference and mixing angle, Δm_M^2 and θ_M . Doing this gives:

$$\Delta m_M^2 = C \Delta m^2 \tag{1.17}$$

$$\sin 2\theta_M = \frac{\sin 2\theta}{C} \tag{1.18}$$

where

$$C = \sqrt{(\cos 2\theta - \frac{2E}{\Delta m^2} V_{CC})^2 + \sin^2 2\theta}.$$
 (1.19)

Accounting for matter effects is then only a case of replacing θ with θ_M in the two-flavour probability equation.

There are a few things to note. First, the oscillation probability now depends on C, which is a non-symmetric function of Δm^2 . This means matter effects allow us to distinguish between the neutrino mass hierarchies. Second, matter effects introduce a difference between the neutrino and antineutrino oscillation probabilities due to the \pm sign in V_{CC} . Finally, θ_M may be maximal even when the vacuum mixing angle is small, if $\frac{2E}{\Delta m^2}V_{CC} = \cos 2\theta$. This is called the Mikheyev-Smirnov-Wolfenstein (MSW) resonance and can lead to large deviations from the behaviour expected in vacuum.

1.4.3 Two-Flavour Approximation

In order to connect the three-flavour neutrino oscillation phenomenology now commonly used in experimental analysis to historical results, it is worthwhile to look at the two-flavour approximations used until recently in the atmospheric oscillations $(\theta_{23}, \Delta m_{32}^2)$ sector.

From Equation 1.9, with $\alpha = \beta = \mu$:

$$P_{\mu \to \mu} = \sum_{i} \sum_{j} U_{\beta i}^{*} U_{\beta j} U_{\alpha j}^{*} U_{\alpha i} \exp(i \frac{L}{2E} \Delta m_{ij}^{2})$$

= 1 - 4 \sum_{i=1,2,3}^{2} \sum_{j < i}^{2} |U_{\mu i}|^{2} |U_{\mu j}|^{2} \sin^{2}(1.27 \Delta m_{ij}^{2} L/E), \qquad (1.20)

where a conversion to the conventional units (L in km, E in GeV, Δm^2 in eV²) has been made. Making the assumption that θ_{13} is zero,

$$U_{\mu 1} = -s_{12}c_{23} \tag{1.21}$$

$$U_{\mu 2} = c_{12} c_{23} \tag{1.22}$$

$$U_{\mu3} = s_{23}.\tag{1.23}$$



Figure 1.4: Comparison of the ν_{μ} disappearance probability in the two- and three-flavour oscillation models. The mixing angles used (in radians) are $\theta_{12} = 1.17$, $\theta_{23} = 0.68$ and $\theta_{13} = 0.16$ and the mass-splittings (in eV²) are $\Delta m_{21}^2 = 7.58 \times 10^{-5}$ and $\Delta m_{32}^2 = 2.41 \times 10^{-3}$.

Under the further condition that $\Delta m_{21}^2 \ll \Delta m_{32}^2$, we obtain the two-flavour ν_{μ} disappearance probability:

$$P_{\mu \to \mu} = 1 - \sin^2(2\theta_{23}) \sin^2(1.27\Delta m_{32}^2 L/E).$$
(1.24)

This function is plotted in Figure 1.4. The depth and position of the first minimum are determined by $\sin^2(2\theta_{23})$ and Δm_{32}^2 respectively. The three-flavour disappearance probability is also shown. Qualitatively, the introduction of non-zero θ_{13} reduces the depth of the minimum, while the effect of Δm_{21}^2 is to slightly increase the oscillation frequency.

1.5 Evidence for Neutrino Oscillations

1.5.1 θ_{12} : Solar Neutrino Oscillations

The first evidence for neutrino flavour oscillations was seen in the θ_{12} sector in experiments looking at neutrinos produced in the Sun.

The Solar Neutrino Problem

In the late 1960s, Ray Davis made the first observation of neutrinos from the Sun in an experiment situated in the Homestake mine [50]. This was a counting experiment based on measuring the amount of radioactive Argon-37 produced in a tank of highly-pure tetrachloroethylene ($\nu_e + {}^{37}\text{Cl} \rightarrow {}^{37}\text{Ar}^+ + e^-$). Davis, studying the flux of solar neutrinos, noticed that his signal was much smaller than predicted by the Standard Solar Model (SSM), the flux predictions for which are shown in Figure 1.5. Understandable suspicions, first of unaccounted-for systematic errors and then of problems with the SSM, were proven unfounded by a succession of later experiments [51,52], leaving the possibility of neutrino flavour oscillations over the distance between the Earth and the Sun.



Figure 1.5: Components of the solar neutrino spectrum as predicted by the BSO5(OP) SSM [5].

SNO

The most powerful early evidence for the existence of solar neutrino oscillations came from the Sudbury Neutrino Observatory (SNO). The experiment consisted of 1000 tons of heavy water (D₂O) in a spherical tank lined with photomultiplier tubes (PMTs), designed to detect Čerenkov radiation. In order to confirm the disappearance of solar ν_e 's, both a ν_e -CC and an NC measurement were made – the NC rate can be used to measure the total number of neutrinos of any flavour passing through the detector.

Firstly, the ν_e flux was measured via the charged-current inverse beta decay channel $\nu_e + d^+ \rightarrow p^+ + p^+ + e^-$, and again observed to be significantly lower than predicted by the SSM (SSM). In addition to this, a measurement of the total neutrino flux from the sun was made by observing NC interactions in the detector of the form $\nu + d^+ \rightarrow p^+ + n + \nu$. An NC flux consistent with the number of predicted solar ν_e s would indicate beyond reasonable doubt that a significant fraction of the ν_e flux was changing flavour between the points of creation and detection. The total solar neutrino flux from the NC channel was measured to be consistent with the SSM. The deficiency in ν_e 's led to a result claiming that around two thirds of the solar flux was reaching the Earth either as tau or muon neutrinos [53].

By this stage, a solid theoretical framework was in place for the study of neutrino oscillations. In 1957, Bruno Pontecorvo suggested that a massive (Majorana) neutrino might undergo $\nu \rightleftharpoons \bar{\nu}$ oscillations, in analogy to the K^0 . The discovery of further generations of fermions and then of flavour mixing in the quark sector led to speculation of similar mechanisms in the lepton sector. A paper by Maki, Nakagawa and Sakata in 1962 [54] first proposed the mixing between the weak eigenstates and 'true' eigenstates of the neutrino that are now believed to give rise to flavour oscillations. Pontecorvo would go on to write several papers extending this work [55–57]. In 1975, he and Samoil Bilenky first wrote about the mixing of Dirac neutrinos, in analogy with recently-discovered quark mixing and the CKM matrix.

The solar oscillation sector is now dominated by reactor experiments – these measure the flux of electron (anti) neutrinos created in nuclear reactors. The current best measurement of the parameters θ_{12} and Δm_{12}^2 comes from combining the results of SNO with those from one such reactor experiment, KamLAND.

KamLAND

KamLAND [58] was a liquid scintillator experiment in Japan designed to measure the energy spectrum of electron antineutrinos created in nuclear reactors across the country. It is conceptually similar to the Cowan-Reines experiment: inverse beta decay is used to detect the neutrinos, whose energy is reconstructed from the scintillation light produced. The neutrinos have energies in the range of 1 to 10 MeV and the average propagation distance is 180 km.



Figure 1.6: Best fit along with 68%, 95% and 99.73% C.L. contours for the parameters Δm_{21}^2 and $\tan^2 \theta_{12}$ from a combined analysis of KamLAND and SNO data [6].

The constraints on θ_{12} and Δm_{12}^2 from a combined analysis of KamLAND and SNO are shown in Figure 1.6, giving the current best estimate of the solar neutrino parameters:

$$\Delta m_{12}^2 = 7.54^{+0.26}_{-0.22} \times 10^{-5} \text{eV}^2 \tag{1.25}$$

$$\sin^2 \theta_{12} = 0.307^{+0.018}_{-0.016} \tag{1.26}$$

1.5.2 θ_{23} : Atmospheric Neutrino Oscillations

The upper atmosphere of the Earth is an isotropic source of electron- and muontype neutrinos. Interacting cosmic rays produce large showers of hadrons including charged pions and kaons, which decay weakly to leptons and neutrinos, mainly ν_{μ} and ν_e . Models of atmospheric neutrino production predict the overall neutrino flux to about 20%, however the ratio of ν_{μ} to ν_e – approximately 2 to 1 – has an uncertainty of only around 2%. Neutrinos are generated with energy of the order of 1 GeV and travel through the Earth relatively unobstructed, giving an experiment at (or near) the surface a large range of L/E to probe, with L varying from 15 to 13000 km.

Kamiokande

The first observations of oscillations in the θ_{23} – or atmospheric – sector was made in 1988 by the Kamiokande experiment in Japan. This was a water Čerenkov detector – a large, shielded tank of water surrounded on all sides by photomultiplier tubes. Kamiokande made separate measurements of the flux of atmospheric ν_e and ν_{μ} . While the ν_e value was in agreement with predictions (in absence of oscillations), a deficit of approximately 60% was seen in the ν_{μ} flux [59]. This result was replicated by several other experiments [60,61] and became known as the Atmospheric Neutrino Anomaly.

Super-Kamiokande

The deficit was shown to be due to neutrino oscillations in 1998 in the Super-Kamiokande experiment, an upgraded version of the Kamiokande experiment, containing 10 times more PMTs and 15 times as much water [62]. By reconstructing the angle of incidence of neutrinos interacting in the detector, an estimate of the propagation length could be found, allowing the L/E dependence of the ν_{μ} deficit – consistent with ν_{μ} to ν_{τ} oscillations – to be determined [63].

More recently, measurements of the atmospheric oscillation sector have been made by beam ν_{μ} disappearance experiments, of which MINOS(+) is one. Another is the Tokai to Kamioka (T2K) experiment in Japan.

T2K

T2K detects neutrinos from a ν_{μ} beam created at the J-PARC facility [64] using two separate detectors. A smaller *near detector* sits just in front of the beam, measuring the energy distribution and flux of neutrinos before any oscillation. The beam then passes through the Super-Kamiokande *far detector*, which observes the oscillated spectrum. The detector can reconstruct ν_{μ} and ν_{e} events, allowing both the ν_{μ} disappearance and ν_{e} appearance probabilities to be probed. T2K is an *off-axis* beam experiment, meaning that its far detector is situated several degrees from the beam axis. Decay kinematics in the beam lead to more monochromatic – though lower intensity – spectrum at this position, providing a greater proportion of events at the L/E point that maximises the oscillation probability and removing high energy events that may cause additional backgrounds.

MINOS

MINOS (Main Injector Neutrino Oscillation Search) is a two detector experiment sitting the NuMI (Neutrinos at the Main Injector) beam at Fermilab, Illinois. The experiment will be described in detail in Chapter 2. MINOS ran from 2005 to 2012, in which time it produced world-leading measurements of the atmospheric mixing parameters, in particular Δm_{23}^2 . The best fits and 68% and 90% confidence limits for the parameters $\sin^2(\theta_{23})$ and Δm_{32}^2 are shown for MINOS and T2K in Figure 1.7.

1.5.3 θ_{13} : Reactor ν_e Disappearance, Beam ν_e Appearance

Measurements of θ_{13} are made either by observing an appeared ν_e signal in a ν_{μ} beam, or the disappearance of $\bar{\nu_e}$ from nuclear reactors. This has traditionally been the most difficult mixing angle to measure due to its small size. Evidence was seen for non-zero θ_{13} by MINOS [65], T2K [66] and Double-Chooz [67] – a reactor disappearance experiment – before the first 5σ non-zero measurement was made at Daya Bay in 2012 [7].


Figure 1.7: Allowed regions for $\sin^2(\theta_{23})$ and Δm_{32}^2 at the 68% (dashed lines) and 90% (solid lines) in MINOS and T2K. Two years of MINOS+ atmospheric data are included in the MINOS contour. The best fits are shown by the red and blue dots for MINOS and T2K respectively.



Figure 1.8: Allowed regions for $\sin^2 2\theta_{13}$ and $|\Delta m_{ee}^2|$ at the 68.3%, 95.5%, and 99.7% confidence level from the Daya Bay experiment. The best estimate of the oscillation parameters is given by the black dot. The adjoining panels show the profile likelihoods for each of the parameters. Taken from [7].

Daya Bay

Daya Bay is a reactor experiment based in China. The experiment consists of two sets of near detectors and one set of far detectors, at average baselines of 470 m, 576 m and 1648 m from the Daya Bay Nuclear Power Plant respectively. The length scale is noticeably shorter than the KamLAND experiment, to avoid effects of atmospheric oscillations. Each detector is a tank containing 20 t of gadolinium-doped liquid scintillator surrounded by 21 t of undoped liquid scintillator for detecting escaped photons, further surrounded by 37 t of mineral oil for shielding. The outer layer is lined with PMTs to detect the scintillation light. Daya Bay's most recent result at time of writing is shown in Figure 1.8.

1.5.4 Mass Hierarchy

The determination of the mass hierarchy has several implications. It will help to better constrain the value of $\delta_{\rm CP}$, as well as potentially giving a lower bound on the neutrinoless double beta decay rate. The most promising method for determining the hierarchy appears to be the use of long-baseline $\nu_{\mu} \rightarrow \nu_{e}$ appearance experiments. These can exploit both matter effects as well as second order dependencies on the sign of Δm_{32}^2 in the $\nu_{\mu} \rightarrow \nu_{e}$ oscillation probability. One such experiment is T2K. A second, recently constructed, is NO ν A.

$NO\nu A$

The NuMI Off-axis ν_e Appearance (NO ν A) experiment [12] is a long-baseline off-axis accelerator neutrino experiment using the Neutrinos at the Main Injector (NuMI) beam at Fermilab. This is the same beam that serves the MINOS+ experiment and it is described in detail in the next chapter. Like T2K, NO ν A measures the disappearance of ν_{μ} and appearance of ν_e in a predominantly ν_{μ} beam. As such, it is sensitive to the atmospheric mixing parameters as well as θ_{13} , δ_{CP} and the mass hierarchy.

NO ν A uses two detectors: a 0.3 kt near detector close to the beam and a 14 kt far detector 810 km from the production point. The detectors are tracking calorimeters, made up of PVC cells filled with liquid scintillator. Light from the scintillator is read out by avalanche photodiodes via wavelength-shifting fibres. The detectors are well-suited to identify electromagnetic showers generated in ν_e CC events, as the radiation length in the detector is many times larger than dimensions of the cells. NO ν A construction was completed in September 2014; Figure 1.9 shows the projected sensitivity of a combination of NO ν A and T2K, with between 1 and 3σ resolution by 2020 depending on the value of δ_{CP} .

1.5.5 δ_{CP}

The level of CP violation currently observed in the universe is not enough to explain its matter-antimatter asymmetry. The measurement of non-zero θ_{13} allows for the possibility of CP violation in neutrino oscillations, which would be a small step



Figure 1.9: Sensitivity to the neutrino mass hierarchy for 6 years of NO ν A running in combination with T2K $\nu_{\mu} \rightarrow \nu_{e}$ data [8].

towards remedying this problem. The violation is of the form [68] $P(\nu_l \to \nu_{l'}) \neq P(\bar{\nu}_l \to \bar{\nu}_{l'})$. The magnitude of this violation is proportional to $\sin \delta_{CP}$. Attempts to measure δ_{CP} are generally based around measurement of the *CP*-asymmetry,

$$A_{CP}^{(\mu e)} = \frac{P(\nu_{\mu} \to \nu_{e}) - P(\bar{\nu}_{\mu} \to \bar{\nu}_{e})}{P(\nu_{\mu} \to \nu_{e}) + P(\bar{\nu}_{\mu} \to \bar{\nu}_{e})} \propto \frac{\sin \delta_{CP}}{\sin \theta_{13}}$$
(1.27)

Both T2K and NO ν A make this measurement. Their sensitivity to δ_{CP} is shown in 1.10. Combined analysis of existing data estimates $\delta_{CP} \approx 3\pi/2$ (shown in Table 1.1) which would maximise the magnitude of CP violation observed.

1.5.6 θ_{23} Octant

The two flavour ν_{μ} disappearance probability is a function $\sin^2 2\theta_{23}$, making it exactly symmetrical in θ_{23} about 45°. In the three-flavour probability, this is an approximate symmetry in $\sin^2(\theta_{23})$ about $1/\cos^2\theta_{13}$. The symmetry is broken by terms containing $\sin^2(\theta_{23})$, which can be found for example in the ν_{μ} disappearance probability in matter or in the ν_e appearance probability.



Figure 1.10: Sensitivity to δ_{CP} for 6 years of NO ν A running in combination with T2K $\nu_{\mu} \rightarrow \nu_{e}$ data [8].

The resolution of this degeneracy will help to constrain other parameters such as δ_{CP} and the mass hierarchy, however it is a particularly difficult measurement. Current global fits [31] give a slight preference for the lower octant $(\sin^2(\theta_{23}) < 1/\cos^2\theta_{13})$, but a definitive measurement is unlikely for many years.

1.6 Summary of Three-Flavour Oscillation Parameters

Table 1.1 summarises the current knowledge of the neutrino oscillation parameters as listed by the Particle Data Group [69].

Parameter	Best Fit Value $(\pm 1\sigma)$
Δm_{12}^2	$7.54^{+0.26}_{-0.22} \times 10^{-5} \mathrm{eV}^2$
$ \Delta m^2 $	$2.43 \pm 0.06 \ (2.38 \pm 0.06) \times 10^{-3} \mathrm{eV}^2$
$\sin^2(\theta_{12})$	0.308 ± 0.017
$\sin^2(\theta_{23})$	$0.437^{+0.033}_{-0.023} \ (0.455^{+0.039}_{-0.031})$
$\sin^2(\theta_{13})$	$0.0234_{-0.0019}^{+0.0020} \ (0.0240_{-0.0022}^{+0.0019})$
$\delta_{\rm CP}/\pi$ (2 σ range)	$1.39_{-0.27}^{+0.38} \ (1.31_{-0.33}^{+0.29})$

Table 1.1: Table of the best-known values of the three-flavour neutrino oscillation parameters. Δm^2 is defined such that the vacuum ν_{μ} disappearance probability is symmetric about $\Delta m^2 = 0$, $\Delta m^2 = m_3^2 - (m_2^2 + m_1^2)/2$. Where the best fit depends on the sign of Δm^2 , the value assuming normal hierarchy is given followed in brackets by the value assuming inverted hierarchy [31–33].

Chapter 2

The NuMI Beam, MINOS Detectors and MINOS+

2.1 Overview

MINOS+ is a continuation of the MINOS experiment, a long-baseline neutrino oscillation experiment designed to make precision measurements of the atmospheric mixing angle and mass-splitting parameters. MINOS did this by measuring the survival probability of approximately 3 GeV accelerator-generated muon-neutrinos as a function of energy over a 735 km baseline. MINOS+ sits in a beam peaked at approximately 6 GeV, probing a region of L/E away from the atmospheric oscillation maximum, but can still be used to constrain the values of $\sin^2(\theta_{23})$ and Δm_{32}^2 .

The experiment consists of three main components. First, muon-neutrinos are created using the NuMI beam at Fermi National Accelerator Laboratory (FNAL or Fermilab), Illinois. The unoscillated energy spectrum of the beam is measured in the Near Detector (ND), approximately 1 km from the neutrino creation point. The Far Detector (FD), situated 735 km from the beam target at the Soudan Underground Laboratory, Minnesota, measures the energy spectrum of the oscillated beam. The geographical layout of the experiment may be seen in Figure 2.1.

In April 2012, the NuMI beam that served the experiment was upgraded to increase its energy and intensity with the purpose of supplying neutrinos to the NO ν A experiment. The MINOS experiment received some upgrades during the shutdown to help its electronics cope with the demands of the new beam configuration,

and began taking data as MINOS+ when NuMI was switched back on in September 2013.



Figure 2.1: Geographical layout of the MINOS+ experiment. The NuMI beam and Near Detector are located at Fermilab in north-east Illinois, the Far Detector at the Soudan Mine in northern Minnesota [9].

2.2 MINOS+ Physics Goals

The design of MINOS facilitated several types of physics analysis, including sterile neutrino, ν_e appearance, and large-extra-dimension searches [70, 71]. Many of these will be continued in MINOS+. The region of L/E probed by the experiment is not well studied and may have potential to uncover as-yet-unseen effects in the neutrino sector. The most-obviously achievable goal of MINOS+ is to further improve MINOS' world-leading measurements of Δm_{32}^2 , as well as $\sin^2(\theta_{23})$. Other goals include constraining active-to-sterile mixing and large extra dimensions searches [72].

Figure 2.2 shows the ranges of L/E probed by MINOS, MINOS+ and NO ν A along with the ν_{μ} disappearance and ν_{e} appearance probabilities for a ν_{μ} beam as a function of L/E. The sensitivity of MINOS+ to atmospheric-scale ν_{μ} disappearance is reduced by the increase in beam energy, with a disappearance probability of 10% expected at the beam peak. Given the increased event rate in MINOS+, the experiment retains a reasonable amount of sensitivity to the atmospheric mixing parameters measured by MINOS.



Figure 2.2: The $\nu_{\mu} \rightarrow \nu_{\mu}$ and $\nu_{\mu} \rightarrow \nu_{e}$ oscillation probabilities are shown as a function of L/E. The shaded bands show the energy range within $\pm 1 \sigma$ of the mean beam neutrino energy in MINOS+, MINOS and NO ν A at their respective Far Detectors. For MINOS+, the mean is taken to be 6.6 \pm 2.5 GeV, for MINOS 3.5 \pm 1.2 GeV and for NO ν A 2.1 \pm 0.5 GeV.

2.3 The NuMI Beam

Neutrinos for MINOS+ and NO ν A are provided by the NuMI beam. NuMI follows the design of a conventional neutrino beam: first, accelerated protons are collided with a graphite target. Next, charged hadrons generated by these collisions are focused and are given time to decay, producing neutrinos.

The design of the NO ν A experiment has necessitated an increase in the energy and intensity of this beam from the levels designed for the MINOS experiment, leading to

a shutdown from April 2012 to September 2013 to perform a range of improvements known as the Accelerator and NuMI Upgrades (ANU) [12], The operation and design of the beam is discussed in the following paragraphs as well as some details of the ANU.

2.3.1 FNAL Accelerator Complex

Protons are provided to the NuMI target by a large accelerator complex at Fermilab that originally served the Tevatron collider experiments. A schematic of the complex is shown in Figure 2.3. Several upgrades have been made as part of the ANU to provide a higher-energy, higher-intensity ν_{μ} beam for the NO ν A experiment. A brief overview of the main stages of the acceleration process is presented, with any upgrades noted.



Figure 2.3: Schematic of the Tevatron accelerator complex. The Tevatron stopped operation in September 2011. Protons are extracted from the Main Injector to the NuMI target [10].

