

Preliminary Measurement Of Neutrino Oscillation Parameters By NuMI/MINOS And Calibration Studies For Improving This Measurement

Philip Andrew Symes

Submitted for the degree of D. Phil.

University of Sussex

November, 2005

Contents

1	INTRODUCTION	1
2	BACKGROUND	4
2.1	Neutrinos in Fermi-Dirac Theory	4
2.1.1	Dirac Particles	4
2.1.2	Weyl Neutrinos	5
2.1.3	Fermi Theory	6
2.1.4	Fermi–Dirac Neutrinos Considered Experimentally	7
2.2	Neutrinos in the Standard Model	7
2.2.1	Weak Interactions With Intermediate Vector Bosons	8
2.2.2	GSW Model	9
2.2.3	SM Neutrinos Considered Empirically (The Solar Neutrino Problem)	10
2.3	Neutrinos (Just) Beyond the Standard Model	13
2.3.1	Quark Sector Flavour Mixing & the CKM Matrix	13
2.3.2	Flavour Mixing in the Lepton Sector	14
2.3.3	Two-Flavour Vacuum Oscillation Probability	16
2.4	Three Flavour Neutrino Mixing	17
2.4.1	The PNMS Matrix	17
2.4.2	Three Flavour Oscillation Probability	19
2.4.3	Solar Neutrino Sector	19
2.4.4	Atmospheric Neutrino Sector	23
2.4.5	Reactor	25
2.4.6	Charge-Parity Violating Phase	25
2.4.7	Majorana Phases: Dirac & Majorana Neutrinos	26
2.5	Neutrino Mass Models & Scenarios	27
2.5.1	Adding Neutrino Mass to the SM Lagrangian	27
2.5.2	Sterile Neutrinos	28
2.5.3	Absolute Neutrino Mass Scale: Experimental Limits	28
2.5.4	Hierarchy Scenarios	29
2.5.5	Summary of Mass Models	30
2.6	The Frontiers of Neutrino Physics	31
2.6.1	Unanswered Questions	31
2.6.2	The Need for MINOS: Scientific Potential of the Experiment	31

3	THE MINOS PROJECT	32
3.1	General Principles	32
3.2	The NuMI Beam	33
3.2.1	Beam Monitoring	35
3.3	The MINOS Detectors	35
3.3.1	Detector Design	35
3.3.2	The Calibration Detector	37
3.3.3	The Near Detector	37
3.3.4	The Far Detector	40
3.3.5	Detector Monitoring	45
4	DATA ANALYSIS AT MINOS	46
4.1	Introduction	46
4.2	Atmospheric Analysis	46
4.3	T -test Analysis Theory	48
4.3.1	Oscillation Probability	48
4.3.2	Event Classification	49
4.4	T -test Analysis Method	51
4.4.1	Data Samples	51
4.4.2	Sample Selection: Cuts & Systematics	51
4.4.3	Fiducial Volume	52
4.4.4	Event Type Identification	54
4.4.5	Optimisation of Selection Cuts	63
4.4.6	Light Injection Cut on Real Data	63
4.4.7	Beam Energy Systematics	69
4.4.8	Expected Number of Events	70
4.5	T -test Results	72
4.5.1	Oscillation Parameter Spaces	72
4.5.2	Discussion	75
4.6	Charged Current Spectrum Analysis	75
5	CALIBRATION OVERVIEW	79
5.1	Calibration Requirements	79
5.2	The MINOS Calibration Chain	79
5.3	Charge Injection	81
5.4	Gain Correction	81
5.5	Light Injection	81
5.6	Strip-to-Strip Calibration	81
5.7	Attenuation Correction	82
5.8	Detector-to-Detector Calibration	83
5.9	Absolute Calibration	83
5.10	Other Calibrations	84
5.10.1	Timing Calibration	84

5.10.2	Magnetic Field Calibration	84
5.10.3	Detector Geometry	84
6	THE LIGHT INJECTION CALIBRATION SYSTEM	85
6.1	LI Calibration Concept	85
6.2	Equipment	87
6.3	Test And Selection of LEDs Used In Calibration	89
6.3.1	Pulse Spectrum & Linearity	92
6.3.2	Pulse Shape & Stability	94
6.3.3	Lifetime	96
6.3.4	Decision to Use Bivar LEDs	97
6.4	Light Injection Gain Curves	97
6.5	Light Injection Drift Method	97
6.5.1	Standard Light Injection Drift Points & Stability	97
6.5.2	Investigation Into Systematic Variations	101
6.5.3	Coherent Drifts At The Far Detector	104
6.5.4	Gain Change Averaging Drift Method	108
6.6	Debugging With Light Injection	117
6.6.1	NearLI Package	117
7	STRIP-TO-STRIP CALIBRATION OF THE FAR DETECTOR	123
7.1	Introduction	123
7.2	Strip-to-Strip Calibration Method	124
7.2.1	General Principles of Cosmic Muon Calibration	124
7.2.2	The Iterative Procedure	126
7.2.3	The Cosmic Muon Sample	128
7.3	Applying Cuts To The Sample	145
7.3.1	Track Cuts	145
7.3.2	Energy Cuts	145
7.3.3	Other Cuts	146
7.4	Applying Corrections To The Sample	146
7.4.1	The Gain Correction	146
7.4.2	The Attenuation Correction	148
7.4.3	The Path-length Correction	155
7.4.4	The Zero Reconstruction	156
7.4.5	The Sparsification Correction	164
7.5	Results	168
7.5.1	Photo-electron Response	168
7.5.2	Agreement of MC at FarDet	169
7.5.3	Data Agreement	177
7.5.4	Stability Of Constants And Detector	181
7.5.5	Caveats	182
7.6	Improvements	183

7.6.1	Truncation Point	183
7.6.2	Convergence Values	184
7.6.3	Attempts At Producing Constants For Pathological Strip-Ends	184
7.6.4	Procedure For Rolling Calibration	188
8	STRIP-TO-STRIP CALIBRATION AT THE NEAR DETECTOR	189
8.1	Introduction	189
8.2	Cosmic Muon Data Sample	189
8.2.1	Track Distributions & Cuts	190
8.2.2	Energy Cuts	191
8.2.3	Other Cuts	195
8.3	Method	197
8.3.1	Gain Correction	197
8.3.2	Attenuation Correction	197
8.3.3	Path-length Correction	202
8.3.4	Zero Correction	202
8.3.5	Sparsification Correction	206
8.4	Results	206
8.4.1	Photoelectron Response	206
8.4.2	Validation	207
9	SUMMARY	218
9.1	Contemporary Neutrino Physics	218
9.2	The MINOS Project	218
9.3	Beam, Beamline and Protons	218
9.4	Detectors	219
9.5	Data Analysis at MINOS	219
9.6	Calibration	220
9.6.1	Light Injection	220
9.6.2	FarDet Strip-to-strip Calibration	221
9.6.3	NearDet Strip-to-strip Calibration	221
9.6.4	Calibration Summary	221
	Bibliography	223
A	The MINOS Collaboration	228

List of Figures

2.1	ν_e-e scattering in Fermi Theory	7
2.2	The Standard Model of Elementary Particles	8
2.3	IVB ν_e-e weak interaction	8
2.4	Energy spectra of solar neutrinos in the SSM [5].	11
2.5	The proton–proton Chain	11
2.6	The carbon–nitrogen Cycle	12
2.7	Mixing in a quark sector interaction	13
2.8	The two right-handed axes show the mass and flavour eigenstates. As long as the mass eigenstates are not degenerate, though they share the same origin as the flavour eigenstates, they mix with the given angles [15].	18
2.9	Current neutrino oscillation parameters.	20
2.10	Fluxes of ^8B neutrinos [5].	21
2.11	Total neutrino rates: B–P SSM vs. experiment	21
2.12	Results from KamLAND	22
2.13	Origins of atmospheric neutrinos	23
2.14	Results of Super-Kamiokande	24
2.15	Results from K2K	25
2.16	Neutrino mass hierarchy	30
2.17	See-saw mechanism interactions	30
3.1	Cross-sectional view of the MINOS baseline	32
3.2	The NuMI target region	33
3.3	The NuMI beamline	34
3.4	The NuMI beam energy configurations	35
3.5	General information diagrams of MINOS	36
3.6	MINOS scintillator module	37
3.7	Photograph of the Calibration Detector	38
3.8	Representation of the Calibration Detector in MINOS software	38
3.9	CalDet response to 2 GeV particles	39
3.10	Photograph of the Near Detector	39
3.11	Diagram of Near Detector layout	40
3.12	Representation of the Near Detector in MINOS software	40
3.13	A NearDet high energy beam spill	41
3.14	The face of the Far Detector	42
3.15	Representation of the Far Detector in MINOS software	43
3.16	The module cross-section of the Far Detector	43

3.17	A Far Detector neutrino event	43
3.18	The moon shadow	44
3.19	Cosmic muon direction spectrum	44
4.1	Parameter space of MINOS atmospheric results	47
4.2	Zenith angle distributions of ν and $\bar{\nu}$ atmospheric events	48
4.3	Long and short events at MINOS	50
4.4	Approximate NearDet event energies in data and MC	52
4.5	FarDet: effect of fiducial volume cut on CC & NC event length distributions	53
4.6	Location of NearDet beam spot	54
4.7	NearDet: effect of fiducial volume and energy cuts on event length distributions	55
4.8	A typical long CC event at FarDet	57
4.9	A typical short CC event at FarDet	58
4.10	A typical NC event at FarDet	59
4.11	A typical long CC event at NearDet	60
4.12	A typical short CC event at NearDet	61
4.13	A typical NC event at NearDet	62
4.14	FarDet event lengths for CC & NC (MC)	64
4.15	NearDet event length distributions of CC & NC (MC)	65
4.16	FarDet strips hit per plane in short sample	66
4.17	Strips per plane cut at NearDet	67
4.18	Comparison of topological and chronological light injection cuts	68
4.19	Approximate event energies in data and MC	69
4.20	CC and NC neutrino event energies as functions of event length	70
4.21	Neutrino event lengths in the FarDet data sample	71
4.22	NuMI protons: integrated intensity	71
4.23	Confidence limits in parameter space for MDC	73
4.24	Confidence limits in parameter space for data	74
4.25	Near detector and far detector neutrino energy spectra	76
4.26	MINOS atmospheric sensitivities	77
4.27	MINOS reactor sector sensitivity	78
5.1	The MINOS calibration chain	80
5.2	Attenuation corrections vs. distance: V-strips, end 0	82
5.3	Stopping power for muons in polystyrene scintillator	83
5.4	An example stopping muon event at FarDet.	84
6.1	Schematic overview of the LI system hardware	86
6.2	A LI pulser box	87
6.3	Diagram of the cone assembly	88
6.4	LI cone and LI fibres	88
6.5	LI fibre connections at the back of the pulser box	88
6.6	LI fibre bundles at NearDet	89

6.7	Ashtray and readout fibre pigtail prototypes	90
6.8	Photograph of different LEDs	90
6.9	Non-linearity of the blue LEDs	91
6.10	Output spectra from blue and UV LEDs.	92
6.11	LED spectra and WLS fibre attenuation length.	93
6.12	LED and wavelength-shifting fibre normalised spectra	93
6.13	Pulse shape of HP blue LED	94
6.14	Pulse shape of Bivar UV LED	95
6.15	PIN diode measured pulse-to-pulse stability	95
6.16	Absolute output power comparison of LEDs	96
6.17	The Bivar UV LEDs were pulsed 10 million times over the period of a weekend, with light output monitored by a PIN photodiode. The diurnal temperature vari- ations over that time can be seen, and these correspond to the variations in light output of around 0.1%.	96
6.18	Gain curve of an example strip-end	98
6.19	Detector-averaged strip-end mean drift over one month	98
6.20	Average strip-end drifts on an example pixel over time	99
6.21	Pattern of problems on the east side of FarDet	100
6.22	Corrected drift in an example pixel	101
6.23	Corrected drift in example PMT showing bad fibres	102
6.24	Gain changes in an example PMT	103
6.25	Discrepancy between PMT and strip-end gains against nominal gain at FarDet . .	104
6.26	Discrepancy between PMT and strip-end gains after the nominal gain correction at FarDet	105
6.27	Effect of small nominal gain correction at FarDet	106
6.28	Nominal gain correction at CalDet	107
6.29	Pixel spot agreement with PMT average across PMT face	108
6.30	Corrected drift in an example PMT at FarDet	109
6.31	All strip-ends: pixel-corrected strip-end drift	110
6.32	Effect of pixel correction on strip-end drift	111
6.33	Pixel mean changes in PMT no. 15517	112
6.34	Strip-end drift mean corrected by PMT drift over time at FarDet	112
6.35	Effect of PMT correction on strip-end drift	113
6.36	Distribution of PMT corrected strip-end gain over time at FarDet	113
6.37	Cross-section of the PMT corrected strip-end gains near the end of the month at FarDet	114
6.38	Relationship between drift and gain at CalDet	114
6.39	Drift and PMT gain variation by drift point	115
6.40	Linear relation between drift and PMT gain at CalDet	115
6.41	Uncorrected drift over time at CalDet	116
6.42	Drift corrected by PMT gain over time at FarDet	116
6.43	Number of pulses and hits in each crate	118

6.44	Number of pulses and hits in each rack (U & V)	119
6.45	LED flashes and correlated hits in spectrometer planes	120
6.46	MINDERS in rack 10	120
6.47	All light and correlated light on an example NearDet plane	121
6.48	Distribution of calculated gains at NearDet	122
7.1	The strip-to-strip correction chain	125
7.2	An example simulated strip-end at CalDet	127
7.3	Ratio of muon distribution mean to fitted Landau peak position for PS muons	127
7.4	Histogram of strip-end convergence values	129
7.5	Spectra of cosmic muon track angles (MC)	130
7.6	Spectra of track directions (data)	130
7.7	Track directions with respect to strip orientations	131
7.8	U-direction cosine spectra as a function of strip no. (MC)	132
7.9	U-direction angle cosine distributions with respect to plane number (MC)	132
7.10	Track direction spectra with respect to strip orientation (data)	133
7.11	Distribution of track lengths (MC)	134
7.12	Distribution of track lengths (data)	134
7.13	Raw track hit distribution (MC)	135
7.14	Raw track hit distribution (data)	135
7.15	Cross-section of muon track hits	136
7.16	Cross-section of muon track hits (data)	136
7.17	Hit response cross-section of detector	137
7.18	Hit response cross-section of detector (data)	137
7.19	MC cross-section of Hit Locations, SM1	139
7.20	Data cross-section of Hit Locations, SM1	140
7.21	MC cross-section of hit locations, SM2	141
7.22	Data cross-section of hit locations, SM2	142
7.23	Vertical position of hits (MC)	143
7.24	Vertical position of hits (data)	144
7.25	Histogram of strip-end gains from the database	147
7.26	Histogram of gain corrected raw hits (MC)	147
7.27	Histogram of gain corrected raw hits (data)	148
7.28	Distribution of attenuation corrections	149
7.29	Attenuation corrections against distance along strip	150
7.30	Effectiveness of the attenuation correction (MC)	151
7.31	Effectiveness of the attenuation correction in normalising response across all strips (data).	152
7.32	Effectiveness of attenuation correction (8 m long strips in MC)	153
7.33	Effectiveness of the attenuation correction in normalising response (8 m strips, data)	154
7.34	Average pathlengths through strips	155
7.35	Relationship between track angle and calculated average path-length	156
7.36	Energy deposited against track angle (data)	157

7.37	Energy deposited against track angle (MC)	157
7.38	Distribution of path-lengths through strips (data)	158
7.39	Distribution of pathlengths through strips (MC)	158
7.40	Effectiveness of the path-length correction (MC)	159
7.41	Effectiveness of path-length correction (data)	160
7.42	Path-length and zero-correction residual	161
7.43	Size of applied zero correction in data (iterative)	162
7.44	Size of applied zero correction in MC (iterative)	163
7.45	Effect of zero reconstruction on average ADC response (MC)	163
7.46	Effect of zero reconstruction on average ADC response (data)	164
7.47	Balance of hits between strip-ends (MC)	165
7.48	Balance of hits between strip-ends (data)	166
7.49	Single photo-electron peak from phototube simulation	166
7.50	Applied weight of the 1-p.e. correction (MC)	167
7.51	Applied weight of the 1-p.e. correction (data)	167
7.52	Histogram of p.e. response of FarDet (MC)	168
7.53	Histogram of p.e. response of FarDet (data)	169
7.54	Statistical uncertainty in photo-electron response (data)	170
7.55	Relative statistical uncertainty in calibration constants (MC)	170
7.56	Typical strip-ends (MC)	171
7.57	Typical strip-ends (data)	172
7.58	Strip vs. plane map of p.e. response (data)	173
7.59	Strip vs. plane map of p.e. response (MC)	174
7.60	Calculated calibration constants vs. true light level	175
7.61	Distributions of calibration constants, true strip variations and the difference between them	176
7.62	Profile of calibration constants as a function of true light level	176
7.63	Convergence to 1.29	177
7.64	Spectrum of photo-electron hits on different light level strip-ends	178
7.65	Discrepancy between truth and calibration constant, east	179
7.66	Discrepancy between truth and calibration constant, west	180
7.67	Sigmap response after calibration, U-view	180
7.68	Sigmap response after calibration	181
7.69	Fractional difference in “siglin” calibration constant	182
7.70	Optimum truncation values on different light-level strip-ends	183
7.71	Spectrum of Photo-electron Hits on 80% Truncated Strip-ends	184
7.72	80% truncation point vs. light level	185
7.73	Systematic error as a function of light level for different convergence values	185
7.74	Optimal convergence value vs. light level	186
7.75	Relationship between calculated calibration constant and the raw hit average	187
8.1	Distribution of track lengths (MC)	190
8.2	Track direction with respect to strip orientations (MC)	191

8.3	Track direction with respect to strip orientations (data)	192
8.4	Track directions (MC)	192
8.5	Strip vs. plane map of the number of entries	193
8.6	Strip vs. plane map of number of entries (data)	194
8.7	Raw track hit distribution (MC)	195
8.8	Spectrum of hits on 5-p.e. light level strip-ends	196
8.9	Spectrum of hits on different light level strip-ends with too many 1-p.e. noise hits	196
8.10	Convergence of iterations	197
8.11	Distribution of strip-end gains	198
8.12	Distribution of deposited light before and after gain correction	198
8.13	Attenuation corrections vs. distance in U & V strips	199
8.14	Uncorrected and corrected hits vs. distance along strip	200
8.15	Attenuation corrections in data	201
8.16	Energy deposited against track angle ($\frac{ds}{dz}$)	202
8.17	Average path-length (ds) against track angle ($\frac{ds}{dz}$)	203
8.18	Path-length corrections applied	203
8.19	Hit vs. average path-length	204
8.20	Size of applied zero correction	205
8.21	Effect of zero correction on strip-end light levels	205
8.22	Size of applied sparsification correction	206
8.23	Strip vs. plane map of photo-electron response (MC)	207
8.24	Strip vs. plane map of p.e. response (data)	208
8.25	Distribution of p.e. calibration constants	209
8.26	Relative statistical uncertainty in calibration constants	209
8.27	Typical strip-end histograms (MC)	210
8.28	Typical strip-end histograms (data)	211
8.29	Agreement between calculated calibration constants and true light level	212
8.30	Residual between calculated calibration constants and true light level	213
8.31	Spectrum of p.e. hits on different light level strip-ends	213
8.32	Strip vs. plane map of discrepancy between truth and calibration constant	214
8.33	Validation of calibration constants with stopping muons	216
8.34	Average raw hit compared to calibration constant (partial planes)	217
9.1	Summary cartoon	222

List of Tables

4.1	$\mathcal{T}$ -test Results	72
6.1	Possible drift points measurement schemes	105
6.2	Comparison between standard LI and the two possible online methods of PMT averaging	108

Declaration

I hereby declare that this thesis has not been submitted, either in the same or different form, to this or any other university for a degree.

Acknowledgements

There are a number of people without whom this work could not have come about in its present form.

Firstly, I would like to thank my supervisor, Dr. P. G. Harris, for his constant guidance and support during my time working on MINOS at Sussex. I would also like to acknowledge Dr. R. Morse for his help when I first joined the experiment and Dr. P. Adamson for guidance during my time at Fermilab.

It has been a pleasure working with the many members of the MINOS collaboration, too numerous to name, that I have had the fortune to meet, both in the UK and USA...see Appendix A for their names.

I would like to thank the following at Sussex for their valuable input whilst preparing this thesis: Drs. C. Bungau, J. V. Dawson, C. B. Smith (now at UCL) and J. Wilson. Dr. Smith deserves a particular mention for his help and advice during the 3rd year of my work on MINOS. Nor would it be fitting to pass by Mr. D. J. Auty, Dr. D. D. Doyle, Mr. J. R. Grozier, Mr. C. Reeve or Mr. J. R. Richardson without recognising the value of their friendship in the group at Sussex.

Finally, I would like to thank my family and all my teachers, to whom I dedicate this work.

Preliminary Measurement Of Neutrino Oscillation Parameters By NuMI/MINOS And Calibration Studies For Improving This Measurement

Philip Andrew Symes

Submitted for the degree of D. Phil.

University of Sussex

November, 2005

Summary

This thesis explains the origins of neutrinos and their interactions, and the phenomenon of neutrino oscillations. Experiments for measuring neutrino oscillations are mentioned and the experiment investigated in this thesis, the “Main Injector Neutrino Oscillation Search”, and its neutrino beam, the Fermi National Accelerator Laboratory’s “Neutrinos At The Main Injector”, are described. MINOS is a long baseline (735 km) neutrino oscillation experiment with a near and a far detector, intended to make precision measurements of the atmospheric sector neutrino oscillation parameters.

A measurement is made of the “atmospheric” neutrino oscillation parameters, Δm^2_{23} and $\sin^2(2\theta_{23})$, using neutrinos from the NuMI beam. The results of this analysis are compared to measurements at MINOS using neutrinos from the atmosphere and with other experiments.

A more detailed method of beam neutrino analysis is discussed, and the extra calibrations needed to perform that analysis properly are described, with special attention paid to two aspects of the calibration, which comprise the bulk of work for this thesis.

The light injection calibration system uses LEDs to illuminate the detector readout and provides a normalisation of the stability of the detector over time. The hardware and different modi operandi of the system are described. There is a description of installation and commissioning of the system at one of the MINOS detectors.

The response normalisation of each detector with cosmic ray muons is described. Special attention is paid to the explanation of necessary corrections that must be made to the muon sample in order for the sample to be used to calibrate each detector to the specified accuracy. The performance of the calibration is shown.

Chapter 1

INTRODUCTION

Since the proposition of a “small neutral particle” by Pauli in 1930, neutrino physics has been at the forefront of research into the understanding of the nature of matter. Originally predicted to interact only weakly, once discovered, neutrinos were found to have curious properties. Not least among these was the phenomenon of neutrino oscillations between different lepton flavours. This process implies that neutrinos have mass, and neutrino oscillations are one of the first “beyond the Standard Model” processes to be studied.

Chapter 2 investigates our current understanding of neutrinos and sets out the background and context of neutrino physics. Neutrinos are considered in and beyond the Standard Model; neutrino interactions and the phenomenon of neutrino oscillations are explained. The chapter concludes by considering the future direction of research into neutrinos and the need for an experiment such as MINOS.

This chapter is a review of current knowledge in this subject by the author.

Chapter 3 describes the MINOS project: the NuMI neutrino beam and the near and far detectors that are used for comparing the spectra of neutrinos. NuMI is the name of the Fermilab project to produce a muon neutrino beam for MINOS. The spectrum of this beam is measured at a near detector, about 1 km from the beam target, and is measured again at a far detector, located 735 km from the near detector.

The experiment is designed to be sensitive to oscillations from $\nu_\mu \rightarrow \nu_\tau$, essentially by a deficit of muon neutrinos compared to the expected number of interactions. This allows MINOS to confirm the neutrino oscillation hypothesis, make precision measurements of the “atmospheric” (2–3) oscillation parameters and also improve the limit on the “reactor” (1–3) parameters (via ν_e appearance).

This chapter is a review of the MINOS project by the author. The project is a collaboration, founded in 1994; the author joined the collaboration in 2002.

Chapter 4 shows how data analysis can be done in MINOS. MINOS can measure charge sign separation of muons from atmospheric neutrinos, and this, the first published analysis by the full MINOS collaboration, is used to provide a limit on CP violation in the neutrino sector.

Work is then shown on a simple method of beam neutrino analysis that uses topological cuts to separate neutrino events into different interaction channels so that the neutrino flavour change

can be measured. The results of this work are shown.

A more complex analysis of the energy dependence of neutrino oscillations is needed in order for MINOS to measure the oscillation parameters to the design sensitivity. This method measures spectral distortion in the number of charged current muon neutrino events as a function of energy. The need for good calibration to reduce the systematic uncertainties in this measurement is explained and this leads on to the latter part of the thesis.

The work on the “ T -test” for measuring the oscillation parameters is by the author. The atmospheric analysis and full spectral beam analysis are the work of the collaboration.

Chapter 5 describes the calibration chain that is used by MINOS in order to achieve the specified level of calibration. There are several steps in this chain, providing corrections at different levels of reconstruction, plus other calibrations and measurements that need to be made in order to minimise systematic uncertainties. The remaining chapters look at some aspects of this calibration in more detail.

This chapter is the author’s review of the calibration chain worked on by the collaboration, in particular the MINOS “Calibration Working Group”.

Chapter 6 describes the light injection calibration system. The system uses LEDs to measure the response of the detector to known amounts of light and corrects variation over time and for PMT non-linearity. The hardware used for this system, and improvements to it, are described. The principles of the calibration and optimisations to the methods are then described. The system can be used to check detector performance and debug problems during commissioning and an explanation is made of a package used for such a purpose during the commissioning of the near detector.

The hardware and calibration concepts described in this chapter is work done by the collaboration, in particular those at the University of Sussex. The work on LEDs, “drift” calibration methods and the debugging package for the Near Detector is by the author.

Chapter 7 describes a method of strip-to-strip calibration using cosmic ray muons (intra-detector response normalisation) and applies this correction in the context of the far detector. The cuts, corrections and methodology are explained in detail, along with methods for boosting the statistical precision and systematic accuracy of the calibration constants. The method is validated by using MC simulations and testing the results using a stopping muon sample.

The corrections used in this section are the work of others in the collaboration. The application and implementation of these methods to the MINOS Far Detector is by the author. The work of other collaborators in this section is cited.

Chapter 8 describes the application of the strip-to-strip calibration to the near detector. The same validation is performed at the near detector as for the far detector and the results are shown to work well.

The application and implementation of the methods described in chapter 7 to the MINOS Near Detector is by the author.

Chapter 9 summarises the content of this thesis and puts the results of the calibrations and analysis into an abstract context. With calibrations such as are described in this thesis, MINOS looks set to make a measurement of the “atmospheric sector” neutrino oscillation parameters to 10% precision from 3 years of results taken with the NuMI beam.

.....

This chapter is a review of the contents of this thesis, i.e. of work by the author and other collaborators.

Chapter 2

BACKGROUND

The story of the neutrino can be taken as far back as the discovery of radioactivity by Henri Becquerel in 1896, but it was not until 1930 that the existence of these small, elusive particles was first postulated by Wolfgang Pauli [1]. Pauli used them to explain the apparent violation of spin angular momentum and energy conservations in beta-decay (such as $n \rightarrow p + e^- + \bar{\nu}_e$). A new spin- $\frac{1}{2}$ particle with no or very little mass and no electromagnetic charge was required, which therefore interacts only through the weak force. Since then, the particles (neutrinos) have been discovered and have been found to possess curious properties.

This chapter discusses these properties and the scientific motivation for the MINOS project. The aim here is to develop the theoretical and experimental background necessary to describe contemporary neutrino physics and the role of MINOS in this field [2]. General references for this chapter include [3],[4],[5],[6].

2.1 Neutrinos in Fermi-Dirac Theory

As fermions, neutrinos can be most simply described using the Fermi–Dirac representation.

2.1.1 Dirac Particles

The Dirac equation for spin- $\frac{1}{2}$ particles is

$$\left(i\gamma^\mu \frac{\partial}{\partial x^\mu} - m\right)\psi(x) = 0 \quad (2.1)$$

where γ^μ are the Dirac matrices, x^μ is space–time, ψ is the wave-function and the mass of the particle is m . This equation is derived and described in many quantum mechanics textbooks. The Dirac matrices are

$$\gamma^0 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \text{ and } \gamma^i = \begin{pmatrix} 0 & \sigma_i \\ -\sigma_i & 0 \end{pmatrix} \quad (2.2)$$

where σ_i are the Pauli matrices. These are

$$\sigma_1 \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \sigma_2 \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \text{ and } \sigma_3 \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (2.3)$$

The particle/anti-particle wave functions are

$$\psi_{p,s}^{[\pm]}(x) = \frac{[u]_{\pm}(p,s)}{2p^0 V} e^{[\pm]ip_\mu x^\mu} \quad (2.4)$$

where the quantities symbolised by p are elements of the particle's momentum 4-vector, the spin polarisation is $s = \pm\frac{1}{2}$, V is the volume (chosen to be normalised to contain one particle) and the particle/anti-particle space-independent spinors are given by u and v , so

$$u(p,s) = \sqrt{p^0 + m} \begin{pmatrix} X_s \\ \frac{\mathbf{s} \cdot \mathbf{p}}{p^0 + m} X_s \end{pmatrix} \quad (2.5)$$

and

$$v(p,s) = \sqrt{p^0 + m} \begin{pmatrix} \frac{\mathbf{s} \cdot \mathbf{p}}{p^0 + m} X_s \\ X_s \end{pmatrix} \quad (2.6)$$

where the two-spinor

$$X_s = \begin{pmatrix} |s + \frac{1}{2}| \\ |s - \frac{1}{2}| \end{pmatrix} \quad (2.7)$$

Massless Dirac particles are the only kind of particle that can have the left-handed spinor $\frac{1}{2}(1 - \gamma_5)\Psi$ as an eigenfunction, where

$$\gamma_5 = i\gamma^0\gamma^1\gamma^2\gamma^3 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad (2.8)$$

They have a simplified equation: the Weyl equation.

2.1.2 Weyl Neutrinos

Neutrinos can be described as Weyl spinors, i.e. Dirac particles that obey the Weyl equation

$$i\dot{\psi} = \boldsymbol{\sigma} \cdot \nabla \psi \quad (2.9)$$

for the two-spinor $\psi = \begin{pmatrix} s_1(x) \\ s_2(x) \end{pmatrix}$ with two independent components, where $\boldsymbol{\sigma}$ is the Pauli matrix 3-vector. These compose the Dirac four-spinor

$$\Psi' = \begin{pmatrix} +\psi \\ -\psi \end{pmatrix} \quad (2.10)$$

A left-handed Weyl two-component spinor satisfies

$$\Psi' = \frac{1 - \gamma^5}{2} \Psi_L = \Psi_L \quad (2.11)$$

and also

$$\frac{1 + \gamma^5}{2} \Psi_L = 0 \quad (2.12)$$

and vice-versa for right-handed spinors.

The free Weyl field is given by

$$\Psi_L(\mathbf{x}) = \int \left[\frac{a_L(\mathbf{p})u_L(\mathbf{p})e^{-i\mathbf{p}\cdot\mathbf{x}} + b_R^\dagger(\mathbf{p})v_R(\mathbf{p})e^{i\mathbf{p}\cdot\mathbf{x}}}{\sqrt{(2ip)^3 2E_p}} \right] d^3\mathbf{p} \quad (2.13)$$

with no sum over the spin, where a and b are right and left handed operators on the spinors.

A Ψ_L annihilates left-handed (L) particles and creates right-handed (R) anti-particles. Through a charge transformation

$$\hat{C}\Psi_L \rightarrow \bar{\Psi}_R \quad (2.14)$$

where, in the Pauli-Dirac representation, $\hat{C} = i\gamma^2\gamma^0$, $\bar{\Psi}_R$ is shown to be the adjoint of Ψ_L .

Neutrinos exist as left-handed spinors whereas charged leptons, such as the electron, can be considered as a projection of a Dirac 4-spinor, $\Psi = \Psi_L + \Psi_R$. We can work with left-handed spinors only in our representation of neutrinos, and with both u_L and v_R for representing other fermions.

2.1.3 Fermi Theory

Fermi Theory describes weak interactions of beta-decay and neutrino scattering as well as quark-sector weak interactions. It was developed in 1934 using charged currents in loose analogy to the theory of quantum electrodynamics. The weak, three-flavour, leptonic, charged, transition current, $J_\alpha^{(L)}(x)$, is defined as the sum over the weak flavours

$$J_\alpha^{(L)}(x) = J_\alpha^e(x) + J_\alpha^\mu(x) + J_\alpha^\tau(x) \quad (2.15)$$

The Hamiltonian of lepton interactions in the theory is

$$H_{int}^{(L)} = \frac{G_F}{\sqrt{2}} \int J_\alpha^{(L)\dagger}(x) J_\alpha^{(L)}(x) d^3x \quad (2.16)$$

where the Fermi coupling constant $G_F = 1.16637(1)10^{-5} \text{ GeV}^{-2}$.

To describe neutrino–electron scattering as shown in 2.1, for example, the weak neutral current is given by

$$J_{(e)}^{\alpha\dagger} = \bar{u}_{\nu_e} \gamma_\alpha (1 - \gamma_5) u_e \quad (2.17)$$

and the effective Lagrangian is thus

$$\mathcal{L} = -\frac{G_F}{\sqrt{2}} J_{(e)}^{\alpha\dagger} J_\alpha^{(e)} \quad (2.18)$$

Weak interactions between different flavour leptons, such as ν_μ – e interactions (substituting the electron-neutrino for a muon-neutrino in figure 2.1), are not allowed in the Fermi theory.




Figure 2.1: ν_e - e scattering in Fermi Theory

Fermi theory breaks down due to, amongst other things, unitarity violations¹ at high energies (when the centre of mass energy is above $E_{CM} > 10^3$ GeV), divergent second-order integrals and non-renormalisability. These limitations of Fermi theory can be circumvented using intermediate vector bosons (IVBs), and this is done to describe the interactions of neutrinos in the Standard Model.

2.1.4 Fermi–Dirac Neutrinos Considered Experimentally

Although hypothesised by Pauli in 1930, neutrinos were not detected directly until the 1956 experiment by Reines and Cowan at the Savannah River nuclear reactor in South Carolina in the USA [7]. They detected electron anti-neutrinos from the nuclear reactions at a short distance from the reactor core using β^+ -decay (inverse beta-decay)

$$\bar{\nu}_e + p \rightarrow e^+ + n \quad (2.19)$$

This was seen using liquid scintillator by looking for the coincidence of the 0.511 MeV photon from the positron annihilation and the subsequent delayed capture of the neutron. The small cross-section of $\mathcal{O}(10^{-38}) \text{ m}^2$ explained why neutrinos had taken so long to be detected.

Along with nuclear reactions, neutrinos have since been detected coming from beta decay, the Sun, the upper atmosphere, cosmogenic origins (e.g. supernovae) and, most recently, specialised particle accelerators. Neutrinos are thus well-established as fundamental particles.

2.2 Neutrinos in the Standard Model

The Standard Model of Particle Physics (SM) is an extremely successful theory for describing particle interactions at low energies, such as can be probed and detected in particle accelerator–collider experiments (figure 2.2).

The Standard Model is, in the syntax of group theory, a $SU(3) \otimes SU(2) \otimes U(1)$ gauge theory with particles and their interactions described by quantum field theory, i.e. in terms of Lagrangians and field equations. We can move from the Fermi description to the SM description by describing weak interactions in terms of IVB.

¹Pure S-wave process unitarity: for an explanation, see [5].

		Model of Elementary Particles					
		Three Generations of Matter(Fermions)			Force Carriers (Gauge Bosons)		
(Name)	Electric Charge	Number of Color Charges	Mass in MeV				
(Symbol)							
Q u a r k s	I			II			
	Up $+\frac{2}{3}$			Charm $+\frac{2}{3}$			
	u	$\frac{3}{3}$	~ 4	c	$\frac{3}{3}$	~ 1300	
	Down $-\frac{1}{3}$			Strange $-\frac{1}{3}$			
	d	$\frac{3}{3}$	~ 7	s	$\frac{3}{3}$	~ 150	
	Top/Truth $+\frac{2}{3}$			Bottom/Beauty $-\frac{1}{3}$			
	t	$\frac{3}{3}$	$>131,000$	b	$\frac{3}{3}$	$\sim 4,800$	
	L e p t o n s	III					
Electron 0				Muon 0			
Neutrino $<15\text{ eV}$				Neutrino $<100\text{ KeV}$			
ν_e				ν_μ			
Electron -1				Muon -1			
e		$\frac{3}{3}$	511	μ	$\frac{3}{3}$	106.76	
Tau 0				Tau -1			
Neutrino $<35\text{ MeV}$							
ν_τ							
τ							

$$\mathcal{L} = \mathcal{L}_{free} + \mathcal{L}_{int} = i\bar{\Psi}\gamma^\mu\partial_\mu\Psi + g\bar{\Psi}\gamma^\mu\frac{\mathbf{A}_\mu\cdot\boldsymbol{\tau}}{2}\Psi \quad (2.20)$$

where $\mathbf{A}_\mu\cdot\boldsymbol{\tau}$ is the scalar product of the vector field and the Pauli vectors in spherical co-ordinates (the generators of the isospin group $SU(2)$). g , the weak coupling, is described by

$$\frac{G_F}{2} \sim \frac{g^2}{8M_W^2} \quad (2.21)$$

where M_W is the W-boson mass.

There are no explicit mass terms in this Lagrangian so that neutrinos and charged leptons are indistinguishable in the SM: the differences between charged lepton and neutrino masses comes from spontaneous symmetry breaking through interaction with the Higgs field.

Although the g is now dimensionless, the theory is still not renormalisable and unitarity conservation has the same high energy limit as the Fermi theory for processes involving an external $W^\pm$. This divergence can be cancelled with the introduction of a third, neutral IVB, now known as the Z (or Z^0): this is the neutral current IVB in this theory.

When the field strength isovectors

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu + gA_\mu \times A_\nu \quad (2.22)$$

and mass terms ($\widehat{M}$) are added, the gauge invariant Lagrangian for interactions in this theory is

$$\mathcal{L} = \mathcal{L}_{free} + \mathcal{L}_{int} + \mathcal{L}_A \quad (2.23)$$

$$= \bar{\Psi}\gamma^\mu(i\partial_\mu + g\frac{\mathbf{A}_\mu\cdot\hat{\boldsymbol{\tau}}}{2})\Psi + \bar{\Psi}\widehat{M}\Psi + \frac{1}{4}\mathbf{F}_{\mu\nu}\cdot\mathbf{F}^{\mu\nu} \quad (2.24)$$

2.2.2 GSW Model

As mentioned earlier, the SM is a $SU(3) \otimes SU(2) \otimes U(1)$ gauge theory and so has gauge symmetries and conserved quantities. When a local gauge symmetry is imposed on a fermion field, a bosonic field term is introduced to preserve this symmetry and associated with this is a conserved quantity. This quantity is known as the “charge” in analogy to the electric charge in quantum electrodynamics. The strength of the coupling between the fields is determined by the coupling constant and the magnitude of the charge.

Electro-magnetic and weak interactions are unified in the SM into the Glashow-Salam-Weinberg (GSW) electro-weak model. In the GSW Model, every fermion generation (e, μ, τ) contains two related left-handed leptons forming an isospin doublet

$$L_i = \frac{1 - \gamma_5}{2} \begin{pmatrix} \psi_{\nu_i} \\ \psi_i \end{pmatrix} \quad (2.25)$$

$$= \frac{1 + \gamma_5}{2} (\bar{\psi}_{\nu_i}, \bar{\psi}_i) \quad (2.26)$$

Since right-handed neutrinos do not exist in this model, the right-handed components are represented by singlet states

$$R_i = \frac{1 + \gamma_5}{2} \psi_i \quad (2.27)$$

The symmetry imposed on this system is $SU(2)_W \otimes U(1)_Y$

The charged weak currents have the form

$$J_{\pm}^{(i)\alpha} = 2\bar{L}_i \gamma^{\alpha} \hat{T}_{\pm} L_i \quad (2.28)$$

with the weak isospin doublets

$$\hat{T}_{\pm} = \hat{T}_1 \pm i\hat{T}_2 \quad (2.29)$$

and

$$\hat{T}_i = \frac{\hat{\tau}_i}{2} \quad (2.30)$$

The electromagnetic (charged) current exists for the charged leptons (i.e. not neutrinos) and is given by

$$J_{EM}^{(i)\alpha} = \bar{L}_i \gamma^{\alpha} \left(\frac{1}{2} - \hat{T}_3 \right) L_i + \bar{R}_i \gamma^{\alpha} R_i \quad (2.31)$$

The weak neutral current has the form

$$J_0^{(i)a} = \frac{\sqrt{2}}{\cos \theta} \left[\bar{L}_i \gamma^{\alpha} \begin{pmatrix} 1 & 0 \\ 0 & -\cos 2\theta \end{pmatrix} L_i + 2 \sin^2 \theta \bar{R}_i \gamma^{\alpha} R_i \right] \quad (2.32)$$

with the neutrino and charged lepton fields written separately.

The complete GSW Lagrangian is constructed as

$$\mathcal{L}_{GSW} = \mathcal{L}_{\text{free}}^{\text{boson}} + \mathcal{L}_{\text{free}}^{\text{fermion}} + \mathcal{L}_{\text{int}}^{\text{lepton-gauge}} + \mathcal{L}_{\text{self}}^H + \mathcal{L}_{\text{int}}^{H\text{-lepton}} + \mathcal{L}_{\text{int}}^{H\text{-gauge}} \quad (2.33)$$

where H denotes the Higgs term, though for the purposes of considering SM neutrinos, we are only interested in the lepton–gauge interaction terms.

2.2.3 SM Neutrinos Considered Empirically (The Solar Neutrino Problem)

The neutrino flux predicted by the Bahcall–Pinsonneault Standard Solar Model [8], is shown in figure 2.4.

Most solar neutrinos, come from the “p–p chain” (figure 2.5), providing 91% of the total flux. The dominant reaction is

$$4p + 2e \rightarrow {}^4\text{He} + 2\nu_e \quad Q = 26.73 \text{ MeV} \quad (2.34)$$

Raymond Davis, Jr. built an experiment in the Homestake Mine in South Dakota in the USA to measure the flux of neutrinos from the Sun [9]. The neutrinos are mainly from the B -branch of the p–p chain (the branch with ${}^8\text{B}$ daughters), but also from the ${}^{13}\text{N}$ and ${}^{15}\text{O}$ reactions of the C–N cycle (figure 2.6), and the additional reaction




Figure 2.4: Energy spectra of solar neutrinos in the SSM [5]. The different colours show the different ranges of energies that different detection methods are sensitive to.




Figure 2.5: The proton-proton chain of reactions in the B-P SSM




Figure 2.6: The carbon–nitrogen cycle of reactions in the B–P SSM



The neutrinos were detected in dry cleaning fluid, C_2Cl_4 , through the reaction



The tank was purged monthly and the number of argon atoms in the tank were counted by the decay rate of the ${}^{37}\text{Ar}$ isotopes in the sample. The experiment took years to run and had large errors due to the low flux, the cosmic ray background, the small amount of target nuclei (only 24% of chlorine is the ${}^{37}\text{Cl}$ isotope) and the inefficiency of the argon detection method. However, Davis measured a distinct deficit of neutrino events compared to solar model predictions.

Subsequent experiments used different detector methods in order to try to confirm Davis' measurement. Experiments which detect neutrinos using ${}^{71}\text{Ga}$, such as SAGE [10], GNO [11] and GALLEX [12] did so via the reaction



and have a relatively low energy threshold of $E > 0.23$ MeV. That means that they are sensitive to the much greater neutrino flux rates of the p–p chain and are also sensitive to different solar reactions. These experiments also measured a ν_e flux deficit, although their results did not agree numerically with Davis'.

The results of further experiments and investigations into solar reactions by other methods, e.g. helioseismology, agree well with the predictions of the B–P SSM, however. In order to explain

Davis' results, and those of the other solar neutrino experiments, we must consider the properties of neutrinos (just) beyond the SM.

2.3 Neutrinos (Just) Beyond the Standard Model

This section will explain neutrino oscillations and why neutrino mass has become a well accepted extension to the SM since oscillations have been observed.

2.3.1 Quark Sector Flavour Mixing & the CKM Matrix

An explanation of the neutrino flux deficit is the flavour oscillation of leptons, which can be considered in analogy to flavour mixing in quarks. The quark sector is divided up into three left-handed doublet states

$$\begin{pmatrix} u_L \\ d_L \end{pmatrix}, \begin{pmatrix} c_L \\ s_L \end{pmatrix}, \begin{pmatrix} t_L \\ b_L \end{pmatrix} \quad (2.38)$$

and 6 right-handed singlet states

$$u_R, d_R, c_R, s_R, t_R, b_R \quad (2.39)$$

Flavour mixing in the quark sector is a SM phenomenon and allows, for example, flavour changing charged currents, such as the s -quark annihilating an $\bar{u}$ -quark in the following interaction (figure 2.7).




Figure 2.7: Mixing in a quark sector interaction: K-meson s -quark interacting with an anti- u -quark. The s -quark effectively mixes with a u -quark first before annihilating the anti- u -quark.

In this interaction, the electric charge (Q), weak isospin (T_3) and weak hypercharge (Y) are all conserved: the only conserved quantity violated is flavour.

This implies that the weak eigenstates of these quarks are not eigenstates of the Hamiltonian (energy) operator, since there is no reason why the weak interaction should distinguish particles by their masses. This mixing between separate doublets, although observed, is “Cabbibo suppressed”,

so that the s -quark– u -quark interactions are about 5 times weaker than the d -quark– u -quark interactions.

The left-handed doublets are adjusted to take account of this mixing by using the form

$$L_{u,d} = \frac{1 - \gamma_5}{2} \begin{pmatrix} u_L \\ ad_L + bs_L \end{pmatrix} \quad (2.40)$$

$$= \frac{1 - \gamma_5}{2} \begin{pmatrix} u_L \\ d'_L \end{pmatrix} \quad (2.41)$$

for the first generation (u, d), for example. In all generations, the lower (primed) component is an arbitrary mixture of mass eigenstates of the d -, s - and b -quarks. This mixing was described by Cabbibo for two-flavour mixing, and expanded by Kobayashi and Maskawa for the three flavour case, by using three Euler angles and a phase difference to formulate a “mixing matrix” [13]. This unitary 3×3 CKM-matrix is defined as

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} c\theta_1 & s\theta_1 c\theta_3 & s\theta_1 s\theta_3 \\ -s\theta_1 c\theta_2 & c\theta_1 c\theta_2 c\theta_3 + s\theta_2 s\theta_3 e^{i\delta} & c\theta_1 c\theta_2 s\theta_3 - s\theta_2 c\theta_3 e^{i\delta} \\ -s\theta_1 s\theta_2 & c\theta_1 s\theta_2 c\theta_3 + c\theta_2 s\theta_3 e^{i\delta} & c\theta_1 s\theta_2 s\theta_3 - c\theta_2 c\theta_3 e^{i\delta} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix} \quad (2.42)$$

where δ is a CP-violating phase, c denotes $\cos$ and s denotes $\sin$. The angles (θ) can be defined arbitrarily, so that in the case of two flavour mixing, only one angle remains: the “Cabbibo angle”. We will consider the qualitative ramifications of this, in terms of flavour mixing in the lepton sector.

2.3.2 Flavour Mixing in the Lepton Sector

No SM processes allow flavour changing for leptons. We can, though, proceed as in the quark case, and define left-handed doublets for the three lepton flavours ($i = e, \mu, \tau$) by:

$$L'_i = \begin{pmatrix} \nu_L \\ i'_L \end{pmatrix} \quad (2.43)$$

We will consider, for simplicity, allowing mixed states between just two flavours, such as electron and muon, via the unitary matrix

$$\begin{pmatrix} e'_L \\ \mu'_L \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} e_L \\ \mu_L \end{pmatrix} \quad (2.44)$$

though the symbols could be replaced to indicate mixing between any two flavours.

We can also introduce completely mixed doublets between these two flavours of the form

$$L''_e = L'_e \cos \theta - L'_\mu \sin \theta \quad (2.45)$$

and

$$L''_\mu = L'_e \sin \theta - L'_\mu \cos \theta \quad (2.46)$$

The interaction Lagrangian from GSW theory is

$$\mathcal{L}_{int}^e = g(\bar{L}_e \gamma^\alpha \hat{\mathbf{T}} L_e) \cdot \mathbf{A}_a - (L + R) B_a \quad (2.47)$$

where $\hat{\mathbf{T}}$ operates on spin and isospin indices. Although we are interested in the first term for our interactions, the second is added for completeness.

The current contribution of these doublets to the total is

$$J = \bar{L}'_e \gamma^\alpha \hat{\mathbf{T}} L'_e + \bar{L}'_\mu \gamma^\alpha \hat{\mathbf{T}} L'_\mu \quad (2.48)$$

$$= \bar{L}''_e \gamma^\alpha \hat{\mathbf{T}} L''_e + \bar{L}''_\mu \gamma^\alpha \hat{\mathbf{T}} L''_\mu \quad (2.49)$$

so it can be seen that the L' and L'' representations are the same. We can now proceed using the L'' forms of the doublets

$$L''_e = \begin{pmatrix} \nu_e \cos \theta - \nu_\mu \sin \theta \\ e'_L \cos \theta - \mu'_L \sin \theta \end{pmatrix} = \begin{pmatrix} \nu'_e \\ e_L \end{pmatrix} \quad (2.50)$$

and

$$L''_\mu = \begin{pmatrix} \nu_e \sin \theta - \nu_\mu \cos \theta \\ e'_L \sin \theta + \mu'_L \sin \theta \end{pmatrix} = \begin{pmatrix} \nu'_\mu \\ \mu_L \end{pmatrix} \quad (2.51)$$

The allowed mixed states are now

$$\begin{pmatrix} \nu'_e \\ \nu'_\mu \end{pmatrix} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} \quad (2.52)$$

$$= \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}^\dagger \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} \quad (2.53)$$

In the quark mixing case of the CKM matrix, only the bottom part of the doublets (d', s', b') was allowed to mix. Here we also allow only the top or bottom part of the doublet to mix, i.e. only charged leptons or neutrinos can mix, but not both. By choosing the L'' forms of the doublets, we have only neutrino mixing, as we require.

The ν_e momentum wave-function is a super-position of the pure (unmixed) electron-neutrino and muon-neutrino momentum wave-functions, i.e.

$$\Psi_{\nu'_e}^{(\mathbf{p})}(x) = \Psi_{\nu_e}^{(\mathbf{p})}(x) \cos \theta + \Psi_{\nu_\mu}^{(\mathbf{p})}(x) \sin \theta \quad (2.54)$$

If neutrinos are massless, applying their wave-function from the Dirac equation yields

$$\Psi_{\nu'_e}^{(\mathbf{p})}(x) = \frac{[u_{\nu_e}(\mathbf{p}, s) \cos \theta + u_{\nu_\mu}(\mathbf{p}, s) \sin \theta]}{\sqrt{2|\mathbf{p}|V}} e^{i(\mathbf{p} \cdot \mathbf{x} - |\mathbf{p}|t)} \quad (2.55)$$

where there are no mixing terms. Therefore, there is either no mixing or lepton sector mixing is in the charged part of the left-handed doublet. Charged lepton mixing is ruled-out experimentally [5] since the branching ratio

$$\frac{W(\mu \rightarrow e\gamma)}{W(\mu \rightarrow e\bar{\nu}_e\nu_\mu)} < 10^{-10} \quad (2.56)$$

However, if neutrinos do have mass, acquired through interactions with the Higgs field (like charged leptons)², the Dirac equation for massive particles gives

$$\Psi_{\nu'_e}^{(\mathbf{p})}(x) = \left[\frac{u_{\nu_e}(\mathbf{p}, s)}{\sqrt{2p_0V}} \cos \theta e^{-ip_0t} + \frac{u_{\nu_\mu}(\mathbf{p}, s)}{\sqrt{2p'_0V}} \sin \theta e^{-ip'_0t} \right] e^{i\mathbf{p}\cdot\mathbf{x}} \quad (2.57)$$

where $p_0 = \sqrt{p^2 + m_e^2}$, etc. Here, there are possible mixing terms.

Therefore, unless the masses of the neutrinos are exactly the same, or the mixing angle $\theta = 0$, energy eigenstates (eigenstates of the Hamiltonian) do not have a defined momentum. Again, as in the quark case, only lepton number is violated by this mixing. This implies that weak and mass eigenstates are not the same, as in the case of quark flavour mixing, and so evolve over space–time for the free neutrinos. It is this mixing that causes the phenomenon of neutrino oscillations.

2.3.3 Two-Flavour Vacuum Oscillation Probability

Let us now consider these oscillations between, for example, ν_μ and ν_τ with mass eigenstates which are therefore ν_2 and ν_3 . In this context, the mixing is

$$\begin{pmatrix} \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} \nu_2 \\ \nu_3 \end{pmatrix} \quad (2.58)$$

Consider the neutrinos as evolving states of definite energy E as a function of position

$$\begin{pmatrix} \nu_2(\mathbf{x}) \\ \nu_3(\mathbf{x}) \end{pmatrix} = \begin{pmatrix} \nu_2(0)e^{i\mathbf{p}_2\cdot\mathbf{x}} \\ \nu_3(0)e^{i\mathbf{p}_3\cdot\mathbf{x}} \end{pmatrix} \quad (2.59)$$

where $p_{2,3} = \sqrt{E^2 - m_{2,3}^2}$.

Using the Hermitian property of the oscillation matrix, we can substitute for the mass eigenstates to get

$$\begin{pmatrix} \nu_\mu(\mathbf{x}) \\ \nu_\tau(\mathbf{x}) \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} e^{i\mathbf{p}_2\cdot\mathbf{x}} & 0 \\ 0 & e^{i\mathbf{p}_3\cdot\mathbf{x}} \end{pmatrix} \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} \nu_\mu(0) \\ \nu_\tau(0) \end{pmatrix} \quad (2.60)$$

With appropriate matrix algebra, we produce

$$\begin{pmatrix} \nu_\mu(\mathbf{x}) \\ \nu_\tau(\mathbf{x}) \end{pmatrix} = \begin{pmatrix} e^{i\mathbf{p}_2\cdot\mathbf{x}} \cos^2 \theta + e^{i\mathbf{p}_3\cdot\mathbf{x}} \sin^2 \theta & (e^{i\mathbf{p}_3\cdot\mathbf{x}} - e^{i\mathbf{p}_2\cdot\mathbf{x}}) \cos \theta \sin \theta \\ (e^{i\mathbf{p}_3\cdot\mathbf{x}} - e^{i\mathbf{p}_2\cdot\mathbf{x}}) \cos \theta \sin \theta & e^{i\mathbf{p}_3\cdot\mathbf{x}} \cos^2 \theta + e^{i\mathbf{p}_2\cdot\mathbf{x}} \sin^2 \theta \end{pmatrix} \begin{pmatrix} \nu_\mu(0) \\ \nu_\tau(0) \end{pmatrix} \quad (2.61)$$

In the NuMI conventional neutrino beam used by MINOS, neutrinos start as essentially $\nu_\mu = 1$ (i.e. $\nu_\tau = 0$) at $x = 0$, so the probability of finding a ν_τ at position x is given by

$$P_{\nu_\mu \rightarrow \nu_\tau} = |\nu_\tau(\mathbf{x})|^2 = |(e^{i\mathbf{p}_3\cdot\mathbf{x}} - e^{i\mathbf{p}_2\cdot\mathbf{x}}) \cos \theta \sin \theta|^2 \quad (2.62)$$

²Neutrino mass models are discussed later, in section 4 of this chapter.

Simplifying the trigonometric identities leads us to just

$$P_{\nu_\mu \rightarrow \nu_\tau} = |\nu_\tau(\mathbf{x})|^2 = \sin^2(2\theta) \sin^2\left(\frac{(p_3 - p_2)x}{2}\right) \quad (2.63)$$

Since the neutrino mass-energy is so small compared to the neutrino kinetic energy, i.e. $E_\nu \ll m_\nu$, we can make the approximation:

$$E \approx p \quad (2.64)$$

so that

$$p_3 - p_2 = \sqrt{E^2 - m_3^2} - \sqrt{E^2 - m_2^2} \quad (2.65)$$

By using the expansion

$$m^2 = (E + p)(E - p) \quad (2.66)$$

and applying the approximation from 2.64, that

$$p = E - \frac{m^2}{E + p} \approx E - \frac{m^2}{2E} \quad (2.67)$$

equation 2.65 becomes

$$p_3 - p_2 \approx \frac{\Delta m_{23}^2}{2E} \quad (2.68)$$

Therefore the two-flavour oscillation probability is

$$P_{\nu_\mu \rightarrow \nu_\tau} \approx \sin^2 2\theta \sin^2\left(\frac{\Delta m_{23}^2 x}{4E}\right) \quad (2.69)$$

The parameter ratio $\frac{x}{E}$ is between the distance travelled by the neutrino and its energy in natural units, but is usually expressed as $\frac{L}{E}$, where L is measure in km and E in GeV. The oscillation probability in these units is

$$P_{\nu_\mu \rightarrow \nu_\tau} \approx \sin^2 2\theta \sin^2\left(1.27 \Delta m_{23}^2 \frac{L}{E}\right) \quad (2.70)$$

Next it is necessary to investigate oscillations between all three active lepton flavours.