First, a radio-frequency quadrupole pre-accelerator system produces 750 keV H⁻ ions, while a chopper is used to set the beam pulse length for the next stage of acceleration [73]. This is a 150 m-long linear accelerator (linac), using oscillating electric fields to accelerate the ions to 400 MeV. The H⁻ are then passed through thin carbon foils to strip their electrons, leaving a 400 MeV proton beam at entry to the next accelerator, the Booster.

The Booster is a circular synchrotron accelerator with circumference 474 m. It accelerates 400 MeV protons from the linac to 8 GeV at a repetition rate of 15 Hz. A total of 96 magnets, grouped into sets of 4 by field strength, steer and focus the beam as it is accelerated by an RF electric field. This field is generated in 18 cavities across the ring, and increases in frequency from 37.9 to 52.8 MHz to match the speed of the beam. After 66.7 ms of acceleration, a $1.58 \,\mu$ s-long *batch* of protons is extracted to the Main Injector.

The Main Injector (MI) is a 3319.4 m-circumference synchrotron which originally served to inject protons into the main Tevatron ring. It can accelerate protons from 8 GeV to a maximum of 150 GeV. The circumference of the accelerator allows it to accept a maximum of seven Booster batches in a single cycle. In practice, six are accepted in order to minimise beam losses during extraction. The MI cycle time has been reduced from 2.2 s to 1.33 s as part of the ANU. This constitutes the majority of the planned increase in beam power and has required several upgrades, namely the installation of two new RF cavities (for a total of 20) to increase the acceleration rate and improvements in the cooling systems of many of the accelerator's components [12].

Since 2005, *slip stacking* [74] has been used in the Main Injector to increase the intensity of the beam. The intensity of a proton batch is doubled by injecting a second batch at a slightly different phase and accelerating frequency, before merging. This technique was originally used to increase the number of protons available for antiproton production for the Tevatron, but has also been used to produce a higher-intensity proton beam for extraction to NuMI. By the end of MINOS running, *multi-batch* slip stacking in the Main Injector allowed acceleration of 11 Booster batches in one MI cycle [75].

Sitting in the same ring as the Main Injector – and with the same circumference – is a third synchrotron: the Recycler. The Recycler was originally used as an anti-proton storage ring for the Tevatron. The ring has been converted to slip stack Booster batches before extraction to the Main Injector. The matching circumferences of the rings make a single-turn (approximately $11 \,\mu$ s) extraction possible. This leaves the Main Injector to simply accelerate the beam to 120 GeV before extraction to NuMI, and facilitates the reduction in its cycle time. Slip stacking in the Recycler is making progress toward the ANU target of 12 batches, with 8-batch slip stacking being introduced in April 2015, shown in the MINOS Near Detector in Figure 2.4.

As in the MINOS era, six batches of 120 GeV protons are extracted from the Main Injector at the end of each cycle. Extraction is performed by a kicker magnet, while five quadrupole magnets guide the beam towards the NuMI target, at an angle of 3.3° from horizontal. These guiding magnets have been replaced with models better able to cope with the increased power and cooling demands from the increased extraction frequency, while the power supply and cooling system for the kicker magnet have also been upgraded [12].



Figure 2.4: Time profile of the the neutrino beam as seen at the MINOS+ Near Detector, April 2015. The first two batches are doubled in intensity by slip stacking.

2.3.2 Neutrino Production

Neutrinos in a conventional beam are primarily produced by focusing positively charged pions and kaons, which may decay into muon-neutrinos by the channels

$$\pi^+ \to \mu^+ \nu_\mu, \quad K^+ \to \mu^+ \nu_\mu. \tag{2.1}$$

These particles are produced by colliding accelerated protons from the Main Injector with a graphite target. Two configurable magnetic horns provide the focusing field, while a long pipe gives the mesons time to decay. A schematic of the NuMI beamline is shown in Figure 2.5.

The new NuMI target was installed in 2013. A technical drawing is shown in Figure 2.6. The target is composed of 12 graphite plates, each 6.4 mm wide and 100 mm long, clamped between water-cooled aluminium plates. The sizes of the MINOS targets and casings were limited to 30 mm at one end by the fact that they had to sit at the same point in the beam as one of the magnetic focusing horns. The increase in beam energy has led to repositioning of the target and horns, making room for a much larger casing [76].

Protons colliding with the target create a shower of hadrons with a wide range of longitudinal and transverse momenta. A pair of magnetic horns is used to focus charged hadrons along the beamline. The energy spectrum of the neutrino beam is very sensitive to the kinematics of the parent hadrons and hence to the lensing strength of the focusing horns. Stronger focusing leads to a beam peaked at higher energy. It is possible to move NuMI's magnetic horns such that the energy peak is anywhere between 3 GeV and 12 GeV. The horns were kept in the lowest-possible energy configuration (LE mode) – with the horns separated by 10 m – for the majority of the lifetime of MINOS, since imperfect knowledge of the atmospheric oscillation sector at the time of design had led to a choice of L/E below the first disappearance



Figure 2.5: Schematic of the NuMI beamline. Protons from the Main Injector hit a graphite target, producing charged hadrons that are focused by a pair of magnetic horns. These hadrons traverse a decay pipe in which they may decay to produce neutrinos, before absorbers remove contaminants such as muons. The horn positions seen here correspond to the low-energy running seen for most of the lifetime of MINOS [11].



Figure 2.6: Drawing of the new NuMI target. Accelerated protons interact with thin graphite plates, which are clamped between two aluminium sheets, cooled by water pipes [12].

maximum. A small amount of data was taken in higher energy configurations. MINOS+ is run in medium-energy (ME) mode, with the horns separated by 23 m.

The current in the focusing horns can be reversed. For the majority of MINOS running, the horns were operated in *forward horn-current* (FHC) mode. In this mode, the field has a focusing effect on positively-charged particles and a defocusing effect on negatively-charged particles. This leads to a beam dominated by neutrinos, with a small antineutrino component. Reversing the current (reverse horn-current or RHC mode), produces a beam with an enhanced antineutrino component. The ability to run in these two modes allowed MINOS to test CPT conservation in neutrino oscillations by separately measuring the disappearance probabilities for ν_{μ} and $\bar{\nu}_{\mu}$ [77] It will also help NO ν A to measure δ_{CP} and the mass hierarchy; equal amounts of FHC and RHC running are expected over the lifetime of NO ν A [12].

The decays of hadrons focused by the horns are facilitated by a steel decay pipe. This pipe is 675 m-long, 2 m in diameter and filled with 0.9 atm of helium. For the first two years of MINOS running, the pipe was evacuated to 1 Torr to minimise the rate of unwanted interactions in the pipe, before degradation in the pipe made this a safety risk. The pipe was designed to be long enough as to not significantly reduce the flux of low-energy neutrinos, while minimising cost [78].

A mixture of mesons, muons and neutrinos exit the end of the decay pipe. The final components of the beamline are designed to remove as many contaminants as possible from the beam before it reaches the Near Detector. An absorber of cooled aluminium is used to remove the mesons. This is followed by layers of steel, concrete and finally 240 m of rock, which attenuate the remaining muons.

2.3.3 Summary of NuMI Upgrades

Table 2.1 summarises the main changes made to the accelerator complex and NuMI as part of the ANU upgrades.

Parameter	Pre-Upgrade	Target
Booster intensity (protons/batch)	4.3×10^{12}	4.3×10^{12}
MI Cycle Time (s)	2.2	1.333
MI Intensity to NuMI (protons per pulse)	$3.7 imes 10^{13}$	$4.9 imes 10^{13}$
Beam Spot Size (mm)	1.2	1.3
NuMI beam power (kW)	400	700
PoT/yr to NuMI	3×10^{20}	6×10^{20}
Peak ν energy (GeV)	3.5	6.6

Table 2.1: Various beam parameters before and after ANU upgrades. Taken from [12].

The MI intensity is currently still limited by the ongoing upgrades in the Recycler. The number of protons on target in 2014 was measured to be 3.1×10^{20} .

2.3.4 Properties of the NuMI Beam

The resultant beam is composed mostly of muon-neutrinos, with a broad energy spectrum whose peak is determined by the configuration of the magnetic horns. The effect of the beam upgrade – including the change in horn configuration to ME mode – on the ν_{μ} energy spectrum is illustrated in Figure 2.7, which compares simulated MINOS, MINOS+ and NO ν A ν_{μ} distributions for an exposure of 6 × 10²⁰ POT. The upgraded beam was designed to produce an L/E spectrum peaked at the atmospheric oscillation maximum for NO ν A, which is located 14 mrad from the NuMI beam axis.

Conventional beams produce neutrinos whose energies are dependent on their angles from the beam axis, with the general effect of a more sharply-peaked, lower-energy spectrum at higher off-axis angles. The MINOS+ Far Detector, which is located along the beam axis, therefore sees a broader, higher-energy spectrum, further from the oscillation maximum (as illustrated in Figure 2.2).



Figure 2.7: Unoscillated Far Detector ν_{μ} energy spectra are shown for 6×10^{20} POT for typical MINOS, MINOS+ and NO ν A running. NO ν A sharply-peaked beam at around 2 GeV due to its position, 14 mrad from the beam axis, while MINOS+, which is on-axis, sees a much higher flux and energy peak [11].

The percentage compositions of the beam in each running mode are shown in Table 2.2 for both MINOS and MINOS+. The only significant beam contaminants are muon neutrinos of the wrong charge-sign and electron neutrinos. The ν_e background is due mainly to muon decay ($\mu^+ \rightarrow \bar{\nu}_{\mu} \nu_e e^+$) and some subdominant meson decays (eg. $K^- \rightarrow \nu_e \pi^0 e^-$). The wrong-sign ν_{μ} background results mainly from mesons of unwanted charge that propagate straight along the beamline. The momenta of these mesons are approximately parallel to the field lines from the horns, resulting in very little defocusing force. These low-transverse-momentum mesons also enhance the high-energy tails of the spectra. The stronger focusing field in MINOS+ leads to a comparative reduction in both of these effects. The overall dominance of ν_{μ} even in RHC mode is a result of two effects: the larger production cross-section of positively-charged mesons over negatively-charged ones [79] and the approximately double cross-section for ν_{μ} -nucleon over $\bar{\nu}_{\mu}$ -nucleon CC interactions [69] at these energies.

Component	FHC MINOS	FHC MINOS+	RHC MINOS
$ u_{\mu}$	91.7%	97.0%	58%
$ar{ u}_{\mu}$	7.0%	1.8%	40%
$ u_e/\bar{ u}_e$	1.3%	1.2%	2%

Table 2.2: Percentage compositions of the beam in each running mode in MINOS and MINOS+. Negatively charged mesons are more strongly defocused by the magnetic horns in the new NuMI configuration, leading to a purer ν_{μ} beam. These compositions are estimated from simulated data.

2.4 The MINOS Detectors

2.4.1 Overview

The MINOS detectors were designed to make precise measurements of the L/Edependent disappearance of muon neutrinos (and antineutrinos) from the NuMI beam. In order to minimise systematic uncertainty associated with the beam (flux, composition etc.), a high-statistics measurement of the unoscillated flux is made by a detector less than 1 km from the neutrino production point, called the Near Detector. A second, larger detector, the Far Detector, measures the oscillated ν_{μ} flux several hundreds of kilometres further along the beamline. The impact of detector-related systematic uncertainties is reduced by constructing the two detectors to be as similar as possible.

As described in Section 2.7, ν_{μ} -CC interactions typically result in a nuclear recoil or shower and a muon track. The MINOS detectors were designed specifically to look for and measure the energy of these types of events, particularly those coming from the direction of the beam. The Far Detector in particular must also have as much mass as possible to maximise the ν_{μ} interaction rate. The detectors are tracking calorimeters, long enough in the beam direction to contain most muons created in ν_{μ} events and wide enough to contain hadronic showers near the vertex. 95% of the target mass is provided by 2.54 cm thick steel planes, upon which several 1 cm thick, 4.1 cm wide strips of scintillator are mounted in parallel. Wavelength-shifting (WLS) fibres set into the strips connect the scintillator to photomultiplier tube (PMT) pixels for readout. A typical mounted plane is shown in Figure 2.8, as well as their arrangement in the detector. The planes are positioned perpendicular to the beam and placed in series, separated by air gaps.



Figure 2.8: General structure of the MINOS detectors. Scintillator strips are mounted onto steel planes ad read out to PMTs. Adjacent planes are aligned orthogonally to one another to enable 3D reconstruction of interactions in the detector [13].

The parallel-strip structure of the scintillator planes allows location of activity in the detector in one transverse dimension, while the plane number allows determination of longitudinal positions (referred to as the z-direction). In order to allow full threedimensional reconstruction, successive planes are aligned with their strips orthogonal to one-another – as illustrated in Figure $2.9 - \pm 45^{\circ}$ from vertical. These define the u and v axes of the detectors.

Each detector is magnetised by means of a current-carrying coil passing longitudinally through its planes. This toroidal field curves the charged particles in a plane defined by the detector's z-axis and the displacement from the coil to the particle. The rate of curvature is determined by the ratio of its charge to its momentum, with the orientation determined by the sign of the charge. This allows the use of track curvature to distinguish ν_{μ} and $\bar{\nu}_{\mu}$ events, as well as to estimate the momentum (and hence energy) of a muon in the detector. In addition to this, the field helps to contain particle tracks of a chosen sign. In normal (FHC) running, negatively charged particles are curved towards the centre of the detector – this can be reversed for antineutrino running.



Figure 2.9: Orientation of scintillator strips on consecutive Far Detector planes. The Near Detector follows the same pattern, with orthogonally aligned strips on consecutive planes [14].

2.4.2 Near Detector

The Near Detector, seen in Figure 2.10, is designed to measure the energy spectrum of the beam before it is modified by atmospheric L/E scale oscillations. It is situated at Fermilab, 1040 m from the NuMI target, at a depth of 110 m. The neutrino flux here is approximately 5×10^5 times greater than at the Far Detector.



Figure 2.10: Schematic and photograph of the MINOS Near Detector. A: The region of the detector in the centre of the beam. B: Entrance point for the coil used to magnetise the detector. C: Electronics racks used to read out the scintillator strips [14].

The Near Detector has a fiducial mass of 0.98 kton. It is made up of 282 octagonal planes, each 3.8 m high and 4.8 m wide. The first 121 planes constitute the *calorimeter region*, with a first plane made from only steel and the rest covered in scintillator. Of

these, only one in five is fully covered in active material, the others being instrumented over a smaller area surrounding the beam spot. The remaining 161 planes make up the *spectrometer region* and are fully instrumented on every sixth plane, with no scintillator on the others. The calorimeter is intended to more precisely measure hadronic showers at the neutrino event vertex, while the spectrometer region is intended to help resolve long muon tracks. The instrumentation of the detector is illustrated in Figure 2.11.



Figure 2.11: ND instrumentation layout. Top: Half-instrumented u and v planes in the calorimeter region. These make up 80% of the calorimeter planes. Bottom: Fully instrumented u and v Near Detector planes. These account for every 5th calorimeter plane and every 6th spectrometer plane. The remainder of the spectrometer planes have no scintillator coverage. The colours distinguish different shapes of scintillator module [14].

The coil used to generate the magnetic field is positioned 2 m from the centre of the detector. This is to ensure that the area of the detector in the centre of the beam – i.e. that samples a similar region of the beam to the Far Detector - (marked A in Figure 2.10) can be fully covered in scintillator and sees a relatively uniform magnetic field. The field has an average strength 1.28 T.

Each scintillator strip in the calorimeter region is read out on one end to one pixel of a Hamamatsu M64 PMT [80]. In the spectrometer region, for the sake of economy, each pixel reads out four separate strips, whose contributions are summed. These strips are separated by around 1 m to prevent ambiguity when reconstructing tracks. A photomultiplier tube is a light detector that produces current in response to extremely small photon flux – even a single photon can produce a measurable signal. The front face of the evacuated tube is lined with a thin conducting layer, which emits electrons due to the photoelectric effect. The electron is accelerated by an electric field and focused onto an electrode, stimulating the emission of further electrons. The electrons are accelerated onto further dynodes, each kept at a larger potential, amplifying the electrical signal, which is read out as a sharp pulse of current.

2.4.3 Far Detector

The Far Detector is designed to measure the spectrum of the neutrino beam after atmospheric-scale oscillations have modified its composition. A schematic and photograph can be seen in Figure 2.12. The detector has a mass of 5.4 kton and sits 705 m underground in the Soudan Underground Laboratory in northern Minnesota, 735 km from the NuMI target.



Figure 2.12: Schematic and photograph of the MINOS Far Detector. A: A detector plane. The whole of the Far Detector is instrumented with scintillator to maximise the active target volume and mass. B: The Far Detector veto shield – two layers of scintillator above the detector that can be used to identify downward-going cosmic muons, the main background at the FD. C: Entrance point for the coil used to magnetise the detector. D: Electronics racks used to read out the scintillator strips [14].

The detector is divided into two *supermodules*, one of 249 and the other 237 planes, with the two separated by 1.2 m. Each plane is an 8 m wide regular octagon.

At this distance from the target the width of the beam is much larger than the detector and as such the entirety of each plane is instrumented. A separate coil runs through each supermodule, providing a field of average strength 1.42 T.

A veto shield is placed above the detector, which can be used to tag downwardgoing cosmic muons – a background mainly to the atmospheric ν_{μ} disappearance analysis. The shield is made up of two layers of the same scintillator modules as used in the detector and is estimated to successfully tag 99.9% of cosmic ray muons that pass through 8 or more Far Detector planes [81].

The Far Detector electronics are based on the Viking VA chip, produced by IDE AS [82]. The strips are read out on both ends to Hamamatsu M16 PMTs, which are functionally similar to the ND PMTs. The Far Detector strips are *multiplexed* – each PMT pixel is coupled to eight strips on the same plane, with their contributions summed. In order to unambiguously *demultiplex* during reconstruction, these eight strips are chosen such that they are separated by approximately 1 m. In order to further reduce ambiguity, a different coupling pattern is used at each end of the detector planes [83].

2.4.4 Scintillator and Readout

The scintillator strips are used to observe energy deposited by ionising particles in the detector. Using individual strips of scintillator allows for the reconstruction of tracks and showers and hence a more thorough understanding of the incident neutrino. Each strip measures 4.1 cm by 1.0 cm in cross-section, with varying length of up to 8.0 m. The strips are made from extruded Dow STYRON polystyrene doped with 1% PPO fluor and 0.03% POPOP. A 2 mm layer of titanium dioxide serves to prevent light loss and protect the strip from external damage.

Light generated from ionisation in the scintillator is collected by wavelengthshifting (WLS) fibres, positioned in a groove along the strip, to be carried to PMTs. The fibres are made from polystyrene doped with Y-11 fluor, absorbing blue scintillator light at a peak wavelength of 420 nm and re-emitting green light isotropically at a peak of 520 nm. This difference in the emission and absorption spectra minimises self-absorption of the scintillator light along the fibres. The



Figure 2.13: Scintillator and wavelength-shifting fibre system. Ionisation in the scintillator creates light, which travels along the fibres to PMTs. A reflective seal prevents light escaping the sides of the scintillator [14].

quantum efficiency of the PMTs at 520 nm is approximately 13%. The scintillator-WLS setup is shown in Figure 2.13.

Scintillator strips are stacked in parallel into planes, which are each sandwiched between two sheets of aluminium providing further light-shielding and structural support. The WLS fibres are connected via clear optical cables to PMTs as shown in Figure 2.14. The PMTs are boxed in steel to block out both external light and the detectors' magnetic fields.

Scintillator signals in the PMTs must be timed, digitised and stored. The vast difference in event rate between the Near and Far Detector lead to necessary differences in the design of the electronics systems used to do this, which can both be found in detail in [84].

The Near Detector is designed to constantly read data without dead time. On each turn of the Main Injector's 53 MHz clock, all ND PMT hits are digitised, with this same 53 MHz used timestamp each hit. A signal derived from the firing of the Main Injector kicker magnet triggers the readout of the Near Detector data for $18.8 \,\mu$ s covering the beam spill.

At the Far Detector, where event rates are much lower, all PMT hits above a threshold charge corresponding to $\frac{1}{3}$ that of a single photoelectron are digitised.



Figure 2.14: WLS readout system. WLS fibres from the scintillator planes are grouped and connected to PMTs via clear optical cables. In the ND, each strip is connected to one of 64 PMT pixels, whereas in the FD, 8 strips are connected to each of 16 PMT pixels [14].

The kicker fire signal is sent to the Far Detector over the internet, triggering PMT readout. Activity-based triggers allow readout outside of the beam spill.

Despite the differences in the readout electronics, both detectors use functionally identical DAQ systems [85]. Minor upgrades were made to the ND data acquisition system to help deal with the increased rates in the upgraded beam.

Further details on the original MINOS design can be found in the MINOS Technical Design Report [84].

2.5 Calibration

The raw data from the detectors comes in the form of ADCs: a PMT pulse-height from a single strip digitised by an analogue-to-digital converter, along with the strip and plane indices. To convert the ADC pulse-height values into a measure of the actual energy deposition in the strip requires a multi-stage calibration process. The result is a time-, position- and channel-dependent correction function for each detector of the form

$$Q_{\rm corr}(i,x,t) = Q_{\rm raw}(i,x,t) \times D(t) \times L(i) \times S(i,t) \times A(i,x) \times M$$
(2.2)

where i is the channel index, x the position of the hit and t the time of the hit [14]. Q_{raw} is the raw pulse-height value and the component correction factors are described in the following paragraphs.

The drift correction, D(t), corrects for changes in detector response over time. This is caused in the long term by the ageing of the scintillator and readout systems and in the short term by daily temperature changes or electronics repairs. A Light Injection (LI) system is used to measure short-term changes [14] – LEDs are flashed in Far Detector 300 times per hour and 1000 times per hour in the Near Detector and the response measured. Cosmic muons are used to detect drift in the scintillator and WLS system in the longer term.

The linearity correction, L(i), accounts for non-linearity in the Far Detector electronics system and the PMTs at both detectors at high light intensity. This is again measured using the LI system, which pulses over a broad range of intensities, allowing parametrisation of the detector response with respect to the light level.

A strip-to-strip correction, S(i, t), is applied to remove systematic differences between different scintillator strips in each detector. These differences are caused by various effects including PMT response and WLS fibre efficiency. The correction is a function of time in order to allow for differences in the ageing of components from strip to strip. Cosmic muons are used to calculate this correction, with the ADC response of each strip being compared to the average across the detector.

An attenuation correction, A(i, x), is applied after event reconstruction has estimated the 3D position of each hit. The pulse-height measured by a scintillator strip varies with the position of the hit along the strip due to the attenuation of light along the WLS fibres.

Having ensured consistency within each detector with respect to strip, position and time, the final task is to attempt to relate each detector's output to a real-world energy scale. This is the absolute energy-scale – or inter-detector – calibration, M. First, the two detectors' responses are compared using the Bethe-Bloch equation. The ionising power of muons with respect to their momentum (measured by range) is equated between the Near and Far Detector to give a shared energy unit, the muon energy unit (MEU).