2.4 Three Flavour Neutrino Mixing

The theory can be extended to three active flavours, such as exist in nature.

2.4.1 The PNMS Matrix

This theory can be expanded from two flavour mixing to three flavours, and a mixing matrix can be defined in analogy to the CKM matrix: the Pontecorvo–Maki–Nakagawa–Sakata matrix [14]. This matrix in 3×3 form is:

$$U_{\text{PMNS}} = \begin{pmatrix} c\theta_{12}c\theta_{13} & s\theta_{12}c\theta_{13} & s\theta_{13}e^{-i\delta} \\ -s\theta_{12}c\theta_{23}-c\theta_{12}s\theta_{23}s\theta_{13}e^{i\delta} & c\theta_{12}c\theta_{23}-s\theta_{12}s\theta_{23}s\theta_{13}e^{i\delta} & s\theta_{23}-c\theta_{13} \\ s\theta_{12}s\theta_{23}-c\theta_{12}c\theta_{23}s\theta_{13}e^{i\delta} & -c\theta_{12}s\theta_{23}-s\theta_{12}c\theta_{23}s\theta_{13}e^{i\delta} & c\theta_{23}c\theta_{13} \end{pmatrix} \quad (2.71)$$

where again, s denotes sin and c denotes cos, so that

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \quad (2.72)$$

where flavour eigenstates are in columns and mass eigenstates in rows. This mixing can be represented diagrammatically (figure 2.8).




Figure 2.8: The two right-handed axes show the mass and flavour eigenstates. As long as the mass eigenstates are not degenerate, though they share the same origin as the flavour eigenstates, they mix with the given angles [15].

This PMNS matrix can be usefully factorised to into three oscillation sectors, and a Majorana neutrino mass matrix ³ as:

$$U_{\text{PMNS}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_{23} & \sin \theta_{23} \\ 0 & -\sin \theta_{23} & \cos \theta_{23} \end{pmatrix} \quad \text{atmospheric sector (23)} \\ \times \begin{pmatrix} \cos \theta_{13} & 0 & \sin \theta_{13}e^{i\delta} \\ 0 & 1 & 0 \\ -\sin \theta_{13}e^{i\delta} & 0 & \cos \theta_{13} \end{pmatrix} \quad \text{reactor sector (13)} \\ \times \begin{pmatrix} \cos \theta_{12} & \sin \theta_{12} & 0 \\ -\sin \theta_{12} & \cos \theta_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{solar sector (12)} \\ \times \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\phi_1} & 0 \\ 0 & 0 & e^{i\phi_2} \end{pmatrix} \quad \text{Majorana phase} \quad (2.73)$$

³Majorana phases are discussed in section 2.4.7.

Experimental neutrino results can be categorised as to which of these sectors they probe.

2.4.2 Three Flavour Oscillation Probability

The three-flavour oscillation probability between two different states, α and β is more complicated than the two-flavour case, so will not be derived here (see [6]). The final result, from algebra with the PNMS matrix, is

$$P_{\nu_\alpha \rightarrow \nu_\beta} = \sum_j |U_{\alpha j} U_{\beta j}|^2 + \text{Re} \left\{ \sum_{i \neq j} U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j} e^{-i(m_i^2 - m_j^2) \frac{L}{2E} \frac{\text{km}}{\text{GeV}}} \right\} \quad (2.74)$$

The current best results and limits of the neutrino oscillation parameters in each sector of equation 2.73 are shown in figure 2.9.

2.4.3 Solar Neutrino Sector

Solar neutrino sector experiments use neutrinos from the Sun, with the flux predicted by the B–P SSM, and the neutrino energy spectrum shown in figure 2.4, as discussed in section 2.3.

The most prominent contemporary solar neutrino experiment is the Sudbury Neutrino Observatory [16, 17], which will soon be decommissioned. SNO is a 1 *kton* heavy-water detector, which measures the ^8B solar neutrino flux. One of the experiments three phases had salt, NaCl , dissolved in the water. The advantage of SNO over the electron neutrino disappearance experiments discussed in section 2.3 is that it can detect neutrinos of all flavours through neutral current (NC) and elastic scattering (ES) interactions, and neutrinos of just electron flavour through charged current (CC) interactions

$$\text{CC:} \quad \nu_e + D \rightarrow e^- + p + p \quad (2.75)$$

$$\text{NC:} \quad \nu_i + D \rightarrow \nu_i + p + n \quad (2.76)$$

$$\text{ES:} \quad \nu_i + e^- \rightarrow \nu_i + e^- \quad (2.77)$$

Thus it can make a measurement of both the total neutrino flux and the flux of electron neutrinos and this means it can prove that the neutrinos from the Sun have oscillated from one type to another, if the flux ratio differs from unity. The SNO results for each reaction channel are summarised in figure 2.10.

Figure 2.11 shows the SNO results (right two data sets), which proves oscillations $\nu_e \rightarrow \nu_x$ and also confirms the predictions of the B–P SSM, solving the solar neutrino problem. The other experiments shown in figure 2.11 are Kamiokande and Super-Kamiokande [18, 19], SAGE [10], GALLEX [12] and GNO [11].

Matter Effects

The oscillation formula (equation 2.74) is appropriate for neutrinos travelling through a vacuum. Like photons, however, neutrinos can change their properties as they travel through matter. The neutrino refractive index can be expressed as




Figure 2.9: The current limits on the measurement of neutrino oscillation parameters in all sectors [5]. ν_X implies any other neutrino flavour.




Figure 2.10: Fluxes of ^8B solar neutrinos from SNO's charged-current, elastic scattering and neutral-current results with salt, along with the B-P SSM prediction. Bands represent 1 sigma uncertainties, and the dashed ellipses show the 68%, 95% and 99% certainty levels [5].




Figure 2.11: Current knowledge of Solar neutrino rates from several of the experiments mentioned in the text [5]. The bars show the discrepancy between the predicted and measured neutrino fluxes. SNO shows that although there is a discrepancy in the ν_e flux, there is no discrepancy in the all-flavour neutrino flux.

$$n = 1 + \sum_{\alpha=n,p,e} \frac{2\pi N_{\alpha}}{p^2} f_{\nu\alpha} \quad (2.78)$$

where N_{α} is the number density of the relevant particle type α and $f_{\nu\alpha}$ is the forward scattering amplitude of the neutrino. All neutrino flavours interact with neutrons and protons, but only electron-neutrinos interact in the CC channel with electrons in ordinary matter. This is known as the Wolfenstein-Mikheyev-Smirnov (MSW) effect [20, 21]. The only remaining forward scattering phase is

$$f_{\nu ee} = \frac{G_F p}{\sqrt{2}ip} \quad (2.79)$$

This is a major effect in many solar sector neutrino oscillation experiments, such as SNO, due to the large matter density inside the Sun. An Earth-based experiment looking for solar-sector oscillations from a terrestrial neutrino source would not see this effect since the Earth has a relatively low matter density and the path-length of the neutrinos is less.

The Kamioka Liquid Anti-neutrino Detector (KamLAND) [22, 23] is a 1 *kton* ultra-pure liquid scintillator detector located at the Kamioka mine in Japan. It uses $\bar{\nu}_e$ from 16 nuclear reactors in Japan and South Korea with an average energy of 4 MeV and an average baseline of 180 km, i.e. covering the solar parameter space. KamLAND was able to confirm neutrino oscillations in this sector (figure 2.12) above other, now obsolete theories and to confirm the results of other experiments in this sector.




Figure 2.12: The $\frac{L}{E}$ survival probability from the KamLAND experiment. The spectral form of the data favours neutrino oscillations over other hypotheses that could have explained the Super-Kamiokande data [23].

The Kamioka detectors, Kamiokande, Super-Kamiokande and the planned Hyper-Kamiokande, also make measurements of the solar neutrino flux and oscillations, but the Super-Kamiokande experiments are also famous for their measurements in the atmospheric neutrino sector.

2.4.4 Atmospheric Neutrino Sector

Atmospheric neutrino experiments look for muon-neutrinos from cosmic ray interactions in the upper atmosphere, such as is shown in figure 2.13. These experiments are sensitive to neutrino direction and compare the muon-neutrino flux coming downward through the atmosphere and neutrinos that have travelled $\mathcal{O}12,000$ km through the Earth and are upward-going in the detector. The flavour change of the the upward-going neutrinos (now known to be mainly due to $\nu_\mu \rightarrow \nu_\tau$ oscillations) had become known as the atmospheric neutrino anomaly. For atmospheric analysis, the only mixing seen was between tau and muon flavour neutrinos, so the MSW effect introduces only a phase difference between the two flavours, and does not affect this result. Examples of experiments in this sector include Soudan 2 [24] and MACRO [25].




Figure 2.13: An interaction in the upper atmosphere such as produces atmospheric neutrinos is shown on the left. The neutrinos have a different path-length through the Earth depending on whether they are upward or downward going (right), allowing them to change flavour [26].

The most prominent atmospheric neutrino experiments, however, are Kamiokande and its successor Super-Kamiokande. These are water-Čerenkov detectors that use ring imaging to distinguish elastic scattering (ES) ν_e events from ν_μ , thus measuring neutrino flavour, energy (E) and path-length (L) through the Earth via the event direction. This meant Super-Kamiokande could measure the oscillation probabilities of the muon-neutrinos to other flavours as a function of the neutrino oscillation parameter $\frac{L}{E}$: a dip in this spectrum (figure 2.14) can be used to determine the neutrino oscillation parameters in the atmospheric sector.

K2K [27] is an upgrade to the Super-Kamiokande experiment in which a beam of muon-neutrinos with a mean energy of around 1.3 GeV from KEK in Japan is aimed at a near detector 300 m away and the Super-Kamiokande detector around 250 km away. This allows K2K to make more precise measurement of the $\frac{L}{E}$ oscillation parameter, hence the atmospheric neutrino oscil-




Figure 2.14: Azimuthal angular distribution of the events seen in Super-Kamiokande [19]. The simulated unoscillated results are shown in red, with the data points in black and the fit to the data in green. The discrepancy between data and unoscillated Monte Carlo can be seen clearly in the μ -like results.

lation parameters in the region of “the dip”(figure 2.15). K2K was also able to prove the neutrino oscillation hypothesis to 3.9σ .




Figure 2.15: The K2K results show a difference between the oscillated and unoscillated spectra and the dip in $\frac{I}{E}$ that is a signature of oscillations [27]. The best fit to the data with a continuous spectrum is shown in red, with the actual data points shown in black and the expected unoscillated spectrum in blue.

2.4.5 Reactor

Experiments in this sector look for the disappearance of $\bar{\nu}_e$ from nuclear fission at short distances from nuclear reactors. The data measured so far are consistent with minimal mixing between mass eigenstates 1 and 3 and the current best limit on oscillations in the reactor sector is given by the Chooz experiment [28]. Based at the Chooz nuclear power station in the Ardennes region of France, the two *Gd*-enriched (gadolinium) liquid scintillator detectors were 1 km from the reactor, 117 m apart and with 300 *MWE* (m of water equivalent) overburden. The typical neutrino energies were around 3 MeV and were detected through the reaction

$$\bar{\nu}_e + p \rightarrow e^+ + n \quad (2.80)$$

where the signature is a coincidence of the prompt e^+e^- annihilation and a delayed neutron capture signal at 8 MeV. The Chooz results are summarised by the red line in figure 2.9.

There are many experiments proposed to improve this limit, including an upgrade to Chooz known as Double-Chooz [29], Daya Bay (China) [30] and Braidwood (IL, USA) [31].

2.4.6 Charge-Parity Violating Phase

CP-violation does not modify the probability that a particular flavour neutrino will oscillate into some other flavour, only that oscillation will occur from one specific flavour to another specific flavour. This makes it difficult to detect in most neutrino experiments, which only look for disappearance or have no neutrino/anti-neutrino sensitivity.

There could be a difference in neutrino/anti-neutrino oscillation rates, but if CPT-conservation holds, then the oscillation probability

$$P(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta) = P(\nu_\alpha \rightarrow \nu_\beta) \quad (2.81)$$

No previous measurement has been made in this sector, but MINOS will be able to measure CP-violation by comparing the oscillation parameters of atmospheric anti-neutrinos and neutrinos. MINOS can distinguish between muons and anti-muons using a magnetic field to distinguish the charge sign of muons in CC events. MINOS will shortly publish results measuring atmospheric neutrinos, which are consistent with no CP-violation [32].

2.4.7 Majorana Phases: Dirac & Majorana Neutrinos

These phases are relevant if neutrinos are Majorana particles, meaning that neutrinos and anti-neutrinos are the same particle. The other possibility is that neutrinos are Dirac particles.

Dirac Neutrinos

The Dirac neutrino mass is generated by the Yukawa coupling of left- and right- handed neutrinos to the Higgs boson. It is known that only left-handed neutrinos couple to the W and Z particles, since this is the only helicity of neutrinos seen experimentally. Right-handed anti-neutrinos could be added to the model as the CPT-conjugate of left-handed neutrinos. If neutrinos are massive, then they must travel slower than c (the speed of light in a vacuum), so it is possible to see a right-handed neutrino. Interaction with the right-handed neutrino (N_R) could yield mass with the Lagrangian of the interaction

$$\mathcal{L}_D = m_D (\bar{\nu}_L N_R + \bar{N}_R \nu_L) \quad (2.82)$$

$$= m_D \nu \bar{\nu} \quad (2.83)$$

where $\nu = \nu_L + N_R$ and m_D is the Dirac neutrino mass.

This interaction conserves lepton number by not allowing any $\nu_L \rightarrow N_L^C$ or $\nu_R^C \rightarrow N_R$ mixing, but does violate weak isospin conservation with $\Delta I_W = \frac{1}{2}$. The mass m_D is generated through the Higgs mechanism

$$m_D = h_\nu \frac{\nu}{\sqrt{2}} \quad (2.84)$$

where $h_\nu = \mathcal{O}10^{-11}$ is the interaction strength, compared to $h_e = \mathcal{O}10^{-5}$, the value for electrons. This Higgs interaction term is calculated from the upper bound on the absolute neutrino mass, $m_\nu < \mathcal{O}1 \text{ eV}$ (see section 2.5.3).

Majorana Neutrinos

The Majorana neutrino mass does not require any new degrees of freedom, but instead a new higher dimensional operator (equation 2.91). The interaction Lagrangian is

$$-\mathcal{L}_T = \frac{m_T}{2} (\bar{\nu}_L \nu_R^C + \bar{\nu}_R^C \nu_L) \quad (2.85)$$

$$= \frac{m_T}{2} (\bar{\nu}_L C \nu_L^T + \bar{\nu}_R^T C \nu_R) \quad (2.86)$$

$$= \frac{m_T}{2} \nu \bar{\nu} \quad (2.87)$$

where

$$\nu = \nu^C C \bar{\nu}^T \quad (2.88)$$

This violates weak isospin with a value of $\Delta I_W = 1$. The mass is generated by the vacuum expectation value of a Higgs triplet.

The K-matrix describes the Majorana phases of neutrinos, which is likely to be a basic property of neutrinos (see equation 2.73). The parameters of the K-matrix

$$K = \begin{pmatrix} 1 & & \\ & e^{i\phi_1} & \\ & & e^{i\phi_2} \end{pmatrix} \quad (2.89)$$

are therefore dependent on the neutrino mass model.

2.5 Neutrino Mass Models & Scenarios

In models based on the SM, neutrinos would acquire mass by interaction with the Higgs field, like other massive particles. However, the details of this mass model are not well understood.

2.5.1 Adding Neutrino Mass to the SM Lagrangian

We can add new physics beyond the SM (BSM) by adding the SM Lagrangian to a Lagrangian describing the BSM interactions suppressed by the scale of new physics ($\frac{1}{\Lambda}$), as proposed by Weinberg [33]

$$\mathcal{L}_{BSM} = \mathcal{L}_{SM} + \frac{1}{\Lambda} \mathcal{L}_5 \quad (2.90)$$

where

$$\frac{1}{\Lambda} \mathcal{L}_5 = \frac{1}{\Lambda} (LH)(LH) \quad (2.91)$$

$$\Rightarrow \frac{1}{\Lambda} (L \langle H \rangle)(L \langle H \rangle) \quad (2.92)$$

$$= m_\nu \nu \nu \quad (2.93)$$

where LH is the interaction between left-handed neutrinos and the Higgs field and $\langle H \rangle$ is the expectation value of the Higgs field.

These effects are suppressed by a factor of $\left(\frac{m_\nu}{E_\nu}\right)^2 < 10^{-18}$, so that neutrino oscillation studies physics at high energy scales. With $\Lambda = \mathcal{O}10^{16}$ GeV as the scale of new physics (Grand Unification, or GUT, scale) from the Minimally Super-symmetric Standard Model (MSSM) gauge coupling at 2×10^{16} GeV, this gives a neutrino mass estimate of

$$m_\nu = \frac{\langle H \rangle^2}{\Lambda} = \mathcal{O}3 \times 10^{-3} \text{ eV} \quad (2.94)$$

Dirac and Majorana neutrinos can both be considered as just Weyl spinors, except with an added mass term.

2.5.2 Sterile Neutrinos

The combined analysis of the “invisible Z -width” (the width of the Z -decay resonance into unseen particles) from ALEPH, OPAL, DELPHI and L3 (the experiments at CERN’s LEP) [34] showed that there are 2.984 ± 0.008 flavours of active (left-handed), light ($M_\nu < \frac{M_Z}{2} = \mathcal{O}40.5 \text{ GeV}$) neutrinos to which the Z -boson could decay via

$$Z \rightarrow \nu_l \bar{\nu}_l \quad (2.95)$$

This does not, however, rule-out the possibility of right-handed neutrinos, which don’t couple to the Z , so would be sterile or heavy neutrinos.

There would not be so much debate about the existence of sterile neutrinos if the Liquid Scintillator Neutrino Detector Experiment (LSND) [35] at Los Alamos had not reported a deficit of $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations from stopped muon decay. They predicted that this fourth neutrino flavour would have a mass-squared difference of

$$|\Delta m_{\text{LSND}}^2| > 1 \text{ eV}^2 \quad (2.96)$$

with a small mixing angle.

Subsequent experiments, such as KARMEN [36] at RAL, have tried to confirm the existence of sterile neutrinos, but have disfavoured the LSND result. There is still some parameter space allowed by both experiments, however, and this will be investigated by MiniBooNE at Fermilab [37]. MiniBooNE will be able to prove conclusively ⁴ the existence or otherwise of sterile neutrinos. Oscillation patterns involving four neutrino flavours will be not be discussed in this thesis.

2.5.3 Absolute Neutrino Mass Scale: Experimental Limits

Experiments have also set about measuring the absolute mass of neutrinos from cosmological and nuclear decay experiments.

The WMAP [38] and SDSS [39] experiments, and analysis of the Lyman- α forest, constrain the neutrino mass through measurements of structure formation in the universe, using their property as hot dark matter candidates. Timing measurements of supernovae neutrinos can also provide a mass limit down to $\mathcal{O}20 \text{ eV}$ [3].

Tritium beta-decay experiments, e.g. Mainz [40], Troitsk [41] and KATRIN [42], make precise measurements of the β -electron energy spectrum to determine the neutrino mass. The experiments need good calibration to measure this energy spectrum to the precision necessary, and KATRIN hopes to bring this limit down to $\mathcal{O}0.2 \text{ eV}$.

⁴MiniBooNE can run with neutrinos and anti-neutrinos.

Another popular method is to look for the possibility of neutrinoless double beta-decay from one of the 35 known isotopes from which two-neutrino double beta-decay is allowed. In neutrinoless double-beta decay, the two Majorana neutrinos annihilate each other (helicity flip) allowing a measurement of their mass by looking for a 0ν decay peak at the 2ν endpoint energy. The most famous double beta-decay experiment was the 11 kg ^{75}Ge Heidelberg–Moscow experiment, which had 0.2% energy resolution. The official result was that no neutrinoless double-beta decay had been observed, but a subset of the collaboration published a re-analysis of the results, claiming a peak that would correspond to a neutrino mass of $\mathcal{O}0.4\text{ eV}$ [43].

2.5.4 Hierarchy Scenarios

Assuming three neutrino flavours, there are several possible patterns of mass hierarchy that are consistent with the current data. The convention used is that $m_1 < m_2$ (the two parameters that are responsible for the solar oscillations) since it is known from matter effects that the solar mass-squared difference is

$$\Delta m_{21}^2 = \mathcal{O}8 \times 10^{-5} \text{ eV}^2 \quad (2.97)$$

with a mixing angle of

$$\sin^2 2\theta_{12} = \mathcal{O}0.8 \quad (2.98)$$

Since atmospheric sector (32) mixing takes place in a relative vacuum, the sign of the atmospheric mass-squared difference is unknown, but

$$|\Delta m_{32}^2| = \mathcal{O}2 \times 10^{-3} \text{ eV}^2 \quad (2.99)$$

with maximal mixing

$$\sin^2 2\theta_{23} > 0.92 \text{ @ } 90\% \text{ C.L.} \quad (2.100)$$

The third angle, from the reactor sector (31) is minimal:

$$\sin^2 2\theta_{13} < 0.03 \text{ @ } 90\% \text{ C.L.} \quad (2.101)$$

The possible hierarchy scenarios are:

1. The normal hierarchy: $m_1 \ll m_2 \ll m_3$, so $\Delta m_{23}^2 > 0$. The lightest neutrino mass is unconstrained, but $m_3 \approx \sqrt{|\Delta m_{23}^2|} \approx 0.03\text{--}0.07 \text{ eV}$ and $m_2 \approx 0.008 \text{ eV}$. This scenario is shown on the left in figure 2.16.
2. The inverted hierarchy: $m_1 \approx m_2 \gg m_3$. Solar neutrino oscillations take place between the higher two masses, so we have no knowledge of m_3 , but we can constrain $m_{1,2} \approx \sqrt{|\Delta m_{23}^2|} \approx 0.03\text{--}0.07 \text{ eV}$ and $\Delta m_{23}^2 = m_3^2 - m_2^2 < 0$. This scenario is shown on the right in figure 2.16. This is also sometimes known as the quasi-degenerate case.
3. The degenerate case: the masses have small splittings ($m_1 \approx m_2 \approx m_3$), so the neutrino masses are large with respect to the mass differences.

4. Interpolations of cases 1–3: interpolations can be made from cases 1 and 2 to case 3 by increasing the lightest neutrino mass.




Figure 2.16: Neutrino mass hierarchy in the normal (left) and inverted (right) scenarios with the mixings also shown [6].

2.5.5 Summary of Mass Models

There have been many different mass models proposed to describe the way that neutrinos acquire mass and it is beyond the scope of this chapter to discuss any of them in detail. Here is a brief summary of their categories:

- See-saw mechanism: a light, active neutrino mixes with a heavy ($\mathcal{O}10^{14} \text{ GeV}$) Majorana neutrino — the heavier the right-handed neutrino, the lighter the left handed neutrino. In this mechanism, GUT-scale particles are exchanged to explain the small Majorana neutrino mass. There are several different types of See-saw model that vary depending on the details of the interaction (figure 2.17).




Figure 2.17: Exchange of GUT-scale particles in the see-saw mechanism to explain small left-handed neutrino mass [6].

- R-parity violation ⁵: once this is violated, there are no longer any quantum numbers distinguishing down-type Higgs doublets and lepton doublets. This allows the generation of hierarchical neutrino masses with large mixing angles.
- Super-symmetry breaking: effective Majorana terms for right-handed neutrinos are generated by super-symmetry breaking at the order of the weak scale, so low left-handed neutrino masses are generated by the see-saw mechanism.
- Texture models: these make specific guesses at the form of the 3×3 mass matrix from current experimental constraints.
- Models that require large extra dimensions.

2.6 The Frontiers of Neutrino Physics

This section considers the yet to be answered questions about the nature of neutrinos, and how MINOS should answer some of them.

2.6.1 Unanswered Questions

We can summarise the questions put to the neutrino physics community and which they will seek to answer within the coming decade:

1. Are neutrinos Dirac or Majorana?
2. What is the neutrino absolute mass scale and hierarchy?
3. Is there maximal mixing in the atmospheric sector (θ_{23})?
4. Is there minimal mixing in the reactor sector (θ_{13})?
5. Is there CP violation in the neutrino sector?
6. Can neutrino oscillations with the LMA solution be confirmed with atmospheric sector neutrinos?
7. Is the LSND result correct? Are there sterile neutrinos?
8. Do neutrinos have a magnetic moment? This is theoretically possible for massive neutrinos and would allow for an electro-magnetic component in, for example, neutrino–electron scattering. The magnetic moment is already constrained to be $\mu_\nu < 10^{-10} \mu_B$.

2.6.2 The Need for MINOS: Scientific Potential of the Experiment

The MINOS long baseline oscillation experiment will investigate items 3, 4, 5 and 6. How MINOS will achieve this, and with what limits, is discussed in the following sections.

⁵R-parity is defined as $R = (-1)^{2S+3B+L}$ where S is spin, B is baryon number and L is lepton number. All SM particles have $R = 1$.

Chapter 3

THE MINOS PROJECT

This section describes the Main Injector Neutrino Oscillation Search: its general principles, beam, beamline and detectors.

3.1 General Principles

MINOS sends muon neutrinos, from the NuMI conventional neutrino beam, 735 km through the Earth from the Fermi National Accelerator Laboratory (Fermilab) near Chicago in Illinois to the Soudan mine in Minnesota. MINOS will be able to set the energy of this beam to one of three configurations, covering a range over the $\frac{L}{E}$ region of interest (section 2.3.3). MINOS will measure the neutrino spectrum with a Near Detector (NearDet) at the Fermilab site and will measure the same spectrum at a Far Detector (FarDet) at the Soudan mine, and infer a change of parameters. Both detectors are similar, magnetised, tracking calorimeters optimised to detect muons in the few GeV range.




Figure 3.1: Cross-sectional view of the MINOS baseline from NearDet to FarDet [47].

The neutrino spectrum is compared between NearDet and FarDet to extract neutrino oscillation parameters in the atmospheric and reactor sectors:

- Precise ($\mathcal{O}10\%$) measurement of Δm_{23}^2
- Investigate whether $\nu_\mu \rightarrow \nu_\tau$ mixing is maximal

- rule out alternatives such as decoherence
- study energy dependence of the oscillation
- Find signs of sub-dominant $\nu_\mu \rightarrow \nu_e$ oscillations
 - look for ν_e appearance
- Measurement of CP-violation by comparing ν with $\bar{\nu}$ oscillations
 - study atmospheric neutrinos, using the magnetic field to identify muon charge sign.

MINOS has been measuring atmospheric neutrinos since 2003, and the beam phase of the experiment started in February 2005.

3.2 The NuMI Beam

Protons with an energy of 120 GeV from Fermilab's Main Injector are fired at a graphite target to produce π^+ and K^+ , which are focused by magnetic horns and then decay into ν_μ in the 675 m decay tunnel (figure 3.2). The total distance from target to detector is about 1 km (figure 3.3).




Figure 3.2: Diagram of the NuMI beamline in the target region

The first neutrinos from NuMI were sent to MINOS in January 2005 and the first neutrino events were seen in January at NearDet and March at FarDet. Apart from 3 weeks of beam “down time” in March and April due to a cooling water leak in the target, the beam has been running consistently since. By July, the proton intensity was at 2.1×10^{13} p/p (protons per pulse) with a spill every 3 s, with the horn in the low energy beam configuration (figure 3.4). The beam-line has been tested to as high as 2.5×10^{13} p/p in February, and the goal of NuMI is to work at this intensity with a beam spill every 2 s by the end of 2005. This beam intensity would make the neutrino interaction rate at NearDet around 1.1×10^4 /day.

The NuMI beam is currently running in low energy (LE) beam configuration (figure 3.4), which is the optimal region for measuring atmospheric sector oscillations, as given by Super-Kamiokande [19] and K2K [27]. By the end of 2005, MINOS will have comparable statistics to K2K and so will be able to re-assess the beam energy required. The design of MINOS is optimised for response in the medium energy (ME) configuration.




Figure 3.3: The NuMI beamline between the Main Injector and the Near Detector. The NuMI extraction area of the Main Injector, magnetic horn, muon absorber and Near Detector are shown.




Figure 3.4: Neutrino energy distributions of the NuMI beam in different energy configurations: low, pseudo-medium (S-ME in blue) and pseudo-high (S-HE in green).

3.2.1 Beam Monitoring

The beam is monitored by various hardware and software mechanisms to ensure the necessary stability and to correct any problems. During beam running, a monitoring system is set up in the control room and data is recorded on the intensity and quality of the beam from hardware monitoring systems, to ensure that the neutrino event spectra are properly understood.

3.3 The MINOS Detectors

The MINOS detectors are tracking calorimeters made of steel planes, to provide the mass for particles interactions, and scintillator planes to detect those interactions. As well as the near and far detectors, there was a calibration detector (CalDet) at CERN in Geneva in Switzerland that has now been decommissioned.