The final conversion, from the Near and Far Detector MEU to the absolute eV scale, is based on measurements made at CERN with the *calibration detector* [86], a small-scale detector using the same scintillator, casing and readout components as the Near and Far Detector. The detector response was recorded in beams of electrons, muons and hadrons of well-known energies, in the range 0.2–10 GeV.

2.6 Monte Carlo Simulation

Simulated data is a vital component of MINOS and MINOS+. In addition to Near Detector data, accurate simulation of the beam is required to make predictions of the Far Detector beam flux, composition and energy spectrum. Further to this, simulation of the detectors is needed to create simulated detector data, used to study various components of physics analyses, such as event selection.

2.6.1 Beam

Single-detector neutrino oscillation experiments rely heavily on models of the beam to make predictions of the neutrino flux and energy spectrum. Experiments of these kinds tend to incur systematic uncertainties on the order of 10 to 30%. At MINOS, the beam flux and energy spectrum is constantly measured by the Near Detector, greatly reducing the overall flux uncertainty. However, as will be described further in Section 3.1.3, the large difference in angular acceptance between the Near and Far Detector leads to differences in energy distribution of the beam. Correcting for this requires accurate modelling of neutrino production in the beam. Beam Monte Carlo is produced for default estimates of the parameters of the beam model, for example hadron production cross-sections and magnetic horn positions. MINOS then uses data from the ND in various beam configurations to constrain these parameters and reweight the simulated data. The result of this is a final FD flux uncertainty on the order of 2%. The effect of the beam reweighing is shown for a MINOS run in 2.15. Several software packages are used for the different stages of the beam simulation. First, the interaction of the proton beam on the NuMI target is simulated, with mesons produced by the FLUKA [87] package. A combination of FLUGG [88] and GEANT4 [89] simulates the propagation of these particles through the magnetic horns and along the beamline until they either decay to produce a neutrino or reach the end of the decay pipe. Each generated neutrino has its energy and trajectory logged. Reweighting is used to avoid unnecessary processing of neutrinos that do not pass near either detector.

In MINOS, the neutrino energy spectrum measured at the Near Detector is poorly modelled by the simulated data from the default beam simulation. There are several uncertainties in the simulation, including hadron production cross-sections and magnetic horn positions, that may modify the predicted beam spectrum. A set of energy-dependent beam weights are applied to correct for the differences between the ND data and simulation, illustrated in Figure 2.15. The weights are calculated by fitting the Near Detector data in various beam configurations with respect to several beam parameters (eg. horn current, pion cross sections).



Figure 2.15: Effect of beam reweighting in the Near Detector in MINOS Run III [15].

2.6.2 Detectors

The detector simulation samples the neutrinos generated in the previous step, before tracking them through the Near and Far Detector as well as the surrounding halls and rock to allow them to interact. Interactions are simulated by **GENIE** [90], which generates quasi-elastic, resonance and deep inelastic scattering interactions between 100 MeV and 100 GeV.

The particles generated in these interactions are then propagated through the detector using a GEANT3 simulation of its geometry as well as the magnetic field. The energy deposited in the scintillator strips is recorded, giving a set of *truth hits*. Finally MINOS software PhotonTransport and DetSim convert these hits into simulated detector output. PhotonTransport simulates the scintillator response to each energy deposit, as well as the propagation of photons through the WLS fibres to the PMTs. DetSim then simulates the conversion of this signal to a pulse-height measurement, producing output of the same format as real data, which can then be fed to the reconstruction software.

2.7 Interaction Topologies

There are three main distinct categories of interaction seen in the MINOS detectors: ν_{μ} CC, ν_{e} CC and NC. Of these, this thesis is primarily concerned with ν_{μ} CC interactions, which can be further subdivided by topology.



Figure 2.16: Muon neutrino CC cross sections and predictions as a function of neutrino energy. Quasi-elastic (QE) interactions dominate at low energy, resonant (RES) interactions and medium energy and deep inelastic scattering (DIS) at high energy. Taken from [4].

MINOS and MINOS+ analyses neutrinos in the energy range of around 1 GeV to 100 GeV. Across this energy range, three types of ν_{μ} -nucleus CC interactions are seen – quasi-elastic (QE), resonance (RES) and deep inelastic scattering (DIS). The cross sections for these interactions are shown as a function of neutrino energy in Figure 2.16. At low energy, QE interactions dominate the total interaction cross-section.

$$n + \nu_{\mu} \to p + \mu^{-}. \tag{2.3}$$

Here, the substructure of the nucleon is not probed and the nucleus absorbs the recoil from the incoming neutrino without breaking apart. From approximately 1 to

10 GeV, RES interactions make a large contribution to the cross-section, for example

$$n + \nu_{\mu} \to \Delta^{+} + \mu^{-}; \qquad \Delta^{+} \to \pi^{+} + n.$$
 (2.4)

Here, the nucleon is excited into a resonance, which quickly decays to another nucleon and one or more other particles, often pions. Above 5 GeV, DIS interactions dominate. One example is

$$d + \nu_{\mu} \to u + \mu^{-}. \tag{2.5}$$

The energy of the neutrino is great enough to probe the structure of the nucleon, interacting with one of its constituent quarks and breaking apart the nucleus, resulting in a shower of strongly interacting particles.

The collection of PMT hits associated with one interaction in the MINOS detectors is called an *event*. *Reconstruction* software is used to separate hits into distinct events before measuring the characteristics of the underlying interaction in order to provide information to later stages of the analysis such as *selection* or energy estimation.

Typical examples of each type of event in the MINOS detectors are shown in $event \ displays - maps \ of hit positions in the detector - in Figure 2.17.$



Figure 2.17: Common event topologies in the MINOS detectors. CC ν_{μ} events are generally identified by a curved track and a shower at the vertex. NC events are characterised by a hadronic shower, while CC ν_e events are usually more compact, EM showers. The true interaction products – marked with arrows – are including for illustration [11].

MINOS was designed to measure CC ν_{μ} interactions, $\nu_{\mu} + N \rightarrow \mu^{-} + X$. These typically produce a hadronic shower at the event vertex from the breakdown of the target nucleus, along with a long track from the generated muon, curved by the detector's magnetic field. The field is chosen such that the direction of curvature is towards the centre of the detector for negative muons travelling along the beam axis, increasing the chance of containing the track in the detector.

NC and CC ν_e events typically display no obvious tracks. The electron created in a CC ν_e interaction is strongly scattered in the detector, leaving its track embedded in the associated hadronic shower and leading to a further compact leptonic shower. NC interactions leave only a hadronic shower, with the incident neutrino escaping the detector.

2.8 Reconstruction

Physics data in the detectors comes in the form of raw ADC hits along with their timestamps, z and u or v positions. At the Near Detector, lists of hits are grouped into files (called *snarls*) by beam spill. These are divided into *slices* – groups of hits approximately contiguous in time – before further analysis by *tracking* and *showering* algorithms, eventually building up a picture of an entire interaction and estimating its properties. At the Far Detector, the slicing step is not performed since each snarl contains approximately contiguous hits and the flux is low enough that most snarls contain no more than one event. The flow of the Near Detector reconstruction is shown in Figure 2.18. Some modifications were made to the ND reconstruction as part of the MINOS+ upgrade in order to cope with increased event pile-up due to the increased energy and flux of the beam.

2.8.1 Slicing

The large neutrino flux at the Near Detector generates multiple neutrino events per beam spill. In order for the later stages of reconstruction to perform optimally, they should be passed a set of hits corresponding to no more than a single event.

Since this algorithm is run before any other reconstruction or analysis step, efficiency is the primary concern – it is preferable to generate slices containing more



Figure 2.18: Near Detector reconstruction chain. Snarls are divided into slices, before clustering algorithms identify potential tracks and showers. Fitters are then used to reconstruct these objects, and an event-builder groups them together and estimate the properties of the underlying neutrino. Taken from [16].

than one event than to wrongly split events. In MINOS, events are sliced by time, with a gap of 100 ns between hits considered enough to start a new slice.

In MINOS+, the increased intensity and event length means that this approach fails to split distinct events in the majority of cases. This is illustrated in Figure 2.19, an example time slice from MINOS+ simulated data. An extension was made to the algorithm to include the position information on each hit. First, the time slicing is performed as before. This produces equal-length slices in each the u and v view. For each view, hits are then partitioned using a near-neighbour search in t, z and u or v, eventually generating a list of distinct clusters of activity in each view. These clusters are then combined across views into slices using their overlap in t and z, with any left over appended to the largest generated slice. The result of applying this algorithm to the time slice in Figure 2.19 is shown in Figure 2.20. A more detailed description of the algorithm is given in [91].



Figure 2.19: A Near Detector time slice in u and v. Three distinct events can be picked out by eye, however they cannot be separated using only timing information.



Figure 2.20: A Near Detector slice in u and v using the new algorithm. A combination of timing and position information has been used to separate the two tracks from the NC-like event. The algorithm is conservative: two CC-like have significant overlap and are not separated. The z-density of hits in these tracks could potentially be used to separate them but this is not attempted here.

2.8.2 Tracking

Each slice is first sent to the track reconstruction. A *track finder* [92] identifies hits that appear on consecutive planes in an approximate straight line, producing a list a list of *seed tracks* in each the u and v view. Efficiency is emphasised at this stage, leaving the job of removing spurious hits to the *track fitter* [93].

The track fitter is implemented with a Kalman filter [94]. The algorithm associates seed tracks across views, deciding which of the components of the seed track will make up the reconstructed track, whilst also providing estimates of some of the properties of the underlying particle, such as its charge to momentum ratio and interaction vertex coordinates.

Where possible, the energy of a track is calculated from its range in the detector. The Bethe-Bloch equation may be used to relate the length of a charged particle in a given medium to its energy (or momentum). This method relies on accurate determination of the thickness and density of the steel in the detectors. It gives a track energy resolution of 4.6% for a 3 GeV muon [86].

In the case that the track's energy may not be calculated by range, it must be estimated by its curvature, given the value of q/p from the Kalman filter. This is the case when a track exits the detector or ends in the coil hole. The energy resolution for these tracks is approximately 11% for a 3 GeV muon.

2.8.3 Showering

The remaining hits in the slice after the tracking are passed to the shower reconstruction. All remaining energy deposits in the slice are grouped into showers in each of the two views using a 2D clustering algorithm [95]. Deposits along a track with energy greater than that expected from ionisation have the track energy subtracted before being added to the shower. These clusters are associated across views by a shower fitter [96], which again estimates various properties of the shower. The shower energy resolutions for hadronic and electromagnetic showers in the MINOS detectors are measured to be $\frac{55\%}{\sqrt{E}}$ [97] and $\frac{20\%}{\sqrt{E}}$ [98] respectively, where E is the shower energy.

Near Detector Modification

A small change was made to the shower reconstruction at the Near Detector, documented in [99]. In the original shower reconstruction, hits can be clustered together despite being separated by large distances. This improves the estimation of the shower energy when only one event is present in a slice. However, even with the use of the updated slicing algorithm, this is rare in the Near Detector in MINOS+. The result is that showers are often seen to 'steal' hits from neighbouring events, reducing energy resolution and reconstruction efficiency. Shower hits separated from their neighbours by more than 20 cm are removed.

Shower Energy kNN

The energy contained in a shower was calculated simply as the sum of the energy deposits in the shower in early versions of the MINOS reconstruction. Improvements in resolution were found by moving to a *k*-nearest-neighbour (kNN) estimator [100] – a simple multivariate analysis technique that can be used for parameter estimation or classification.

The kNN algorithm proceeds in general terms as follows: a *feature space* of arbitrary dimension is populated with a large number of events, in our case events from our Monte Carlo simulated data. This is sometimes called *training* the kNN. The nearest k neighbours to a *test event* (a data event) are then found according to some distance metric. The properties of the test event can then be estimated as the mean of the properties of its neighbours.

The shower energy is calculated using a three-dimensional feature space in total calorimetric shower energy, calorimetric shower energy within 1 m of the track vertex and transverse shower length, using the Euclidean distance metric. These inputs must be normalised such that each variable contributes a similar amount to the calculated distance. This is done simply by scaling by the width of the distributions of these variables in the simulated data. The value k = 400 was seen to optimise the shower energy resolution. The kNN is trained separately for each MINOS and MINOS+ beam run.

2.8.4 Near Detector Reconstruction Performance in MINOS+

As well as the changes to the showering and slicing algorithms, several bugs were fixed in the track reconstruction. All of the changes made between MINOS and MINOS+ are documented in [17].

The general performance of the reconstruction can be measured with two variables: the reconstruction efficiency, η_{reco} , defined as the fraction of true neutrino interactions that are successfully reconstructed into events. Reconstruction of a true simulated event is considered successful if a reconstructed event matches its vertex time and position. Figure 2.21 shows the reconstruction efficiency of the final version of the MINOS reconstruction and the first version of the MINOS+ software for MINOS+ simulated data. An increase in efficiency is seen for the new software above 1 GeV.



Figure 2.21: Reconstruction efficiency vs. energy for MINOS+ simulated data using the final version of the MINOS reconstruction ("Dogwood") and the first version of the MINOS+ reconstruction ("Elm"). A modest improvement in efficiency is seen above 1 GeV [17].

No significant difference is seen in track and shower energy resolutions when moving to the new version of the reconstruction. The vertex position resolution is also unaffected. The resolution of x and y position is 5 cm in MINOS and MINOS+. The z position resolutions are 14 cm and 19 cm respectively [17].
Slicer

Analysis of the slicer centres on the *slice purity* and *completeness*. The slice purity is defined as the fraction of the hits in the slice belonging to the largest true event in that slice (in terms of calorimetric energy), and is low for slices containing multiple events. The slice completeness is defined as the fraction of the hits in the largest event in the slice that are contained within that slice, and is low for events that have been split by over-aggressive slicing. Table 2.3 compares the distributions of these variables for MINOS+ simulated data using the old and new slicing algorithms. As might be expected, the more aggressive new algorithm significantly increases the fraction of high-purity slices, at the cost of a slight reduction in completeness. The fraction of events with both high completeness and purity shows clear improvement.

	Old Slicing	New Slicing
$\overline{\text{Completeness} > 50\%}$	96.9%	92.7%
Completeness $> 80\%$	94.8%	84.5%
$\overline{\text{Purity} > 50\%}$	87.1%	95.3%
Purity $> 80\%$	47.7%	65.4%
$\overline{\text{Completeness and Purity} > 50\%}$	84.6%	88.7%
Completeness and Purity $> 80\%$	46.2%	76.5%

 Table 2.3: Table showing the distribution of slice completeness and purity for MINOS+ simulated data using the old and new slicing algorithms.

Chapter 3

Three-Flavour Neutrino Oscillations at MINOS

MINOS and MINOS+ can constrain the oscillation parameters Δm_{32}^2 , $\sin^2(\theta_{23})$, $\sin^2(\theta_{13})$ and δ_{CP} by observing both the disappearance of ν_{μ} and the appearance of ν_e as a function of L/E. Three categories of analysis were performed in MINOS, each using distinct techniques. These analyses focused on ν_{μ} disappearance in the NuMI beam, atmospheric ν_{μ} disappearance and ν_e appearance in the NuMI beam respectively. Since 2014, the are combined to produce a single three-flavour result constraining Δm_{32}^2 and $\sin^2(\theta_{23})$ (and δ_{CP} to a lesser extent) [27]. θ_{13} is treated as a nuisance parameter.

In this chapter, an overview of the techniques used in each of the three analyses is given, along with the MINOS results. Particular focus will be given to the ν_{μ} disappearance analysis, the techniques of which will be reused in the MINOS+ ν_{μ} disappearance analysis presented in Chapter 4.

3.1 Beam ν_{μ} Disappearance

The broad outline of a neutrino oscillation analysis is as follows: a sample of *signal* events are *selected* from the data based on their reconstructed properties. A certain fraction of unwanted events – *background* – pass into the selected sample. The fraction of selected events that are part of the signal is defined as the *purity* of the

selection, while the *efficiency* is defined as the fraction of true signal events that pass the selection.

The neutrino energy distribution of the selected sample of events can be plotted in order to search for energy-dependent oscillation effects in the flux. A statistical analysis is performed to estimate the parameters of the oscillation model from the data. In order to do this, oscillation-dependent flux-energy predictions are generated and compared to the data. In general, these predictions depend on estimates of various properties of the experiment as well as physics models and are susceptible to associated *systematic errors*. Much of the analysis involves mitigating and understanding these errors.

In this analysis, the energy-dependent reduction in ν_{μ} (and $\bar{\nu}_{\mu}$) flux at the MINOS Far Detector, due to oscillations on the atmospheric scale, is studied. CC ν_{μ} events are selected against a background primarily consisting of NC events. The ν_{μ} disappearance equation (in vacuum) is given in Equation 1.20 and is dependent to first order on $\sin^2(2\theta_{23})$ and $|\Delta m_{32}^2|$. Matter effects and higher-order terms introduce weaker dependencies on the octant of θ_{23} , the mass hierarchy and the other oscillation parameters.

Systematic uncertainties introduced in the predictions of the Far Detector flux are mitigated by the two-detector design of the experiment, which is common in neutrino oscillation experiments using man-made sources. The distribution of events in one detector is used to predict the distribution in the other. At MINOS, the Near Detector observes the beam with negligible flavour-change from three-flavour oscillations, greatly reducing uncertainty on the flux and energy distribution of the beam. Transformations are applied to the Near Detector data in order to generate Far Detector predictions. This process is called *beam extrapolation* (often referred to simply as *extrapolation*).

The reconstruction and selection of events in both detectors is kept as similar as possible. This, along with the similar design of the ND and FD, ensures that many single-detector systematic uncertainties are correlated between both detectors and cancel out in the extrapolation process.



Figure 3.1: Rate of protons delivered to NuMI per week over the lifetime of the MINOS experiment. RHC (orange) and high-energy (red) running are shown separately to standard low-energy FHC running (green). The integrated POT count over time is shown by the blue line [18].

3.1.1 Data

Data taken during MINOS and MINOS+ beam running is divided into *runs*, with each run analysed separately in the fit. A new run is started whenever the beam conditions change significantly, necessitating the production of new simulated data taking these changes into account. The amount of data taken is expressed in terms of the estimated number of protons incident on the NuMI target (protons-on target or POT).

3.1.2 Selection

The three main types of beam interaction in the MINOS detectors are discussed in Chapter 2. This analysis attempts to select a signal of CC-interacting muon neutrinos and antineutrinos originating in the beam. The dominant background is from neutral-current events, which occur at a similar rate to the signal in the $\nu_{\mu}/\bar{\nu}_{\mu}$ -dominated beam Atmospheric neutrinos and cosmic muons are additional, small backgrounds.

The criteria used to discriminate this signal from the background are referred to as the *analysis selection*. A *preselection* is first applied in order to remove data likely to be unreliable, for example data taken during periods when the detector was known to be misfunctioning.

Preselection

Several cuts are made to eliminate data taken under either poorly-understood experimental conditions or inconsistent conditions compared to the rest of the run:

- Good beam: cuts are made on the alignment and width of the proton beam on the target as measured by the beam monitors. The current in the horns is further required to fall between an acceptable range.
- Light injection: data taken around the time that the light injection system was running in either detector is removed to avoid errant high-energy hits.
- Magnetic field: the coil current at each detector is required to be of the correct sign to contain ν_µ or ν
 _µ, depending on the beam configuration.
- Detector performance: Several conditions are required at each the Near and Far Detector to ensure good detector performance. For example, problems with the readout electronics may lead to consistently low or high readings in certain strips. Constant monitoring of the detectors allows for quick correction of these types of faults as well as simple removal of affected data.

Next, cuts are made to remove certain classes of background events:

- **Require track**: At least one track must be reconstructed. This removes a large number of NC events.
- **Track fit success**: Removes events for which the Kalman filter has failed to fit the track, which are regarded to be misreconstructed.
- Muon direction: the track fit gives an estimate of the direction of the muon generated in the event. With θ as the angle between the initial track direction and the beam direction, events for which $\cos(\theta) < 0.6$ are removed as they are likely to be from atmospheric sources.

• Spill coincidence: a further cut is made to remove cosmic muons – removing events that are not in time with the beam spill. The reconstructed event time must fall within 2 μ s of the 10 μ s-long spill.

Fiducial Volumes

Finally, geometrical cuts are made in each detector, defining a *fiducial volume*, intended to eliminate events in poorly-understood regions of the detectors, or those unlikely to have all of their energy contained. The reconstructed vertex of each event must lie within this volume. A cross section of each the ND and FD fiducial volumes is shown in Figure 3.2. Track containment is not required since the energy of exiting muons can be estimated by their curvature in the detector's magnetic field.



Figure 3.2: Front-on view of fiducial volumes in the Near (left) and Far Detector (right). Preselection eliminates events with vertices outside of these volumes [13].

At the Near Detector, the fiducial volume is a cylinder of radius 80 cm centred on the beam spot. This cylinder is angled 3° downwards to follow the direction of the beam. The z-span of the volume is from the 14th to the 68th plane. The size and shape of the ND fiducial volume is designed such that only events from the portion of the beam likely to reach the Far detector are sampled, while ensuring that as much of the energy in the event as possible is contained, with showers being contained in the calorimeter.

At the Far Detector, as much of the detector is sampled as possible to maximise the target mass. The volume consists of two annular cylinders – one for each supermodule – with an inner radius of 40 cm an outer radius of 3.74 m. This ensures that any event closer than 35 cm to the outer edge of the detector or the inner coil-hole are removed. The first 4 planes of each detector are excluded, along with the last 8 planes of the first supermodule and the last 20 of the second supermodule. This eliminates events originating outside the detector as well as those whose showers are likely to leak outside the detector.

Analysis Selection

Most ν_{μ} CC interactions are easily identified by eye due to their long muon track. Early MINOS analyses used just a few simple cuts to select events with clear muon tracks. However, some NC events can mimic low-energy CC events, for example if a small track is reconstructed among the shower deposits of an NC event. Later analyses implemented two kNN particle-identification discriminators (PIDs), one specifically tuned for low-energy events and the other applied across all energies. In MINOS+, the low-energy discriminator is dropped, with the analysis selection performed using the second kNN, named *roID* [101]. A description of the low-energy kNN can be found in [102].

roID kNN Discriminator

A brief overview of the kNN algorithm is given in Section 2.8.3. For this discriminator, a 4D feature space is used in the following variables, chosen for maximum separation between signal and background. The variables are normalised by width.

- Number of active planes in the track: In a ν_μ CC event, the characteristic muon track tends to lead to a large number of activated planes in the event, while NC events are likely to be more compact, with more energy spread in the transverse dimension.
- Transverse profile: This is the fraction of the energy in the last 50% of active planes in an event that is contained within ± 2 strips of its reconstructed track. Signal events, away from the vertex shower, are likely to have the vast majority of their energy deposits distributed in and around the track.

- Mean track pulse-height per plane: This exhibits low values for signal events, since muons are minimally ionising particles, depositing only a small amount of energy per plane. Hadrons are likely to leave significantly larger energy deposits. In order to exclude the vertex shower, only the last 70% of active planes are considered.
- Track energy fluctuation: Defined as the ratio of the least energetic 50% of the hits in the track to the most energetic 50%, this quantifies the variation in energy deposited along the track. Again, since muons are minimally-ionising, they should deposit approximately the same amount of energy in plane, whereas hadronic activity identified as a track may exhibit large variations in energy.

The output of the kNN is a single value equal to the fraction of the nearest neighbours that are true ν_{μ} CC events. The value of k = 80 was chosen as the value that maximises the separation between signal and background. Events with output value ≤ 0.3 are cut from the analysis, with this cut position optimised for sensitivity to the oscillation parameters in fake data studies.

3.1.3 Far Detector Prediction

Data taken at the Near Detector can be used to generate a prediction of the Far Detector energy spectrum for an arbitrary set of oscillation parameters. The procedure for doing so has been developed in previous MINOS analyses [20] and is shown in the flow chart in Figure 3.3. The three main steps are the estimation of the true energy spectrum and flux of the beam at the Near Detector from the reconstructed ND spectrum (*detector unfolding*), followed by the conversion to a Far Detector flux, including the effects of oscillations (*beam extrapolation*) before applying the first step in reverse to convert the FD flux into a prediction of the reconstructed FD spectrum.

Detector Unfolding

For the purposes of the extrapolation, the reconstructed energy spectrum at the Near Detector, is represented as a column vector, N_{ND} of dimension n. The elements of this vector are N_1, N_2, \dots, N_n , with N_i as the number of events in each reconstructed



Figure 3.3: Flow chart illustrating the beam extrapolation procedure. Taken from [19].

energy bin. The aim of the first step of the extrapolation is to estimate the true flux at the Near Detector, ϕ_{ND} , a column vector in m bins of true neutrino energy. To do this, we use the following transformation:

$$\phi_{\rm ND} = \eta_{\rm ND} \mathbf{M}_{\rm RT} \rho_{\rm ND} \mathbf{N}_{\rm ND}, \qquad (3.1)$$

where ρ_{ND} and η_{ND} are, respectively, $n \times n$ and $m \times m$ diagonal matrices, called the *purity correction* and *efficiency correction*. \mathbf{M}_{RT} is the $m \times n$ Near Detector *reco-to-true* matrix. In the MINOS software, the vectors and diagonal matrices used in the extrapolation are represented as 1D histograms with non-diagonal matrices represented as 2D histograms.

First, a *purity correction* is applied to remove background events that have passed the selection. This is achieved by means of a histogram in which the MC purity is recorded for each bin of reconstructed energy in the ND spectrum. The purity ρ_i for bin *i* is given by the number of signal events in that bin that pass the selection divided by the total number of selected events in that bin,

$$\rho_{\rm ND} = diag\{\rho_1, \rho_2, \cdots, \rho_n\}; \quad \rho_i = \frac{N_{i, \text{ signal}}}{N_{i, \text{ signal}} + N_{i, \text{ background}}}.$$
 (3.2)

Next, this reconstructed energy distribution is converted to a true energy distribution. This process is referred to as *unfolding*. The reconstructed energy distribution is a function of detector effects, such as imperfect resolution, that act to smear out the the true energy distribution. The generation of a reconstructed energy distribution from a true distribution is performed with the following transformation,

$$\mathbf{N}_{\text{reco}} = \mathbf{A} \mathbf{N}_{\text{true}},\tag{3.3}$$

where \mathbf{A} is an $n \times m$ matrix – the transition matrix describing the smearing of true energy across reconstructed energy bins, such that A_{ij} is the probability of an event in true energy bin j having reconstructed energy i. \mathbf{A} can be constructed from the distribution of true and reconstructed energies for a set of simulated events passing the specified selection. One technique for unfolding is the inversion of \mathbf{A} , since $\mathbf{N}_{true} = \mathbf{A}^{-1}\mathbf{N}_{reco}$. This solution, however, can lead to numerical instabilities – even worse, these are dependent on the exact form of \mathbf{A} , which is model-dependent. Though various techniques exist for solving this problem, in MINOS the Far Detector prediction turns out to be relatively insensitive to the specific unfolding procedure, due to the lack of sharp feature in the beam and the re-smearing that happens when converting back from true to reconstructed energy at the Far Detector [103], and instead a *reco-to-true* transition matrix, $\mathbf{M}_{\rm RT}$ is used – the transverse of matrix \mathbf{A} , normalised columnwise – producing a smeared but unbiased estimate of the true energy distribution.

Next, an efficiency correction factors in CC ν_{μ} events that were removed by the selection. The transformation is analogous to the purity correction mentioned above, with the transformation defined as

$$\eta_{\rm ND} = diag\{\eta_1, \eta_2, \cdots, \eta_m\}; \quad \eta_i = \frac{N_{i, \text{ accepted signal}}}{N_{i, \text{ accepted signal}} + N_{i, \text{ rejected signal}}}.$$
 (3.4)

Finally, the total beam flux is reconstructed by multiplying out the detector mass and neutrino interaction cross-section.

Beam Extrapolation

The flux of the beam at the Near Detector must now be extrapolated to the Far detector. In addition to the $1/r^2$ reduction in the flux, angular acceptance differences between the two detectors must be accounted for, as illustrated in Figure 3.4. At the Near Detector, neutrinos up to $(4.6 \times 10^{-2})^{\circ}$ from the beam direction may reach the detector, while at the Far detector, only the inner $(2.9 \times 10^{-4})^{\circ}$ is sampled.



Figure 3.4: Diagram demonstrating the difference in angular acceptance between the Near and Far Detectors (not to scale) [20].

Decay kinematics of pions and kaons lead to a non-trivial angular dependence (and therefore energy dependence) in the beam's neutrino flux, such that

$$\frac{dN}{d\cos\theta} = \frac{1}{2\gamma^2(1-\beta\cos\theta)^2},\tag{3.5}$$

where γ and β are the Lorentz factor and speed of the meson respectively and θ is the angle between the direction of the meson and neutrino, all in the laboratory frame.

The general effect is that the Near Detector sees a slightly higher flux from unfocused, low-energy mesons than the Far Detector. The differences are illustrated in Figure 3.5. Another matrix transformation is applied to correct for this.

$$\phi_{\rm FD} = \phi_{\rm ND} \mathbf{M}_{\rm beam},\tag{3.6}$$

where ϕ_{FD} and ϕ_{ND} are the Far and Near Detector fluxes respectively and \mathbf{M}_{beam} is an $m \times m$, non-diagonal matrix called the *beam matrix*.

This beam matrix is again populated from a large number of Monte Carlo simulated events. The output from the beam simulation – described in Section 2.6.1 – includes a list of the properties of all neutrino parents in the beam. These parent



Figure 3.5: Example ν_{μ} energy spectra in the Near and Far Detector. Hadrons that produce events of one energy at the Near Detector will produce events of different energies at the Far Detector due to acceptance effects, as illustrated by the hatched bands [21].

mesons are grouped according to the energy of their daughter neutrinos at the Near Detector. The energy distribution of neutrinos from kinematically identical mesons that pass through the Far Detector is then calculated.

 $\mathbf{M}_{\text{beam }i,j}$ is equal to the probability that an FD event in energy bin j was created from a parent kinematically similar to an ND event in energy bin i, multiplied by the probability of an event passing through the ND reaching the FD, such that multiplying the ND flux by this matrix gives the FD flux.

FD Prediction

This flux is then converted into an FD predicted spectrum by applying the ND unfolding process in reverse – specifically, applying the transpose of the transformation in Equation 3.1 – using Far Detector simulated data. First, the effect of oscillations on the FD flux is accounted for. This can be considered an additional multiplication by an $m \times m$ diagonal matrix, \mathbf{P}_{osc} , whose elements P_j correspond to the ν_{μ} disappearance probability calculated at the centre of bin j. The full transformation is

$$\mathbf{N}_{\rm FD} = \rho_{\rm FD}^{-1} \mathbf{M}_{\rm RT, FD}^{\top} \eta_{\rm FD}^{-1} \mathbf{P}_{\rm osc} \phi_{\rm FD}.$$
(3.7)

Separate spectra are produced for events reconstructed as ν_{μ} and $\bar{\nu}_{\mu}$, which are fitted separately. For each of these spectra, the extrapolation is performed individually for each of its true components $-\nu_{\mu}$, $\bar{\nu}_{\mu}$ and ν_{τ} CC events and NC events, with the appropriate oscillation probabilities, purities, efficiencies and cross sections applied. The final Far Detector prediction is calculated as the sum of each of these extrapolated components.

Systematic uncertainties are accounted for by changing the parameters used to generate simulated data in the extrapolation. For example, the track energy scale may be shifted by $\pm 1\sigma$ and the matrices used in the extrapolation recalculated. This is generally achieved by reweighting the simulated data already created. Fitting the data requires the ability to generate extrapolated predictions for an arbitrary point in some pre-defined systematic and oscillation parameter-space. This can be done exactly, by applying the appropriate oscillations during the extrapolation and using simulated data with the appropriate systematic shifts applied, but it is very computationally intensive. Instead, a large set of predictions is generated in a grid of oscillation and systematic parameter-space, and interpolation is performed between points during the fit. This is described further in Section 3.4.2.

3.1.4 Systematics

A number of effects may cause systematic errors on the Far Detector prediction. In general, uncalibrated systematic differences between the two detectors give rise to an error. The transformations applied in the previous section are based on simulation and implicitly use best-guess measurements of various parameters. As an example it can be easily seen from the discussion in the previous section that a constant underestimation of the FD selection efficiency in simulation will lead to a systematic deficit in the predicted FD spectrum. However, if the ND efficiency has the same effect, the deficit will cancel perfectly. If the effects are identical but have an energy dependence, small differences will appear.

Any systematic uncertainties with the potential to significantly alter the values estimated for Δm_{32}^2 and $\sin^2(\theta_{23})$ in the fit must be included as model parameters. A description of the four systematic uncertainties included in the MINOS and MINOS+ fit is given below.

Relative Normalisation

Several uncertainties are approximated to have the effect of introducing a scale-factor uncertainty on the Far Detector spectrum. The major contributors to this uncertainty are differences between the selection and reconstruction at the two detectors – if they are not identical, the relative number of events cut can only be simulated. Other factors include the uncertainty on the fiducial masses and live-times of the detectors.

The size of this uncertainty was estimated by hand-scanning real and simulated events in each detector, using people to cross-check the reconstruction and selection by identifying interactions in event displays. The difference between the average number of events selected by hand and by the MINOS software is taken as the per-detector normalisation uncertainty, while the difference between the values for the ND and FD is taken as the relative normalisation uncertainty.

NC background

The NC background component of the ν_{μ} -CC-selected ND data is estimated by simulation. The background is extrapolated separately to the rest of the data, with no oscillations applied. Uncertainty in the simulation, due mainly to uncertainty in the modelling of hadronic showers, gives rise to an uncertainty in the final prediction on the size and shape of the true NC background, which is represented by an energydependent scale uncertainty on the predicted NC component of the extrapolation.

The uncertainty is estimated by first examining the difference between data and simulation for events rejected by the CC selection. This constitutes an NC-rich sample, and as such any NC mismodelling is more easily seen. The 1σ uncertainty is taken to be the ratio between data and simulation in each energy bin. A second study involves the removal of tracks from CC events in data and simulation [104], creating a fake NC data set that is then passed through the reconstruction and selection and allowing an estimate the NC efficiency to be made. The ratio between this and the nominal NC efficiency is taken as the 1σ uncertainty on the efficiency. The total NC background uncertainty per bin is then given by the quadrature-sum of these two uncertainties.

Track Energy Scale

This is the uncertainty on the overall track energy scale calibration measured by CalDet. It is estimated by comparing CalDet data to simulation and incorporating energy-loss uncertainties [105], giving a $\pm 2\%$ shift on the track energy at 1σ , 100% correlated between the Near and Far Detector.

Shower Energy Scale

The uncertainty on the absolute shower energy scale comes from various sources. Cosmic muons are used to perform this absolute calibration, with uncertainty introduced in the modelling of their interactions. This uncertainty is estimated by varying the model parameters used in the simulation (cross sections etc.) within reasonable limits [106]. An energy dependent uncertainty is found, with 8.2% scale < 0.5GeV and approximately 3.5% scale > 6GeV.

A further, flat 5.7% scale uncertainty is calculated from uncertainty on the CalDet measurements. This includes a 5% uncertainty on the energy deposited by individual hadrons, a 2% uncertainty on the CalDet beam energy and a 1.4% uncertainty on the energy of stopping muons. A further 0.9% uncertainty is included, due to systematic differences observed in the detector response for cosmic and beam muons [107].

The flat 5.7% uncertainty is added in quadrature to the energy-dependent uncertainty, with a final, multiplicative systematic uncertainty on the energy of a shower given as $6.6\% + (3.5\%) \times e^{-E_{\rm shw}/1.44}$.

Additional Systematic Uncertainties

Several additional systematic uncertainties that have a smaller effect on the fit have been modelled in previous MINOS analyses. Their effects on the estimation of oscillation parameters are small compared to the four parameters listed above and are neglected when performing combined analyses, in which the number of systematic parameters is large [108].

Neutrino interaction cross sections are poorly constrained in general, and are estimated in MINOS by comparing simulations to Near Detector data. Uncertainties of around 3% for DIS interactions and around 10% for QE and single-pion resonance events at the energies probed by MINOS and MINOS+ are seen. The effect of cross section systematics in the Near and Far Detector are highly correlated, cancelling to a large extent in the Far Detector predictions.

Uncertainty on the beam flux is also greatly mitigated by the use of an identical Near and Far Detector. However, uncertainty in the beam simulation that is used to convert between Near and Far Detector flux generates a small energy-dependent uncertainty on the Far Detector prediction. This uncertainty is evaluated from the uncertainties of the parameters of beam model fits to the Near Detector data.

3.2 Atmospheric ν_{μ} Disappearance

Neutrinos are created in the upper atmosphere in a manner similar to that utilised by a neutrino beam. Cosmic rays interact with nuclei to produce a hadronic shower including pions and kaons, which commonly decay to neutrinos. Neutrinos are produced roughly isotropically, with around double the number of ν_{μ} generated compared to ν_{e} .

Experiments near the surface of the Earth can detect atmospheric neutrinos over a large range of L/E, covering the atmospheric oscillation maximum. This again allows constraint of the atmospheric oscillation parameters Δm_{32}^2 and $\sin^2(\theta_{23})$. In addition to this, neutrinos created on the other side of the Earth are sensitive to the the MSW effect and MSW resonance. This adds sensitivity to the θ_{13} sector, also allowing constraints of the mass hierarchy, θ_{23} octant and δ_{CP} .

This analysis aims to select muon neutrinos and antineutrinos at the Far Detector. Reconstructing the angle of the neutrino in the detector allows an estimation of the distance between production and detection so that the L/E shape of the spectrum can be compared to predictions.

3.2.1 Selection

There are two selection paths in the atmospheric analysis – an event may be selected either because it is upward-going or because the highest point on its track is contained within the Far Detector's fiducial volume. Both selections are designed to remove the dominant background to the atmospheric ν_{μ} signal, which are cosmic or atmospheric muons.

First, data quality cuts are made to ensure good detector performance. These are outlined in section 3.1.2. Next, beam events are removed by applying a time window cut, excluding snarls beginning within 15 μ s of a beam spill. Preselection cuts are then made to ensure a track is reconstructed and remove potentially poorly-reconstructed events, before the selection forks into two.

Upward-going muons are identified mostly by timing. Further cuts are made on the timing reconstruction quality of the event as well as the event's topology, the event timing itself and position. Events passing these cuts are then separated into one of two categories – those whose tracks are completely contained in the FD fiducial volume, called 'Fully-Contained Upward-Going' and those that exit the fiducial volume, called 'Partially-Contained Upward-Going'. Contained-highest-end events are selected by a range of containment cuts, and event-topology cuts, both aimed at eliminating cosmic muons from the selection sample. Events are again divided into 'Fully-Contained' and 'Partially-Contained samples', while those with sufficiently high-quality timing-reconstruction may be defined as 'Fully-Contained Downward-Going'.

Due to the difference between the ν_{μ} and $\bar{\nu}_{\mu}$ oscillation probabilities in the presence of matter effects, these two samples are treated separately in the fit. These samples are further divided into four bins by the estimated resolution on their reconstructed L/E, improving sensitivity to oscillations. One further sample is included in the analysis: events consisting of contained-vertex showers – mainly NC and ν_e -CC – that help constrain the overall atmospheric neutrino flux in the fit. A full description of the Atmospheric selection criteria can be found in [109].

3.2.2 Far Detector Prediction

The single-detector nature of this analysis makes the generation of predictions less complex than in the beam disappearance analysis, though more vulnerable to systematic uncertainties. Predictions of the neutrino flux at the Far Detector are based on calculations by Barr *et al.* [110]. Oscillations can then be applied to this flux

according to energy, flavour and propagation distance, with matter effects applied assuming a four-layer model of the Earth's electron density [111] based on the PREM model [112]. Finally, the detector simulation described in Section 2.6.2 is used to convert this oscillated flux into Far Detector predicted spectra. As in the beam analysis, systematic effects are modelled by varying the parameters used in this prediction process to generate new 'shifted' predictions. The predictions are again generated on a grid of oscillation and systematic parameter-space, to be interpolated between during the fit.

3.2.3 Systematics

The lack of a Near Detector measurement in this analysis leaves the predictions of the Far Detector flux highly model-dependent. This leads to the necessary inclusion of several systematic parameters in the analysis, mostly relating to the properties and flux of neutrinos in the atmosphere. In total, 11 systematics are included in the analysis. A full discussion of the source and estimation of these uncertainties can be found in [111].

- Flux Uncertainties: Contained ν_{μ} Normalisation, Rock ν_{μ} Normalisation, Contained ν_{μ} up/down ratio, Contained ν_{e}/ν_{μ} ratio, Contained $\nu_{\mu}/\bar{\nu}_{\mu}$ ratio, Rock $\nu_{\mu}/\bar{\nu}_{\mu}$ ratio, Contained NC/CC ratio, ν Spectrum, $\bar{\nu}$ Spectrum
- Energy Scales: Track Energy, Shower Energy.

The flux systematics (all except the track and shower energy scales) are chosen to cover differences between different flux models. The dominant uncertainties here are in the overall scale of the neutrino flux, with a 15% uncertainty on the contained ν_{μ} normalisation. Rock events have a higher energy on average, leading to a 25% uncertainty on the rock ν_{μ} normalisation [111].

3.3 Beam Neutrino Appearance

Whilst MINOS was designed to measure the atmospheric sector parameters, it may also be used to observe electron-neutrino appearance in the ν_{μ} -dominated NuMI beam, allowing constraint of the parameters θ_{13} and δ_{CP} . The approximate equation for the relevant oscillation probability is

$$P(\nu_{\mu} \to \nu_{e}) \approx \sin^{2}(2\theta_{13}) \sin^{2}\theta_{23} \frac{\sin^{2}(A-1)\Delta}{(A-1)^{2}} + 2\alpha \sin^{2}\theta_{13} \cos \delta \sin 2\theta_{12} \sin 2\theta_{23} \frac{\sin A\Delta}{A} \frac{\sin(A-1)\Delta}{A-1} \cos \Delta \qquad (3.8) - 2\alpha \sin^{2}\theta_{13} \sin \delta \sin 2\theta_{12} \sin 2\theta_{23} \frac{\sin A\Delta}{A} \frac{\sin(A-1)\Delta}{A-1} \sin \Delta,$$

where

$$\Delta = \frac{\Delta m_{32}^2 L}{4E} \qquad A = \frac{G_F n_e L}{\sqrt{2}\Delta} \qquad \alpha = \frac{\Delta m_{21}^2}{\Delta m_{32}^2}.$$
(3.9)

It can be seen from Equation 3.8 that the number of appearing electron neutrinos in the MINOS beam will be a function not only of θ_{13} and δ_{CP} but of a number of oscillation parameters. Of particular note is Δm_{32}^2 , on which the appearance probability depends asymmetrically, giving sensitivity to the neutrino mass hierarchy.

The observation of ν_e appearance in MINOS is challenging for several reasons. The first is that $\nu_{\mu} \rightarrow \nu_e$ oscillation probability is only approximately 10% at the MINOS L/E (see Figure 2.2). The second is that the MINOS detectors were not built to look for ν_e events, and as such their reconstruction is difficult. This both increases background contamination and reduces energy resolution.

3.3.1 Data

It is very difficult to select electron-neutrinos in the MINOS detectors. Given the low number of events selected in this analysis, the MINOS data are combined into two runs: one FHC and one RHC. Poor resolution on the energy of the incident neutrino smears out small differences between the different beam configurations in the beam runs used in the beam disappearance analysis.

3.3.2 Selection

Electron neutrino charged-current events result in a nuclear recoil and the creation of an energetic electron. The small mass of the electron means that it is scattered more readily than a muon, generating bremsstrahlung photons and further e^+/e^- pairs. The result is a compact electromagnetic shower. In MINOS, the relatively sparse instrumentation optimised for long muon tracks leads to EM showers with very little recognisable structure in the detector, as shown in Figure 2.17. This makes it difficult to separate ν_e -CC events from the neutral-current background.

Preselection

The same data quality, cosmic and fiducial cuts as for the ν_{μ} disappearance analysis, described in 3.1.2, are first applied to the data. Additional cuts aim to increase the purity of the sample before the labour-intensive selection stage.

- Reconstructed Energy: Events with E_{reco} < 1GeV and E_{reco} > 8GeV are removed. At low energy, NC events dominate and the small event sizes make it harder to pick out ν_es. At high energy, the probability P_{νµ→νe} vanishes a cut at 8 GeV removes some of the intrinsic background of ν_es created in the beam. Figure 3.6 illustrates the effect of this cut.
- Track Cuts These aim to remove obvious ν_{μ} CC events. Events with tracks longer than 24 detector planes are cut, as well as events with more than 15 'track-like' planes planes containing only hits from reconstructed tracks.
- Shower Requirement At least one shower is required to have been reconstructed in the event. Both NC and ν_{μ} CC events may leave a track with no shower, whereas ν_e CC events will always contain an EM shower.
- Contiguous Planes At least 4 consecutive planes in the event are required to contain energy deposits of at list 0.5 MIPs. EM showers tend to be more compact than hadronic ones, meaning that this cut will remove many NC events.