3.3.1 Detector Design

Each detector is built from planes comprised of 1 in thick steel and 1 cm thick solid scintillator, so that the thickness of each plane is 6 cm, including an air gap between planes that exists to accommodate warping and variations in the steel thickness. The scintillator is divided into 4.1 cm wide strips, co-extruded with reflective TiO_2 coating, to give spatial resolution in one direction. Alternating planes have strips oriented orthogonal to each other provide both x and y resolution (figure 3.6). Light is collected by wavelength-shifting optical fibres running in a groove along the strips and is read out by Hamamatsu M16 (FarDet) or M64 (NearDet) PMTs (figure 3.5).




Figure 3.5: Various pieces of MINOS apparatus, showing a M16 PMT, a diagram of the Far Detector, of particle interactions in the steel–scintillator sandwich design, of light output from the scintillator and how this is collected by the green wavelength-shifting fibres.



Figure 3.6: The inside of a FarDet scintillator module during construction. The strips can be seen running along the length of the module. The scintillator modules are then encased in 0.5 mm thick aluminium and attached to the steel planes.

3.3.2 The Calibration Detector

CalDet was 60 planes long with 24 vertical and horizontal strips in each plane, so it was physically 3.6 m long and 1 m wide (figures 3.7 and 3.8). The detector was placed in the T7 and T11 test beams on the PS ring at CERN in Geneva.

The detector was designed to characterise the response to various particles at different energies (figure 3.9). This is used for the final stage of calibration to get the absolute energy from an interaction from the light level and particle identification. The absolute calibration and the need for CalDet are described in section 5.9.

The detector was also useful as a prototype to test MINOS hardware and software, and to develop protocols for installing and running the other detectors.

3.3.3 The Near Detector

The NearDet is 282 planes or 17.5 m long with different numbers of strips in each plane. These strips are oriented at 45° to vertical in U and V rather than x and y planes, in order to allow measurement of atmospheric events and calibration with cosmic ray muons (the pathlengths of vertical muons will be distributed the same in U and V strip orientations), and to ease the installation of the readout system at the lower end of the strip (vertical strips would require readout underneath the detector). The detector has a 1.5 T magnetic field, centred at the coil hole (figure 3.10), offset from the 0.5 m diameter beam spot. The total mass of NearDet is 0.96 kton.

In NearDet, only every 5th plane is fully instrumented, i.e. 96 strips across. In the forward 120 planes of the detector (the calorimeter) the planes between full planes are partially instrumented, i.e. have 64 strips. In the back 160 planes (the spectrometer) there is no readout between fully instrumented planes. The purpose of the spectrometer is to measure the momenta of high-energy




Figure 3.7: A photograph of CalDet during operation in the CERN test beamline. The vertical and horizontal strip orientations can be seen.




Figure 3.8: Representation of the Calibration Detector in MINOS software, showing orientation and strip-number starting position in the U-view (left) and V-view (right) [49].




Figure 3.9: The response of CalDet to e^- , π^+ , μ^- and p at 2 GeV. These responses can be characterised to identify different reactions at the NearDet and FarDet.




Figure 3.10: A photograph of NearDet taken shortly after construction. The coil hole is offset from the beam spot; the latter is at the centre of the scintillator.

beam muons from curvature in the magnetic field. Figures 3.11 and 3.12 show this diagrammatically.




Figure 3.11: The forward calorimeter section (left) and backward spectrometer section (right) have different scintillator layouts, as described in the text.

The detector has 19 ns timing resolution between planes to separate events. In order to prevent beam muons being mis-tagged as neutrino events, the planes in the first 0.5 m comprise a “veto region” and those in the next 1 m the “target region” for neutrino events; together these comprise the trigger region.

Because of the high neutrino flux from the NuMI beam, beam spills at NearDet have multiple neutrino events (figure 3.13). These events are classified as a “snarl”, defined on the basis of timing, and “event slicers” are used to split up the events in offline reconstruction software.

There have been problems in NearDet with “ghost events”, a phenomenon that sees low energy neutrino showers being reconstructed in the detector soon after real events have died away. The problem is caused by residual space charge in the PMT after an event that has deposited large amounts of light into a PMT pixel, i.e. gaseous heavy ions moving slowly between the dynodes of the PMT. Offline software can be used to remove these events during reconstruction, based on event energy and topology.

3.3.4 The Far Detector

FarDet is 485 planes long, and each plane has 192 strips. The detector is 8 m across and 31 m long, and it has a mass of 5.4 kton. The magnetic coil produces a field of 1.5 T and runs through




Figure 3.12: Representation of the Near Detector in MINOS software, showing orientation and strip-number starting position for full and partial planes in the U-view (left) and V-view (right)




Figure 3.13: A NearDet high energy beam spill with 8 events. The high energy muon tracks can be seen ranging-out into the spectrometer section. Showers from high energy events can be reconstructed as “ghosts”.

the centre of the detector (figure 3.14).



Figure 3.14: The face of FarDet, showing the last steel plane where the neutrino beam exits. The veto shield can be seen covering the top of the detector. [Photo: Jerry Meier.]

The strips of FarDet are also oriented at 45° to vertical (figure 3.15). The scintillator strips running across the detector are divided into groups of 20 or 28, encased within modules (figure 3.16). There is also a veto shield with 86% efficiency over the top of the detector to reduce the cosmic ray background.

FarDet sees atmospheric neutrinos at a rate of about 1 every 5 days, as well as beam-induced neutrino events at a rate of at least 1 event per day at nominal beam intensity. The NuMI beam has diverged to around 1 km wide by the time it reaches FarDet (figure 3.1).

A partially contained beam ν_μ event candidate at the Far Detector is shown in figure 3.17, as an example of a neutrino event in the detector. The neutrino has interacted in the rock and the muon has ranged into the detector; it has been classified as a beam event since it is travelling along the beam axis in coincidence with a beam spill. The background of stopping muons are distinguished from partially-contained atmospheric neutrino events using timing.

The detector has been operational since 2003 and has been taking atmospheric neutrino and cosmic muon events since then. FarDet has an angular resolution of 1° to cosmic muons and since around 10^7 cosmic muons have been detected, FarDet has been able to observe the shadow of the moon in the cosmic muon spectrum (figure 3.18). When only high-momentum muons are selected, since these curve less in the magnetic field, the dip can be seen due to lunar attenuation in the cosmic ray direction spectrum (figure 3.19).




Figure 3.15: Representation of the Far Detector in MINOS software, showing orientation and strip-number starting position in the U-view (left) and V-view (right)




Figure 3.16: The number of strips in each FarDet module, using the U-plane orientation as an example.




Figure 3.17: A FarDet partially contained beam neutrino event from an online event display.




Figure 3.18: A shadow is formed by the pattern of high energy cosmic rays absorbed by the moon.




Figure 3.19: The cosmic muon direction spectrum with respect to the moon shows the lunar absorption.

3.3.5 Detector Monitoring

A monitoring system and a detector control system run online during data taking at both detectors. These systems measure and record hall temperatures, electronics singles rates (single photoelectron noise), information from the calibration systems (discussed in section 5) and a myriad of other data about the runs being taken. These systems are used to compliment the calibration systems by allowing any problems seen during data taking to be corrected, and to validate data integrity offline.

Chapter 4

DATA ANALYSIS AT MINOS

This chapter describes the way that neutrinos events are analysed in MINOS and how the neutrino oscillation parameters can be extracted from the data.

4.1 Introduction

The collaboration’s first published results, with atmospheric neutrinos at FarDet, are discussed at the beginning of this chapter. These data can then be compared to the beam data that is the main focus of MINOS.

In this thesis, the first seven months of MINOS beam results are analysed using a simple $\mathcal{T}$ -test method to separate events into different physics channels: neutral current and charged current. Neutral current events can be equally of all neutrino flavours, whereas charged current events come mainly from muon neutrino interactions. The $\mathcal{T}$ -test uses this characteristic to measure the difference in the $\nu_\mu:\nu_\tau$ ratio at each detector. This method is relatively insensitive to many systematic errors and does not require good energy calibration.

Beam events reconstructed by the MINOS collaboration are separated into two data sets based on an energy algorithm, which masks any spectral distortions in the energy spectrum that could occur due to oscillations. This allows the collaboration to perform a “blind analysis” in order to add confidence to the final results. The analysis shown here is not blind, and combines results from both the blinded and un-blinded data sets.

An improved method of beam analysis is then expounded, which investigates the change in the charged current spectrum as a function of energy, in a similar way to Super-Kamiokande and K2K (section 2.4.4). This method produces more precise results, but it is more sensitive to systematic errors and requires accurate energy calibration.

4.2 Atmospheric Analysis

FarDet is 2070 *MWE* deep, giving good shielding from the cosmic ray muon background, and it is large enough to measure a significant atmospheric neutrino event rate. Since August 2003, FarDet has collected 420 days of atmospheric neutrino data, equivalent to 6.18 *kT.y*. Atmospheric

data taking will continue between spills during beam running. These data have been advertised and have been submitted for publication.

Between March 2003 and March 2005, a total of 107 fully contained and partially contained events have been seen, 77 of which have been reconstructed with good information about the neutrino direction. An upward-going muon analysis has been performed using neutrino interactions in the surrounding rocks from atmospheric neutrinos that have travelled through the Earth. MINOS looked at the ratio, R , between upward and downward going events. The ratio of the R ratios between data and non-oscillation MC was found to be

$$\frac{R_{\frac{up}{down}}^{data}}{R_{\frac{up}{down}}^{no\ osc}} = 0.62 \pm 0.14 \text{ (stat.)} \pm 0.02 \text{ (sys.)} \quad (4.1)$$

The data are consistent with the oscillation parameters measured by Super-Kamiokande [19], K2K [27], MACRO [25] and Soudan 2 [24], although the statistics are much poorer. The oscillation parameter fit is shown in figure 4.1.




Figure 4.1: The MINOS parameter space shows the parameters measured with 77 neutrinos from the MINOS atmospheric neutrino analysis. The results are consistent with the other experiments shown. [A. Mann]

FarDet is unique as a large underground neutrino detector in that it is magnetised, so can measure charge separation between muon events in the detector with 99% purity. This means MINOS can tell whether a neutrino event is a neutrino or anti-neutrino — MINOS measured 34 ν and 18 $\bar{\nu}$ (figure 4.2). Thus MINOS can make a measurement of the CP violating parameter δ by looking for differences between ν_μ and $\bar{\nu}_\mu$ event rates and hence oscillations. I.e., a ratio, R , is calculated between anti-neutrino event rates and all neutrino event rates

$$R = \frac{\bar{\nu}_\mu}{\bar{\nu}_\mu + \nu_\mu} \quad (4.2)$$

MINOS measured the ratio of R between data and MC with no CP violation to be

$$\frac{R_{\text{data}}}{R_{\text{MC}}} = 0.98 \pm 0.19 (\text{stat.}) \pm 0.06 (\text{sys.}) \quad (4.3)$$

which is consistent with no CP violation.




Figure 4.2: Upward- and downward-going ν (left) and $\bar{\nu}$ (right) atmospheric events at FarDet. The oscillation parameters do not differ between the two samples, with the current statistics.

4.3 $\mathcal{T}$ -test Analysis Theory

For the purposes of a simple beam data analysis, we can disregard some aspects of neutrino oscillations and concentrate only on those most easily visible in MINOS. This section discusses the oscillation probability for two-flavour maximal mixing (i.e. atmospheric sector) applied to MINOS and the $\mathcal{T}$ -test that can be used to distinguish different neutrino flavours and measure the neutrino oscillations.

4.3.1 Oscillation Probability

It is known from the Chooz θ_{13} result [28] that $\nu_\mu \rightarrow \nu_\tau$ is the dominant oscillation, so we can use two-flavour mixing and derive an oscillation probability for this. We use the defined neutrino mass eigenstates $\nu_{2,3}$ for consistency, which oscillate with respect to the flavour eigenstates $\nu_{\mu,\tau}$ via the oscillation matrix (equation 2.73). The two-flavour oscillation probability (from equation 2.69) is

$$P_{\nu_\mu \rightarrow \nu_\tau} \approx \sin^2 2\theta_{23} \sin^2 \left(\frac{\Delta m_{23}^2 x}{4E} \right) \quad (4.4)$$

In the context of MINOS, the length L is 735 km, and by converting to appropriate units (c.f. equation 2.70) the measurable oscillation probability is

$$P_{\nu_\mu \rightarrow \nu_\tau} \approx \sin^2 2\theta_{23} \sin^2 \left(933 \frac{\Delta m_{23}^2 (\text{eV})}{E (\text{GeV})} \right) \quad (4.5)$$

For the purpose of this analysis, we are expecting mixing mainly between tau and muon flavour neutrinos, so the MSW effect (section 2.4.3) introduces only a phase difference between the two flavours, and does not affect this result ¹.

4.3.2 Event Classification

Events in MINOS can be classified in terms of neutral current (NC) and charged current (CC) events, i.e. in terms of the underlying weak interaction. The difference between the two events is in whether the interaction involves the exchange of $W^\pm$ (CC) or Z^0 (NC) boson, i.e.

$$\text{NC:} \quad \nu_l + N \rightarrow \nu_l + X \quad (4.6)$$

$$\text{CC:} \quad \nu_l + N \rightarrow l^- + X \quad (4.7)$$

where N is a nucleon, X is a hadronic final state and l is any specific lepton flavour (e, μ, τ).

This is a useful separation, since all flavours couple equally with the NC interaction, which gives a measure of the total neutrino flux. The vast majority of CC events, however, will be from ν_μ , which gives a measure of the flavour of the neutrinos seen at each detector.

Experimentally, the difference between NC and CC events can be measured by looking for short and long events in the detectors (figure 4.3). Long events have a track due to the muon in the final state, short events have a shower. Short events are comprised of NC events and short CC events, the latter having an electron or low energy muon in their final state. Long events are ν_μ CC events, with some small contamination from a fraction of ν_τ CC events above the 3.5 GeV energy threshold where the τ decays leptonically to a μ . (The τ -decay modes are 17.84(6)% to electronic final states, 17.36(6)% to muonic final states and the rest to hadronic showers [5], so around 83% of τ -events will be short.)

The $\mathcal{T}$ -test is the ratio of ratios of events in each classification between the two detectors (near and far). The test distinguishes between short and long events and measures the ratio between them for each detector. The formula used for the $\mathcal{T}$ -test is

$$\mathcal{T} = \frac{\left(\frac{N_S}{N_L}\right)_{\text{near}}}{\left(\frac{N_S}{N_L}\right)_{\text{far}}} \quad (4.8)$$

where N_S are short events and N_L long events.

The advantage of using a method such as this for analysis is that there is no sensitivity to different rates (fluxes) at each detector, since the test uses ratios. The method is also relatively insensitive to energy calibration, since only the topology of events is used, rather than any measure of the light deposited by each interaction.

The main disadvantage is that the method has poor discrimination between lepton flavours, and the uncertainty will be dominated by the low number of NC events in the short sample at FarDet. The method will also have systematic errors from energy-related systematics in the beam cross-section (NearDet and FarDet see different areas of the beam cross-section), event selection criteria, data to Monte Carlo simulation (MC) disagreement and near/far event systematics.

¹Even when mixing into electron flavour is introduced, the small mass density of the Earth (compared to the Sun, for example) does not modify the effective oscillation parameters greatly, so the MSW effect will not affect the validity of the maximal mixing analysis.




Figure 4.3: Long and short events in MINOS from NC and CC interactions with different neutrino flavours.

4.4 T -test Analysis Method

4.4.1 Data Samples

Five different data and MC samples were available for this analysis. There are MC samples for each detector, with an additional oscillated MC sample for FarDet. The data samples come from the NuMI beam data, and are April–October for NearDet and March–October for FarDet, with beam monitoring data. The samples used are:

- NearDet MC: a simulated sample of neutrino events;
- FarDet Un-oscillated MC: a sample of neutrino events with no oscillation, so that systematic effects in the reconstruction (e.g. event energies) can be quantified between NearDet and FarDet. This sample can also be oscillated with fixed parameters and compared with data to find the best fit (and confidence limits) of the oscillations;
- FarDet Tau Event MC: ν_μ events from the un-oscillated MC sample can be replaced with equivalent ν_τ events from this sample in order to mimic oscillations;
- FarDet Oscillated “mock data challenge” (MDC) MC: an oscillated data sample, such as we would expect for real data. The neutrino oscillation parameters used for this sample are $\Delta m_{23}^2 = 2.1233 \times 10^{-3} \text{ eV}^{-2}$ and $\theta_{23} = 0.60958$. This sample can also be used to determine how well we can measure the parameters with this method;
- NearDet Data: the NearDet in-spill data sample of neutrino-like events;
- FarDet Data: both the “blinded” and “un-blinded” data are combined for this analysis, making it the first full analysis of NuMI beam data with MINOS.

4.4.2 Sample Selection: Cuts & Systematics

There are two types of cut that need to be made in this analysis: cuts to ensure the quality of the data and their freedom from systematic biases and cuts made as part of the selection criteria for different types of events — these second type of cuts will be discussed in section 4.4.6.

In this analysis, cuts of the first type include:

- number of planes triggered by the interaction: at least 3 planes must be read out in order for the event to be classified as a neutrino event. This helps to remove background events from radioactive decays in the detector halls;
 - proper reconstruction: this ensures that the reconstruction software is confident that there are no errors in its de-multiplexing², vertex finding, etc.;
- timing: the event must start within the NuMI spill time. This cut, along with other procedures applied during reconstruction, filters out most noise and also “ghost events”³ at NearDet;

²At FarDet and in the spectrometer section of NearDet, more than one strip-end readout fibre is read out by the same piece of electronics, which is a design known as “multiplexing”. “De-multiplexing” refers to the process of reconstructing which physical strip an electronic event corresponds to.

³Ghost events have been seen to occur at NearDet. After an authentic neutrino event, particularly a high energy shower, has died away, PMT afterpulsing in the channels hit by the event causes another event to be reconstructed as a neutrino event with the same topology as the first, but with lower energy.

- fiducial volume: this ensures that the event is sufficiently contained within the detector (see section 4.4.3);
- proximity to LI flashes: at FarDet, a LI flash can coincide with the beam, and this could potentially cause the event to be reconstructed as a long shower. In order to prevent this, a timing cut is placed around to ensure that no LI triggers coincided with the event (see section 4.4.4);
- event energy: events with low pulse height have a high probability of being ghost events that need to be cut from the sample (see figure 4.4). This also removes some of the lowest energy neutrinos. A maximum event energy cut is also placed on the samples, to ensure the samples used are the same between each detector. The event energy is only known approximately, and the approximate range used is 1–25 GeV at each detector. This introduces only a weak dependence on the energy calibration.




Figure 4.4: The sub-5 GeV event energy spectra of data and MC at NearDet. The shapes of the distributions agree above about 1 GeV, but there are significantly more sub-GeV events reconstructed in data than MC.

4.4.3 Fiducial Volume

The fiducial volume cut ensures that data is collected from a region far enough inside the detector that there are minimal systematics resulting from different vertex locations. Neutrino-induced rock muon events, for example, could be counted in the CC sample; but NC events from the same vertex location and same energy could not be measured, because the events would not reach the

detector. Also important is that long events must be contained long enough inside the measurement region of the detector to be classified as “long”, if they leave the detector before having travelled 50 planes, they would not be placed in the long sample.

To ensure that this target region is the same for all event topologies, the cuts placed on FarDet events are:

- the event vertex must be at least 10 planes (0.6 m) inside the detector from the beam direction, to ensure that the event really began inside the detector;
- the event vertex must allow the event at least 50 planes (3 m) to travel before leaving the end of the detector, so that the event can be classified as “long”;
- the event vertex must be within a cylindrical volume of radius 3 m: since the planes are hexagonal and 8 m across, there needs to be a cut to ensure that the vertex is within the detector, and that it is not so close to the edge of the detector for an event that would range over 50 planes and coming from the beam direction to range-out before it has travelled those 50 planes.

The effect of these cuts on the FarDet sample is shown in figure 4.5.




Figure 4.5: The effect of the fiducial volume cut on event length distributions at FarDet, showing the patterns before and after. The ratio of data removed from each sample is similar for NC and CC.

For NearDet, the cuts are made differently, but the same logic applies: the events measured at NearDet must look the same as those that would be measured at FarDet if there were no oscillations. MC can be used to ensure that this is the case, and to make a correction if it is not.

One particularly important systematic is the difference in the beam cross-section measurable by each detector: if the same beam radius that can be seen at FarDet were seen at NearDet, the cut would be at 4 mm around the beam axis. This is too tight to get enough data, even from the large NearDet sample, and a cut this tight is probably unnecessary. Instead, a cut of 0.25 m is placed around the centre of the beam axis, in both the x - and y -views (figure 4.6).




Figure 4.6: Neutrino event vertices in NearDet in the x - and y - views. The bulk of events have vertices in the partially-instrumented region of the detector.

The z -direction cut ensures that events are sufficiently inside the detector, so that the vertex farthest upstream in the beam-line is at 10 planes, i.e. 0.6 m. The downstream cut is at 4.9 m, so that short events (<35 planes) are fully contained within the calorimeter section, whilst long events can range out in the spectrometer section.

The effect of these cuts in the NearDet sample is to reduce the number of events in the sample dramatically (figure 4.7), but since the flux at NearDet is at 10^4 neutrino-induced events per day (several orders of magnitude higher than the FarDet flux), systematic effects at NearDet are more important than statistics.

4.4.4 Event Type Identification

In order to perform the analysis, long and short events are classified by their length in the detector. The number of strips hit per plane can also be used for events classified as short in order to get better discrimination between NC and CC events in the short sample. The event length cut is made first, and the strips per plane cut applied to the short event length sample.

The cut below which events are classified as short is at 36 planes. There is then an ambiguous region between 36 and 50 planes, events above 50 planes being classified as long. Any events in




Figure 4.7: The effect of the fiducial volume cut on event length distributions at NearDet, showing the patterns before and after. The cut is stringent so that only the central part of the beam is used, to reduce systematics. NB: The merlons on the long end of the CC sample are from the 5-plane resolution of the spectrometer section.

the short sample with more than 6 hits per plane are classified as ambiguous. An initial study with loose cuts allows the locations of these cuts to be found approximately. These event classification criteria are optimised and justified later in this section.

Figures 4.8, 4.9 and 4.10 show typical examples of the different event classifications at the FarDet, and figures 4.11, 4.12 and 4.13 the same at NearDet. The plots are taken from an event display running on MC samples.

Since short CC events and NC events are similar, there is no way to separate them in the short sample, but by measuring the ratio of short to long CC events and CC to NC events from MC, the NC:CC ratio can be calculated.




Figure 4.8: A typical long CC event at FarDet seen in a MC event display in various different views. Red denotes hits that are part of a track, yellow events are part of a shower, green events are coincident noise hits and blue events are de-multiplexing uncertainties. Hits considered part of the physics event (snarl) are ringed.




Figure 4.9: A typical short CC event at FarDet seen in a MC event display. The key is as for figure 4.8.




Figure 4.10: A typical NC event at FarDet seen in a MC event display. The key is as for figure 4.8.




Figure 4.11: A typical long CC event at NearDet seen in a MC event display. The key is as for figure 4.8.




Figure 4.12: A typical short CC event at NearDet seen in a MC event display. The key is as for figure 4.8.




Figure 4.13: A typical NC event at NearDet seen in a MC event display. The key is as for figure 4.8.

4.4.5 Optimisation of Selection Cuts

After the hard cuts placed to remove systematics have been made, the samples can be studied to find the best places to make cuts to distinguish the short and long samples from MC. The optimisation is made on the number of NC events in the FarDet sample, as this is likely to be the sample with the fewest entries. In order to attain this, a trade-off must be made between the purity of the NC events in the short sample and the total number of NC events in the short sample (figure 4.14). The uncertainty in the number of NC events in the short sample is the measure of the optimal place to cut on event length, and this is at 36 planes at FarDet. The “shape” of the error on the NC sample is defined as uncertainty of the NC content of the short sample, normalised to fit in the range 0–1.

The same distributions are shown for NearDet in figure 4.15. The optimal place to cut on ND events to maximise the certainty on the NC sample would be around 25 planes, but the same cut is placed on both NearDet and FarDet samples to get a consistent spectrum⁴. The larger amount of data at NearDet means that a poorly placed cut has less effect here than at FarDet, nor does the uncertainty increase greatly between samples cut at 25 and 36 planes.

The other selection criterion used to distinguish NC and CC events is the strips-per-plane cut placed on the sub-36 plane sample. The cut is on the maximum number of strips hit per plane, which is lower in NC showers than CC showers. The purpose of this cut is to further optimise the amount of NC events in the short sample. The number of hits per plane from sub-36 plane NC and CC events are shown for FarDet in figure 4.16 and for NearDet in figure 4.17. Again, the optimisation should be made for the FarDet NC sample, since this has fewest entries in real data. The optimum place to cut is between 6–7 strips per plane at both detectors.

4.4.6 Light Injection Cut on Real Data

The real data has additional data quality cuts that need to be applied on top of those cuts placed on the MC and MDC samples. This accounts for detector systematics and other issues not sufficiently modelled in MC. At NearDet, an additional timing cut was put in place to remove ghost events from the sample. At FarDet, an additional cut is used to remove light injection events from the neutrino candidate sample.

LI calibration runs continuously at FarDet, and LI events can correspond to a time window coincident with a beam pulse. In this case, the LI event gets tagged as a neutrino candidate event. Since these events are often long (40–80 planes), they would contaminate the long sample in the FarDet data and bias the result.

LI events can be removed from the sample in two ways. Firstly, a timing cut can be placed on candidate neutrino event, which insists that the event not be in the same time window as a LI event. Alternatively, the LI events can be removed spatially by making cuts on the number of strips and planes that a LI event would have compared to a neutrino event. The two methods are compared in figure 4.18. The timing cut shows candidate events inside and outside a LI time window, and these two samples are topologically distinct. This implies that either cut could be used to remove LI events and the timing cut was used for the purpose in this analysis.

⁴In fact, the FarDet sample used here includes neutrino oscillations. Since we are expecting neutrino oscillations (see chapter 1), it makes sense to optimise the cuts for an oscillated FarDet sample.




Figure 4.14: Event lengths for NC and CC events in the FarDet MC sample. The bottom plot shows the purity and efficiency of the short sample for including NC events above the CC background. The shape of the error on the NC sample is the uncertainty of the NC content of the short sample, normalised to fit in the range 0–1. This latter parameter is minimised at 36 planes.





Figure 4.15: Event lengths for NC and CC events in the NearDet MC sample. The bottom plot shows the purity and efficiency of the short sample for including NC events above the CC background. The error on the NC sample is minimised at 25 planes.

FarDet Strips Hit Per Plane in Short Sample (MC)



FarDet Strips Hit Per Plane Cut: Efficiency & Purity





Figure 4.16: The number of strips hit per plane in sub-36 plane CC and NC events at FarDet. The purity and efficiency of the cut in including NC events are shown on the bottom plot, as well as a normalised measure of the uncertainty on the NC sample with cuts at different values.

NearDet Strips Hit Per Plane in MC Sample 1



NearDet Strips Hit Per Plane Cut: Efficiency & Purity




Figure 4.17: The number of strips hit per plane in sub-36 plane CC and NC events at NearDet. The purity and efficiency of the cut in including NC events are shown on the bottom plot, as well as a normalised measure of the uncertainty on the NC sample with cuts at different values. Compare with figure 4.16.




Figure 4.18: The effect of making a LI cut based on timing compared to one based on the size and shape of LI events. The dots each represent a neutrino candidate event, before fiducial and energy cuts. Red dots are coincident with a LI flash whereas blue dots have no coincidence: the topology of the events is markedly different. This shows that either timing, or cuts on numbers of strips and planes hit, can be used to cut LI events from the sample.

4.4.7 Beam Energy Systematics

The probability of neutrino oscillations is dependent on the $\frac{L}{E}$ parameter, which, in terms of this analysis, is a function of the neutrino energy (equation 2.72). The approximate event energies is shown for data at NearDet and MC in both detectors in figure 4.19, above the minimum energy cut at 1 GeV. The shape of the spectra are qualitatively the same, with the similar means and spreads, so that the average neutrino energy can be considered the same at both detectors, as necessary. The surfeit of low energy events in the NearDet data sample of figure 4.19 are caused by "ghost events" (see section 4.4.2).




Figure 4.19: The super- GeV reconstructed energy spectra for neutrino events in different samples. The event energies are approximate and are based on summed event pulse heights, without a validation of the calibration. In each case the distributions are qualitatively similar, about 9% of the distribution is above the 25 GeV limit, the mean energy is at about 6.5 GeV, with an energy spread of about 5.5 GeV. The means of the NearDet data and MC distributions differ by about 5%. The approximate energy spectrum of FarDet data neutrino events is similar.