LEM Selection

In order to maximise signal and background separation, a computationally-intensive multivariate analysis algorithm, called Library Event Matching (LEM), is used for the analysis selection. The LEM algorithm aims to improve particle identification by identifying events based on their raw hit positions, times and pulse-heights, as opposed to usual methods involving analysis of reconstructed variables like track



Figure 3.6: Figure illustrating the ν_e analysis preselection cuts on reconstructed energy. Below 1 GeV, the sample is poorly reconstructed and dominated by NC events, while above 8 GeV, the signal drops to almost zero [22].

and shower energy, maximising the information available to identify the event. The method was developed over several years and is documented in [113, 114].

The algorithm works by comparing each event to a large library of simulated events, identifying the N best matches. Variables are calculated from the truth properties of these N events, which are then used as inputs to a neural network. This network is trained against simulated data and produces a single output indicating the probability of the input event being a ν_e CC event.

In MINOS, the number of matches was chosen to optimise the performance of the discriminator as 50. The variables used as input to the NN are:

- **f50**: The fraction of matches which are ν_e CC
- **y50**: The mean y of the ν_e CC Matches
- **q50**: The average fractional matched charge for ν_e CC matches

3.3.3 Far Detector Prediction

This analysis uses data-driven methods to decompose the ND data into its components, which are extrapolated separately into FD predictions [70]. NC, ν_{μ} CC, and beam ν_{e} CC events are extrapolated by multiplying the ND data with the ratio of FD MC to ND MC for each bin, i.e.

$$F_i = \frac{f_i}{n_i} N_i, \tag{3.10}$$

where F_i and N_i are the FD prediction and ND data respectively for bin i; f_i and n_i are the Monte Carlo simulated samples for the FD and ND respectively.

The extrapolation of appearing components $-\nu_e$ and ν_{τ} – is described in [115], and is largely similar to the matrix method outlined in Section 3.1.3.

3.3.4 Systematics

The ν_e appearance analysis is subject to many of the same uncertainties as the beam ν_{μ} disappearance analysis. The presence of appearing flux components leads to some extra uncertainty given that these can not be measured in the Near Detector. For many of the systematics listed, the beam-based backgrounds and appearing components are shifted by different amounts [115].

The following ten systematic effects are accounted for in the ν_e fit [116]:

- Normalisation Appearing ν_e signal, combined backgrounds.
- Cross-talk Mismodelling
- Calibration Uncertainties: Relative ND and FD Energy Scale, ND and FD Gains, Linearity Correction, Attenuation Correction, Strip-to-Strip Correction, Absolute Energy Scale, Hadronic Energy Scale

A 2% normalisation systematic is applied to the background F/N ratios, while 2.4% systematics are given to the appearing components to include additional uncertainty on the beam ν_{μ} flux (eg. on the size of the NC background at the Near Detector). The use of the scintillator hits in the LEM selector cause sensitivity to cross-talk mismodelling in the detectors. The remainder of the systematics come from calibration uncertainties, which may affect the LEM PID values of events as well as their energies.

3.4 Fitting the MINOS and MINOS+ Data

A statistical analysis must be performed to estimate the parameters of the threeflavour neutrino oscillation model, combined with experiment-specific modelling from the measured data in each analysis. In MINOS and MINOS+, this is done by *maximum-likelihood estimation*.

3.4.1 Maximum-Likelihood Estimation

More specifically, the data are fitted to a statistical model based on the oscillation parameters and the systematic parameters chosen in the analysis. These parameters are represented as a vector $\boldsymbol{\theta} = \theta_1, \theta_2, \dots, \theta_m$. Let the number of events in analysis bin *i* predicted by the model be $\lambda_i(\boldsymbol{\theta})$. Given a set of independently distributed data points, $\boldsymbol{x} = x_1, x_2, \dots, x_n$ (in this case the number of events observed in each analysis bin), the *likelihood function* is defined as

$$L(\boldsymbol{x};\boldsymbol{\theta}) = \prod_{i=1}^{n} f(x_i;\boldsymbol{\theta}), \qquad (3.11)$$

where $f(x_i; \boldsymbol{\theta})$ is the probability of observing exactly x_i events in this bin for the model parameters specified. Since in MINOS the data are poisson distributed, $f(x_i; \boldsymbol{\theta}) = \frac{\lambda(\boldsymbol{\theta})_i^x e^{-\lambda(\boldsymbol{\theta})}}{x_i!}.$

The maximum-likelihood estimate of $\boldsymbol{\theta}$ is then the value of $\boldsymbol{\theta}$ that maximises the likelihood of the data. For practical purposes, the minimum of $-2 \ln L(\boldsymbol{x}; \boldsymbol{\theta})$ is equivalently found.

$$-2\ln L(\boldsymbol{x};\boldsymbol{\theta}) = 2\sum_{i}^{n} \lambda(\boldsymbol{\theta}) + \ln x_{i}! - x_{i}\ln\lambda(\boldsymbol{\theta})$$
(3.12)

Using Stirling's approximation, $\ln n! \approx n \ln n - n$, gives

$$-2\ln L_0 \equiv -2\ln L(\boldsymbol{x};\boldsymbol{\theta}) = 2\sum_{i}^{n} \lambda(\boldsymbol{\theta}) - x_i + x_i \ln x_i - x_i \ln \lambda(\boldsymbol{\theta})$$
(3.13)

This is the form of the (unconstrained) likelihood function employed in all of the MINOS and MINOS+ three-flavour oscillation analyses. Computational methods described in Section 3.4.2 are used to perform the minimisation.

The Likelihood-Ratio Test

The comparison of hypotheses within the model of three-flavour neutrino oscillations and the calculation of *confidence intervals* is achieved using likelihood-ratio tests. The likelihood ratio is defined as

$$\Lambda = \frac{L(\boldsymbol{x}; \boldsymbol{\theta}_0)}{L(\boldsymbol{x}; \boldsymbol{\theta}_1)},\tag{3.14}$$

where $\boldsymbol{\theta}_0$ is called the *null hypothesis* and $\boldsymbol{\theta}_1$ the *alternative hypothesis*.

The test statistic Λ is itself a random variable sampled from a distribution $f(\Lambda)$. The null hypothesis is rejected at a significance, S, if the fraction of tests expected to have Λ smaller than that measured is greater than the chosen significance of the test. In other words, the cumulative distribution function, $F(\Lambda)$ gives the *p*-value for the measured value of Λ .

Nuisance Parameters and the Profile Likelihood

In the MINOS analyses, we aim to constrain only one or two oscillation parameters – the *parameters of interest*, $\boldsymbol{\theta}_{I}$ – from a model with multiple other oscillation and systematic parameters, known as *nuisance parameters*, $\boldsymbol{\theta}_{N}$.

The profile likelihood ratio, λ is calculated as

$$\lambda = \frac{L(\boldsymbol{x}; \boldsymbol{\theta}_{I}, \boldsymbol{\hat{\theta}}_{N})}{L(\boldsymbol{x}; \boldsymbol{\hat{\theta}})}, \qquad (3.15)$$

where $L(\boldsymbol{x}; \boldsymbol{\theta}_{I}, \hat{\boldsymbol{\theta}}_{N})$ is the null hypothesis corresponding to $\boldsymbol{\theta}_{I}$, calculated by maximising the likelihood with respect to the nuisance parameters. $\hat{\boldsymbol{\theta}}$ is the maximum likelihood estimator for $\boldsymbol{\theta}$, or the global maximum for the likelihood. Again, the cumulative distribution of λ gives the *p*-value of the hypothesis.

Confidence Intervals

The results of an analysis are generally presented as one- or two-dimensional confidence intervals created from a profile likelihood. Before this can be done, the distribution of λ must be known. The distribution can be approximated by repeating multiple fake experiments – this computationally intensive approach is required in some cases, for example in the MINOS sterile neutrino analysis [71]. More commonly, the fact is used that for a large sample size – under certain conditions $[117] - 2 \ln \lambda$ is distributed as χ^2 , with the number of degrees of freedom defined by the number of parameters of interest. This holds in a wide range of scenarios and is used in the MINOS and MINOS+ analysis, allowing us to calculate *p*-values from standard tables of the cumulative χ^2 distribution.

Penalty Terms

The model used to fit the MINOS and MINOS+ analyses includes both neutrino oscillation parameters and systematic parameters. In the case of systematics in particular, we may want to incorporate prior knowledge of the value of a parameter into the likelihood function. This is done by adding a *penalty term* – an additive term in the log-likelihood function that increases with distance from its constrained central value. This term can have any form, but is most commonly chosen to approximate a Gaussian prior likelihood. For P independent parameters θ_k with Gaussian penalty terms, this gives a constrained log-likelihood

$$-2\ln L(\boldsymbol{x};\boldsymbol{\theta}) = -2\ln L_0(\boldsymbol{x};\boldsymbol{\theta}) + \sum_k^P (\theta_k - \mu_k)^2 / \sigma_k^2, \qquad (3.16)$$

where $L_0(\boldsymbol{x}; \boldsymbol{\theta})$ is the unconstrained likelihood defined in 3.13, μ_k is the measured value of θ_k and σ_k is the uncertainty on that measurement.

3.4.2 GhostFitter

MINOS and MINOS+ use a software package called GhostFitter to fit the oscillation and systematics model to the data. GhostFitter is a *template fitter*, using the ROOT package MINUIT [118] to minimise the likelihood function. GhostFitter is used to simultaneously fit the MINOS beam appearance, disappearance and atmospheric disappearance data, along with the MINOS+ beam disappearance in the next chapter. Upgrades made to the software for this thesis now allow for a more thorough treatment of the correlation between systematic parameters in these analyses. **GhostFitter** greatly speeds up the fitting process by decoupling the extrapolation and oscillation of Far Detector predictions from the parameter minimisation. It does this by pre-computing a large set of oscillated Far Detector predictions, called *templates* on a grid in parameter space. These are then loaded by the fitter and interpolated to generate a prediction at an arbitrary point.

Templates

Far Detector predictions are first generated on a grid of points in oscillation parameter space. For the MINOS combined and MINOS+ analysis, a 4-dimensional grid in Δm_{32}^2 , $\sin^2(\theta_{23})$, $\sin^2(\theta_{13})$ and $\delta_{\rm CP}$ is used. Oscillation parameters not included in the grid must be kept fixed during the fit. Specifically, in each of the analyses, $\sin^2(\theta_{12})$ is fixed at 0.307 and Δm_{12}^2 is fixed at 7.54⁻⁵eV². MINOS has very little sensitivity to either of these parameters.

For each point on the oscillation grid, ± 1 and 2σ shifted predictions are stored for each systematic, as well as a single unshifted prediction. A Far Detector spectrum is generated for an arbitrary point in parameter space by interpolating between points on the grid and in systematic space. First, multilinear interpolation is applied to reach the required point in oscillation space – performing a weighted sum of the corresponding templates in each of the files neighbouring the oscillation point (16 files in a 4D grid). From here, weighted linear interpolation – described further in [13] – is used to sum the systematically shifted predictions.

The result is that Far Detector predictions can be generated on a continuous surface of parameter (oscillation and systematic) space, as is required by the MIGRAD algorithm. MIGRAD can then be used to minimise over the parameter space to produce the maximum-likelihood-estimator and profile likelihoods.

External Data

In the case where template files do not exist for an analysis, it can be incorporated into a fit by importing its externally-generated likelihood grid – profiled over its systematic parameters. **GhostFitter** interpolates the likelihood from the grid for a given oscillation point and adds the value to the likelihood. This is the method employed to include the MINOS ν_e appearance data in the combined MINOS and MINOS+ fit.

Chapter 4

MINOS+ Beam Disappearance Analysis

This chapter documents the analysis of the first year of MINOS+ beam data according to the standard three-flavour neutrino oscillation model. The analysis uses the techniques of the MINOS beam disappearance analyses, as outlined in Section 3.1, with some small changes. The data is analysed both as a distinct MINOS+ set as well as in combination with the MINOS beam and atmospheric data.

This analysis is unique for several reasons. It is the first physics analysis of data from the upgraded high-energy NuMI beam. In terms of techniques, this is also the first analysis to use the updated ND reconstruction described in Section 2.8. The selection used in the 2014 MINOS analysis is also modified.

4.1 Data

MINOS+ began data taking in September 2013. The rate of accumulation of POT in MINOS and MINOS+ shown in Figure 4.1. The beam upgrade is still in progress and the intensity of the proton beam to NuMI over the data taking period is at a similar level to that at the end of MINOS running. However, the increase in beam energy, along with minimal down-time, has allowed MINOS+ to analyse around a quarter as much data as was collected in seven years of MINOS running.



Figure 4.1: Rate of protons delivered per week to MINOS and MINOS+ from start of running to 4 May 2015. Protons from the upgraded NuMI beam used in MINOS+ are shown in pink, which has only been operated in FHC mode. FHC MINOS data is shown in green, with RHC data shown in orange. Special, short, high-energy MINOS runs are marked in red [23].

 2.99×10^{20} protons-on-target were analysed for this thesis, taken in the period September 2013 to September 2014. This is known as *Run XI*. Runs I to X make up the MINOS beam data set, which is analysed in combination with MINOS+ Run XI at the end of this chapter. Figures 4.2 and 4.3 show the ratio of selected Near and Far Detector events to recorded POT against time over the run period. These data in these plots are consistent with a constant neutrino rate across the run, indicating good performance in the detectors and beam. Drift and/or large fluctuations in the neutrino rate may indicate detector malfunction or target degradation.

4.2 Selection

The selection of CC ν_{μ} events in MINOS+ is modified only slightly from that used in the last MINOS analysis, as described generally in Section 3.1.2. The low-energy MINOS PID is no longer used, while the kNN for the general PID has been retrained for MINOS+ data. In addition to this, a near-detector specific geometric cut around the magnetic coil has been relaxed.



Figure 4.2: Rate of selected events per 10^{16} POT in the Near Detector during Run XI [24].

4.2.1 roID retuning

The kNN PID described in Section 3.1.2, roID, was retrained on MINOS+ simulated data to help optimise performance in the higher-energy beam [26]. Figure 4.4 illustrates the improvement in sensitivity in MINOS+ when using the retrained kNN PID.

4.2.2 ND Coil Hole Cut

In the last MINOS beam ν_{μ} disappearance analysis, a cut was introduced at the Near Detector to remove events crossing or ending near the Near Detector coil hole. Specifically, tracks that are reconstructed as having crossed or ended within 60 cm of the coil hole are removed. This cut was introduced to improve agreement between ND data and simulation, which was strained by poor understanding of the magnetic field near the coil as well as truncation of muon tracks that end in the hole.



Figure 4.3: Rate of selected events per 10^{18} POT in the Far Detector during Run XI. The first bin is approximately 2.4σ from the mean, but is consistent with the statistical expectation for a 70-bin histogram [25].



Figure 4.4: Fake data sensitivities to the parameters of the 2-flavour approximate disappearance model in MINOS+ for the old (blue) and new (red) kNN. A third line (black) shows the sensitivity when using perfect PID information. Taken from [26].

This cut has a significant impact on the selection efficiency at the Near Detector, particularly for high-energy events, which are longer and more likely to cross the coil. The increase in beam energy in MINOS+ means that this cut removes a large fraction of Near Detector events. This, along with the fact that the agreement between real and simulated data for events ending in the coil hole has improved since MINOS, has led to a relaxation of the cut.

For tracks ending in the spectrometer region, the cut geometry remains the same. Tracks ending in the calorimeter, however, are required only to have not crossed the coil hole, removing the radial 60 cm cut.

4.3 Fitting

The fit is performed with respect to four oscillation parameters: Δm_{32}^2 , $\sin^2(\theta_{23})$, $\sin^2(\theta_{13})$ and $\delta_{\rm CP}$, and four systematic parameters: normalisation, NC background, shower energy and track energy. Penalty terms are applied to the systematics to represent our prior estimates of their values. A further penalty term is applied to $\sin^2(\theta_{13})$, constraining it to 0.0242 ± 0.0025 at 1σ , based on an average of Daya Bay, Double-Chooz and RENO measurements. $\sin^2(\theta_{13})$ is treated as a nuisance parameter.

4.3.1 Binning

The MINOS+ beam data is divided into two samples: one made up of events with tracks reconstructed as positively-charged, and one with tracks reconstructed as negatively-charged. Due to the increase in energy of the NuMI beam, MINOS+ has relatively few events at low energy. Assuming the best fit MINOS oscillation parameters, approximately one event is expected below 1.5 GeV.

For this reason, MINOS+ uses only one reconstructed bin from 0 to 1.5 GeV, replacing three 0.5 GeV MINOS bins. The rest of the binning scheme remains the same as used in [27], with 0.25 GeV bins from 1.5 to 20 GeV; 1 GeV bins from 20 to 30 GeV; 2 GeV bins from 30 to 50 GeV and one bin from 50 to 200 GeV. The separation of the data into a positively-charged and a negatively-charged sample gives a total of 196 analysis bins.

4.3.2 Systematics

The set of systematics considered in the analysis remains unchanged from MINOS. The impact of the four systematic parameters analysed in MINOS were studied along with a fifth parameter representing the total uncertainty on various measurements of the properties of the new NuMI beam that are assumed in the extrapolation procedure, for example the exact position of the magnetic horns. Figure 4.5 shows the shift in the MINOS+ two-flavour best fit point caused by fixing each of these systematic parameters to $\pm 1\sigma$, providing an estimate of their impact on the parameters of interest. This shows that the beam flux error remains small compared to the four dominant systematics.



Figure 4.5: Shift in the MINOS+ two-flavour best fit point caused by $\pm 1\sigma$ shifts to five systematic parameters.

Of the four parameters included in the fit, the size of the shower and track energy scale uncertainties are based on the same measurements as made for MINOS, as is the relative normalisation uncertainty. Differences in the selection and reconstruction used in MINOS and MINOS+ are seen to be small. Reanalysis of the NC background uncertainty in MINOS+ leads to an increase in the magnitude of this uncertainty from 20% to 50% [119]. Note that this value corresponds to the uncertainty of the

absolute scale of the NC background, and that this scale is reduced with respect to the total number of selected events due to the increase in beam energy in MINOS+.

4.4 Results

The selected ND energy spectrum can be seen in Figure 4.6, along with the spectrum predicted by simulation. The difference between data and simulation is covered by the 1σ systematic uncertainty on the simulated spectrum.

Given the Near Detector spectrum, and assuming oscillations at the MINOS best fit point, 1969.2 events are expected at the Far Detector, while 2208.5 events are expected in the absence of any oscillations. The number of FD events selected in the analysis is 1889.

Figure 4.7 shows the energy distribution of the selected FD data, along with predictions for no oscillations, at the MINOS best fit oscillations and at the MINOS+ best fit point, which is specified in Table 4.1. Note that throughout this chapter, the distinction between normal and inverted hierarchy is specified by the sign of Δm_{32}^2 (positive values indicate normal hierarchy). This plot illustrates the consistency between the MINOS and MINOS+ data, in spite of the large changes to the neutrino beam. The ratios of data and oscillated predictions with respect to the MINOS+ prediction in the absence of oscillations is shown in Figure 4.8, showing the classic oscillation dip.

The MINOS+ best fit point is found at $\Delta m_{32}^2 = -2.89 \times 10^{-3} \text{eV}^2$, with a reduced χ^2 of 1.04, indicating a high goodness of fit. Note that there is no sensitivity to the hierarchy and an essentially equally likely best fit point exists for positive Δm_{32}^2 . This is some distance from the MINOS best fit of $-2.37^{+0.11}_{-0.07} \times 10^{-3} \text{eV}^2$. Two-dimensional 68% and 90% confidence intervals in Δm_{32}^2 and $\sin^2(\theta_{23})$, constructed using profile likelihood ratios as specified in Section 3.4.1, are shown for the MINOS+ data in Figure 4.9. The MINOS best fit point can be seen to lie within the MINOS+ 68% confidence interval. The difference in $-2 \ln L$ between the MINOS+ best fit point and the best fit at the MINOS oscillation parameters is 1.12, corresponding to a *p*-value of 0.57. In other words, assuming the true oscillation parameters to lie at the



Figure 4.6: Selected MINOS+ Near Detector data binned in reconstructed energy, along with ND simulated data. The difference between data and simulation is covered by the 1σ systematic error on the simulation [25].



Figure 4.7: Selected MINOS+ Far Detector data binned in reconstructed energy. Predictions at the MINOS+ best fit point and a prediction assuming the best-fit oscillation parameters from the MINOS analysis are also shown.

MINOS best fit point, MINOS+ data more consistent with this best fit than those observed would be expected on average 43% of the time.

There are four degenerate minima in the likelihood surface, reflecting the approximate mass hierarchy and mixing-angle octant symmetry in the ν_{μ} disappearance equation. Given that matter effects have very little impact over the relatively short



Figure 4.8: The number of selected MINOS+ events per reconstructed energy bin over the number expected given no neutrino oscillations is shown. The ratio of the MINOS+ best-fit prediction and MINOS prediction with respect to the no-oscillation prediction are also shown.

Parameter	Best Fit Value
$\overline{\Delta m^2_{32}}$	$-2.89 \times 10^{-3} \mathrm{eV}^2$
$\sin^2(\theta_{23})$	0.264
$\sin^2(\theta_{13})$	0.02421
$\delta_{ m CP}$	0.99π
Normalisation	$+0.15\sigma$
NC Background	$+0.12\sigma$
Shower Energy	-0.65σ
Track Energy	-0.17σ

Table 4.1: Best-fit parameters from fit to MINOS+ data. The negative value of Δm_{32}^2 indicates a preference for the inverted hierarchy, however MINOS+ has no real sensitivity to the mass ordering. The best fit values of all nuisance parameters are also given. Systematic parameter values are given as shifts from the nominally estimated value in standard deviations.

distance between the NuMI facility and the Far Detector, this degeneracy is almost perfect in the MINOS+ data. The $-2 \ln L$ values for the four minima are shown in Table 4.2.


Figure 4.9: Likelihood contour in Δm_{32}^2 - $\sin^2(\theta_{23})$ for MINOS+.

Mass Hierarchy	θ_{23} Octant	$-2\Delta \ln L$
Normal	Lower	0.005
Normal	Upper	0.001
Inverted	Lower	0.000
Inverted	Upper	0.010

Table 4.2: $-2\Delta \ln L$ for the four degenerate minima in the MINOS+ Δm_{32}^2 - $\sin^2(\theta_{23})$ likelihood surface with respect to the global minimum.

Finally, the one-dimensional profile likelihoods for Δm_{32}^2 and $\sin^2(\theta_{23})$ are shown in Figure 4.10 Profiles are given for the normal and inverted mass hierarchy hypotheses in each plot. One-dimensional confidence intervals are constructed from these profiles as described in Section 3.4 to give

$$\Delta m_{32}^2 = +2.96^{+0.54}_{-0.46} \times 10^{-3} \text{eV}^2 \text{ (normal hierarchy) and}$$

$$\Delta m_{32}^2 = -2.89^{+0.46}_{-0.50} \times 10^{-3} \text{eV}^2 \text{ (inverted hierarchy) (68\% C.L.)}$$

 $0.16 < \sin^2(\theta_{23}) < 0.87$ (normal hierarchy) and $0.16 < \sin^2(\theta_{23}) < 0.86$ (inverted hierarchy) (90% C.L.)



Figure 4.10: Likelihood profiles over Δm_{32}^2 (left) and $\sin^2(\theta_{23})$ (right) are shown for the MINOS+ data. Separate profiles are presented assuming normal and inverted mass hierarchy.

4.5 Combined MINOS+ and MINOS Analysis

The MINOS+ data analysed above can be combined with the full MINOS data set to improve the MINOS constraints on the three flavour oscillation parameters. Combination of the three MINOS analyses was first performed in [27].

4.5.1 Data

The amount of data used from each analysis in this combination is as follows:

- 2.99 × 10²⁰ POT FHC MINOS+ beam ν_{μ} disappearance.
- 10.71×10^{20} POT of FHC and 3.36×10^{20} POT of RHC MINOS beam ν_{μ} disappearance and ν_e appearance.
- 48.67 kT-yr Atmospheric ν_{μ} disappearance.

Beam ν_{μ} Disappearance

A summary of the MINOS and MINOS+ beam data runs used in this analysis is shown in Table 4.3. The distributions of the MINOS and MINOS+ beam data are plotted together in Figure 4.11, illustrating the relative amounts of data taken by each. MINOS+ contributes a significant fraction of the data, but mostly away from the approximately 3 GeV oscillation maximum.

Name	Beam Configuration	POT ($\times 10^{20}$)
Run I	low-energy FHC	1.27
Run I	pseudo-high-energy FHC	0.15
Run II	low-energy FHC	1.94
Run III	low-energy FHC	3.89
Run IV	low-energy RHC	1.71
Run V	low-energy FHC	0.46
Run VI	low-energy FHC	0.57
Run VII	low-energy RHC	1.24
Run IX	low-energy RHC	0.41
Run X	low-energy FHC	1.47
Run XI	medium-energy FHC	2.99

Table 4.3: Summary of MINOS and MINOS+ runs used in this analysis. A total of 10.71×10^{20} POT of FHC data and 3.36×10^{20} POT of FHC data was takenduring MINOS running.

Atmospheric ν_{μ} Disappearance

The Far Detector continued to take atmospheric data during the NuMI shutdown. In June 2014, the 10.79 kT-yrs of data taken since the analysis published in [27] was analysed in combination with the MINOS data [1]. The rate of atmospheric data taken over the lifetime of MINOS and MINOS+ is shown in 4.12.

ν_e Appearance

No MINOS+ data has yet been analysed for ν_e appearance. The number of events selected in this analysis is small, and the MINOS beam runs are combined into one 10.71×10^{20} POT FHC sample and one 3.36×10^{20} POT RHC sample. The selected data along with best-fit predictions from the appearance-only analysis are shown in Figure 4.13.



Figure 4.11: Selected MINOS+ and MINOS data binned in reconstructed energy. Predictions at the MINOS+ best fit point and a prediction assuming the best-fit oscillation parameters from the MINOS analysis are also shown.



Figure 4.12: Number of events selected for atmospheric analyses over time in MINOS and MINOS+. The blue points show the data analysed for the last published MINOS analysis [27]. An analysis including the MINOS and MINOS+ data has since been performed [1].

The ν_e appearance data are included in the fit using a four-dimensional likelihood surface in the oscillation parameters. A two-dimensional profile of this surface can be seen in Figure 4.14, which illustrates the effect of the data on the combined analysis,



Figure 4.13: The reconstructed energy distributions for three α_{LEM} ranges. The events with energy greater than 5 GeV are combined into a single bin for the fits. The vertical bars through the data points denote statistical uncertainties. The signal predictions assume $\sin^2(2\theta_{13}) = 0.051$, $\Delta m_{32}^2 > 0$, $\delta = 0$, and $\theta_{23} = \pi/4$. The plots in the left column correspond to data collected in the ν beam mode. The plots in the right column correspond to data collected in the $\bar{\nu}$ beam mode [28].



showing asymmetry in the likelihood with respect to the octant of $\sin^2(\theta_{23})$ and the mass hierarchy.

Figure 4.14: A two-dimensional $(\Delta m_{32}^2 - \sin^2(\theta_{23}))$ profile of the log-likelihood surface from the MINOS ν_e appearance analysis used in the combined MINOS and MINOS+ result. Note that the surface is asymmetric in the octant of θ_{23} and the sign of Δm_{32}^2 . This surface contributes the majority of the small sensitivity to the octant and hierarchy in the MINOS and MINOS+ combined analyses.

4.5.2 Systematic Parameter Correlation

Care must be taken in the treatment of systematics when combining the MINOS and MINOS+ analyses in a fit. The discussion in Section 4.3.2 indicates that the values of the track and shower energy scales, as well as the relative normalisation, are presumed to be 100% correlated between MINOS and MINOS+, and they are treated as such in this combined fit. The relationship between the MINOS and MINOS+ NC

background uncertainties is less clear, and they are conservatively assumed to be completely uncorrelated.

Correlation between the track and shower energy scales between the beam and atmospheric disappearance analyses is accounted for in the MINOS combined fit [27]. In this analysis, the corresponding MINOS+ and MINOS beam and MINOS atmospheric parameters are all treated as correlated – i.e. all three are modelled with the same fitted parameter. The inclusion of the ν_e appearance data as a profiled likelihood surface precludes the correlation of any of the appearance systematics.

4.5.3 Results

The best fit point for this combined analysis is found at $\Delta m_{32}^2 = -2.45 \times 10^{-3} \text{eV}^2$, with a reduced χ^2 of 1.04. The best fit parameters are given in Table 4.4. Figure 4.15 shows the 68% and 90% confidence intervals in Δm_{32}^2 and $\sin^2(\theta_{23})$ obtained from this analysis. The values of $-2\Delta \ln L$ for the four degenerate minima are given in Table 4.5. The inverted hierarchy, lower octant hypothesis is very weakly preferred, with the normal hierarchy, upper octant hypothesis disfavoured at $-2\Delta \ln L = 1.66$.

Figure 4.16 compares the contour in Figure 4.15 to those from each the MINOSonly analysis and from T2K. The MINOS+ data move the MINOS best fit point to larger $|\Delta m_{32}^2|$ and away from maximal θ_{23} . The T2K best fit point lies just inside the 90% confidence limit for this analysis.

The one-dimensional Δm_{32}^2 and $\sin^2(\theta_{23})$ profiles are shown in Figure 4.17. Again, one-dimensional confidence intervals are constructed to give measurements of the two parameters:

$$\Delta m_{32}^2 = +2.45^{+0.08}_{-0.11} \times 10^{-3} \text{eV}^2$$
 (normal hierarchy) and
 $\Delta m_{32}^2 = -2.42^{+0.09}_{-0.11} \times 10^{-3} \text{eV}^2$ (inverted hierarchy) (68% C.L.)

 $0.35 < \sin^2(\theta_{23}) < 0.65$ (normal hierarchy) $0.35 < \sin^2(\theta_{23}) < 0.66$ (inverted hierarchy) (90% C.L.).

This is an approximate 5% reduction in the size of the uncertainty on Δm_{32}^2 in the MINOS-only analysis presented in [1] and an approximate 9% reduction compared to the last published MINOS analysis [27].

Parameter	Best Fit Value
$\overline{\Delta m^2_{32}}$	$-2.46 \times 10^{-3} \mathrm{eV^2}$
$\sin^2(\theta_{23})$	0.404
$\sin^2(\theta_{13})$	0.0244
$\delta_{ m CP}$	2.00π
Normalisation	$+1.12\sigma$
MINOS NC Background	$+0.11\sigma$
MINOS+ NC Background	$+0.23\sigma$
Shower Energy	-0.17σ
Track Energy	0.69σ

Table 4.4: Best-fit parameters from a combined fit to MINOS and MINOS+ data. The
MINOS and MINOS+ NC background uncertainties are treated as uncorrelated.

Mass Hierarchy	θ_{23} Octant	$-2\Delta \ln L$
Normal	Lower	0.06
Normal	Upper	1.61
Inverted	Lower	0.00
Inverted	Upper	0.66

Table 4.5: $-2\Delta \ln L$ for the four degenerate minima in the combined MINOS+ and MINOS Δm_{32}^2 - $\sin^2(\theta_{23})$ likelihood surface with respect to the global minimum.



Figure 4.15: 68% and 90% confidence limits in $\Delta m_{32}^2 - \sin^2(\theta_{23})$ for the MINOS, MINOS+ combined analysis. The best fit point – marked with the star – is found at $\Delta m_{32}^2 = -2.46 \times 10 - 3 \text{ eV}^2$, $\sin^2(\theta_{32}) = 0.404$.



Figure 4.16: 68% and 90% confidence limits in $\Delta m_{32}^2 - \sin^2(\theta_{23})$ for the MINOS, MINOS+ combined analysis are shown along with 90% confidence limits for the MINOS-only analysis and the 2014 T2K ν_{μ} disappearance analysis [29].



Figure 4.17: Likelihood profiles over Δm_{32}^2 (left) and $\sin^2(\theta_{23})$ (right) are shown for the combined MINOS and MINOS+ analysis. Separate profiles are presented assuming normal and inverted mass hierarchy.

Chapter 5

Measuring Neutrino Speed at MINOS

5.1 Introduction

Neutrinos are known to have very little mass. Given the laws of special relativity, even neutrinos at the lowest energies detectable in current experiments should have speed indistinguishable from c, the speed of light in vacuum. Relativity states that no object may exceed this speed in any frame of reference.

By far the most stringent experimental constraint on the speed of neutrinos was obtained from observations of supernova SN1987A [120–122]. A total of 24 low-energy (approximately 10 MeV) electron-antineutrinos were observed across three experiments, Kamiokande II, IMB and Baksan. This observation led to limits on various neutrino properties, including a velocity constraint of $|(v - c)|/c < 2 \times 10^{-9}$. MINOS published its first measurement of the speed of ν_{μ} in the NuMI beam in 2007 [123], reporting $(v - c)/c = (5.1 \pm 2.9) \times 10^{-5}$.

In 2011, the OPERA collaboration measured

$$(v-c)/c = (2.37 \pm 0.32 \,(\text{stat.}) \pm 0.29 \,(\text{sys.})) \times 10^{-5},$$
 (5.1)

using a ν_{μ} beam with $\langle E_{\nu} \rangle \approx 17 \,\text{GeV}$ across the approximately 730 km baseline from CERN to Gran Sasso (CNGS) [124]. This result, now known as the *faster-than-light* neutrino anomaly, excludes $v \leq c$ at the 6.2σ confidence level.

Soon after the OPERA measurement was made public, a number of long-baseline neutrino experiments began work on independently reproducing it. OPERA announced two potential unnoticed sources of error in February 2012, the largest being a loose GPS cable. In summer 2012, Icarus [125], Borexino [126], LVD [127] all announced new results from the CERN to Gran Sasso (CNGS) beamline, consistent with neutrinos travelling at the speed of light. OPERA published a second paper correcting the error in their 2011 result and reporting v consistent with c [128]. MINOS performed two analyses by the end of 2012, which form the subject of this chapter. These analyses are distinguished from the CNGS measurements by the use of a different beam in a distinct energy region, $\langle E_{\nu} \rangle \approx 3$ GeV, as well as the use of two neutrino detectors along the beamline.

The first 2012 MINOS analysis was designed as a repeat of the 2007 analysis with added data and re-evaluation of several systematic uncertainties. It is referred to here as the Retrospective Time of Flight (RTOF) analysis, and described in full in Section 5.3.

The second measurement, named here as the High-Resolution Time of Flight analysis (HRTOF), is described in Section 5.4. This analysis utilises several new pieces of timing equipment, using the last two months of MINOS data taken after their installation. This analysis results in the most precise measurement of the flight-time of a neutrino beam ever made, and the second-most-precise measurement of the speed of a neutrino beam, and is published in [2].

5.2 MINOS Time of Flight Analyses

The measurement of neutrino speed is conceptually simple, requiring a measurement of the time taken for neutrinos in the beam to travel between two fixed points (the Near and Far Detector) – the neutrino *time of flight* – and of the distance between those points. The distance measurement follows standard geodetic techniques, complicated by the location of the detectors beneath the ground. Measurement of the time of flight requires precise knowledge of the neutrino creation time and the interaction times in each detector, whose clocks must be very accurately synchronised over a large distance. Systematic timing offsets present another challenge to the analyses. These can arise in various places in the timing system and, if left unmeasured, can lead to systematic shifts on the result of magnitude far greater than its statistical uncertainty. The RTOF analysis was designed primarily to better characterise the delays of the old timing system, reducing the systematic uncertainty on the result. The HRTOF analysis improves on this with constant monitoring of several significant delays that may change with time.

5.2.1 Timing in the NuMI Beam

The creation time of individual neutrinos cannot be measured, and can only be known definitively up to the approximately $10 \,\mu$ s width of the beam spill. We can however construct PDFs of the creation time based on the known time-structure of the beam in order to reduce the level of ambiguity. The time-structure of the beam is discussed in the following paragraphs.



Figure 5.1: Illustration of the main time structure of a NuMI beam spill, showing Booster batches and RF bunches. Taken from [2].

During MINOS running, the MI was extracted once every 2.2s to NuMI. The acceleration process described in Section 2.3.1 leads to spills consisting of 6 Booster batches. During Tevatron running, one of these batches would be sent to the main ring for collisions, otherwise every batch is extracted to the NuMI target. Each batch lasts $1.58 \,\mu$ s and consists of 81 *bunches*. This finer structure is caused by protons collecting around the same point in the phase of the MI's RF field, whose 53.1 MHz

accelerating frequency leads to an 18.83 ns period in the bunch structure. The width of each bunch is approximately 0.8 ns. Figure 5.1 illustrates the main features of the beam. In the 2007 and RTOF analyses, the event time resolution at the Near and Far Detector was enough to resolve the Booster batch structure. In the HRTOF analysis, the RF bunch structure can be resolved, facilitating an order-of-magnitude reduction in the statistical error of the experiment.

The extraction of protons from the Main Injector to NuMI uses a kicker magnet. A 'kicker fire signal' derived from the current firing this magnet is sent to the ND to be used as a trigger. This signal arrives several μ s before the beam, with a near-constant delay and as such is used as the reference for the ND and FD event times in the 2007 and RTOF analyses.

5.2.2 Timing in the MINOS Detectors

The Near and Far Detector are connected to identical, synchronised GPS receivers, with an estimated random error on the synchronisation of 100 ns. At the Near Detector, the GPS timestamps the arrival of the kicker fire signal, which triggers ND data taking using its own internal clock. PMT hits are measured relative to the kicker fire signal using a 53.1 MHz voltage-controlled crystal oscillator.

The ND's kicker fire timestamp is sent to the Far Detector over the internet to be used to tag data in the beam spill window [14]. At the Far Detector, the GPS receiver is directly used for event timing. A 40 MHz reference is distributed from the receiver to the FD electronics and used to timestamp PMT hits.

In addition to this, a system called the 'timing fiducial' is in place to measure the latency of the FD electronics. Every second, a signal derived from the PPS is sent into the FD electronics, mimicking a PMT signal. The time assigned to this fake PMT hit gives the response time of the detector electronics and is subtracted from the event time for that second.

At each detector, the event time relative to the internal clock is reconstructed using the Kalman-filter-based track fitter mentioned in Section 2.8.2. The estimated resolution on the event time is calculated using Monte Carlo simulations to be approximately 1.5 ns at each of the detectors. The Far Detector resolution is diminished by the synchronisation error when measuring times relative to the kicker fire signal.

5.3 Retrospective Time of Flight Analysis

The first of the two MINOS time of flight analyses performed in 2012 was made using the original Near and Far Detector timing systems, as described above. Event times are recorded at each detector relative to the kicker fire signal, which is assumed to arrive a constant time before the beam. The distribution of event times at the ND is used to generate predictions of the distribution of event times at the FD as a function of the overall neutrino time of flight.

The synchronisation of the Near and Far Detector is subject to several sources of systematic uncertainty, and this dominates the 2007 MINOS analysis. The primary aim of this analysis was to reduce some of the largest of these uncertainties, which are listed in Table 5.1, with improved measurements of cable and electronics delays. Some of these improved measurements are then reused in the HRTOF analysis described in Section 5.4.

Description	Uncertainty $(\pm 1\sigma)$
Distance between detectors	2.3 ns
ND antenna cable delay	$27\mathrm{ns}$
ND electronics latency	$32\mathrm{ns}$
FD antenna cable delay	$46\mathrm{ns}$
FD electronics latency	$3\mathrm{ns}$
GPS and transceivers	$12\mathrm{ns}$
Detector readout differences	$9\mathrm{ns}$
Total (sum in quadrature)	64 ns

Table 5.1: Table of systematic uncertainties for 2007 time of flight analysis.

5.3.1 Antenna Cable Delays

The ND and FD GPS receivers are connected to antennae on the surface by long cables. Several instruments were used to measure the average delay along each cable for the 2007 analysis, with the systematic error taken as the RMS of the different measurements. The Far Detector antenna cable delay was the largest single uncertainty on the 2007 TOF measurement. Two independent measurements of the FD antenna cable delay were made, using a Vector Network Analyser and oscilloscope respectively and improving the uncertainty to 1 ns [129]. A similar measurement at the Near Detector yields an uncertainty of 0.1 ns [130].

5.3.2 Detector Latencies

The next largest contribution to the systematic uncertainty from 2007 was the measurement of the relative Near and Far Detector electronics latency. The time taken for activity in the scintillator of the detector to be travel through the detector electronics may be different in the ND and FD, and contributes an offset to the relative time of flight. At the Far Detector, the timing fiducial system reduces this uncertainty by measuring the delay between a PMT hit and the DAQ timestamp, but some uncertainty remains in the scintillator response and WLS fibre delays.

The 2007 uncertainty was based on individual test stand measurements of the various components of the detectors. A more integrated approach is taken for this analysis, with the use of identical *auxiliary detectors* (ADs), one at each detector. The relative detector-AD latency is measured for each the ND and FD, with the difference providing an estimate for the ND-FD latency with very little systematic error.

The Auxiliary Detectors

The ADs are simple counters, each consisting of two 1 m^2 orthogonal scintillator planes, made from spare MINOS detector planes. The scintillator is read out to PMTs attached to a coincidence module; the detector registers an event if the two planes see activity at the same time. Coincidences in the ADs are timed by two independent clocks. Each clock is synchronised by an external, stable PPS, while a high-resolution internal clock is used to timestamp each coincidence.

Each AD is positioned against back face of the detector, such that muons from beam neutrino events in the detector may pass through the AD. Coincidences in the AD are matched to events in the main detectors by using the transverse positions of tracks that exit the detector – which must match the location of the AD – and by requiring that the event times measured by the ND and the AD match within $1 \mu s$.

A distribution of the difference between AD and detector timestamps can then be built up for matched events. A correction is made to the detector timestamps, since smaller time differences will be recorded for events occuring at the back of the detector, near the AD. The detector times are extrapolated to the front face of the detector using the event vertex position and assuming speed of light neutrinos. Finally, the relative AD-detector latency is taken as the mean of this time-difference distribution.

The relative FD-ND latency is measured to be

$$L_{\rm FD} - L_{\rm ND} = (-27.8 \pm 0.4) \,\mathrm{ns.}$$
 (5.2)

5.3.3 Selection

The 2010 beam ν_{μ} disappearance analysis selection was used to select CC- ν_{μ} events. The criteria are similar to those described in Section 3.1.2. A description can be found in [131]. A total of 5584 events were selected at the Far Detector.

5.3.4 Fitting

Maximum-likelihood and likelihood-ratio tests are again used in this analysis. Far Detector events are taken to be distributed in time according to a probability-density function (PDF) generated from the ND data with two model parameters, τ , the neutrino time of flight and σ_t , the Far Detector timing resolution. The data and predictions are binned in time, with bin width 1.88 ns. Pearson's chi-squared test is used to compare hypotheses:

$$\chi^2 = \sum_{i}^{N} \frac{(O_i - E_i(\tau, \sigma_t))^2}{E_i},$$
(5.3)

where O_i is the observed number of events in time-bin *i* and $E_i(\tau, \sigma_t)$ is the expected number given the model parameters. The predictions $E_i(\tau, \sigma_t)$ are generated by scaling the ND data, shifting by the time of flight τ and convolving with a Gaussian of width σ_t .

The observed and predicted event times are calculated relative to the kicker fire signal, i.e. they are effectively binned in event time modulo the 2.2s between beam spills. This increases the sensitivity of a chi-squared fit, with the possible drawback that the time of flight itself can only be definitively measured modulo 2.2s. It is also possible to bin the plots in time modulo the duration of a batch, increasing the effective sample size fivefold, to estimate the time of flight modulo approximately 1.58 μ s. This is referred to as the 'wrapped batch' fit, and is performed along with the standard 'full spill' fit.



Figure 5.2: Comparison of the "wrapped batch" and "full spill" fit results for each of 27 data subsets. The results vary over time by more than expected from the statistical error on the fit, due to a problem with the ND-FD GPS connection.

The data was divided into 27 sets according to date. Figure 5.2 shows the results of each of the two fits for each data set. This shows a few things: that the two approaches are consistent and that the wrapped batch approach does indeed have the smaller statistical uncertainty. However, it is clear that the spread of the points is too large compared to their error bars. This was discovered to be caused by a problem with the GPS, left unnoticed in the 2007 analysis.

5.3.5 GPS Reset Instability

The Near and Far Detector GPS 10 MHz references were observed to pick up a random phase-change with respect to the PPS every time the GPS unit was reset, causing otherwise unexplained jumps in the measured TOF between resets. Tests on spare GPS receivers showed the jumps to be normally distributed with a standard deviation of 60 ns.

5.3.6 Result

The GPS reset problem was incorporated into the statistical uncertainty of the result by taking the mean of the 27 data sets. Note that this then became the dominant statistical uncertainty, negating the improved statistical power of the wrapped batch method. The results for the two methods are shown below; the data are consistent with v = c.

Full spill:

$$\Delta t = 18 \pm 11 \,(\text{stat.}) \pm 29 \,(\text{syst.}) \,\text{ns} \tag{5.4}$$

Wrapped batch:

$$\Delta t = 11 \pm 11 \,(\text{stat.}) \pm 29 \,(\text{syst.}) \,\text{ns}$$
 (5.5)

5.4 High-Resolution TOF Analysis

The installation of new timing equipment allows the HRTOF analysis to drive down both the statistical and systematic uncertainty of the RTOF analysis. Highly stable, precise atomic clocks are installed at the Near and Far Detector to calibrate their internal timing systems, as well as new synchronisation systems, bringing the absolute ND and FD event time resolution to $\mathcal{O}(1 \text{ ns})$ compared to $\mathcal{O}(100 \text{ ns})$ for the FD in previous analyses. This provides sensitivity to the fine time-structure of the beam, greatly improving the statistical precision of the analysis. A resistive-current wall monitor (RWCM) [132] is installed, providing detailed imaging of the beam. The neutrino creation-time distribution can now be measured precisely for each spill instead of relying on models of the beam. The time of flight of the beam is measured separately between the RWCM and each the Near and Far Detector. Taking the difference between the two measurements gives the ND-FD neutrino flight-time.