In order to perform event separation in an unbiased way, it must be shown that event lengths are not dependent on neutrino energy. If this were not the case, an event length cut would also be a cut on neutrino energy, so the two samples would be dominated by different oscillation probabilities. Figure 4.20 shows true energy of the neutrino in FarDet events as a function of event length (note: true neutrino energy and not reconstructed event energy). The cuts are placed at below 36 planes for the short sample and above 50 planes for the long sample. The NC events have a roughly linear increase of neutrino energy with event length, but this is not important since there are few NC events outside the short sample. More important is that the neutrino energies of CC events in the short and long samples are similar, i.e. that there is no obvious difference in neutrino energies

below 36 planes and above 50 planes. This seems to be the case, with the profile for each at about 10 GeV on average over the sample range.




Figure 4.20: True neutrino event energy profiles at different event lengths from FarDet MC. There is a linear relationship of event energy with event length for the NC sample, but the CC sample is flatter.

4.4.8 Expected Number of Events

Since the FarDet data sample has fewest entries, it is here that the expected number of entries is most important. Of the 75,977 events from the “blind” sample, there are 69,313 candidate events in the sample after the 3 plane minimum cut. Once the LI timing cut is applied, only 1,746 candidate events remain; and a further 47 were removed by the LI topology cut. A beam quality cut was applied to all real data, after which 1,437 events remained. When the fiducial volume cut (section 4.4.2) was made, only 391 events remained, and a further 98 were removed by the energy cuts (section 4.4.3). This left 293 neutrino events in the FarDet data sample (figure 4.21).

This number is consistent with the number of protons on target used for the data samples (around 1×10^{20} POT, figure 4.22), compared with 25×10^{20} POT for the MDC. The LI and beam quality cuts do not apply the MDC sample, but the effect of the cuts is the same for both samples.




Figure 4.21: The event length distribution of real data events at FarDet. In all 213 events passed all the cuts from around 1×10^{20} protons on target when both LI cuts were used.




Figure 4.22: Integrated intensity and horn configuration (beam energy) of the NuMI beam in 2005. About 0.8×10^{20} POT, mainly in the low energy horn configuration, were used for the FarDet and NearDet data in this analysis.

4.5 $\mathcal{T}$ -test Results

The data samples are processed by applying the cuts and selection criteria described in section 4.4. For each data sample, the number of short and long events are compared to make a parameter

$$\mathcal{T} = \frac{\frac{\text{Short}}{\text{Long}}(F)}{\frac{\text{Short}}{\text{Long}}(N)} \quad (4.9)$$

that can be used as a comparable measure of the ratio of events in each sample. The results of the analysis are summarised in table 4.1.

Data Sample	Entries	Short	Long	$\frac{\text{Short}}{\text{Long}}$	$\mathcal{T}$
NearDet MC	16,974	7,317	7,781	0.940(15)	
NearDet Data	70,482	29,176	33,776	0.8638(69)	
FarDet Unosc. MC	74,175	32,913	31,858	1.0331(81)	1.099(18)
FarDet MDC	8,943	3,520	4,014	0.877(26)	0.933(34)
FarDet Data	293	144	112	1.286(162)	1.49(18)

Table 4.1: The results of the different data samples analysed using short and long events. The $\mathcal{T}$ values compare FarDet MC to NearDet MC and FarDet data to NearDet data.

A measurement can be made of the discrepancy between the un-oscillated and the data by using the “ratio of ratios”:

$$\mathcal{R} = \frac{\mathcal{T}_{\text{data}}}{\mathcal{T}_{\text{MC}}} \quad (4.10)$$

where the measured ratio is $\mathcal{T}_{\text{data}}$ and the expected ratio from the un-oscillated MC is $\mathcal{T}_{\text{MC}}$. The ratio $\mathcal{R}$ would be 1 in the case of no oscillations, and significant deviation from this proves that there has been a distortion to the event length spectrum.

For beam data, the value is $\mathcal{R} = 1.36 \pm 0.16$ with a chi-square value of 6.13. This proves that neutrinos have changed from one flavour to another in real data to over 98.4% certainty, implying neutrinos in the beam have oscillated between the near and far detectors (assuming none of the more exotic explanations, such as described in [6]).

4.5.1 Oscillation Parameter Spaces

These oscillations can be measured and represented on a parameter space, showing the best fit of the Δm^2 and $\sin^2(2\theta)$ parameters and the confidence limits on the result. This also allows the results of this analysis to be compared with other experiments.

The parameters are measured by taking a sample of un-oscillated FarDet MC, which contains the same proportion of neutrino flavours as seen by NearDet, i.e. almost exclusively ν_μ . Each point on the parameter space has oscillation parameters Δm^2 and $\sin^2(2\theta)$, and the un-oscillated events are read in. These events are then randomly changed from ν_μ to ν_τ according to their energy dependent oscillation probability for that particular point on the parameter space (equation 7.2). The events that have been tagged as having changed flavour are replaced in the sample by a ν_τ event from the tau-neutrino file (section 4.4.1).

The ratio of short to long events in the simulated oscillation sample is calculated, and the T -value of the simulated FarDet sample is calculated by comparing it to the NearDet MC sample. The chi-square value between the calculated T -value and the T -value of the data is plotted on the parameter space, and contours connect the lines of chi-square that correspond to integer values of sigma in two-dimensions ($1\sigma = 2.3$, $2\sigma = 4.6$, $3\sigma = 9.2$). The method was tested by oscillating the MC with certain parameters, and checking that any discrepancy between the best fit and the inputted values was merely statistical.

The confidence limits in parameter space for the mock data challenge are shown in figure 4.23, with the input parameter shown to be within the 1σ confidence region. This proves that the analysis method described can measure neutrino oscillation parameters with reasonable accuracy, so that the systematic errors are not greater than the statistical errors with 25×10^{20} POT.

Confidence Limits in Parameter Space for MDC




Figure 4.23: Confidence limits in parameter space for the MDC sample. The input oscillation parameters to the mock data is shown, along with the 1, 2 and 3 sigma confidence limits. This shows the viability of the method, which includes the real result to within 1 sigma from 25×10^{20} POT. The chi-square value of the best fit region is around 1.

The confidence limits in parameter space for data are shown in figure 4.24. Table 4.1 shows

that, although both the data and MDC samples show oscillations near maximal mixing compared to the MC, these oscillations are fundamentally different between the two samples. In the MDC sample, the ratio of short to long events decreases, whilst in data it increases. Therefore, whilst the MDC challenge includes the input parameter and the best fit regions of figure 4.1 within its 90% confidence limit, the real data excludes these.

The fit to the data, although it has an overly good chi-square value, is poor in excluding regions of parameter space. Apart from excluding to 3σ the world average, the fit has poor exclusion of small $\sin^2(2\theta)$ values to 90% and cannot pin-down any confidence on Δm^2 to over two orders of magnitude.

Confidence Limits in Parameter Space for Data




Figure 4.24: Confidence limits in parameter space for data. The best fit oscillation parameter is shown, along with the 1 (turquoise), 2 (green) and 3 (red) sigma confidence limits. The areas enclosed by the red line is similar in this figure and figure 4.23, except the contour is exclusive in this figure and inclusive in the MDC fit. The chi-square value on the best fit is $\sim 10^{-5}$.

The results shown in this section are the work of an attempt to perform an independent analysis of the MINOS beam data using the data, simulations and reconstruction algorithms available at the time of analysis. Therefore the numbers presented here should not be considered as official

MINOS results and the methods used should not be considered as necessarily being similar to those used for official MINOS results. Furthermore, no part of this work should be taken as representative of the current state of data quantity, detector simulation or reconstruction efficiency of the MINOS project.

By the time of publication, an official review and analysis of the current MINOS beam data had been published at <http://www-numi.fnal.gov/talks/results06.html>.

4.5.2 Discussion

The difference in the long:short event ratios for data and MC show a distinct difference, implying neutrino oscillations with over 98.4% certainty. The expected oscillation parameters were not measured, however, probably due to a discrepancy between aspects of the data and MC. More work should be done on understanding the relationship between data and MC before this result can be trusted.

The data fit may be hampered by poor neutrino energy calibration in the MC. Although the long:short ratio is relatively independent of the event energy spectrum for CC events (figure 4.20), a systematic error in the energy tuning of the MC with respect for data will affect the oscillation probabilities. In such a case, the MC will model a particular set of oscillation parameters to have a wrong T -value, and this will skew the shape of the chi-square contours. This would not be the case for the MDC, which uses the same energy tuning as the MC, and so would not be seen in tests of the method and MC with MDC. This may be an explanation why the results of MC and MDC differ from that of real data.

It is therefore clear that more work must be put into:

- the correct neutrino event energy tuning of the MC;
- understanding the reconstruction related dependence of the oscillation fit to the MC;
- an alternative method, using an analysis technique based on energy bins.

These will yield a more reliable result. The MINOS analysis of the charged current spectrum is an example of such a method, whilst also being capable of higher precision than the T -test method described here.

4.6 Charged Current Spectrum Analysis

The beam analysis in MINOS compares the un-oscillated NearDet energy spectrum with the oscillated FarDet energy spectrum as a function of neutrino energy (figure 4.25). The NearDet spectrum is extrapolated to FarDet using Monte Carlo simulations of the different oscillation parameters, and a Feldman–Cousins analysis predicts the best fit near the physical boundaries in parameter space with the smaller number of entries in each energy bin. Good knowledge of parent hadron production, neutrino interaction physics and the energy calibration of detectors is necessary to reduce systematics in this analysis.

The oscillation parameters can be measured to 30% over the 3-year run plan using muon energy alone, with the limiting factor being the uncertainty in the fraction of neutrino energy




Figure 4.25: Comparison of neutrino energy spectra from CC interactions between NearDet and FarDet with oscillations between the two given by the parameters specified.

transferred to the muon in the 1–10 GeV range. On average, the muon has around 65% of the observable energy and the hadronic component around 30% with the rest being electromagnetic. Thus, a better measurement, i.e. to 10%, requires the summation of energy in the muon track (E_μ) and in the showers (E_{shw}) to find the total visible neutrino energy

$$E_\nu \approx E_\mu + E_{\text{shw}} \quad (4.11)$$

The difference between the two spectra is plotted as a function of neutrino energy for different numbers of protons on target on the left in figure 4.26, where the plots assume the Super-Kamiokande best fit parameters: $\Delta m_{23}^2 = 2.5 \text{ meV}^2$ and $\sin^2 2\theta_{23} = 1$ [19]. This is a powerful test of the oscillation hypothesis. The position of the dip in the ratio between the two spectra gives the value Δm_{23}^2 , whilst the size of the dip gives the oscillation parameter $\sin^2 2\theta_{23}$ (shown in blue on the top left plot).

The plots on the right in figure 4.26 show MINOS sensitivity in atmospheric parameter space with each number of protons. In particular, the bottom right plot shows MINOS' sensitivity from 3 years of running. It can be seen that MINOS limit improves significantly on that of Super-Kamiokande in this case.

Using event topology and energy deposition, MINOS can distinguish between hadronic and electromagnetic showers. MINOS can also measure the oscillation parameters in the reactor sector by looking for ν_e appearance in the ν_μ beam. If the parameters are not close to the Chooz limit [28], MINOS can halve this limit towards minimal mixing (figure 4.27).

Absolute calibration is important in this analysis for determining the position of the dip, and relative calibration is important for the comparison of the energy spectra between the detectors. MINOS aims to calibrate to 5% absolute and 2% relative uncertainty, and this is discussed in the next chapter. A good understanding of the beam is also necessary for this precision as FarDet sees a smaller cross-section of the beam than NearDet, and this difference in the beam profile measured could cause systematic errors in the analysis.




Figure 4.26: MINOS sensitivity to the atmospheric neutrino oscillation parameters from different numbers of protons on target. The plots on the left show the difference between the oscillated and unoscillated spectra as a function of neutrino energy. The position of the dips shows the parameter Δm_{23}^2 , whilst the depth (in blue) shows the magnitude of θ_{23} . The plots on the right show MINOS' limits in parameter space from these results.




Figure 4.27: MINOS sensitivity to the reactor sector oscillation parameters with different numbers of protons on target compared to the Chooz limit. The plot assumes minimal mixing, below the MINOS limits.

Chapter 5

CALIBRATION OVERVIEW

Calibration is required to normalise the response across the detectors and to convert light from hadrons, electrons and muons into energy. This section gives an overview of how calibration is done in MINOS. Two calibration stages, strip-to-strip calibration and light injection calibration, will be discussed in more detail in the proceeding sections.

5.1 Calibration Requirements

The MINOS detectors should be calibrated to 2% relative precision between NearDet and FarDet and 5% precision of the absolute energy scale.

With the 5% absolute precision, the systematic uncertainty will be smaller than the statistical uncertainty on the measurement of Δm_{23}^2 with 2.5×10^{21} protons on target (the 3 year timescale when the parameter is accurate to 10%).

The 2% relative calibration precision was selected since a shift in energy scales between the two detectors affects measurement of both oscillation parameters. This precision is necessary to ensure that the systematic effect from such a shift:

- on the best fit parameter is small compared to the statistical error
- does not distort the energy spectrum, so there is not a large chi-squared value on the fit parameters (i.e. any systematic effect on a bin-by-bin basis is less than the statistical error).

The other major sources of uncertainty between NearDet and FarDet are in the neutrino fluxes. The results of the study of the NuMI beam by the Main Injector Particle Production Experiment (MIPP) [62] will reduce this uncertainty to 2%.

5.2 The MINOS Calibration Chain

The MINOS calibration chain has several links required to ensure MINOS can work as a precision experiment (figure 5.1). These eventually convert a raw amount of light read as digitised charge (analogue-to-digital converter or ADC counts) from a PMT into an energy measurement. The various links in the calibration chain are discussed in turn in the proceeding sections.




Figure 5.1: The MINOS calibration chain applies corrections successively

5.3 Charge Injection

The first stage in the calibration chain is the electronics linearisation and pedestal subtraction. Both of these calibrations are done on the hardware level, with no DAQ read-out.

To ensure linear electronics read-out, known amounts of charge are injected into all the electronics channels at the beginning of each run sequence and stored in an onboard “look-up table”. Pedestal subtraction is done using the background noise on the PMT readout with the high voltage off.

This stage is important, since a pedestal set too low would swamp the DAQ with noise and one set too high would fail to readout data. Charge injection also linearises the PIN photo-diode scale, which are used to linearise the PMTs in the light injection calibration [63].

5.4 Gain Correction

The gain-corrected link is only used for intermediate calibrations. It converts the charge collected in ADC counts back to the value in photo-electrons seen at the PMT photo-cathode using the pre-measured PMT gain. It is added here for completeness and does not need to be discussed further [64],[65].

5.5 Light Injection

This stage of the calibration uses LEDs to inject light into each strip of the detector. This amount of light is recorded by PIN diodes and the PMT response to the light is measured. This done to fulfill the light injection system’s three main requirements:

- to correct for the non-linearity of PMT response with incident light (known as “gain curves”), monthly;
- to correct for the change of gain over time between gain curves (known as “drift points”), about every 3 hours;
- to measure absolute PMT gain.

The LI system is discussed in more detail in the next chapter.

5.6 Strip-to-Strip Calibration

This stage of the calibration uses cosmic ray muons as a standard candle to normalise the response of all strip-ends within each detector. The main variations corrected for are:

- differences in clear fibre lengths;
- differences in read-out fibre pigtails;
- differences in PMT gain;
- differences in scintillator response;

- differences in read-out fibre collection and connector efficiencies.

This calibration is discussed in detail for the FarDet in chapter 7 and NearDet in chapter 8. Once this calibration is applied, there should be no more calibrations that can be applied to a hit without knowing where along a strip a hit occurred.

5.7 Attenuation Correction

The attenuation correction takes account of light lost in the green readout fibres within the scintillator module. Therefore it is necessary to know where along a strip a hit occurred, so a track-fitter or shower finder must be applied to a hit before this correction can be applied. The light readout can then be corrected to what the equivalent amount of light seen would have been had the hit occurred at the centre of the strip.

The attenuation corrections were measured using a radioactive source to measure the double-exponential (i.e. two decay modes) attenuation curve along each end of each strip. The mapper corrections were then parameterised.




Figure 5.2: The mapper attenuation corrections applied to a sample of cosmic muons, for a particular strip orientation and read-out side, against position along the strip. Most strips have mapper correction parameterisations along the red band. High correction values, away from the swathe at the far end of the strip, are a sign of a damaged read-out fibre within the strip.

Figure 5.2 shows the mapper corrections applied to hits in V-strips at FarDet seen from the East end from a sample of cosmic muons. The response is normalised to 1 at the centre. Most mapper responses follow the red band; but some fibres have a much lower response at the far end, for example, if they are broken. These lower response strip-ends can be seen as indigo lines on the left of the plot.

5.8 Detector-to-Detector Calibration

The calibration between detectors is done by comparing the response of stopping muons in each detector. An example stopping cosmic ray muon at FarDet is shown in figure 5.4.




Figure 5.3: The energy lost by muons in the scintillator as a function of muon momentum at CalDet. The theoretical Bethe-Bloch curve, full MC and data samples are compared [66].

The muons' energy is known from range and the amount of light deposited in each plane is known from the Bethe-Bloch curve (figure 5.3). The response of the muons in a track window starting several planes from the end of the track is characterised as the minimum ionising particle (MIP) response for each detector. This value is also known as the muon energy unit (MEU) and is a value in ADC counts. The relative value of the MEU in the NearDet and FarDet compared to CalDet is used as the relative calibration [66].

5.9 Absolute Calibration

The absolute calibration is the final stage of the calibration in which light (now in MEU) is converted to energy in GeV. Data from particle interactions measured at CalDet (see chapter 2) is compared to the data at NearDet and FarDet. It was necessary to use a separate detector for absolute calibration, since the strip width and plane thickness makes the spatial resolution too poor to measure hadronic response by reconstructing the invariant π^0 mass, and because the detector event live time is too short to reconstruct electron response from stopping muon decay.

At CalDet, the responses to different particles at different known energies were measured, and this was used to tune the MC simulation. The different energy depositions of electrons [67] and hadrons [68],[15] compared to muons, and their different event topologies, identify the different particles and allows calorimetry with real data.




Figure 5.4: A stopping muon at the far detector, such as would be used for inter-detector calibration. Hits in the cosmic ray “veto shield” can also be seen.

5.10 Other Calibrations

Other calibrations are carried out on the MINOS hardware that are not part of the main calibration chain, but provide constants used in the latter stages of event reconstruction.

5.10.1 Timing Calibration

The timing calibration uses cosmic ray muons at FarDet to measure the timing difference of signals readout from different strip-ends in the detector from a muon track. A measurement of the “time walk” of each strip-end in the detector and veto shield allows MINOS to measure muon track direction through timing.

This measurement is useful in beam analysis to ensure that tracks and showers come from neutrinos travelling along the beam axis. It is also needed in atmospheric analysis to determine whether a track began inside the detector (neutrino) or outside the detector (stopping cosmic ray muon) [69].

5.10.2 Magnetic Field Calibration

The values of the magnetic field must be known and gaussing/degaussing must be performed so that the field inside the detector is well understood. Good knowledge of the magnetic field is necessary for proper momentum reconstruction of muons, especially from long CC events. Any known variations in the field over space and time can be used to correct the data, and can also be used to improve MC simulations. Work on the magnetic field is continuing on MINOS.

5.10.3 Detector Geometry

A record is kept of various other aspects of the detector, such as plane separation, steel thickness and strip alignment. It is useful to have this information in order to make MC simulations as realistic as possible and in order to reconstruct events accurately.

Chapter 6

THE LIGHT INJECTION CALIBRATION SYSTEM

This chapter discusses the MINOS Light Injection System [70],[71],[72],[73],[74],[65], what it does, how it is used and some of the improvements that have been made. The section will start with the concept of LI calibration, i.e. its *raison d'être*. Next, the hardware of the LI system will be discussed, with some particular attention being put on the LEDs, which are the source of light in the system. The section will then move on to discussing the two modes of LI: “drift” calibration and “gain curves”. Finally, use of the LI system as a debugging tool is investigated, particularly in the context of NearDet.

6.1 LI Calibration Concept

The MINOS light injection system is used for measuring differences and monitoring changes of the photo-multiplier tubes (PMTs) and readout electronics at all three MINOS detectors (Calibration, Far and Near), as mentioned in previous sections. There are three main requirements for the light injection system:

- to measure the non-linearity of PMT response with incident light (known as “gain curves”);
- to monitor the change of gain over time between gain curves (known as “drift points”);
- absolute gain measurement.

The LI system is also a very effective debugging tool. Since known relative amounts of light can be shone onto any optical read-out channel in the detector, readout holes and broken channels can be found and fixed. It was used extensively for this purpose during the construction and early operation of all three detectors [75],[76]. The LI corrections are applied to data after electronic non-linearity corrections have been applied using charge injection.

In order to measure the gain of the PMTs as a function of incident light level, known amounts of light illuminate the PMT face via optical fibres from light emitting diodes (LEDs), as shown

in figure 6.1. Twenty points are taken over the dynamic range of the PMT and the relative measurement of light from the LEDs is measured by PIN photodiodes ¹. The resulting PMT vs. PIN photodiode response curve can be parameterised to allow the charge fall off with respect to incident light at the high end of the scale. This allows for the measurement of the saturation at high light levels. These curves are currently taken monthly at the Far and Near Detectors (FarDet and NearDet) and were also taken during normal running at the Calibration Detector (CalDet).




Figure 6.1: A schematic diagram of the LI hardware. Light from LEDs, monitored by PIN diodes, is shone onto many fibres. Those fibres in turn illuminate a number of readout fibres, which are readout by the PMTs. Since the light levels are known, this can be used to calibrate the PMTs and readout system [65].

It is necessary to track the read-out response between gain curves, so light is shone at a default light level point every 3 hours. This corrects mainly for diurnal temperature changes and drifts of gain throughout the month.

The light injection system therefore acts effectively as a link in the calibration chain (figure 5.1) since it deals with all the corrections that can be performed without dealing with any real data readout, i.e. everything downstream from the scintillator in the readout chain.

¹p-type, intrinsic, n-type diode

6.2 Equipment

The LI system uses special hardware connected to the detector. The first piece of hardware to discuss is the pulser box, and an example from NearDet is shown in figure 6.2 (the FarDet pulser boxes look similar). The pulser box hold the electronics that control the LEDs and the LEDs themselves. The optical fan-out at the back of the pulser box contains optical fibres and the cone assemblies that surround the LEDs.




Figure 6.2: An example of a LI pulser box from NearDet. The figure shows the electronics compartment of the pulser box: the control card can be seen in the middle, the power supply at the top and the RS-232 cable for communicating with the DAQ is in green on the left. Harder to see, but still visible, are the driver cards, mounted horizontally on shelves, and it is to these that the legs of the LEDs are attached.

Such a cone assembly is shown in diagrammatically in figure 6.3, with a photograph of the cone and LI fibre “pigtail” shown in figure 6.4. The pigtail connects into the wide end of the cone, with the “snout” that carries the light to the PIN photodiodes facing the LED. The many fibres in the pigtail were routed to allow each LED to contribute light to LI fibres on different bundles.

In each LED fibre bundle, each fibre carries light routed from a different LED. The optical connection between pigtails and LI fibres is made at the back of the pulser box (figure 6.5). Each LI fibre bundle feeds one plane with light and each LI fibre illuminates one set of 8–10 strips. LI fibre bundles at NearDet, where there are 6 fibres in a bundle for a partial, are shown in figure 6.6.

The LI fibres illuminate the read-out fibres in the modules in a construction known as an “ashtray” (figure 6.7). Light from each LI fibre illuminates 8–10 consecutive strips by shining over them in the ashtray. The routing is done so that no LED illuminates more than one group of strips in any one plane. A further constraint is that at FarDet and the spectrometer section of NearDet, which are multiplexed, strips which are read out by the same readout channel cannot be illuminated by the same LED. Thus gains can be calculated and drifts monitored for each




Figure 6.3: Diagram showing the cone assembly in which the LED shines on the LI fibres running to the PMTs and PIN photodiodes [65].



Figure 6.4: A photograph of a prototype cone and LI fibre “pigtail”. The legs of the LED can be seen poking out of the back of the cone. The snout of the pigtail connector fits into the open end of the cone.



Figure 6.5: LI fibres connecting to the back of a NearDet pulser box. Black tape and RTV were used to prevent light leaks.



Figure 6.6: Sets of LI fibres at NearDet, where the fibres were connected to the pulser boxes before installation to the planes (this was done the other way around during FarDet installation). Orange rubber protects the optical output end of the fibres, which shines on the readout fibres.

individual strip-end by knowing which LED was firing when a pixel was read out.

This knowledge of which LED is flashing is controlled by the “LI Master” process running on the DAQ. This process tells the pulser box to flash a particular LED with particular intensity settings and also performs online summary making to compress the data from 1,000 pulses into a mean and RMS, to reduce disk storage.

There are three online modi operandi for the pulser boxes. When there is no beam, the onboard controller card flashes the LEDs and simultaneously flashes a trigger PMT (tPMT). Hits from PMT channels that were flashed coincident with the tPMT hit are recorded as “correlated” hits and are used for the calibration (uncorrelated hits are considered to be noise once the system is properly configured). When the beam is running, an inhibit stops the pulser box from flashing at in-spill times (at NearDet). Vice-versa, the third mode works by telling the pulser box to flash at a particular time, but this is rarely used. The tPMT is still needed in the latter two cases.

6.3 Test And Selection of LEDs Used In Calibration

This section describes the selection of LEDs for use in the LI system from three possible contenders: blue HP LEDs, UV Bivar LEDs and UV Nichia LEDs (figure 6.8).

The system was originally designed to use the blue HP LEDs, but a non-linearity discovered at CalDet proved them unsuitable for the purpose (figure 6.9). The non-linearity was between the light measured from clear fibres by the PIN photodiodes and that measured by PMTs, from secondary light produced in the green wavelength-shifting fibres.



Figure 6.7: Prototypes of the ashtray and readout fibre pigtails in a FarDet module. The LI fibre connects into the ashtray (circled in red) and illuminates the readout fibres on the module (green) [65].




Figure 6.8: The Nichia UV LED in metal casing is shown on the left. This casing and the flat face of the LED could cause mechanical difficulties. The blue HP LED is shown on the right. The Bivar UV LED has the same shape as the HP LED.




Figure 6.9: The blue HP LEDs that the system was originally designed for proved unsuitable because of the non-linearity between light collected through clear fibres, seen as blue, and light collected through wavelength shifting fibres, seen as green. This graph is taken using a LED and data from CalDet, showing all channels (hence the scatter due to the multiple M16 value for each PIN value). The graph is normalised to 1 at the highest output on each channel. Against this line is plotted a line showing true linearity. This graph also shows integrated spectrometer data (blue squares), which fits with the non-linearity curve from CalDet.

6.3.1 Pulse Spectrum & Linearity

The cause of this non-linearity was postulated to be a spectrum shift of the light output by the blue LEDs between low and high output intensities. UV LEDs were considered as a replacement since they only have a very slight spectrum shift between a pulse height of 100 and 1000, and what change can be seen does not result in a broadening of the spectrum (figure 6.10). The spectrum shift alone does not explain the non-linearity, however [74].




Figure 6.10: Normalised output spectra from the blue and UV LEDs. “UV” are the Nichia LEDs, “blue” the HP LEDs and “bivar” the Bivar LEDs. Compared to the relative output from the blue LEDs’ between pulse heights of 100 and 1000, the UV spectra have a negligible change.

Figure 6.11 shows the spectra for the UV and blue LEDs superimposed onto the graph of absorption length for the wavelength-shifting (WLS) fibre. This line shows the length of fibre needed to absorb light of a particular wavelength and shift it into green light to travel down the fibre. The low edge on the blue spectrum, which broadens at high currents, sits on a rapidly-changing edge in the absorption spectrum: this is the underlying cause of the non-linearity. The absorption length is changing rapidly (over two orders of magnitude) over the range at which the LED spectrum broadens.

The WLS fibre is much more efficient at absorbing light at UV wavelengths than blue. This means that there will be less bouncing around of fibre in the ashtray (where light is absorbed due to imperfect reflection) in the UV case than for the blue LEDs, so the amount of green light is boosted when UV is used.

This assertion is confirmed by figure 6.12, on which it can be seen on the spectrum of blue light (the PIN diode output) that there is a distinct spectrum shift to the left as current through the LED increases, causing a small relative change in the green light output. This change causes the feature shown in the inset: a difference in WLS fibre output spectrum shape between low and high LED pulse heights.

The UV LEDs have neither the same spectrum shift as the blue LEDs, nor is the WLS fibre




Figure 6.11: The normalised LED output spectra of figure 6.10 superimposed onto the graph of wave-length shifting fibre attenuation length (in green). The blue LED spectrum shifts over the range where the attenuation length rises by two orders of magnitude.