Systematic uncertainty is minimised by using measurements from the RTOF analysis, as well as the installation of interval timers to continuously measure various delays in the system. Redundancy is employed in the ND-FD link to provide estimates of the size of systematic effects in the synchronisation.

5.4.1 Resistive Wall Current Monitor

The RWCM measures the current passing through the extraction line through its effect on the beam pipe. Protons travelling inside the pipe produce an image current, equal and opposite to the proton current, along the pipe. A section of the pipe is cut and replaced with a ceramic resistive gap. The voltage across this gap is measured with a waveform digitiser with analogue bandwidth 1.5 GHz.

The arrival of the kicker fire signal at the RWCM triggers approximately $16 \,\mu s$ of continuous data taking, in which the output of the RWCM is digitised by an oscilloscope at 2.5 GHz. This provides resolution of 0.4 ns. The oscilloscope is synchronised by a 10 MHz signal from a local rubidium clock, and the time of the trigger with respect to this clock is recorded by an interval timer.

Combining these measurements allows reconstruction of a 0.4 ns-resolution time profile of the beam current for each spill with respect to the local rubidium clock, as well as the time of the kicker fire signal.

A demonstration of the utility of the RWCM is shown in Figure 5.3. Here, the time of the first proton bunch relative to the kicker fire is shown, testing the assumption from the 2007 and RTOF analyses that this difference is constant. This can be seen to be false – there is a discrete random distribution of start times in buckets of 18.83 ns. This effect will have contributed to the statistical error in previous analyses. It is however eliminated from the HRTOF analysis by using the RWCM to make this per-spill measurement of the beam.



Figure 5.3: Distribution of the time of the first proton bunch relative to the kicker fire signal, measured by the RWCM.



Figure 5.4: Diagram of the main components of the RWCM-ND-FD timing system. Original timing equipment is shown below the detectors, with equipment installed for this analysis shown above. Connection points (marked "C"), show points where equipment is connected temporarily, e.g. travelling GPS units and the TWSTT equipment. Taken from [2].

5.4.2 Synchronisation

Figure 5.4 illustrates the main features of the new timing system. The original MINOS ND and FD timing systems are described in Section 5.2.2, and are retained for this analysis. New timing equipment supplements the existing timing system at each of the detectors, allowing precise characterisation and offline correction of times

recorded by these systems without interfering with data acquisition. Caesium clocks are installed at the Near and Far Detector. At the Near Detector, the time of the kicker fire trigger is measured relative to this clock by an interval timer, and the ND's internal clock is used to record the event time relative to this signal. At the Far Detector, each PPS from the FD GPS is measured relative to the Cs PPS using an interval timer.

In order to measure the RWCM-ND and RWCM-FD time of flight, the three reference timers (the Rb clock at RWCM and Cs clocks at the ND and FD) need to be synchronised. In the previous MINOS analyses, the GPS system used for FD event timing was used to synchronise it to the ND. This synchronisation was the limiting factor on the resolution of FD event times, which were measured relative to the arrival of the kicker fire signal at the ND, with an estimated random synchronisation error of 100 ns from second to second and a further approximately 20 ns reset error.

Synchronisation in MINOS is now achieved using a multi-step process, involving several independent methods using equipment newly installed for this analysis. Constant monitoring of the difference between each of the time references in the system allows a stability of less than 1 ns to be achieved in the synchronisation between the three detection points. This monitoring is performed with two-way time transfer (TWTT) over optical fibres, two-way satellite time transfer (TWSTT) and the Global Positioning System (GPS), with some measurements made using more than one technique to check for consistency and quantify systematic error.

Two-Way Time Transfer

Time transfer is a simple technique of synchronising two clocks. In one-way timetransfer, the time of one clock is sent along some communication line to an observer. If the observer takes the difference between the time on his clock to the received time, allowing synchronisation. The problem with this is that any delay caused by the communication is unaccounted for. If the time at two clocks is t_A and t_B respectively, and the communication delay is C, the time difference measured is

$$\Delta t = t_A - t_B + C. \tag{5.6}$$

Two-way time-transfer solves this problem by transferring time in both directions along the line, such that

$$\Delta t_1 = t_A - t_B + C, \qquad \Delta t_2 = t_B - t_A + C, \tag{5.7}$$

allowing computation of the communication delay,

$$\frac{\Delta t_1 + \Delta t_2}{2} = C,\tag{5.8}$$

and hence absolute synchronisation of the two clocks once their measurements are shared.

Fibre-Based TWTT

Reference clocks at the Near and Far Detector are synchronised to the surface using TWTT over optical fibres. TWTT between the RWCM and the ND surface reference also allows synchronisation of the ND and RWCM clocks. This fibre-based TWTT works as specified above, with signals transferred along optical cables. Two cables in the same bunch to are used to ensure identical delays for the two transfers.

Satellite-Based TWTT

One of the methods of synchronising across the long distance between the ND and FD is to used two-way satellite time transfer. This works on the same principal as fibre-based TWTT, using radio-frequency satellite communications for the transfer. US Naval Observatory (USNO) personnel maintained TWSTT links directly between the ND and FD as well as between each the ND and FD and a USNO facility in Washington DC, providing two means of performing the ND-FD synchronisation.

Global Positioning System

MINOS employs six identical GPS receivers for this analysis, with two at each detection site. A further two similar, newer units are used as 'travelling' receivers. The GPS system is described in detail in [133, 134].

At each detector, the difference between the local atomic clock, t_A and the GPS time t_{GPS} is recorded by each receiver once per second, such that

$$\Delta t_{A-GPS} = t_A - t_{GPS}.$$
(5.9)

The GPS time calculations were performed by members of the National Timing and Standards Institute (NIST), using precise point positioning (PPP) algorithms [135]. The use of two receivers at each point allows the estimation of systematic effects in the receivers. The two travelling receivers spend three days at each detector, as well as at the NIST facility in Boulder, Colorado. This allows systematic differences between the receivers at each site to be measured and corrected for.

5.4.3 Validation of the Timing System

The synchronisation system was built with as much redundancy as possible, in order to compare independent measurements and establish a level of systematic uncertainty on the synchronisation. For example the RWCM and ND can be synchronised using either GPS or fibre-based TWTT. Equations 5.10 and 5.11 show the two calculations. From 5.4.2, the calibration using TWTT is

$$\Delta t_{\rm TWTT} = t_{\rm ND} - t_{\rm RWCM} = \frac{\Delta t_{\rm ND-RWCM} - \Delta t_{\rm RWCM-ND}}{2}.$$
 (5.10)

Now from 5.9, the GPS synchronisation is given by

$$\Delta t_{GPS} = t_{ND} - t_{RWCM} = \Delta t_{ND-GPS} - \Delta t_{RWCM-GPS}.$$
(5.11)

Given perfect synchronisation, $\Delta t_{\text{GPS}} = \Delta t_{\text{TWTT}}$. In practise, Δt_{TWTT} is expected to provide the more accurate synchronisation, since there are fewer potential sources of systematic error. The stability and accuracy of the GPS synchronisation can therefore be measured by plotting $\Delta t_{\text{GPS}} - \Delta t_{\text{TWTT}}$. This *double-difference* calculation is performed for each pair of GPS receivers, and between synchronisation methods.

Independent measurements of the ND-FD synchronisation are provided by the GPS system and the temporary ND-FD TWSTT performed for two days in April 2012. Three double differences calculated from four independent measurements are



Figure 5.5: Differences between four measurements of the ND-FD synchronisation are plotted as a function of the Modified Julian Date over the duration the USNO TWSTT test (April 18-19, 2012). The four measurements are direct ND-FD TWSTT (TWdirect); indirect ND-USNO-FD TWSTT (TW); and synchronisation using each pair of ND and FD GPS receivers (GPS₂₋₄ and GPS₈₋₃). Taken from [2].

plotted in Figure 5.5: direct ND-FD TWSTT; indirect ND-USNO-FD TWSTT; and synchronisation using each pair of ND and FD GPS receivers. The four measurements can be seen to be stable to within 800 ps. The distribution of the means is used to calculate the systematic uncertainty on the ND-FD synchronisation of 480 ps.

5.4.4 ND-to-FD Distance Survey

The distance between the front faces of the Near and Far Detector was re-evaluated for this analysis. The work is described in detail in [136]. GPS location was used to measure the position of control marker on the surface at the two detector sites to approximately 1 cmprecision [14]. The distance from the surface reference to the detector is easily calculated for the ND using standard optical survey methods to millimetre accuracy [137].

At the Far Detector, atmospheric effects prevent accurate optical surveys. In addition to this, there is no direct plumb line from the surface to the Soudan Underground Laboratory. Instead, the position of the detector relative to the surface reference is determined using an Inertial Navigation Unit (INS), containing three gyroscopes and three accelerometers. The INS travels between the surface and the lab in the mine's elevator, the vibration of which limits the accuracy of the survey. This measurement in turn limits the overall accuracy of the ND-FD distance measurement, which is calculated to an uncertainty of 70 cm or 2.3 ns. This constitutes the dominant uncertainty in the final measurement of neutrino speed.

5.4.5 Systematic Uncertainties

Systematic uncertainties in the analysis come in the form of direct shifts on the predicted time of flight. Table 5.2 lists the dominant systematic uncertainties on the final measurement. The largest single uncertainty is on the Near-to-Far Detector distance measurement, described above. The 0.5 ns GPS synchronisation uncertainty is a conservative estimate based on the double-differences described in 5.4.3, shown in Figure 5.5. The relative detector electronics latency measurement is described in 5.3.2. Finally, smaller contributions come from the measurements of the various cables connecting components of the timing system, the largest of which is the TWTT link between the FD and its GPS units on the surface.

Description	Uncertainty $(\pm 1\sigma)$
Distance between detectors	$2.3\mathrm{ns}$
ND-FD relative latency	$1.0\mathrm{ns}$
FD-surface TWTT accuracy	$0.6\mathrm{ns}$
GPS synchronisation accuracy	$0.5\mathrm{ns}$
Total (sum in quadrature)	2.6 ns

Table 5.2: Dominant systematic uncertainties in the 2012 high-resolution TOF analysis.

5.4.6 Selection

The same selector is used to select $CC-\nu_{\mu}$ events as in the RTOF analysis. A further cut on the track fit quality is made in order to exclude events with poor

time resolution. At the Far Detector, the sample size is increased by including an additional *rock and anti-fiducial* (RAF) event sample.

RAF events are partially reconstructed interactions whose vertices are located outside the detector's fiducial region or in the surrounding rock. Their selection is described in [131]. These events are characterised by the lack of a reconstructed vertex inside the detector volume, with a muon track entering either through the front face or the sides of the detector. A partial shower may also be observed at the edge of the detector. The lack of a detectable vertex increases the uncertainty on the interaction time of a RAF event. In addition to this, a delay is observed with respect to contained events, caused by the average increase in path-length for muons created off the beam axis.

5.4.7 Fitting

The time of flight is extracted by minimising the following log-likelihood function, with confidence intervals constructed as in the oscillation analysis in the previous chapter, and with the likelihood profiled over a single nuisance parameter, σ_t :

$$-\ln L = \sum_{i}^{N} -\ln P_{i}(t_{i};\tau,\sigma_{t}), \qquad (5.12)$$

where N runs over the number of selected events. P_i is the detection time PDF for event *i*, which is constructed from the RWCM data for each spill, and τ is the neutrino time of flight.

Given the highly-degenerate likelihoods produced from the symmetries of the beam, a grid search is preferred to the gradient-descent method favoured in other analyses. A window of approximately $6 \,\mu s$ around the expected time of flight is used as the search space.

Detection Time PDFs

In order to form likelihoods for a set of events x_i , the distributions of their parent population must be known. For example, in the previously described analyses the data x_i are event counts sampled from a Poisson parent distribution, with its mean calculated using Near to Far Detector extrapolation. In this analysis, x_i are event times measured relative to some reference. For a single neutrino, this could be approximated as sampling from a Gaussian distribution with width corresponding to the event time resolution, σ_t and centred on the overall time of flight, T. This is complicated by lack of knowledge about the exact neutrino creation time. Using the RWCM, a creation time distribution can be constructed for each beam spill and hence each neutrino.



Figure 5.6: A single spill (top-left) as seen by the Resistive Wall Current Monitor, with a close-up of the first few bunches (bottom-left). The final PDF generated from this spill is also shown (top-right and bottom-right). Note that the last three batches of this spill are slip-stacked: the unsmeared measurement resolves two peaks for each slip-stacked bunch, while the smeared PDF sees one peak of double intensity.

The left half of Figure 5.6 shows the time-current profile of the beam for a single spill. The beam current is proportional to the proton flux in the beam pipe and therefore to the resulting neutrino flux, allowing us to estimate a neutrino creation-time PDF relatively simply. Systematic effects in the baseline current are removed, as well as low-level noise, and the distribution is normalised to unity. The resulting PDF is shown in the right half of Figure 5.6. For an event i, we name this PDF $P_{B,i}(t)$.

The final, detection-time PDF is calculated by shifting the creation-time PDF by the overall time of flight, τ , and including the effects of timing resolution by

convolving with respect to time with a Gaussian of width σ_t , $\mathcal{N}(t, \sigma_t^2)$, as in previous analyses. This PDF is given by

$$P_i(t_i;\tau,\sigma_t) = P_{B,i}(t-\tau) * \mathcal{N}(t,\sigma_t^2).$$
(5.13)

For a RAF event, the additional delay and smearing on the event time must be accounted for. Monte Carlo simulation was used to estimate the distribution of delays on RAF events compared to their contained counterparts. This distribution is shown in Figure 5.7. Normalising this to one and convolving with the the PDF in Equation 5.13 gives the detection time PDF for a RAF event.



Figure 5.7: Distribution of the difference between reconstructed and true RAF neutrino interaction times [30].

5.4.8 Results

RWCM-ND

Final synchronisation between the three detector sites is performed using the GPS data, such as to minimise systematic effects in the synchronisation process when taking the difference between the two measurements. The other synchronisation techniques can be used to measure the relative stability of the GPS system. At the ND, the fibre-based TWTT system should provide more stable synchronisation than the GPS.



Figure 5.8: Distribution of reconstructed Near Detector event time with respect to the start of a beam spill, modulo the bunch length (18.83 ns). The ND event time resolution is demonstrated by the clear emergence of this peak.



Figure 5.9: Best fit time of flight, along with 1σ uncertainty, for one-day subsets of the HRTOF dataset. The fit is performed both using two-way time-transfer and GPS for ND-FD synchronisation.

We select 1.67×10^4 events at the Near Detector. We can test the resolution of the RWCM-ND timing system before performing the analysis fit by plotting the ND event time modulo 18.83 ns, as shown in Figure 5.8. Fibre-based TWTT is used for the RWCM-ND synchronisation in this plot. The clear appearance of the RF bunch structure in this plot demonstrates the high level of resolution achieved: the width of the bunch is approximately 1.6 ns, indicating a total resolution in the ND-RWCM system of 1.4 ns. Given the large number of events selected at the Near detector, we can divide the data by date in order to measure the stability of the ND TOF. The data are divided into runs each covering a single day, and the fit is performed for each run individually. This is done for each the GPS and TWTT synchronisation procedures and the results are plotted in Figure 5.9.

The TWTT-based measurement displays the greater stability, with RMS approximately 50 ps. However, the GPS measurement is also stable to approximately 200 ps. The difference between the means of the two measurements is as little as 300 ps. The GPS-based measurement is used in the final result, and gives a RWCM-ND time of flight of

$$4\,621.10 \pm 0.05\,\mathrm{ns}\,(\mathrm{stat.\ only})$$
 (5.14)

compared to an expected TOF, given the surveyed RWCM-ND distance, of

$$4\,622.7 \pm 4.0\,\mathrm{ns.}$$
 (5.15)

RWCM-FD

We select 193 CC and 177 RAF events at the Far Detector in this analysis. Figure 5.10 shows the equivalent of Figure 5.8 for the Far Detector, the event time relative to the kicker fire modulo 18.83 ns. In this plot, TWSTT was used for the synchronisation. The slight delay and smearing of the RAF distribution due to the increased path length can be seen.

Figure 5.11 shows the profile RWCM-FD TOF likelihood over the entire search space, with a closer view of the best fit point and surrounding minima shown in Figure 5.12. From this the RWCM-FD time of flight can be seen to be

$$2\,453\,935.0\pm0.1\,\mathrm{ns}\,(\mathrm{stat.\ only}).$$
 (5.16)



Figure 5.10: Distribution of RWCM-FD event time modulo 18.83 ns. The RAF distribution is shown uncorrected, and hence has a slight delay and smearing with respect to the CC distribution. The clear appearance of this structure demonstrates the resolution achieved on the event time and RWCM-FD synchronisation, which was performed using TWSTT for this plot.



Figure 5.11: Likelihood profile in the RWCM-FD time of flight over a range of $\pm 3 \mu s$ from the v = c prediction.



Figure 5.12: Likelihood profile in the RWCM-FD time of flight, showing the best fit point at $2\,453\,935.0\pm0.1$ ns and surrounding local minima. The early and late minima are excluded at 8.15σ and 4.5σ respectively.

Neutrino Speed

Subtracting the ND and FD results presented gives the ND-FD neutrino time of flight,

$$\Delta t_{\rm ND-FD} = 2\,449\,313.9 \pm 0.1\,(\text{stat.})\,1.2\,\text{ns}\,(\text{syst.}). \tag{5.17}$$

This is the most precise measurement of the flight-time of a neutrino beam ever made.

The predicted time of flight given the surveyed ND-FD distance is

$$\Delta T_{\rm ND-FD} = 2\,449\,316.3 \pm 2.3\rm{ns.} \tag{5.18}$$

The difference between the observed and predicted flight-time is therefore

$$\delta = \Delta T_{\rm ND-FD} - \Delta t_{\rm ND-FD} = 2.4 \pm 0.1 \,(\text{stat.}) \pm 2.6 \,(\text{syst.}) \text{ns}, \tag{5.19}$$

leading to a neutrino speed, expressed as the fractional difference from the speed of light, of

$$(v-c)/c = \frac{\Delta z_{\rm ND-FD}/c}{\Delta t_{\rm ND-FD}} - 1 = (1.0 \pm 1.1) \times 10^{-6},$$
 (5.20)

consistent with v = c.

Chapter 6

Summary and Conclusion

MINOS has produced the tightest constraints on Δm_{32}^2 for many years. MINOS+, despite its reduction in sensitivity to three-flavour oscillations, continues to improve this measurement. This thesis documents the first neutrino oscillations analysis of data from MINOS+, resulting in a world-leading measurement of Δm_{32}^2 . The precise measurement of the atmospheric oscillation parameters can help provide constraints for other experiments in the field. In addition, a world-leading measurement of the time of flight of a neutrino beam is described. This analysis is published in [2]. In addition to these analyses, my work has included contributions to the MINOS+ reconstruction software and a MINOS sterile neutrino analysis.

Neutrino physics remains an exciting field. Daya Bay has made very precise measurements of a large θ_{13} , paving the way for the potential discovery of CP violation in the neutrino sector. T2K and NO ν A are poised to begin making real progress in this area, as well as in the determination of the neutrino mass hierarchy. Further afield, there are still several orders of magnitude to search for the absolute neutrino mass scale: KATRIN will place new limits on the ν_e mass in the next few years. Majorana neutrinos, favoured by many beyond-the-Standard-Model theories, remain elusive: SuperNEMO will hope to catch sight of one when it begins data taking this year. Finally, sterile neutrinos refuse to go away. MINOS+'s L/E configuration, though non-optimal for traditional oscillation physics, gives the experiment unprecedented sensitivity to large regions of the sterile neutrino parameter space and work is currently underway on the first MINOS+ sterile neutrino analysis.

Neutrinos have proven themselves to be unorthodox characters. I hope that they can surprise us further.

Bibliography

- MINOS+, MINOS Collaboration, A. B. Sousa, First MINOS+ Data and New Results from MINOS, AIP Conf. Proc. 1666 (2015) 110004, arXiv:1502.07715 [hep-ex].
- [2] MINOS Collaboration, P. Adamson et al., Precision measurement of the speed of propagation of neutrinos using the MINOS detectors, Phys. Rev. D92 (2015) no. 5, 052005, arXiv:1507.04328 [hep-ex].
- [3] SLD Electroweak Group, DELPHI, ALEPH, SLD, SLD Heavy Flavour Group, OPAL, LEP Electroweak Working Group, L3 Collaboration, S. Schael et al., *Precision electroweak measurements on the Z resonance*, Phys. Rept. 427 (2006) 257-454, arXiv:hep-ex/0509008 [hep-ex].
- [4] J. Hewett, H. Weerts, R. Brock, J. Butler, B. Casey, et al., Fundamental Physics at the Intensity Frontier, arXiv:1205.2671 [hep-ex].
- [5] J. N. Bahcall, A. M. Serenelli, and S. Basu, New solar opacities, abundances, helioseismology, and neutrino fluxes, Astrophys.J. 621 (2005) L85-L88, arXiv:astro-ph/0412440 [astro-ph].
- [6] KamLAND Collaboration, S. Abe et al., Precision Measurement of Neutrino Oscillation Parameters with KamLAND, Phys. Rev. Lett. 100 (2008) 221803, arXiv:0801.4589 [hep-ex].
- [7] Daya Bay Collaboration, C. Zhang, Recent Results from the Daya Bay Experiment, AIP Conf. Proc. 1666 (2015) 080003, arXiv:1501.04991 [hep-ex].
- [8] J. Bian, The NOvA Experiment: Overview and Status, in Meeting of the APS Division of Particles and Fields (DPF 2013) Santa Cruz, California, USA, August 13-17, 2013. 2013. arXiv:1309.7898 [physics.ins-det].
https://inspirehep.net/record/1256026/files/arXiv:1309.7898.pdf.