Figure 6.12: LED and wavelength-shifting fibre spectra at high and low LED intensity. This graph has been normalised so that the blue spectra of high and low current both have equal area. The blue light output spectrum at pulse height 1000 is obviously wider and shifted to shorter wavelengths compared to the spectrum at pulse height 100. The green spectra have been normalised using the same relative factors. The inset shows the difference between the WLS spectra. The interesting feature of this plot is the peak at around 490 nm. Although the data points on this graph look sporadic (because the low intensity curve is at the sensitivity limit of the spectrometer), this feature can still be recognised as corresponding to the difference between the spectra shown as blue and yellow on the main chart.




Figure 6.13: The pulse shape of the HP blue LED measured by an oscilloscope. There is a long drop-off tail to the pulse.

as sensitive to the shifts at UV wavelengths. Therefore switching to one of the UV LEDs (Nichia or Bivar) would solve the non-linearity problem, but it must be shown that those LEDs meet other requirements, such as intensity, pulse shape, pulse-to-pulse stability and lifetime.

6.3.2 Pulse Shape & Stability

The shape of the pulse is important since it must be intense enough to be measured by the PMT. The shorter the pulse, the better it simulates real scintillator response. The pulse shapes of the blue HP and UV Bivar LEDs were measured by an oscilloscope (figures 6.13 and 6.14). Both pulses are square shaped and have a short (~ 10 ns) rise time, but the UV LED also has a 10 ns drop-off (maybe dominated by capacitance in the readout), whilst the blue LED's tail is twice as long.

The pulse-to-pulse stability is a measure of the variation of output intensity from the same input current (set by the pulser box) and is important since it is necessary to have a deterministic light output from the LED for given input settings. This measure for the three types of LEDs tested is shown in figure 6.15, where it can be seen that the response from all the LEDs is sufficiently stable.

A measure also needs to be made of the absolute light output of the UV LEDs to ensure that they are sufficiently bright to illuminate the PMTs across their optical range and perform the calibration, as the blue LEDs originally used were able to do. The Bivar UV LED is shown in figure 6.16 compared to the HP blue LED, and compares favourably with it, having more light seen by the PMT. This can be explained by figure 6.11, which shows that the WLS fibre absorbs light from UV wavelengths more efficiently than blue wavelengths. An advantage of this is that the UV LEDs can be run at a shorter “pulse width” (15 ns compared to 35 ns for the blue LEDs — see figures 6.13 and 6.14), so the pulse is even shorter and simulates the scintillator response better. A disadvantage of this, though, is that there is relatively less light to the PIN photodiodes.




Figure 6.14: Pulse shape of Bivar UV LED measured by an oscilloscope at maximum width. The pulse is squarer and has a shorter tail than the HP blue LED shown in figure 6.13.




Figure 6.15: Pulse-to-pulse stability of the different LEDs: HP, Nichia and Bivar (labelled as "new 3mm"), measured with a PIN photodiode. All the LEDs are stable enough to be used.




Figure 6.16: Comparison of light intensity between the blue HP LED and the Bivar UV LED (labelled as “new 3mm”). Much more light is seen at the PMT from the UV LEDs than the blue LEDs (see figure 6.11).

6.3.3 Lifetime

It is also finally necessary to ensure that the LEDs chosen will be able to work without any degradation in performance over the lifetime of the experiment. A set of Bivar LEDs were pulsed 10 million times, equivalent to 10,000 drift points or 1,250 live days, to detect whether there was any change in output with lifetime. Figure 6.17 shows that there were no ill effects and that the output variation over that period is merely a function of temperature.




Figure 6.17: Variation of LED performance over 10 million pulses

6.3.4 Decision to Use Bivar LEDs

The decision was made to switch from blue LEDs to UV LEDs to solve the non-linearity problem. The Bivar UV LEDs were chosen since there were still several difficulties with the Nichia UV LEDs:

- The Nichia UV LEDs had a lower light output than the blue LEDs, so, amongst other upgrades, the circuitboard resistors would have to have been changed on the LED driver boards.
- The Nichia LEDs have a metal casing that would have to be insulated from the pulser box, causing mechanical difficulties during installation (figure 6.8).
- The flat front face of the Nichia LEDs causes problems with the cone geometry since it does not illuminate all fibres on the pigtail optical connector uniformly.
- There were manufacturing complications with Nichia so LEDs in the intensity bin required by MINOS were not in production, so were expensive (around \$5 (US) per LED).

6.4 Light Injection Gain Curves

In order to measure the linearity of the PMTs with incident light, runs known as “gain curves” are taken. LEDs are tuned to illuminate the PMTs with light at 20 points over the full range of the PMT, i.e. from the pedestal at around 20 ADC to the readout maximum at 16,000 ADC. This light is monitored by PIN photodiodes² and the response of the PMT can be measured at different incident light levels (figure 6.18). This response is linear at low light levels, but flattens near the top of the dynamic PMT range due to saturation.

During normal operation, a gain curve is taken monthly for each strip-end. Figure 6.18 in fact shows two gain curves, taken one month apart. The second gain curve is mapped back onto the first by using the ratio of drift point values taken at the time of the respective gain curves. The next section describes how drift is measured by the LI system.

6.5 Light Injection Drift Method

6.5.1 Standard Light Injection Drift Points & Stability

The modus operandi of standard LI drift points is to flash PMTs with 1,000 pulses. The mean value of these flashes is corrected for by the PIN diode mean in order to account for temperature dependent changes in LED brightness. The largest cause of variation in drift point values is from the diurnal temperature variation (figure 6.19), where the peak-to-peak variation is up-to 1%. The drift points nullify this effect and other temporal variations on a strip-end-by-strip-end basis.

The relative value of this mean compared to its value at the time of the last gain curve is then used to correct the response of each strip-end. The statistical uncertainty of this method is 1.0% and has been shown to work at CalDet [63].

This section investigates weaknesses in the standard mean-based approach. These problems affect either the PIN diode response or the amount of light travelling through the LI fibre and

²the PIN photodiode response is linear and the readout electronics is linearised by charge injection




Figure 6.18: An example gain curve showing the original gain curves in red and blue taken one month apart. The blue curve is mapped back onto the original red curve using the ratio of the drift points, so the drift points monitor the change in gain curve mean.




Figure 6.19: Average strip-end mean drift over one month at FarDet. Note the diurnal pattern showing 28 days of data from August 2004. The full range of the diurnal variation is $\mathcal{O}1\%$.

so affect the mean response of the channel, although the gain of the channel does not change. Regimes based on averaging over pixel or PMT responses are considered instead of the standard drift points.

Problems Related To PIN Diodes

The PIN diodes are used to measure the LED light output, which may vary at up to the few percent level. The PIN value is used as the measure of injected light, rather than the arbitrary pulser box “pulse height” setting. Any changes in the mean response must then be due to gain changes of the strip-end readout (PMT and electronics). It sometimes happens that PIN diodes are not read out properly, due to noise, or have a high uncertainty on their mean, due to damaged electronics, cross-talk, etc. If there is no PIN readout for a particular drift point, the PIN value for the previous drift point is used, but if PIN diodes are not reading out the correct value, this can be difficult to monitor and there can be large effects on the drift result.

At FarDet during August 2004, the magnetic coil was turned off and temperatures consequently dropped. Figure 6.20 shows the corresponding change in mean response of $0.5\%/^{\circ}\text{C}$ in the average response of an example pixel. It was considered that the temperature fluctuation led to a small change in LED brightness and a larger change in phototube and electronics response. If this change in response was not picked-up by the PIN photodiodes, then an erroneous (0.4 %) drift point value would be recorded. The problem is that the PIN photodiode readout was also affected by changes in the environment when the magnetic field was switched-off, so this mean change was seen as a genuine drift.




Figure 6.20: The mean response of pixel 2 on PMT 0 in mux box 2100 during August and September 2004. This corresponds to work on the detector. The PIN diodes corresponding to these channels were not reading out correctly, so a response dip of 4% can be seen in the mean over this time.

Figure 6.21 shows the predominant pattern of problems correlating to PIN diodes in indigo on a strip vs. plane map. The bad PIN diodes led to a “LED-shaped” pattern (a pattern of strips and planes corresponding to the same LED and hence the same PIN) of erroneous drift values. These

problems can be associated with a bad PIN diode value.




Figure 6.21: A strip vs. plane map of the East side of FarDet showing where the drift has changed from the initial value by more than 2%. “LED-shaped” patterns can be seen (indigo), implying that problems seen in these channels are due to PIN-diodes not monitoring LED response changes correctly. The remaining patterns mostly correspond to light injection modules (red).

When known PIN diode problems are removed, there is some scatter of problem strip-ends (shown in red in figure 6.21), defined here as having drifts that diverge by more than 2% from the average drift of their respective pixel more than 8 times in that month. However, a residual pattern of errors remains: groups of 8-10 neighbouring strips showing a discrepancy. These correspond to light injection fibres that illuminate the readout fibres.

Problems Relating To Light Injection Fibre

The patterns of 8–10 adjacent strips correspond to LI fibres, which run from the pulser box to illuminate the read-out fibres from the scintillator strips. Hence, the other possible cause of problems for the standard drift method is that any changes to the light transmission properties in the LI fibre can not be corrected for by any monitoring device. Problems can be found by analysing patterns in strip–plane plots, such as figure 6.21, where sets of adjacent problematic strips imply a LI fibre problem.

Over a two month period between August and September 2004, a sudden change of 5% in the response of some of the strip-ends was observed. This corresponds to a period at the FarDet when cables were being moved, loosened and re-tied in order to reduce tension on the optical fibres. Figure 6.22 shows an example of a strip that had its light injection fibre loosened in this time — the response can be seen to jump sharply at a particular date. The plot is normalised to 1 over the course of the period, so that it starts around 0.98 and drags the pixel-average down, then jumps to 1.02 and pulls that average up.

Figure 6.23 shows what happened to all the pixels on the same PMT at that time. The red




Figure 6.22: Example of a strip-end showing anomalous behaviour (red) in pixel 5 of PMT 1 in mux box 3203. This plot shows the PIN-corrected strip-end means read-out by this pixel, with each of the strip-ends represented by a different colour. The response anomaly of the strip-end shown in red is related to a LI fibre problem.

fibres, which correspond to the same LI fibre, clearly show a coherent change in mean consistent with the checkerboard wiring of strip-ends to PMT pixels (designed to allow multiplexing and reduce cross-talk [65]).

Further evidence that this shift is caused by a difference in incident light, and is not a systematic of the PMT or readout electronics, can be seen in figure 6.24. Rather than measuring the response change using the mean of the pulses, photo-electron statistics³ can be used to calculate the gain [77]. Figure 6.24 shows this gain for those channels over the same period, averaged over 8 drift point runs to improve statistics and remove diurnal fluctuation, and normalised to the starting value rather than the average, so that the lines can be seen separately. The gains are not sensitive to LI fibre light output and since no large shifts in response are evident, the response difference must be upstream of the PMT.

6.5.2 Investigation Into Systematic Variations

Nominal Gain Dependence

There is a small dependence of the relative drift measured by both mean and gain with respect to the nominal gain of the strip-end. Channels with a higher nominal gain drift less than strip-ends with a lower nominal gain, shown for FarDet in figure 6.25. The gradient of the line is $\mathcal{O} - 5 \times 10^{-5} \text{ ADC/pe}$. This allows a small correction to be made if averaging over a whole PMT with higher gain strip-ends given a slightly smaller drift compared to the average.

When this nominal gain correction is applied, there is no longer any dependency of strip-ends drift dependent on nominal gain (figure 6.26). This correction narrows the distribution of PMT mean drift from a high precision drift point run by 3.2% (figure 6.27), showing that a systematic improvement has been made.

³Gain measured by $(\frac{\bar{x}}{\sigma})$, as described later




Figure 6.23: As figure 6.22, but showing all the pixels showing all the stripends on PMT 1 of mux box 3203. The pattern of the anomalous strip behaviour (red) is shown repeating on the other pixels shone on by the same LI fibre (checkerboard pattern).




Figure 6.24: As figure 6.23, but this shows the gain for the strip-ends on PMT 1 of mux box 3203 instead of the mean. The gains of each strip-end (by $\frac{\bar{g}}{\sigma}$) are offset by the uncertainty at the first drift point (5.1%) and are average over 1 day (8 drift points brings the uncertainty down to 1.8% so that 5% changes of gain are obvious). There are no jumps in the gain, so the anomalies are due to a change of light level arriving on the PMT face.




Figure 6.25: How the discrepancy of strip-end drift to PMT averaged drift varies with nominal strip-end gain at FarDet.

There is also a nominal gain dependence on drifts in gain measured with photo-electron statistics, shown here for CalDet in figure 6.28. The gradient of the line is of the same order of magnitude, but with poorer statistics.

Pixel-spot Dependence

It is conceivable that there could be a difference in the agreement between strip-end and PMT-averaged drifts with respect to the strip-end's pixel-spot location on the face of the M16 PMT. Figure 6.29 shows a cross-section of the pixel-spot positions on the PMT face showing the summed value of strip-end drift corrected for by PMT-averaged drift. The overall range between indigo and red (i.e. the full scale) is 0.13% and it shows no pattern of pixels or pixel-spots has a significantly higher or lower value compared to the average. Therefore, no correction needs to be made for pixel-spot location.

6.5.3 Coherent Drifts At The Far Detector

Ideas For Improvement

There are several conceivable methods to improve the way that light injection drift points are taken if it can be shown that the majority of systematic gain changes seen by strip-ends are also seen by the other strip-ends the same pixels or PMTs. In this case, the average pixel or PMT response can be used instead of the individual strip-end response. This also brings in the possibility of using gains calculated using photo-electron statistics to measure the drift, instead of the mean. That is, the possibilities are to average over pixel or PMT and to use either PIN-regulated mean or gain by photo-electron statistics.

Each method has its benefits and draw-backs in terms of number of LED flashes needed, size of




Figure 6.26: There is no systematic variation of discrepancy between strip-end drift and PMT drift after a correction has been applied.

the offline PULSERDRIFT database files ⁴, dependence on PIN photodiodes (and other hardware), need for new online and offline software, detector dead time (where no physics events can be read out) and the statistical precision of the method. The possible options for taking LI drift points are summarised in the table below:

Option	Online Load	Offline Size	PIN Ind.	Dead Time	Stat. Prec.
Standard LI	1	1	No	0.1%	1%
Pixel mean avg	0.15	0.125	No	0.01%	1%
PMT mean avg	0.15	0.008	No	0.01%	0.2%
Pixel gain avg	6	0.125	Yes	0.5%	1%
PMT gain avg (1)	1	0.008	Yes	0.1%	1%
PMT gain avg (2)	0.3	0.008	Yes	0.02%	1%

Table 6.1: Possible drift points measurement schemes compared. The columns denote the possible drift point regime with its relative load on the detector in terms of the amount of data that needs to be collected, size of database files (stored offline), dependence on PIN-photodiodes, detector “dead time” (time that the detector cannot be used for taking other data) as a fraction of total potential detector “live time” and the statistical precision of the method.

Drift Within Pixel

The first possible averaging method is over strip-ends multiplexed into the same pixel. If there is no coherent drift pattern within the pixel over the period of about one month, then the dominant variation is in the individual strip-end response and no averaging is possible. Figure 6.30 shows the variation of strip-ends in all 16 pixels of an example PMT over a period of 200 drift points. The uncertainty of each drift point is 1.0%, which accounts for the nominal spread, and sinusoidal variations with a period of one day can also be understood as correcting for diurnal temperature variations, as discussed earlier.

⁴PULSERDRIFT is the name of the database file used to apply the LI drift correction. When a value is needed for




Figure 6.27: The distribution of PMT drifts between 2 high precision drift point runs at FarDet is normalised to 1 by applying the dependence on nominal strip-end gain.

In this case, there does appear to be coherent drift, but at a similar or smaller scale than the statistical uncertainty. At the general case, over the whole detector, the spread of the fit to the pixel-corrected strip-end drift distribution over a selection of this range is 0.92%, and has relatively few strip-ends more than 3% from the mean (figure 6.31).

When the distribution of strip-end drift is compared with that of the pixel corrected strip-end drift, the latter distribution can be seen to be noticeably narrower: figure 6.32. This is expected as the non-pixel corrected drift is the convolution of the statistical uncertainty of the strip-end mean with the systematic changes in strip-end response. The pixel corrected drift, on the other hand, reflects only the statistical variations (strip-end and pixel). This implies that the spread of 1.10% is well understood just from knowledge of the uncertainty of the corrected strip-end mean from 1,000 LED flashes (1.0%) which predicts a spread of 1.06% when the uncertainty on pixel mean is taken into account.

This proves the effectiveness of pixel averaging as a viable alternative to measuring the response of each individual strip-end. The systematic variations and strip-end uncertainties would be removed by this method and replaced by only the statistical uncertainties in the average pixel mean drift, $\mathcal{O}0.3\%$. An even better averaging method, however, would be to average over the whole PMT.

Drift Within PMT

Firstly, it can be seen from, for example, figure 6.30, and in other figures, that pixel drift changes and PMT drift changes are highly correlated. This is expected as the major sources of variation in response of individual pixels and whole PMTs should be the same.

This means that it is now sensible to bypass pixel-averaging for strip-ends and move directly each strip-end at each detector every 3 hours, this file becomes unmanageably large.




Figure 6.28: The discrepancy of strip-end gain drift (from photo-electron statistics) corrected by their PMT-average, with respect to nominal strip-end gain at CalDet.

to PMT-averaging. Figure 6.34 shows the effect of PMT-averaged drift being used to correct strip-end drift, where figure 6.30 is the variation of strip-end drifts within an example PMT. Figure 6.34 shows the agreement between the two drifts over time and shows that there are no gross systematic changes in the shape of the agreement over time, even on a log scale.

Figure 6.35 shows additional error in PMT-averaging compared to pixel-averaging (figure 6.32), which is due to differences in pixel response of pixels on the same PMT. The important point is whether the drifts of the strip-ends and PMTs are effectively the same, i.e. whether this 0.37% higher uncertainty varies with drift. This will effect the viability of PMT-averaging as a possible drift method.

The alternative method of drift correction uses gains instead of means to calculate drift (see table 1). The same histogram of agreement over time can be drawn, using gains calculated with photo-electron statistics instead of drift-point mean, as figure 6.36. The statistical uncertainty of the strip-end gains calculated with this method is higher, so the spread of the distribution is much wider. Again, no systematic variations can be seen using this method over the same period of about 4 weeks.

A cross-section of figure 6.36 was taken near the end of the month of drift points, and a Gaussian fitted for comparison of statistical errors (figure 6.37). The width of the fit to the peak is 6.4%, in agreement with the expected width from the statistical uncertainties in the gain calculation from the strip-end and PMT, but the spread is wider in the tails of the distribution. This will be studied in more detail in the next section.




Figure 6.29: The detector averaged discrepancy in strip-end drift with respect to its PMT average. The variation is shown as a function of pixel-spot on a full spectrum scale. There are 8 spots on each pixel and pixel boundaries are shown with black lines. No discernable patterns can be seen in terms of a difference between, for example, corner and centre pixel-spots.

6.5.4 Gain Change Averaging Drift Method

In order to see whether the pixel or PMT gain drift is a viable alternative to standard LI stripend mean drift, it is necessary to see whether it works over a drift range of 2–3%, as this is the maximum range expected during normal running at the NearDet and FarDet. In order to reduce reliance on extra hardware, in particular PIN photodiodes (that have been seen to sometimes cause problems in averaging methods) and on other light yield effects (such as in LI fibres) the better method of improving drift point measurements is by using average calculated gain. In particular, since pixel-averaging gain requires a lot more data to be taken in order to get a calibration precise to 1.0%, the preferable method is to use PMT gain drift averaging. The two options for taking LI drift points in this method are summarised in the table below:

Option	Online Load	Offline Size	PIN Ind.	Dead Time	Stat. Prec.
Standard LI	1	1	No	0.1%	1%
PMT gain avg (1)	1	0.008	Yes	0.1%	1%
PMT gain avg (2)	0.3	0.008	Yes	0.02%	1%

Table 6.2: Comparison between standard LI and the two possible online methods of PMT averaging. The columns are as defined in Table 6.1.

The effective differences in how the gain-averaging is done are:

Standard LI this is the current drift point method. For standard LI, the PIN-corrected mean response, d , of the strip-end is used to calculate the drift. The average mean drift on the abscissa of figures 6.39, 6.38 and 6.40 are therefore given by:




Figure 6.30: The drift of the individual strip-ends in PMT 1 of mux box 3203 are shown after cuts and corrections have been made for effects understood to cause anomalies. The stripends follow the drift of the others in the pixel well over 2 months.




Figure 6.31: The projection of strip-end mean corrected for by pixel mean. The distribution follows a normal distribution drop-off qualitatively well for two orders of magnitude from the peak at the high end and 4 orders of magnitude at the low end.

$$\frac{\sum_{i=n}^{i=0} d}{n} \quad (6.1)$$

where n is the number of stripends in the detector and i is a counter over those strip-ends.

Gain averaging the average relative gain change, g , is given similarly by:

$$\frac{\sum_{i=n}^{i=0} g}{n} \quad (6.2)$$

in the ordinates of figure 6.38.

PMT gain averaging for the PMT gain averaging method, the factor for the averaging is:

$$\frac{\sum_{p=q}^{p=0} g}{q} \quad (6.3)$$

where q is the number of strip-ends in the same PMT and p is a counter over them. So the formula for the line labelled “PMT gain” in figure 6.39 and the ordinates of figure 6.40 is:

$$\frac{\sum_{i=n}^{i=0} \frac{\sum_{p=q}^{p=0} g}{q}}{n} \quad (6.4)$$

The relative gain change of the PMT can be used for each strip-end in that PMT, and this can be compressed offline to one entry per PMT rather than one entry per strip-end. In order




Figure 6.32: When the pixel-average is used to correct the strip-end drift, the distribution narrows. This implies that the pixel-averaged drift accounts for most of the same variations observed in the strip-end drift.

to show that PMT average gain drift is a viable alternative to standard LI data and corrects all the systematic effects in monitored by the per strip-end drift, data was taken over a 30 hour time period at CalDet. Figure 6.38 shows that corrected strip-end mean drift and strip-end gain drift are exactly proportional to each other with a gradient of unity.

Figure 6.39 shows the variation of standard LI (drift) and PMT average gain variation (PMT gain) over this period. In particular, the two are shown to match each other's systematic diurnal change.

It only remains to show that the standard LI and PMT-averaged drifts are directly proportional. This is shown by figure 6.40, where the average values of PMT-averaged drift and per strip-end drift at each drift point are plotted against each other and yield the required gradient of unity.

Figure 6.41 shows the uncorrected spread of strip-end drift for the same CalDet runs. The shape of the distribution profiled in figure 6.39 can be seen, showing the sinusoidal diurnal response.

When the PMT-averaged drift is used to correct the strip-end drifts on a one-by-one basis, the systematic effect that the light injection is designed to take out is seen to be removed from the distribution figure 6.42 (the distribution widens due to poorer statistics on the PMT average gain drift than on the strip-end mean drift).

This shows that a method of measuring the average PMT gain change can be used as a viable alternative for standard LI. This allows for a reduction in the PULSERDRIFT database table as well as reducing the reliance of the system on extra hardware (as described earlier). The statistics are sufficient at FarDet to use the current standard LI online method to continue to take LI drift points, so that this method can run in parallel with standard LI drift. Since only 40,000 pulses are required




Figure 6.33: The variation of all pixels in an example PMT over the course of 2 months. The pixel means vary coherently, especially with the diurnal temperature changes.

Strip-end Drift Mean Corrected By PMT Drift over Time for the Whole Detector




Figure 6.34: The variation of pixel mean corrected for by PMT mean over the course of 1 month. The mean and spread do not vary over the course of the month.




Figure 6.35: Most of the systematic effects measured variations by the strip-end mean drift are also measured with the PMT averaged drift

Distribution of PMT-Corrected Strip-end Gain over Time for the Whole Detector




Figure 6.36: The variation of strip-end gain corrected by its PMT over the course of 1 month. The mean varies diurnally and spread is dominated by the uncertainty on the strip-end gain measurement. Neither varies over the course of the month.




Figure 6.37: The projection of strip-end relative gain corrected for by PMT mean. The distribution roughly follows a Gaussian shape well out to 10% from the mean.




Figure 6.38: The strip-end mean drift and strip-end gain changes agree with each other within standard errors so the linear relationship between the two has no offset and a gradient of 1.




Figure 6.39: The variation of strip-end mean drift (grey dots with red spread) and PMT average gain change (by $\frac{\bar{x}}{\sigma}$, green dots) are shown against drift point number over a 30 hour period at CalDet during October 2003. The blue spread on the gain represents $\frac{1}{4}$ of the statistical uncertainty. The two follow each other well over the course of the period, apart for some anomalous drift points between drift points 40-50.




Figure 6.40: The average strip-end mean drift and average PMT gain changes also agree. This proves that PMT gain averaging is a viable method for calculating strip-end drift.




Figure 6.41: The spread of drift points over time shows the distinct diurnal pattern due to temperature changes at CalDet. The gaps in the spread are when drift point LI was suspended so that LI gain curves could be taken.




Figure 6.42: The stripend drifts corrected for by their PMT gain averages. The distribution loses its systematic diurnal shape as this is nullified by the PMT averaging. This distribution widens as the uncertainties on the PMT gains are much greater than those of the drift points.

to calculate the PMT gain to a precision of 1.0%, instead of 128,000 when for all the strip-ends on a PMT using standard LI, the online load can be reduced by a factor of 3. This reduces deadtime and data file size, as well as reducing the proportion of overall time that LI flashes for further than the recent increase in LI flashing rates [72].

6.6 Debugging With Light Injection

The light injection system is a very powerful tool for checking the health of the MINOS detectors. For example, during commissioning it can show-up such problems as cabling errors and afterwards it can find faulty equipment. The “NearLI” package was used for this purpose during NearDet commissioning.

6.6.1 NearLI Package

NearLI uses functions specific to the physical and readout geometry of NearDet to find problems with individual strip responses in the detector and tries to form patterns out of these problems in order to diagnose specific problems. The input for the program is the summary of LI data from the DAQ, described earlier, which is then run through an algorithm to convert raw channel IDs to physical and electronics locations by using the plex (the database that converts electronics addresses in the readout to physical locations in the reconstruction). The output comes in the form of a webpage.

The aim of the package is to give an output that shift crew can use, rather than just LI or electronics experts. The specific diagnoses made the program allow the problems to be fixed faster by determining particular pieces of hardware to be checked. The diagnoses are backed-up by graphical output for every issue, so that if the diagnosis is wrong, the hardware does not need to be checked.

Several criteria are used to flag strips as “bad”, for example, if they have a mean less than 10% of the detector average or if they record fewer than 10% of LED flashes, and a list is made of the bad strips. The pattern finder tries to associate poorly strips with various possibly pathologies by comparing the patterns seen with patterns for those pathologies. The algorithm works in a cascading system, so that pathologies that would be associated with a lot of strip-ends not working, such as an electronics crate not working properly, are checked for first. The order is then electronics rack, various electronics problems (e.g. readout cards), plex problems, miswirings and individual strip or pixel related problems. When a potential cause is found for a bad strip, the strip is removed from the list that is then passed for matching patterns to smaller problems.

Crate Health

The first thing to check for is that all the crates are turned on. Crates could be off because they have been turned off or are not connected, or because a circuit breaker has tripped. Figure 6.43 shows the NearLI output for an example drift point run (3612) from 2004. All the crates are working properly in this run.

Rack Performance

For all crates in use during the run, the program then looks to see whether all the electronics racks or “MASTER” readout cards were working — it is not effectively possible to distinguish which




Figure 6.43: The total number of hits in each crate during an example NearDet LI drift point run is shown in red. The number of correlated hits (i.e. excluding noise) is shown in blue. This shows that all the crates are working.

of the two is not working if there is a problem. The number of hits and number of flashes is compared for each rack, although different racks have a different number of strips read out by them, e.g. between the calorimeter and spectrometer sections. If the program suggests that the rack is bad, this can be cross-checked by looking at the graph shown like figure 6.44, which shows no problems in rack performance from the run.

Plex Problems

The program then looks for problems that could arise with strips of a certain number and orientation being ignored in the plex. The program compares the number of times a correlated hit is seen by a strip to the number of times it was flashed by the LED. The potential problem is that physical numbers and plex numbers have become mixed up in a one-off way, or in a pattern.

Figure 6.45 shows the results in the spectrometer section from an example run for strips numbered 6. All strips with strip number 6 see a healthy ratio of correlated LI hits to flashes from this run.

Dead LED or PIN Photodiode

The program checks that all the LEDs are working properly and that all the PIN photodiodes that monitor them are also working properly. The LEDs are monitored by high-gain and low-gain PIN diodes, so that a problem with either PIN diode or with the LED can be distinguished. There have not yet been any problems with LEDs or low gain PIN diodes at NearDet.

Problems at “MINDER” Level

Bad strips not yet associated with another problem are checked for patterns on the MINDER level (i.e. MASTER channel). Figure 6.46 shows all the channels for an example MASTER from run 3612, showing the number of flashes and the number of correlated LI hits seen, so that any problems diagnosed by the program can be cross-checked.




Figure 6.44: The number of LI pulses (red) and the number of correlated LI hits (blue) for each U and V electronics rack at NearDet from run 3612.




Figure 6.45: The number of LED flashes (red) compared to the number of correlated LED hits (blue) in the spectrometer section of NearDet from run 3612.




Figure 6.46: Numbers of flashes and correlated hits in each MINOS Near Detector Electronics Readout Card (MINDER) in rack 10 from run 3612.

The program also checks for the health of individual MINDER channels, known as “MENUs”, in the same way.

Modules & Miswirings

A potentially big problem with the LI is a “bad” modules (a problem with the readout fibres, LI “ashtray” or optical connector) or damage to the LI fibre leading to the readout module. These can also be confused with another potentially common problem: miswiring by the shift crew of the right LI fibre to the right set of readout fibres (i.e. the right ashtray).

Figure 6.47 shows a dearth of injected light in two adjacent “module sized” (i.e. 10-strip wide) spaces in both correlated and total hits. If enough uncorrelated light was seen, but substantially less correlated light, this would imply a miswiring. However, since there is also very little correlated light, this implies that there is no light shining on either module, so the problem is with a damaged LI fibre or poor optical connection.




Figure 6.47: Miswirings are found by comparing the all light seen on each strip with the correlated light seen by that strip. If LI fibres have been plugged into the wrong modules or ashtray, then the total amount of light seen by the strips will be right, but this won’t be correlated with the right LED.

Other Problems

Bad strips that cannot be correlated to any others, i.e. are not part of a pattern, are listed at the end. As other problems in the detector get fixed, the occasionally broken fibres become the biggest sources of problems — little can be done about these.

The program can also find other problems, such as bad PMTs or “Ainer boxes” (the light-tight boxes that house the PMTs), but no such problems were found by this system, mainly because this part of the electronics were tested with singles rates first. Other tests that can be done by the LI system, such as measuring PMT gains (figure 6.48), can also be used to summarise the health of the detector.




Figure 6.48: The LI system can be used to calculate the gains of each strip, and a distribution of these values at NearDet is shown here from run 3612. Problems could occur if strips have gains that are too high or too low.

The NearLI package is effective at finding many problems that occur during commissioning so that they can be fixed, but it does have limitations, meaning that other systems must also be used. For example, since there is only single ended readout, the program cannot check for broken strips — a low singles rate with otherwise efficient readout and electronics was the signature for this during commissioning.

Chapter 7

STRIP-TO-STRIP CALIBRATION OF THE FAR DETECTOR

Several methods of measuring the strip-to-strip response using cosmic ray muons are compared for use at the Far Detector. Systematic response differences and corrections to overcome them are discussed. An iterative method is shown to work accurately at FarDet by Monte Carlo and data tests. This method gives a consistent way of normalising the response of all three detectors. Improvements to the method at FarDet show that it is a viable calibration method to give constants accurate to within 2% from 3 months' worth of cosmic ray neutrino data. This method is suitable for rolling calibration using the procedure described.

7.1 Introduction

The strip-to-strip correction is the link in the MINOS calibration chain (figure 4.1) and deals with normalising the strip-end response differences observed with cosmic ray muons. Factors affecting strip-end response include clear fibre length, readout fibre pigtails, gain, scintillator light output, wavelength-shifting fibre collection efficiency and connector efficiency. All of these effects can be normalised at the strip level without the need for events to be reconstructed further. Once this calibration, and the subsequent attenuation correction, have been applied to a hit, the response should be the same at each point within each detector. Sufficient cosmic muons are required to normalise the response within each of the detectors to within the 2% specification.

In order to properly normalise the detector, the data sample used for each strip-end must be the same. There are several factors that have to be corrected for in the data before this is the case. The main corrections that need to be applied are an attenuation correction (to account for wavelength-shifting fibre length differences along the strips) and a path-length correction so that tracks coming-in at different angles with respect to the strip orientation give the same response. The performance of attempts to model these corrections are discussed in section 7.5.

There have been several possible methods of strip-to-strip calibration considered and these vary in the way that they model the response of a strip to “zeros”: when a muon has crossed the strip, but no photoelectrons have been detected at the PMT photocathode. This is an important

systematic since it accounts for around 6% of hits for an average strip-end.

The first method is to ignore such hits and to concentrate on the light seen, using this as a measure of light yield. The average light seen per muon is then normalised to the detector average, and the constant reflects this comparative response. The other methods try to reconstruct in some way where there are likely to be "zeros", and then correct for them.

This chapter explains why zeros are so important and will describe how a strip-to-strip calibration can be performed for response normalisation, what corrections are necessary and what effect these have. The performance of the chosen method in terms of how well it can normalise the detector is shown using Monte Carlo simulation and real data. This will then allow further ideas of how to boost the precision of the calibration constants produced to be discussed. Several methods will be discussed and the most successful described in detail, as it is used for the strip-to-strip calibration.

7.2 Strip-to-Strip Calibration Method

This section describes the chosen method used to calculate the strip-to-strip calibration constants.

7.2.1 General Principles of Cosmic Muon Calibration

The calibration is performed by taking reconstructed cosmic muon tracks travelling through the detector and looking at the distribution of light (known as "siglin" in MINOS reconstruction and in figure 7.1) from many such tracks at each strip-end. It requires around 1,200 track-hits to build-up a "strip-end histogram" with a statistical accuracy of 2%, as is necessary for this calibration. At FarDet, using the cuts described in this section, this requires around 3 months of data.

The response of muons through each strip is characterised, but complications arise since, for example, the path-lengths of the muons through the strip and the hit location of the muon along the strip vary. In order to get a cosmic muon sample which is the same in all strips, these must be corrected for. The number of zeros seen by each strip-end is an exponential function of the light level and it is vital that these are taken account of to perform a viable calibration.

There are several ways that zeros can be accounted for:

1. Use single-ended hits at one end of the strip to reconstruct zeros at the other end;
2. Use the tracker to identify strip-ends that a track has passed through, whether or not a hit has readout;
3. Assign all the charge in each plane from a muon track to the strip in that plane with the highest hit;
4. Estimate the number of zeros based on the light level of the strip-end in an iterative procedure.

The first method relies on having an equally efficient double-ended readout, and this is only the case at the FarDet¹, although readout fibre-lengths differ between each strip-end even here. Nor does this take into account the increased probability that, if a hit is not seen at one end, that

¹This method would not be valid at the calibration or near detectors.




Figure 7.1: Schematic diagram of the iterative method showing corrections made.

hit will not be seen by the other end, either ². For these two reasons, this method has been shown to not be viable at CalDet [79],[80].

The second method relies on an accurate prediction of where the missing strips are, and it is found that the tracker efficiency is not accurate enough to reconstruct zeros through this method.

The third method is biased against low light level strips and has been shown not to work at the CalDet.

The fourth method must be iterative, since an estimation of the light level is needed before a zero correction can be applied. This method was chosen for FarDet, and was also applied at CalDet. This means that the same method can be used at all the detectors, reducing detector-to-detector systematics. This iterative technique requires various stages of systematic corrections to be made.

7.2.2 The Iterative Procedure

For an initial, conceptual understanding of the procedure, consider the following. Each strip-end is either nominally set to a light level of 4 p.e. at the beginning of the iteration, when no prior assumptions can be made of the strip-end's light level, otherwise a previous set of calibration constants can be used as the starting point. This light level can be used to estimate the number of unseen strip hits, for example. After an iteration, the mean of each strip-end histogram is calculated in p.e. and divided by the previous light level used to estimate the number of zeros in the strip-end histogram. If this ratio is about unity, then the light-level used can be assumed to be correct.

In reality, several effects complicate this. The goal of the iteration is to make the ratio for each strip-end tend to 1.29 [80]. The value of 1.29 is not arbitrary: it is an estimate of the ratio of the Landau most probable value (MPV) to the (truncated) mean. The assumption being that for zero reconstruction it is better to tune the light levels to the most probable value rather than the mean since $P(0; \lambda)$ is non-linear with light level. The value of 1.29 was measured at CalDet by fitting to through-going beam muons [81] (figures 7.2 and 7.3), whereas a value of 1.25 was predicted by MC simulation. Most strips converge after 3-4 iterations, but in order to allow the strip-ends with light-levels at the extrema to converge to their appropriate value, up to 10 iterations are needed.

For each hit, the zero and single p.e. probabilities are calculated based on a prescription in [80]. These probabilistic corrections are applied as weights on a hit-by-hit basis, since they are dependent on the incident angle of the muon. Each strip-end histogram is filled with:

- Value: p , weight: $\left(1 - \langle P(0) \rangle_{xy} - f \langle P(1) \rangle_{xy}\right)$
- Value: 0, weight: $\langle P(0) \rangle_{xy}$
- Value: $t/2$, weight: $f \langle P(1) \rangle_{xy}$

where p is the fully corrected ADC-like hit value, $t/2$ is the mean of the sparsified part of the single p.e. peak ³ and f is the fraction of the single p.e. peak that is sparsified. The mean of the histogram is corrected for zeros and sub-sparsification hits, without the zero and sub-sparsification

²Such hits are likely to have a short path-length through the scintillator.

³The area of the distribution below the electronic readout threshold.

Example CalDet Strip-end (Plane 31, Strip 5)




Figure 7.2: An example strip-end from CalDet (plane 31, strip 5). Results from the horizontal muons from the PS beam are shown. The strip-end response in ADC is on the x-axis and the the number of hits at each count on the y-axis [81].

Ratio of Muon Distribution Mean to Fitted Landau Peak Position for PS Muons




Figure 7.3: Histogram of ratios of peak to mean (optimal convergence values) for all strip-ends at CalDet from horizontal PS muons. The ratio is different for horizontal and vertical strips (odd and even planes) [81].

bins being added explicitly. The shape of the sub-sparsification part of the strip-end histogram is approximated to be triangular, so sub-sparsification hits are added back in at $\frac{2}{3}$ the sparsification value (typically 0.2–0.3 p.e.).

If the calculated ratio is not 1.29, an estimate of the true light level can be made by using

$$\bar{x} = \frac{\sum_{i=0}^N (1 - \beta_i(\lambda)) n_i}{N\lambda} \quad (7.1)$$

$$\hat{x} = \frac{\sum_{i=0}^N (1 - \beta_i(\hat{\lambda})) n_i}{N\hat{\lambda}} \quad (7.2)$$

and

$$\frac{\bar{x}}{\hat{x}} = \frac{\lambda \sum_{i=0}^N (1 - \beta_i(\lambda)) n_i}{\hat{\lambda} \sum_{i=0}^N (1 - \beta_i(\hat{\lambda})) n_i} \quad (7.3)$$

and the approximation

$$\left(\frac{d \langle P(0) \rangle_{xy}}{d\lambda} \right)_i \approx \left(\frac{\Delta \langle P(0) \rangle_{xy}}{\Delta \lambda} \right)_i = \frac{\beta_i(\hat{\lambda}) - \beta_i(\lambda)}{\hat{\lambda} - \lambda} \quad (7.4)$$

where x is the corrected mean in arbitrary “muon energy units”, λ is the calibration constant in p.e., β_i is the probability that hit i with n_i p.e. is not seen, N is the number of entries and the symbols bar ($\bar{}$) and circumflex ($\hat{}$) represent the calculated and true values respectively. In this case

$$\hat{\lambda} \approx \lambda \frac{\bar{x} + \frac{1}{N} \sum_{i=0}^N n_i \frac{d \langle P(0) \rangle_{xy}}{d\lambda}}{\hat{x} + \frac{1}{N} \sum_{i=0}^N n_i \frac{d \langle P(0) \rangle_{xy}}{d\lambda}} \quad (7.5)$$

Thus, during each loop the quantity on the right in equation 7.5 is calculated and the next iteration uses the new light level. Figure 7.4 gives the convergence values attained by this calibration on FarDet data. Strip-ends that do not converge properly can be “fixed” using one of several methods, described later.

7.2.3 The Cosmic Muon Sample

FarDet is described in section 3.3.4 as having 484 active planes, divided into 2 supermodules (SM1 is planes 1–248, SM2 is planes 250–485, with no readout on the boundary plane, 249) with 192 double-ended readout strips on each, so there are 185,856 strip-ends to calibrate. The single cosmic muon flux is 0.42 Hz which gives about 400 hits per strip per month [2]. Through-going muons are used since there is only a 10% difference in energy deposition between 1–10 GeV [66].

Each muon produces 12 p.e. summed at centre of strip, which is seen as typically 4 p.e. at each end when read out. The light is lost through attenuation in the wavelength shifting fibres and readout fibres.

The data sample used for this analysis are taken from 3 months of running from August to October 2004, reconstructed with MINOS-soft release R1.11, giving a sample of around 1200 hits at each strip-end. The simulated muon sample used for this analysis is cosmic ray Monte Carlo, reconstructed on the batch processing farm using release R1.14. This MC is designed to simulate all known systematics and expected behaviour in the Far Detector from almost 1 month of normal




Figure 7.4: Convergence value of strip-ends in Far Detector data. The entries are inversely weighted by RMS. Strip-ends furthest from 1.29 have a high RMS. The majority of strip-ends converge closely to 1.29.




Figure 7.5: Spectra of cosmic muon track angles in the x-, y- and z- directions in the MC sample. The distributions are very similar for data (figure 7.6) and MC.




Figure 7.6: Spectra of cosmic muon track angles in the x-, y- and z- directions in the data sample. Compare with figure 7.5.

running at FarDet. Further cuts and corrections are made to these data samples and these are described in the following sections.

Figures 7.5 and 7.6 show the track directions for cosmic muons entering FarDet for MC and data. The main difference between these figures, other than one is shown on log axes and the other on linear axes, is in the red y-distribution. There are more upward-going (positive) y -events in data than in MC. This is caused by poor timing calibration in the data that was not modelled in the MC. Between the data and MC reconstruction times, there was a considerable effort to improve the timing calibration at FarDet. This effort was spurred on by the need to reduce the stopping muon background for the atmospheric neutrino analysis, described in section 4.2, by improving the discrimination between upward and downward going events.




Figure 7.7: Track directions in MC with respect to strip orientations (U and V), offset by 45° from horizontal and vertical. The spectra seen by each orientation is the same.

It can be seen from figures 7.7–7.10 that there is no large systematic difference in track orientations, hence path-lengths through strips, for any potential pattern: plane number, strip number, orientation or position along the strip. This was not the case at CalDet, where the x and y (as opposed to U and V) strip orientations caused a large systematic difference in pathlengths between different strip orientations. Any differences in average muon track pathlengths through strips at FarDet can be removed with a pathlength correction without any higher order residual effects causing systematic biases.

The distributions of track lengths are shown for the MC sample in figure 7.11 and data in figure 7.12. The mean track length is about 8 m in MC, but only 4 m in data. There is a deficit of short events in the MC, so the distributions look different. This may be due to a slightly different

U-Direction Cosine Spectra As A Function of Strip No. (MC)




Figure 7.8: U-direction cosine track angle spectra as a function of strip number (MC). The V-direction cosines look similar, as do the spectra for data.

U-Direction Angle Cosine Distributions With Respect To Plane Number (MC)




Figure 7.9: U-direction track angle cosine distributions with respect to plane number (MC). The spectra look similar in the V-direction and in data.




Figure 7.10: Track direction spectra with respect to strip orientation (data). The distributions are the same for both strip orientations.

cut or correction in the MC sample compared to the data sample, or that fewer short events were generated in the MC than are present in data.

This should not affect the calibration since the procedure described in this chapter aims to correct any differences seen between different strip-ends that may be due to different muon spectra being seen. I.e. the corrections described in this section act to convert each muon hit to what it would be if it were a horizontal muon travelling through the centre of the strip. Therefore, any differences seen should be due to just strip-end light level output, and would be consistent within each sample, which is all that is required. Furthermore, if the issue is merely due to track length differences, this should have no impact on the light levels seen by each strip-end.

The raw track hits vary between data and MC at the high end (figures 7.13 and 7.14), but this is 3 orders of magnitude from the peak, so is a small effect. The mean track hit energy varies by only about 1%, so the MC can still be considered to represent the data well.

The cross-section of hit locations show that there are no systematic patterns in hit locations, since figures 7.15 and 7.16 shows this cross-section is only a function of the shape of the detector. The energy deposited by these hits in profile (figures 7.17 and 7.18) does vary as a function of position, but this is a genuine strip-to-strip variation (readout fibre length) and is one of the effects that the strip-to-strip calibration corrects for.

When this cross-section of hit locations is plotted in terms of vertical and horizontal positions, rather than radial position, no distinct patterns can be seen either in separate U and V-views (figures 7.19 and 7.20) or combined (figures 7.21 and 7.22). Patterns can, however, be seen corresponding




Figure 7.11: Distribution of the length of tracks passing basic cuts in the MC sample. The distribution for data in figure 7.12 has a lower mean track length and a longer tail.




Figure 7.12: Distribution of the length of tracks passing basic cuts in the data sample. Compare the distribution to that for MC (figure 7.11), which has a higher mean and shorter tail. There is a small peak at around 7 m that is also present in the MC.




Figure 7.13: The distribution of track hits light output in MC. The shape of this distribution differs from data at the high end.




Figure 7.14: The data distribution of light from track hits. Compare with the MC distribution in figure 7.13.




Figure 7.15: The octagonal cross-section of the Far Detector causes this radial pattern in hit locations for MC over the whole detector. The dip in the centre is due to the coil hole. The pattern is repeated in data (figure 7.16).




Figure 7.16: The octagonal cross-section of the Far Detector causes this radial pattern in hit locations for data over the whole detector. The dip in the centre is due to the coil hole. The pattern is repeated in MC (figure 7.15).




Figure 7.17: The variation in corrected hit response as a function of radial distance from the centre of the whole Far Detector. This plot is for MC, but the pattern is repeated in data (figure 7.18). Note the response dips at module boundaries.





Figure 7.18: The variation in corrected hit response as a function of radial distance from the centre of the whole Far Detector in the data sample. Compared to MC, (figure 7.17), the distribution has higher peak structure around the centre of the detector, due to differences in clear fibre lengths.

to module boundaries, reflecting the spatial resolution of muon tracking.

The y -position (vertical position) of hits in these samples (figure 7.23 for MC and 7.24 for data) are relevant since subsequent calibrations, such as the relative calibration that uses stopping muons, have a strongly asymmetric y -distribution of hits in FarDet. In these samples, there is a 5% bias towards higher- y hits, but this should not cause any significant systematic effects as it only affects the number of entries and not the light deposited.

Any strip-to-strip differences in the samples either do not directly affect the p.e. response of the strip-ends (e.g. only affect the number of entries) or can be removed by subsequent corrections. Therefore, these data and MC samples are appropriate for response normalisation and strip-to-strip calibration.

Cross-section of Hit Locations - U-strips, SM1



Cross-section of Hit Locations - V-strips, SM1




Figure 7.19: Cross-section of hit locations in SM1 the Far Detector from the cosmic muon Monte Carlo sample. The plot is divided into U and V views to show any differences in number of entries that could be dependent on strip orientation. The top plot also has a profile fitted showing that the average vertical hit location is at the centre of the detector. Some hit vertices are reconstructed to be outside the detector.




Figure 7.20: Cross-section of hit locations in the FarDetector from cosmic muon data. Compare with figure 7.19 where the dips in the number of entries at module boundaries can also be seen.

Cross-section of Hit Locations - Both Orientations, SM2




Figure 7.21: Cross-section of hit locations in SM2 the Far Detector from the cosmic muon Monte Carlo sample. The plot sums the U and V views to show clearly any differences. Patterns that could occur due to position along a strip are reduced so that other systematics can show up more clearly. There do not appear to be any such systematic patterns. Compare with figure 7.22.




Figure 7.22: Cross-section of hit locations in SM2 the Far Detector from data. No systematic variations in hit location can be seen across the detector. Compare with figure 7.21.




Figure 7.23: The octagonal cross-section of the Far Detector causes this vertical pattern in hit locations for MC over 4 different orientations of the detector. The dip in the centre is due to the coil hole. This pattern for through-going muons is different to that for a stopping-muon sample, which would be sharply skewed towards more hits at high- y values.




Figure 7.24: The octagonal cross-section of the Far Detector causes this vertical pattern in hit locations for cosmic muon data in different views of the detector. These plots show data, compare with figure 7.23 (MC).

7.3 Applying Cuts To The Sample

There are some standard cuts applied to the data from the data acquisition system (DAQ) and standard reconstruction as well as those mentioned in this section, such as the 4/5 online plane trigger cut on the DAQ determining what events will be written out to file. None of these cuts are likely to remove any significant number of genuine muon tracks from the sample.

7.3.1 Track Cuts

A loose cut is placed on tracks so that tracks must be at least 8 planes long and have at least 6 track-like planes (hits in contiguous planes that look like a track). This is in order to get as much data as possible, especially in the corners of the detector, where there are proportionally fewer long tracks. For this calibration, unlike alignment, for example, accurate knowledge of the track fitting is not necessary. This is only used to provide an attenuation correction for the light lost in the wavelength-shifting fibres running along the strips.

The tracks must enter and leave the detector, which cuts down the fluctuations due to stopping muons, which deposit different amounts of light in each plane. A steepness cut on track angle is set at 86° , so that a strip entering the top of the detector and leaving the bottom must traverse at least 8 planes. The average pathlength through the strip for a particular set of angles is limited to 2.5 cm through the strip. This reduces problems which arise at high path-lengths from the Landau effect of light deposited as a function of path-length through a medium.

The only other purpose of using tracks and requiring a cut is so that we know that our events are real muons. There must be no showers along the track, only the first track in any snarl is used (for reconstruction reasons).

Cuts are also applied to individual strip hits to determine whether they are included in the sample. The scintillator is not covered with the co-extruded reflective TiO_2 coating at the ends of the strips. Light is lost at the ends of the strips due to this lack of reflective coating. To prevent any systematic effects due to the different size of this effect between different size strips, a fiducial cut is made against hits within 15 cm of the edge of the strip or the coil hole (see figures 7.21 and 7.22).

7.3.2 Energy Cuts

A cut is also placed on the amount of light seen in each hit. Any hits below a sparsification value of 0.3 p.e. is rejected: a correction is applied later for the amount of the single photo-electron peak not seen due to sparsification, and this requires this value to be set to the same value in terms of p.e.s for all strip-ends.

Any hits above 30 p.e. are rejected as having a light level not consistent with a muon hit. The expected light deposition follows a Landau distribution that peaks at around 4 p.e., as described in the next section. This cut prevents hits from Bremsstrahlung entering the sample, since these do not follow the Landau distribution of energy loss, and will not be the same for all strip-ends.

A high-end cut is set at 8 times the light-level of the strip, and this gives the same truncation proportion (>99%) for all light-level strip-ends. This reduces the problematic effect of a shower hit being mistagged as part of the track.

7.3.3 Other Cuts

A small number of other cuts are made when needed, for example, to prevent floating point exceptions in the analysis program. These cuts are rarely made so have negligible side-effects. (Data collected from strip-ends with hardware problems are dealt with in the usual way: sanity checking of the calculated calibration constants is described in a later section.)

7.4 Applying Corrections To The Sample

Corrections are needed to ensure that the sample of muons seen at every strip-end is the same. A gain correction is used to convert ADC-like hits into p.e. in order for the p.e.-based zero and sparsification corrections to work — the gain is convolved back in later. Path-length (the average distance the muons travel through the strip for a given angle) and fibre attenuation corrections are made based on muon track information.

7.4.1 The Gain Correction

When a hit registers as an ADC count, we need to know the number of p.e. seen at the photocathode so that we can use Poisson statistics to predict zeros from the light seen, so the “siglin” hit ⁴ in ADC counts is converted into a hit in p.e.. After the procedure, since we need an ADC like value for the constants, the same numbers are used to correct the constants back from p.e. to ADC, so this strip-to-strip calibration corrects gain differences between strip-ends. The distribution of PMT gains at FarDet is shown in figure 7.25.

The distribution of hits added to the sample after gain correction is shown in figure 7.26 for MC and 7.27 for data.

⁴Linearised light using the “gain curve” of PMTs and their drift correction over time, see section 6.4.




Figure 7.25: Histogram of the gain values used in the calibration. This distribution is the same for data and the MC sample since it is the distribution of strip-end gains in the database. The default value, when none has been written in the database, is 67 ADC/p.e..




Figure 7.26: The gain-corrected distribution of hits on an ADC-scale from the MC sample. The width of this distribution is dominated by the natural Landau width of energy deposition, and attenuation and pathlength variations.




Figure 7.27: The gain-corrected distribution of hits on an ADC-scale for data. Compare with figure 7.26.

7.4.2 The Attenuation Correction

The attenuation correction is applied in the standard reconstruction chain after track fitting, but it is necessary to take account of the amounts of attenuation due to different length strips. The charge is normalised as if the hit had been at the centre of the strip. If the attenuation correction to be applied is not a sensible value, the hit is cut from the sample.

Figure 7.28 shows the distribution of attenuation corrections applied to hits in the samples. There are more hits closer to the readout end (hits from the far end have a higher chance of being attenuated below sparsification) and there is a dip in the centre, since there are fewer hits around the coil hole. Figure 7.29 shows the value of the attenuation correction for all strip hits in different views at FarDet, mapping-out the correction as a function of position along the strips.

The attenuation along the longest strips are corrected from up to 600% differences to within 5%, the residual difference is shown in figures 7.30, 7.31 7.32 and 7.33 (i.e. both data and MC, with zero correction and for strips of the same length).

A similar attenuation under-correction is seen with stopping muons [66], and an extra empirical correction to the attenuation correction was applied in this work on relative calibration. Such a correction could be applied here, but has not been since the attenuation correction may be modified in the future (i.e. to “future-proof” the method) and the current under-correction will have only a small effect on the resulting calibration.




Figure 7.28: The distribution of attenuation corrections applied to the samples. There are more corrections below 1 as there is a greater chance of a hit at the near end of a strip being above sparsification (the other end would be a zero-hit). The dip in the number of events at 1 is due to the coil. This is the distribution of attenuation corrections in the database.




Figure 7.29: The profile of attenuation corrections applied against strip position in data, in East, West, U and V views. Damaged strip-ends can be seen by purple lines away from the main swathe of values. The profiles are the same in MC.




Figure 7.30: Effectiveness of the attenuation correction in normalising response across for all strips in the MC sample. The black line shows the profile. Note that the residual effect shows several convolved dependencies as the strips are different lengths, c.f. figure 7.32.




Figure 7.31: Effectiveness of the attenuation correction in normalising response across the length for all strips in data. The effective undercorrection is 10% in this view, c.f. 7.31.




Figure 7.32: Effectiveness of the attenuation correction in normalising response across the length of 8m long strips in the MC sample. The plot is split into 4 views showing the response against the length of the strips before and after the correction. Two of the corrected plots are on full scales for comparison with the original and two on zoomed scales to show the residual effect. Note that the residual effect is almost linear across the length of the strip for strips of the same length. The profile is flattened at the high (read-out) end by the exclusion of hits above the axis maximum into the profile. The sharp drop-off at each end is a separate effect.




Figure 7.33: Effectiveness of the attenuation correction in normalising response across the length of 8 m long strips in data. Compare to figure 7.32.

7.4.3 The Path-length Correction

The path-length correction used is that described in reference [79]. The path-length correction takes account of the effect from muons entering the strips at different positions along the cross section of the strip and at different angles. The improvement of this method over a simple geometrical calculation of length through an infinitely wide (δy) strip is that this method takes account of corner-clipping (figure 7.34). Thus the actual average path-length given (ds) is the average path-length through the strip for a given angle when all possible entry locations are integrated over. This means the path-length is less than $\frac{ds}{dz}$, the path-length of an infinitely long and wide strip (figure 7.35). In the application of this correction, the strips are taken to be 1 cm thick, 4.1 cm wide and 8 m long.




Figure 7.34: Different possible track angles through a strip and their resulting path-lengths [79].

The energy deposited as a function of average path-length through the strip is shown in U and V-views is shown for data in figures 7.36 and 7.37, which shows that this correction needs to be done, since the amount of light deposited climbs so sharply with track angle.

The formula used to calculate the average path-length through the strip based on the track angle through the strip is

$$\langle ds \rangle_{xy} = \frac{ds}{dz} \frac{\Delta x \Delta y \Delta z}{\Delta x \Delta y + \left| \frac{dx}{dz} \right| \Delta z \Delta y + \left| \frac{dy}{dz} \right| \Delta z \Delta x} \quad (7.6)$$

The distribution of number of tracks at each angle shows few path-lengths above 2.5 cm in data (figure 7.38).




Figure 7.35: The average pathlength (ds) as a function of track angle ($\frac{ds}{dz}$) in data. This displays the output of the function which calculates average pathlength from track angles $\frac{dU}{dz}$ and $\frac{dV}{dz}$ so the pattern is the same for all data samples (i.e. MC).