- [9] U. H. M. Group, Geographical Overview of the MINOS Experiment, http://www.hep.ucl.ac.uk/minos/images/minosmap.jpg.
- [10] R. L. Dixon, Farewell to the Tevatron, CERN Cour. 51N9 (2011) 28-31, arXiv:1209.6229 [physics.acc-ph].
- [11] MINOS Collaboration, J. Evans, The MINOS experiment: results and prospects, Adv. High Energy Phys. 2013 (2013) 182537, arXiv:1307.0721 [hep-ex].
- [12] NOvA Collaboration, D. Ayres et al., The NOvA Technical Design Report, tech. rep., 2007.
- [13] J. S. Mitchell, Measuring ν_{μ} Disappearance with the MINOS Experiment. PhD thesis, 2011.
- [14] MINOS Collaboration, D. Michael et al., The magnetized steel and scintillator calorimeters of the MINOS experiment, Nucl. Instr. & Meth. A596 (2008) 190–228.
- [15] M. Dorman, Beam Fit Position Paper, MINOS-doc-7146-v1, 2010.
- [16] L. Whitehead, Near Detector Beam Reconstruction Notes, MINOS-doc-9368-v1, 2012.
- [17] A. Aurisano et al., *Elm reconstruction*, MINOS-doc-10535-v3, 2014.
- [18] P. Adamson, Protons Delivered to MINOS, http://adwww.fnal.gov/userb/pic/capture/www/trish.png.
- [19] A. Radovic, Measuring the Disappearance of Muon Neutrinos with the MINOS Detector. PhD thesis.
- [20] J. J. Evans, Measuring Antineutrino Oscillations with the MINOS Experiment. PhD thesis, 2008.
- [21] C. J. Backhouse, Measuring neutrino oscillation parameters using ν_{μ} disappearance in MINOS. PhD thesis, 2011.
- [22] MINOS Collaboration, P. Adamson et al., New constraints on muon-neutrino

to electron-neutrino transitions in MINOS, Phys. Rev. **D82** (2010) 051102, arXiv:1006.0996 [hep-ex].

- [23] P. Adamson, Protons Delivered to NuMI, http://adwww.fnal.gov/userb/pic/capture/www/allbig.png.
- [24] A. Timmons, MINOS+ DQM Plots, MINOS-doc-10595-v1, 2014.
- [25] A. Blake, MINOS+ Standard Oscillation Blessing Package (Run 11), MINOS-doc-11102-v1, 2015.
- [26] J. Huang, Sterile Neutrino Searches in MINOS/MINOS+ Experiment. PhD thesis, U. Texas, Austin (main), 2015. http: //lss.fnal.gov/archive/thesis/2000/fermilab-thesis-2015-06.pdf.
- [27] MINOS Collaboration, P. Adamson et al., Combined analysis of ν_μ disappearance and ν_μ → ν_e appearance in MINOS using accelerator and atmospheric neutrinos, Phys.Rev.Lett. **112** (2014) 191801, arXiv:1403.0867 [hep-ex].
- [28] MINOS Collaboration, P. Adamson et al., *Electron neutrino and antineutrino appearance in the full MINOS data sample*, Phys. Rev. Lett. **110** (2013) no. 17, 171801, arXiv:1301.4581 [hep-ex].
- [29] T2K Collaboration, K. Abe et al., Precise Measurement of the Neutrino Mixing Parameter θ₂₃ from Muon Neutrino Disappearance in an Off-Axis Beam, Phys. Rev. Lett. **112** (2014) no. 18, 181801, arXiv:1403.1532
 [hep-ex].
- [30] J. O'Connor, Far Detector TOF combination, MINOS-doc-9531-v1, 2012.
- [31] F. Capozzi, E. Lisi, and A. Marrone, Neutrino mass hierarchy and electron neutrino oscillation parameters with one hundred thousand reactor events, Phys.Rev. D89 (2014) no. 1, 013001, arXiv:1309.1638 [hep-ph].
- [32] M. Gonzalez-Garcia, M. Maltoni, J. Salvado, and T. Schwetz, Global fit to three neutrino mixing: critical look at present precision, JHEP 1212 (2012) 123, arXiv:1209.3023 [hep-ph].
- [33] M. Gonzalez-Garcia, M. Maltoni, and T. Schwetz, Updated fit to three neutrino mixing: status of leptonic CP violation, Journal of High Energy Physics 2014

(2014) no. 11, . http://dx.doi.org/10.1007/JHEP11%282014%29052.

- [34] W. Pauli, Dear radioactive ladies and gentlemen, Phys. Today **31N9** (1978) 27.
- [35] J. Chadwick, Possible Existence of a Neutron, Nature 129 (Feb., 1932) 312.
- [36] E. Fermi, An attempt of a theory of beta radiation. 1., Z.Phys. 88 (1934) 161–177.
- [37] H. Bethe and R. Peierls, *The "neutrino"*, Nature **133** (1934) 532.
- [38] C. Cowan, F. Reines, F. Harrison, H. Kruse, and A. McGuire, Detection of the free neutrino: A Confirmation, Science 124 (1956) 103–104.
- [39] S. H. Neddermeyer and C. D. Anderson, Note on the Nature of Cosmic-Ray Particles, Phys. Rev. 51 (May, 1937) 884-886. http://link.aps.org/doi/10.1103/PhysRev.51.884.
- [40] G. Danby, J. Gaillard, K. A. Goulianos, L. Lederman, N. B. Mistry, et al., Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos, Phys.Rev.Lett. 9 (1962) 36–44.
- [41] DONUT Collaboration, K. Kodama et al., Observation of tau neutrino interactions, Phys.Lett. B504 (2001) 218-224, arXiv:hep-ex/0012035 [hep-ex].
- [42] T. Lee and C.-N. Yang, Question of Parity Conservation in Weak Interactions, Phys.Rev. 104 (1956) 254–258.
- [43] C. Wu, E. Ambler, R. Hayward, D. Hoppes, and R. Hudson, Experimental Test of Parity Conservation in Beta Decay, Phys.Rev. 105 (1957) 1413–1414.
- [44] B. Kayser, F. Gibrat-Debu, and F. Perrier, *The Physics of massive neutrinos*. World Scientific, 1989.
- [45] Planck Collaboration, P. Ade et al., Planck 2015 results. XIII. Cosmological parameters, arXiv:1502.01589 [astro-ph.CO].
- [46] Troitsk Collaboration, V. Aseev et al., An upper limit on electron antineutrino mass from Troitsk experiment, Phys.Rev. D84 (2011) 112003, arXiv:1108.5034 [hep-ex].

- [47] KATRIN Collaboration, A. Osipowicz et al., KATRIN: A Next generation tritium beta decay experiment with sub-eV sensitivity for the electron neutrino mass. Letter of intent, arXiv:hep-ex/0109033 [hep-ex].
- [48] P. Guzowski, A combined limit for neutrinoless double-beta decay, arXiv:1504.08285 [hep-ex]. Prospects in Neutrino Physics Conference, 15 -17 December, 2014, Queen Mary University of London, UK.
- [49] Z. Maki, M. Nakagawa, and S. Sakata, Remarks on the Unified Model of Elementary Particles, Progress of Theoretical Physics 28 (Nov., 1962) 870–880.
- [50] J. Davis, Raymond, D. S. Harmer, and K. C. Hoffman, Search for neutrinos from the sun, Phys.Rev.Lett. 20 (1968) 1205–1209.
- [51] A. Abazov, O. Anosov, E. Faizov, V. Gavrin, A. Kalikhov, et al., Search for neutrinos from sun using the reaction Ga-71 (electron-neutrino e-) Ge-71, Phys.Rev.Lett. 67 (1991) 3332–3335.
- [52] GALLEX Collaboration, P. Anselmann et al., Solar neutrinos observed by GALLEX at Gran Sasso., Phys.Lett. B285 (1992) 376–389.
- [53] SNO Collaboration, Q. Ahmad et al., Measurement of the rate of ν_e + d → p + p + e⁻ interactions produced by ⁸B solar neutrinos at the Sudbury Neutrino Observatory, Phys.Rev.Lett. 87 (2001) 071301, arXiv:nucl-ex/0106015 [nucl-ex].
- [54] Z. Maki, M. Nakagawa, and S. Sakata, *Remarks on the unified model of elementary particles*, Prog.Theor.Phys. 28 (1962) 870–880.
- [55] S. M. Bilenky and B. Pontecorvo, Quark-Lepton Analogy and Neutrino Oscillations, Phys.Lett. B61 (1976) 248.
- [56] S. M. Bilenky and B. Pontecorvo, Again on Neutrino Oscillations, Lett.Nuovo Cim. 17 (1976) 569.
- [57] S. M. Bilenky and B. Pontecorvo, Lepton Mixing and Neutrino Oscillations, Phys.Rept. 41 (1978) 225–261.
- [58] KamLAND Collaboration, S. Abe et al., Precision Measurement of Neutrino Oscillation Parameters with KamLAND, Phys.Rev.Lett. 100 (2008) 221803,

arXiv:0801.4589 [hep-ex].

- [59] Kamiokande-II Collaboration, K. Hirata et al., Experimental Study of the Atmospheric Neutrino Flux, Phys.Lett. B205 (1988) 416.
- [60] R. Becker-Szendy, C. Bratton, D. Casper, S. Dye, W. Gajewski, et al., A Search for muon-neutrino oscillations with the IMB detector, Phys.Rev.Lett. 69 (1992) 1010–1013.
- [61] W. Allison, G. Alner, D. Ayres, W. Barrett, C. Bode, et al., Measurement of the atmospheric neutrino flavor composition in Soudan-2, Phys.Lett. B391 (1997) 491-500, arXiv:hep-ex/9611007 [hep-ex].
- [62] Super-Kamiokande Collaboration, Y. Fukuda et al., The Super-Kamiokande detector, Nucl.Instrum.Meth. A501 (2003) 418–462.
- [63] Super-Kamiokande Collaboration, Y. Fukuda et al., Evidence for oscillation of atmospheric neutrinos, Phys.Rev.Lett. 81 (1998) 1562-1567, arXiv:hep-ex/9807003 [hep-ex].
- [64] T2K Collaboration, K. Abe et al., The T2K Experiment, Nucl.Instrum.Meth. A659 (2011) 106-135, arXiv:1106.1238 [physics.ins-det].
- [65] MINOS Collaboration, P. Adamson et al., Improved search for muon-neutrino to electron-neutrino oscillations in MINOS, Phys. Rev. Lett. 107 (2011) 181802, arXiv:1108.0015 [hep-ex].
- [66] T2K Collaboration, K. Abe et al., Indication of Electron Neutrino Appearance from an Accelerator-produced Off-axis Muon Neutrino Beam, Phys. Rev. Lett. 107 (2011) 041801, arXiv:1106.2822 [hep-ex].
- [67] Double Chooz Collaboration, Y. Abe et al., Reactor electron antineutrino disappearance in the Double Chooz experiment, Phys. Rev. D86 (2012) 052008, arXiv:1207.6632 [hep-ex].
- [68] V. D. Barger, K. Whisnant, and R. Phillips, CP Violation in Three Neutrino Oscillations, Phys.Rev.Lett. 45 (1980) 2084.
- [69] Particle Data Group Collaboration, K. Olive et al., *Review of Particle Physics*, Chin.Phys. C38 (2014) 090001.

- [70] A. P. Schreckenberger, Electron Neutrino and Antineutrino Appearance in the MINOS Detector. PhD thesis, University of Minnesota, 2013.
- [71] MINOS Collaboration, P. Adamson et al., Active to sterile neutrino mixing limits from neutral-current interactions in MINOS, Phys.Rev.Lett. 107 (2011) 011802, arXiv:1104.3922 [hep-ex].
- [72] MINOS+ Collaboration, G. Tzanankos et al., MINOS+: a Proposal to FNAL to run MINOS with the medium energy NuMI beam, 2011.
- [73] C. D. Curtis, G. M. Lee, C. W. Owen, C. W. Schmidt, and W. M. Smart, LINAC H- BEAM OPERATION AND USES AT FERMILAB., in 8th Particle Accelerator Conference San Francisco, California, March 12-14, 1979, vol. 26, pp. 3760–3762. 1979.
- [74] K. Seiya, B. Chase, J. Dey, P. Joireman, I. Kourbanis, and J. Reid, Slip Stacking, in CARE-HHH-APD Workshop on Finalizing the Roadmap for the Upgrade of the CERN & GSI Accelerator Complex (BEAM'07) Geneva, Switzerland, October 1-5, 2007, pp. 66-70. 2007. http://lss.fnal.gov/cgi-bin/find_paper.pl?conf-07-650.
- B. C. Brown, P. Adamson, D. Capista, W. Chou, I. Kourbanis, D. K. Morris, K. Seiya, G. H. Wu, and M.-J. Yang, *Fermilab main injector: High intensity operation and beam loss control*, Phys. Rev. ST Accel. Beams 16 (Jul, 2013) 071001. http://link.aps.org/doi/10.1103/PhysRevSTAB.16.071001.
- [76] M. Martens, Target and Horn Configuration for SNuMI, NOvA, and MINERvA, FERMILAB-BEAMS-DOC-2392, 2006.
- [77] MINOS Collaboration, P. Adamson et al., Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS, Phys. Rev. Lett. 110 (Jun, 2013) 251801.
 http://link.aps.org/doi/10.1103/PhysRevLett.110.251801.
- [78] K. Anderson, B. Bernstein, D. Boehnlein, K. R. Bourkland, S. Childress, et al., *The NuMI Facility Technical Design Report*, tech. rep., 1998.
- [79] NA49 Collaboration, C. Alt et al., Inclusive production of charged pions in p+C collisions at 158-GeV/c beam momentum, Eur.Phys.J. C49 (2007) 897-917, arXiv:hep-ex/0606028 [hep-ex].

- [80] N. Tagg, A. De Santo, A. Weber, A. Cabrera, P. Miyagawa, et al., Performance of Hamamatsu 64-anode photomultipliers for use with wavelength-shifting optical fibres, Nucl.Instrum.Meth. A539 (2005) 668-678, arXiv:physics/0408055 [physics].
- [81] P. Adamson, T. Alexopoulos, W. Allison, G. Alner, K. Anderson,
 C. Andreopoulos, M. Andrews, R. Andrews, C. Arroyo, S. Avvakumov, et al., *First observations of separated atmospheric* ν μ and ν⁻ μ events in the *MINOS detector*, Physical Review D **73** (2006) no. 7, 072002.
- [82] 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 (2004) 2193–2195.
- [83] B. J. Rebel, Neutrino Induced Muons in the MINOS Far Detector. PhD thesis, Indiana U., 2004. http://lss.fnal.gov/cgi-bin/find_paper.pl?thesis-2004-33.
- [84] MINOS Collaboration, I. Ambats et al., The MINOS Detectors Technical Design Report, tech. rep., 1998.
- [85] A. Belias, G. Crone, E. Harris, C. Howcroft, S. Madani, T. Nicholls, G. Pearce, D. Reyna, N. Tagg, and M. Thomson, *The MINOS data acquisition system*, in *Nuclear Science Symposium Conference Record*, 2003 IEEE, vol. 3, pp. 1663–1667 Vol.3. Oct, 2003.
- [86] P. Adamson, G. Crone, L. Jenner, R. Nichol, R. Saakian, et al., *The MINOS calibration detector*, Nucl.Instrum.Meth. A556 (2006) 119–133.
- [87] F. Ballarini, G. Battistoni, F. Cerutti, A. Empl, A. Fasso, et al., Nuclear models in FLUKA: Present capabilities, open problems and future improvements, AIP Conf.Proc. 769 (2005) 1197–1202.
- [88] M. Campanella, A. Ferrari, P. Sala, and S. Vanini, *First Calorimeter Simulation with the FLUGG Prototype*, tech. rep., 1999.
- [89] R. Brun, F. Carminati, and S. Giani, GEANT Detector Description and Simulation Tool, tech. rep., 1994.
- [90] H. Gallagher, The NEUGEN neutrino event generator, Nucl. Phys. Proc. Suppl.

112 (2002) 188–194.

- [91] M. Medeiros and J. O'Connor, New Near Detector Slicing, MINOS-doc-9781-v8, 2013.
- [92] J. Marshall, Documentation for the New Track Finding Package, MINOS-doc-2104-v1, 2006.
- [93] J. Marshall, A New Implementation of the Kalman Filter, MINOS-doc-1248-v1, 2005.
- [94] R. Fruhwirth, Application of Kalman filtering to track and vertex fitting, Nucl.Instrum.Meth. A262 (1987) 444–450.
- [95] J. Urish, Event Reconstruction in the Near Detector, MINOS-doc-917-v1, 2003.
- [96] C. Smith, The CandSubShowerSR Package, MINOS-doc-2820-v1, 2007.
- [97] M. A. Kordosky, Hadronic interactions in the MINOS detectors. PhD thesis, Texas U., 2004. http://lss.fnal.gov/cgi-bin/find_paper.pl?thesis-2004-34.
- [98] P. L. Vahle, Electromagnetic interactions in the MINOS detectors. PhD thesis, Texas U., 2004. http://lss.fnal.gov/cgi-bin/find_paper.pl?thesis-2004-35.
- [99] J. O'Connor, Shower Reconstruction Position Paper, MINOS-doc-10531-v1, 2014.
- [100] C. Backhouse, Estimating shower energies using a kNN-based method -Position Paper, MINOS-doc-6868-v14, 2010.
- [101] R. Ospanov, A measurement of muon neutrino disappearance with the MINOS detectors and NuMI beam. PhD thesis, 2008.
- [102] J. S. Ratchford, Identifying Muons for Neutrino Oscillation and Cross Section Experiments. PhD thesis, 2012.
- [103] R. Nichol 2016.
- [104] J. Marshall, NC Background Update, MINOS-doc-3313-21, 2007.

- [105] R. Hatcher et al., Range Curvature Task Force Position Paper, MINOS-doc-3134-v2, 2007.
- [106] H. Gallagher, M. Kordosky, and S. Dytman, Shower Energy Scale Uncertainty For the Run I+II CC Analysis, MINOS-doc-4287-v1, 2008.
- [107] M. C. Group, 2009 Calibration Position paper on Runs I-II-III, MINOS-doc-6717-v2, 2009.
- [108] MINOS Collaboration, P. Adamson et al., Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS, Phys. Rev. Lett. 110 (2013) no. 25, 251801, arXiv:1304.6335 [hep-ex].
- [109] J. D. Chapman, Atmospheric neutrino observations in the MINOS far detector. PhD thesis, 2007.
- [110] G. Barr, T. Gaisser, P. Lipari, S. Robbins, and T. Stanev, A Three dimensional calculation of atmospheric neutrinos, Phys.Rev. D70 (2004) 023006, arXiv:astro-ph/0403630 [astro-ph].
- [111] MINOS Collaboration, P. Adamson et al., Measurements of atmospheric neutrinos and antineutrinos in the MINOS Far Detector, Phys.Rev. D86 (2012) 052007, arXiv:1208.2915 [hep-ex].
- [112] A. Dziewonski and D. Anderson, Preliminary reference earth model, Phys.Earth Planet.Interiors 25 (1981) 297–356.
- [113] M. Thompson, A new approach to the MINOS electron neutrino search., MINOSdoc-1398, 2006.
- [114] J. P. Ochoa, The Monte Carlo Nearest Neighbors (MCNN) Selection., MINOSdoc-5031, 2008.
- [115] R. B. Toner, Measuring θ_{13} via Muon Neutrino to Electron Neutrino Oscillations in the MINOS Experiment. PhD thesis, 2011.
- [116] M. Betancourt et al., Systematic uncertainties for the third nue box opening, MINOS-doc-7976-v1, 2011.
- [117] S. S. Wilks, The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses, Annals Math. Statist. 9 (1938) no. 1, 60–62.

- [118] F. James, MINUIT Function Minimization and Error Analysis: Reference Manual Version 94.1, 1994.
- [119] L. Whitehead, Normalisation and NC Background systematics for CC2014, MINOS-doc-10510-v1, 2014.
- [120] Kamiokande-II Collaboration, K. Hirata et al., Observation of a Neutrino Burst from the Supernova SN 1987a, Phys.Rev.Lett. 58 (1987) 1490–1493.
- [121] R. Bionta, G. Blewitt, C. Bratton, D. Casper, A. Ciocio, et al., Observation of a Neutrino Burst in Coincidence with Supernova SN 1987a in the Large Magellanic Cloud, Phys.Rev.Lett. 58 (1987) 1494.
- [122] E. Alekseev, L. Alekseeva, I. Krivosheina, and V. Volchenko, Detection of the Neutrino Signal From SN1987A in the LMC Using the Inr Baksan Underground Scintillation Telescope, Phys.Lett. B205 (1988) 209–214.
- [123] MINOS Collaboration, P. Adamson et al., Measurement of neutrino velocity with the MINOS detectors and NuMI neutrino beam, Phys.Rev. D76 (2007) 072005, arXiv:0706.0437 [hep-ex].
- [124] OPERA Collaboration, T. Adam et al., Measurement of the neutrino velocity with the OPERA detector in the CNGS beam, arXiv:1109.4897v2, 2012.
- [125] ICARUS Collaboration, M. Antonello et al., Measurement of the neutrino velocity with the ICARUS detector at the CNGS beam, Phys.Lett. B713 (2012) 17-22, arXiv:1203.3433 [hep-ex].
- [126] Borexino Collaboration, P. Alvarez Sanchez et al., Measurement of CNGS muon neutrino speed with Borexino, Phys.Lett. B716 (2012) 401-405, arXiv:1207.6860 [hep-ex].
- [127] LVD Collaboration, N. Y. Agafonova et al., Measurement of the velocity of neutrinos from the CNGS beam with the Large Volume Detector, Phys.Rev.Lett. 109 (2012) 070801, arXiv:1208.1392 [hep-ex].
- [128] OPERA Collaboration, T. Adam et al., Measurement of the neutrino velocity with the OPERA detector in the CNGS beam, JHEP 1210 (2012) 093, arXiv:1109.4897 [hep-ex].
- [129] R. Nichol, J. Meier, and A. Sousa, In-Situ Measurement of the MINOS Far

Detector GPS Antenna Fibre Transit Time, MINOS-doc-9078-v3, 2012.

- [130] B. Pahlka, Fiber Optic Cable Delay Measurements, MINOS-doc-9032-v2, 2012.
- [131] MINOS Collaboration, P. Adamson et al., Measurement of the neutrino mass splitting and flavor mixing by MINOS, Phys.Rev.Lett. 106 (2011) 181801, arXiv:1103.0340 [hep-ex].
- [132] J. Crisp and B. Fellenz, Tevatron Resistive Wall Current Monitor, Journal of Instrumentation 6 (2011) no. 11, T11001. http://stacks.iop.org/1748-0221/6/i=11/a=T11001.
- [133] P. Adamson, N. Ashby, R. Bumgarner, and N. Ashby, Measurement of the Velocity of the Neutrino with MINOS, arXiv:1408.6267 [physics.acc-ph].
- [134] S. Römisch et al., Synchronization between remote sites for the MINOS experiment, in 44th PTTI Meeting. 2012.
- [135] P. Tétreault, J. Kouba, P. Héroux, and P. LEGREE, CSRS-PPP: an Internet service for GPS user access to the Canadian Spatial reference Frame, Geomatica 59 (2005) no. 1, 17–28.
- [136] V. Bocean, Geodetic Distance Determination between the MINOS Detectors for the Neutrino Time of Flight Measurements, in 44th PTTI Meeting. 2012.
- T. Soler, R. H. Foote, D. Hoyle, and V. Bocean, Accurate GPS orientation of a long baseline for Neutrino Oscillation Experiments at Fermilab, Geophysical Research Letters 27 (2000) no. 23, 3921–3924.
 http://dx.doi.org/10.1029/2000GL011539.