The energy deposited as a function of average path-length is shown with and without the correction in figure 7.40 for MC. There is a 10% under-correction between 1–2.5 cm (the range used in this study).

The histogram of fully corrected light deposited is shown as a function of track angle ($\frac{ds}{dz}$) in figure 7.42. A fit is made in the 1.0–2.5 cm region for data, showing the further under-correction that would be made if $\frac{ds}{dz}$ alone was used for this correction, rather than ds .

7.4.4 The Zero Reconstruction

This correction is also applied from [79]. Along with the sparsification correction, the zero correction takes account of muons travelling through the scintillator but not depositing any light. The several possible methods of applying this correction were discussed earlier in this chapter.

With the light measured by a PMT in photoelectrons, the probability of seeing a certain amount of light at the PMT face is given by Poisson statistics. The Poisson probability from the photoelectron spectrum of seeing 0 p.e. is given by

$$P(0) = e^{-\lambda \cdot ds} \quad (7.7)$$

where λ is the number of photo-electrons produced per cm and ds is the path-length through the strip. When integrated over all path-lengths, an equation can be derived that predicts the average




Figure 7.36: Strip-end ADC response as a function of track angle ($\frac{ds}{dz}$) in data. The features are similar in the MC sample (figure 7.37).




Figure 7.37: Strip-end ADC response as a function of track angle ($\frac{ds}{dz}$) in MC. The features are similar in the data sample (figure 7.36).




Figure 7.38: The distribution of path-lengths through strips in the data. The distribution peaks at 1.25 cm. The maximum pathlength is cut at 2.5 cm so the number of tracks cut out of the sample can be seen. Compare with figure 7.39, for the distribution in MC.




Figure 7.39: The distribution of path-lengths through strips in the MC sample. The distribution peaks at 1.25 cm. The maximum pathlength is cut at 2.5 cm, the same as for data (figure 7.38).




Figure 7.40: The effectiveness of the pathlength correction for the MC sample, figure 7.41 shows the same for data. The detector is split into 4 views showing the deposited energy distribution as a function of pathlength (ds) before and after the correction. The residual dependence on pathlength is an over-correction of 10% over the range of pathlengths used in the calibration (1–2.5 cm), similar to data.




Figure 7.41: The effectiveness of the pathlength correction for the data sample, the same is shown for MC in figure 7.40. The detector is split into 4 views showing the deposited energy distribution as a function of pathlength (ds) before and after the correction. The residual dependence on pathlength is an over-correction of 10% over the range of pathlengths used in the calibration (1–2.5 cm), the same as for MC.




Figure 7.42: The residual relationship of pathlength, zero- and 1-p.e. corrected response on track angle for data in 4 views. There is also a 10% effect here (seen in the zoomed top left plot) as in figure 7.40.

probability of getting a zero, given the strip light level and muon angles [79]. This formula can be simplified to

$$\langle P(0) \rangle_y = \frac{2 \left| \frac{dy}{dz} \right|}{\lambda \frac{ds}{dz} \left(\Delta y + \left| \frac{dy}{dz} \right| \Delta z \right)} \quad (7.8)$$

(where x , y and z represent the U, V and z directions respectively) by taking the limit where x , the strip length (of order several metres) is large with respect to y , the strip width (4.1 cm).

The size of the zero correction is typically around 5%, and can be seen for data and MC in figures 7.43 and 7.44. The shape of the distribution is asymmetric since strips at high light levels have a low probability of getting a zero, e.g. 1%, whereas low light level strip-ends have a much higher proportion, up to 20%, and the difference has exponential factors (equation 7.7), although the dominant factor is λ^{-1} . The mean of the distribution varies slightly since the light-level tuning of the data (in terms of muon energy units, known as MEU or MIP) is 370 ADC in data and 358 ADC in MC.




Figure 7.43: The distribution of applied weights of the iterative zero correction in data. The mode correction is 4% and there are larger corrections in data than MC (figure 7.44). This is due to a higher light level tuning in the MC than is seen in the data.

Since zero hits are applied by reweighting the light seen, and keeping the number of entries the same, rather than by explicitly adding-in the 0th bin, which would reduce the mean. The difference of the distributions in data and MC is shown by figures 7.45 and 7.46.

Another effect of the zero reconstruction is that the difference in the number of hits seen at each end of the same strip is reduced, in both MC and data (figures 7.47 and 7.48). Whilst both distributions in each figure are centred around 0, the difference reduces by $\frac{1}{3}$ for both MC and data.




Figure 7.44: The distribution of applied weights from the zero-correction in the MC sample using the iterative zero correction. The mode correction is just below 4%.




Figure 7.45: A direct comparison between the zero-corrected strip-end response means with and without zero correction for the MC sample. The means clearly shift to lower values once the zero hits have been added back in.




Figure 7.46: A direct comparison between zero-corrected strip-end response means with and without zero correction in data. As in MC (figure 7.45), adding in the zero hits reduces the mean and width of the distribution.

7.4.5 The Sparsification Correction

This correction accounts for the part of the one photo-electron peak not seen due to sparsification [79]. This is the correction for the part of the 1-p.e. peak below the sparsification threshold, that is therefore not read out (figure 7.49).

This effect is different from the zeros described since it is part of the 1-p.e. peak and not the 0-p.e. peak, but the effect (of seeing no light from a muon hit) is the same. Using the same derivation as before, the probability of getting a single p.e. ($P(1)$) can be calculated for a particular light level and set of muon angles

$$\langle P(1) \rangle_y = \frac{\left(\frac{ds}{dz} \lambda \Delta z e^{-\lambda ds} \left(\Delta y - \left| \frac{dy}{dz} \right| \Delta z \right) \right) + 2 \left| \frac{dy}{dz} \right| \left(\frac{1 - e^{-\lambda ds}}{\lambda \frac{ds}{dz}} - \Delta z e^{-\lambda ds} \right)}{\Delta y + \left| \frac{dy}{dz} \right| \Delta z} \quad (7.9)$$

where symbols are as previously defined (see [79] for more details).

Since only about 4% of the 1-p.e. peak is below this threshold of 0.3 p.e. (although this can be higher for lower light-level strip-ends), and the 1-p.e. peak is only a small part of the light level spectrum (especially at high light levels) this correction is only important for low light level strip-ends. For most strip-ends the size of this correction is $\mathcal{O}0.2\%$, over 20 times smaller than the zero correction. The amounts added back in are shown in figures 7.50 and 7.51 for MC and data.




Figure 7.47: Most zeros are light seen at one end and not the other. Once the zero correction has been applied, there should be an even number of zeros at each end. The difference between the number of hits at each strip-end is lower after calculated zeros have been added in, and the distribution is narrower, as can be seen in the MC sample. This implies that the correction is working in the right way. There is a small offset since one set of strip-ends has a lower light level than the other.




Figure 7.48: The difference in the number of hits between different ends of the same strip in data is wider than that in MC because there are more hits per strip-end in the data sample. Once the zero correction is applied, the difference in the number of hits seen at each end is reduced (c.f. figure 7.47).




Figure 7.49: The fraction of the single photo-electron peak below sparsification is about 4% from a phototube simulation [79].




Figure 7.50: The mode weight of the 1-p.e. correction is 0.1% in the MC sample. The effect of the correction is small but there could be up to 1% discrepancies between strip-ends if it is not applied.




Figure 7.51: The mode weight of the 1-p.e. correction added back in is 0.2% in data. This is a larger correction than the MC (figure 7.50) because the light level in the MC is tuned higher than in data.

7.5 Results

The results of the calibration describe the response and performance of the detector.

7.5.1 Photo-electron Response

The p.e. response of the detector shows the effects of everything except PMT gain. Since the gain variation (figure 7.25) is a major strip-to-strip difference, and can be measured by light injection “gain curves”, it can be useful to deconvolve its effects on the strip-to-strip calibration constants in order to highlight other response patterns.

Figures 7.52 and 7.53 show the distribution of strip-end p.e.-response means for data and MC. The light-level tuning of the MC (i.e. the number of p.e. per muon energy unit) is different to the data, so the means are different. The relative widths of the two distributions is not important since this just depends on the arbitrary strip responses set as truth in the MC. For data, the distribution is qualitatively Gaussian, which does not imply that there are any causes of strip response differences not described in the introduction (although there may be systematics much smaller than the spread of the distribution).




Figure 7.52: Histogram of photo-electron light level for all strip-ends in the detector in the MC sample. These values are also known as the calibration constants in pe units. The MC mean at 4.4 p.e. is 0.5 p.e. higher than the data (figure 7.53).

The statistical uncertainty in p.e.-response (figure 7.54) reflects the number of entries and the shape of the strip-end histogram. Systematic uncertainty will be discussed later in this section by comparison with MC. The number of entries in the data sample is around 1200 per strip-end, and the absolute uncertainty in the calibration constants averages at 0.08 p.e., corresponding to 2.1%. For MC, figure 7.55 shows the relative uncertainty, and is around 4% from around 250




Figure 7.53: Histogram of photo-electron light level for all strip-ends in the detector in data. The mean and spread in photo-electron responses are smaller in data than MC (figure 7.52).

entries per strip-end. The distribution shows that few strip-ends have large statistical errors, i.e. over twice the average; the RMS of the distribution is 22%, showing that strip-ends tend to have similar uncertainties.

Typical strip-end responses for the data and MC samples in terms of corrected ADC response show what the data filling the strip-end histograms actually look like from the given number of entries (figures 7.56 and 7.57). With around 1000 entries, the underlying shape of the Landau distribution can be seen, but it is difficult to see so clearly with only around 250 entries. This is one reason why the mean of the distributions is used in this analysis (truncated at 8 times the light level), since it is difficult to accurately fit to find the peak.

Figures 7.58 and 7.59 show how the p.e.-response varies with respect to the position and end of the strip in the detector. Response variation with gain can be seen in figures 7.45 and 7.46.

7.5.2 Agreement of MC at FarDet

For the MC sample, the truth is known, as well as the calibration constant, so the two can be compared and the accuracy of this method can be quantified. Figure 7.60 shows the scatter plot of the agreement between calibration constants produced by the method and the MC truth values in which the width of the main swathe of results is statistically dominated. Figures 7.61 and 7.62 show different profiles of this, the gradient of the latter (unity) shows that the analysis method can fully normalise the detector for all strip-to-strip response differences.

Using the iteration to the convergence value of 1.29 justified in section 2.2, figure 7.63 shows the relative discrepancy ($\frac{A}{B}$) between truth and calibration constant as a function of light yield.




Figure 7.54: The absolute statistical uncertainty in photoelectron response in data. The average relative uncertainty is 2.1% for the sample.




Figure 7.55: The relative statistical uncertainty in calibration constants (MC). The average uncertainty is around 4%, with a tail to higher uncertainties from strip-ends with fewer entries.




Figure 7.56: Six typical strip-ends in the MC sample from different parts and orientations of the detector, and with different numbers of entries.




Figure 7.57: Six typical strip-ends from the data in different parts and orientations of the detector, and with different numbers of entries. The same strip-ends are used in the MC equivalent (figure 7.56).




Figure 7.58: Detector photoelectron response map in strip vs. plane view. Topological differences in response can be seen from this map without being smeared by differences in gain. Newer strip-ends (with higher plane numbers) tend to have a higher response, as do lower strip numbers compared to high. This latter effect is due to clear fibre length which is corrected for by this calibration.




Figure 7.59: Detector map in strip vs. plane view of the calibration constants from the MC sample.

Calculated Calibration Constants vs. True Light Level




Figure 7.60: Spread of calculated calibration constants compared to true light level from MC with 280 entries per strip-end. The spread is statistically dominated. Both axes are normalised to 1 at their mean value.

Distribution of Calibration Constants, True Strip Variations and the Difference Between Them




Figure 7.61: The distribution of strip-end light level variations is shown in blue and the distribution of calculated calibration constants in red. The difference between the two values on a strip-end-by-strip-end basis is in black, with a width of 5.8%.

Profile of Calibration Constants as a Function of True Light Level




Figure 7.62: There is a linear relationship with a gradient of unity between the calculated calibration constants and the true light level from the strips in MC.

This shows that the systematic error is $<1.0\%$ for the majority of strip-ends, with calibration constants in the range 0.7–2.0 times the detector average (compare with figures 7.45 and 7.46). There is a larger discrepancy of upto 5–6% at the low end and possible ways to correct this is discussed in section 6.2. The cause is that the shape of the strip-end histograms (figure 7.64) change with the light-level of the strip, causing the optimum strip-end convergence values to change (the strip-end shape changes with all light levels, not just low ones).




Figure 7.63: When the iteration converges to 1.29, the majority of the strip-ends have a systematic discrepancy below 1%.

Looking at the discrepancy between truth and constant as a function of position in the detector is another powerful way to look for systematic errors in this method. This is done for each read-out side of the detector in strip vs. plane maps in figures 7.65 and 7.66. No discernable patterns (other than those of broken electronics) are discernable in the range shown, implying that the discrepancy is dominated by statistical scatter.

7.5.3 Data Agreement

Confirmation that the calibration works on real data comes from analysing the energy deposition of stopping muons, where the energy is known from range and the minimum ionising region of the Bethe–Bloch energy loss curve.

A properly working strip-to-strip calibration reduces the response difference between different regions of the detector. As an example, figure 7.67 shows the effect of this normalisation on the scale of the original scatter, i.e. when the data is calibrated using this strip-to-strip correction and the attenuation correction, the detector appears normalised.

When the scale is zoomed to the scale of the remaining variations, only statistical scatter and the residual variation along the strips from the attenuation under-correction, discussed in earlier in




Figure 7.64: Normalised spectra of hits on different light level strip-ends (MC). The underlying shape of the strip-end histograms may change with light level, which causes a difference in optimal convergence values as a function of light level.




Figure 7.65: The strip vs. plane map of discrepancy between calibration constant and truth on the East side of FarDet shows no obvious patterns.




Figure 7.66: The strip vs. plane map of discrepancy between calibration constant and truth on the West side of FarDet shows no obvious patterns.




Figure 7.67: Response cross-section of FarDet with the strip-to-strip and attenuation corrections applied in the U-view [66]. The strip orientation is shown by the black line running across the detector and the position of the readout by the magenta line on one side of the detector cross-section.

this chapter, remain (see figure 7.67).




Figure 7.68: Response cross-section of FarDet with the strip-to-strip and attenuation corrections applied [66]. Compare with figure 7.67.

7.5.4 Stability Of Constants And Detector

The difference between the calibration constants from two data samples, taken 6 months apart, is shown in figure 7.69. The width of this distribution is dominated by statistical uncertainty and there is no systematic change in performance (mean) over that period. This confirms that it is feasible to use 3 months or more of cosmic muon data to perform the strip-to-strip calibration, since there is a negligible change in detector performance over that time scale. If there is any seasonal change in response due to temperature changes, this is already calibrated-out by the light injection system (chapter 6).




Figure 7.69: The difference in calibration constant taken from 2 runs, 6 months apart. There is a negligible shift in the mean and the spread is statistically dominated.

7.5.5 Caveats

The iterative method used for this calibration comes with caveats, since it makes a number of assumptions:

1. That all strips have perfect 1.0×4.1 cm cross-section;
2. That the gain is well known;
3. That the attenuation correction is applied correctly;
4. That the hits processed consist only of track hits (and comprise all track hits for zero reconstruction);
5. That there is no contamination from cross-talk or shower hits, since this will affect both the visible mean and the prediction of zeros;
6. The integrals assume a uniformly distributed entry point into the strip, i.e. corner clippers should be included;
7. That the convergence value is correctly tuned to get the absolute strip light levels.

These effects introduce biases with light level — the MC test indicates some sensitivity to these.

7.6 Improvements

7.6.1 Truncation Point

Thus far in this analysis, truncation to the strip-end histograms have only been made at the high end to prevent non-track-like hits from entering the sample. The resulting strip-end histograms for various light levels have long tails, which have a large effect on the mean (figure 7.64).

By truncating these histograms at the same proportion of the number of entries, effectively cutting off the tail, the precision on the mean can potentially be increased. Even though it is the peak of the distribution that is used to calculate the calibration constant; it is the mean that is used to determine where this peak is. There are not enough events in a typical strip-end histogram of 1200 entries to determine the peak directly.

Applying truncation therefore reduces the RMS without changing the value of the calibration constant. A study was carried out to measure the statistical precision on the calibration constant as a function of the amount of data truncated away (figure 7.70), for different light-level strip-ends with 1200 entries. For all strip-ends, there is an optimum truncation value with around 80% of the number of entries, i.e. 20% of the data is removed from the tail. The resulting distributions are shown in figure 7.71.




Figure 7.70: Optimum truncation values on different light-level strip-ends. All the strip-ends have a minimum at 20–25% truncation.

This 80% truncation point is linear as a function of light-level (figure 7.72). This feature could potentially be used to apply the truncation to a sample of strip-ends when the light-level is known (i.e. at the end of the iteration). With the improvement of using truncation means, the statistical




Figure 7.71: Strip-end histograms with 80% truncation in terms of p.e.. The scale is zoomed to give the best resolution on the 2–8 p.e. histograms.

precision is 1.4% from strip-ends from 3 months of data, and allowing up to a further 1.4% from systematic error, calibration constants can be calculated to the required 2.0%.

7.6.2 Convergence Values

So far in this analysis, it had been assumed from the work done at CalDet that the ratio of mean to MPV, the convergence value of the iteration, was 1.29 (as discussed earlier in the chapter). The shape of the strip-end histogram changes with light-level, so this value may not be optimal for all light level strip-ends. Figure 7.73 shows the agreement between truth and calculated constant for different light-level strip-ends and convergence values.

Converging to different values as a function of light-level may not lead to improved results, since the iterations can get “stuck” at local maxima/minima, so values of 1.25–1.29 give the best accuracy. This is also shown by figure 7.74, which shows the optimum convergence value as a function of light-level. Most strip-ends are in the range 3–6 p.e., where the value of 1.29 produces constants accurate to within 1.0%.

7.6.3 Attempts At Producing Constants For Pathological Strip-Ends

Strip-ends that have too few entries or will not converge to a stable value in the iteration can be “fixed” to produce a best guess of the true light-level. The different attempts at such fixes are:

1. If both ends of a strip had sufficient entries, but neither end converges, then the strip-end




Figure 7.72: The light (in p.e.) at which different light level strip-ends would be truncated at 80%.




Figure 7.73: Systematic error as a function of light level for different convergence values. The values of 1.25 and 1.29 give the flattest response over the range with the most strip-ends.




Figure 7.74: The ratio of corrected mean to true light level (MPV) for different light level strip-ends. A third-order empirical fit highlights the trend.

can be fixed using adjacent planes. The number of muons passing through the strip can be estimated by taking the average of the number of muons passing through the strip with the same number on the planes either side. The correct number of zero hits are then added to the histogram and a new light level can be calculated.

2. If both ends of a strip had sufficient entries but one does not converge, then the strip-end can be fixed using the other end of the strip. The number of muons passing through the strip is known from the other end, and the correct number of zero hits are then added to the histogram.
3. If the strip-end had insufficient entries but other end of the strip had enough entries, then the strip-end can be fixed using the other end of the strip.
4. If neither end of the strip converged, adjacent planes are used to fix the strip-end.

If the strip-ends are dead and no other information is available, setting the light-level to the detector average value gives the best estimate of light level: the strip-ends may be subsequently fixed (in the case of faulty electronics, for example).

Figure 7.75 shows the scale of this problem in real data. Strip-ends with calibration constants away from the main swathe of values have been fixed by using one of the above methods to produce better calibration constants.




Figure 7.75: The raw hit average as a function of calculated calibration constant. Strip-ends away from the main body of hits have been “fixed”. The spread of the swathe also highlights the need for a full analysis, such as the one described here, rather than just a raw hit average.

7.6.4 Procedure For Rolling Calibration

The procedure used for truncation is very similar to a potential procedure which can update the calibration constants frequently, such as daily, for greater accuracy (a “rolling calibration”). This procedure can be summarised as:

1. The light level of the strip is known approximately from a previous calibration, such as one performed by the iterative method described in this note.
2. A special file type holding data for each strip-end and for each of the previous 100 days is read in to the strip-end histograms.
3. New data can be made for that day with zero correction applied, since the light-level is known approximately.
4. Data from this day and the previous days are added together for each strip-end and truncated to get a calibration constant.
5. This value is then put into offline database for this strip-end to be used for calibration.

A further advantage of this method is that it can, with light injection, track detector bugs and performance changes quickly, thus improving the health of the detector. A rolling calibration should be implemented as the next improvement to strip-to-strip calibration.

Chapter 8

STRIP-TO-STRIP CALIBRATION AT THE NEAR DETECTOR

This chapter describes the strip-to-strip calibration of NearDet using a method based on that used for the FarDet calibration. The results of the calibration are verified in the same way.

8.1 Introduction

The procedure for strip-to-strip calibration has been shown to work at CalDet [80] and FarDet ([82] and chapter 7), and the work in this chapter follows on from chapter 7. The success of strip-to-strip calibration at these detectors is the incentive for testing the applicability of the iterative method to NearDet. Cosmic ray muon MC is used to look for potential systematic biases in this procedure, then cosmic ray data is used to perform a strip-to-strip calibration.

The geometry of NearDet is more complicated than that of CalDet and FarDet, and is described in section 4.3.3. The aim is to calibrate only the calorimeter section of NearDet (front 120 planes) with cosmic ray muons using this method, since these are the only planes in which energy measurements will be made. The first 5 planes of the spectrometer section will also be calibrated, but beyond this it is not possible to collect enough cosmic ray muons to perform a calibration¹ — beam muons must be used.

8.2 Cosmic Muon Data Sample

There are two data sets used in this analysis: simulated MC and real cosmic muon data. The real data was taken from early July 2005 and there are, on average, 6,000 muons per strip from this sample. The MC simulation has around 240 entries per strip and an average, corrected detector response of 707 ADC.

¹Cosmic ray muons come from above and all NearDet tracks must start in the calorimeter section in reconstruction, so there are relatively few track hits outside of the first few planes in the spectrometer section.

8.2.1 Track Distributions & Cuts

The distribution of track lengths in the MC is shown in figure 8.1. The shape of the distribution agrees with data. A cut is placed so that the minimum track length used in the analysis is 8 planes long and with a minimum of 6 track-like planes (planes that have the number of strip hits and energy deposition consistent with only a muon track having passed through it). Furthermore, to improve track quality, the track must start in the calorimeter section, there must be no showers along the track and only the first track in a snarl is used. Since light is lost at either end of the scintillator strips, a fiducial volume cut of 15 cm is placed around either end of strips.




Figure 8.1: Distribution of track lengths through the detector in the MC sample. Most tracks are 3–4 m long.

The track angle distributions with respect to strip orientation and in Cartesian co-ordinates are shown in figures 8.2, 8.3 and 8.4. Although a path-length correction is performed, track angle cuts are also made. There is a minimum track angle cut around the beamline so that no beam muons enter the sample and a maximum track angle, set so that the average path-length through strips is less than 2.5 cm from a given set of $\frac{dU}{dz}$ and $\frac{dV}{dz}$ angles. This is described in more detail later in section 8.3.3.

In figure 8.2 there is a bias in the U-coordinate with respect to the V-coordinate. This is caused by the bump in the tails of the distribution because more U than V angles are mis-reconstructed as upward rather than downward going. These distributions are just the U and V angles (cosines) for all tracks, so the number of tracks is the same in each sample.

For an explanation of this, consider figure 8.5, for example, and note that planes 1 and 121, the spectrometer boundaries (where all tracks start in reconstruction), are both U-planes. Also note that the lowest three strips in partial planes are lower for U-planes than V-planes. Tracks corner-

clipping the detector low down on either face travel through fewer planes and, on average, travel through more U-planes than V-planes. These tracks have the least efficient timing calibration and are therefore most likely to be mis-reconstructed as upward going. E.g. a track could go through 6 planes, but only two V-planes, therefore its U-angle has a high likelihood of being mis-reconstructed.




Figure 8.2: Distribution of track directions with respect to strip orientations from the MC sample. Strips in both orientations see the same direction spectrum. The high-end bump is due to mis-reconstruction. The distribution looks similar for data (figure 8.3).

The geometry of NearDet is such that different strips see different numbers of entries, and strip-plane maps show how many entries each strip sees from the muon samples (figures 8.5 and 8.6). The calorimeter section and the first few planes of the spectrometer are shown in the plots. Although the strip numbering convention makes it unclear, strips in the centre of the detector in full planes see the most hits. Strips at the edges, e.g. where there is no overlap between U and V oriented planes, and at the corners of the detectors see fewer hits. Strips in the spectrometer section only see tracks that have ranged in from the calorimeter, and since most muons are downward-going, these planes see few hits.

8.2.2 Energy Cuts

Cuts are also placed on a hit-by-hit basis on the light deposited. The light deposited by track hits from all tracks in the sample is shown in figure 8.7. There is an exponential drop-off in hits from the peak at around 700 ADC.

There is a tracking cut of 200 ADC placed on hits as a criterion for them to be used by the tracker. Hits with light levels below this cut-off are added back to the track later during track fitting. Furthermore, a minimum light level cut of 0.3 p.e. is applied to hits in this sample, since




Figure 8.3: Distribution of track directions with respect to strip orientations from the data sample. Strips in both orientations see the same direction spectrum. The high-end bump is due to mis-reconstruction. The distribution looks similar for MC (figure 8.2).




Figure 8.4: Track direction spectra with respect to the x , y and z axes. Most tracks come from above, since they are cosmic, and there is a cut around the beam axis that excludes horizontal muons. The high-end bump on the y distribution is due to mis-reconstruction.




Figure 8.5: The number of entries seen by each strip-end in the MC sample in a strip vs. plane map representation of the calorimeter section and the first few planes of the spectrometer. The central regions of the detector see the most hits. The strip numbering convention causes the offsets between U and V planes.




Figure 8.6: The number of entries seen by each strip-end in the data sample in a strip vs. plane map representation of the calorimeter section and the first few planes of the spectrometer. Compare with figure 8.5.




Figure 8.7: Distribution of hits associated with a track in raw ADC counts from the MC sample. Note the exponential drop-off in the high-end tail.

light levels below this are sparsified away and are reconstructed by the separate sparsification correction in this analysis (section 8.3.5). Figure 8.8 shows the distribution of hits in p.e. on 5-p.e. light level strips, and the tracking cut at around 2 p.e. can be seen, although it is smeared by the applied gain correction.

A maximum light level cut of 8 times the strip light-level (in p.e.) is placed on hits in the sample, since any hits greater than this are unlikely to have come from a minimum ionising muon. It can be seen from figure 8.8 that this cut removes little data from the tail. When analysing data with few hits, a hit a long way into the tail of the distribution increases the RMS of the distribution disproportionately and can also skew the mean of the distribution to values that are too high. Therefore it is necessary to make a maximum value cut, and this cut must be proportional to the strip response. (A better method would be to use a mean truncated by a fraction of the entries from the high end.)

8.2.3 Other Cuts

Once the cuts already mentioned have been applied, the response histograms for strips with different light levels have the form of those in figure 8.9. It is clear from this figure that there are too many single photo-electron noise hits in the distribution, skewing the distributions to too low values. Therefore a further cut at around 1-p.e. must be placed to remove the extraneous hits in the single p.e. peaks. It was not necessary to apply such a cut at FarDet.




Figure 8.8: The distribution of light deposited at strip-ends with a calibrated light level of 5-p.e.. The shape of the underlying Landau distribution of energy loss in the scintillator is convolved with photo-electron statistics. The sharp cut-off at around 3 p.e. is due to a 200 ADC tracking cut in the reconstruction.




Figure 8.9: The distribution of hits on different light level strip-ends, without a cut to remove single p.e. noise from the sample. The 1-p.e. peak can be seen to dominate the distributions. The blue histogram is at 2-p.e. and the black histogram at a light level of 9-p.e., with the other histograms taking the integer p.e. intervening values.

8.3 Method

The method used here is based on that used for FarDet (section 7.2). The only difference (apart from the cuts applied) is that the value used for the first iteration is a nominal detector average of 5 p.e., rather than 4 p.e. at FarDet. Figure 8.10 shows the convergence of the iterations to the value of 1.29 for NearDet strips.




Figure 8.10: Actual values of strip-ends in the MC sample to the convergence value of 1.29. The iterations converge successfully, even with the limited amounts of data available in the MC sample. This shows that the iterations converge properly.

8.3.1 Gain Correction

The gain correction converts the data from ADC-like units to p.e.-like units so that Poisson statistics can be applied to the light seen, and the p.e.-based zero and sparsification corrections can be applied. The distribution of strip-end gains used for this conversion is shown in figure 8.11.

As at FarDet (section 7.4.1), the gain correction narrows the distribution of hits as it acts as part of the strip-to-strip calibration (figure 8.12), i.e. gain differences are part of the strip-to-strip calibration. After the iterative procedure, in order to get an ADC-like value for the final calibration constants, the same gain numbers are used to correct the constants back from p.e. to ADC, i.e. this strip-to-strip calibration corrects gain differences between strip-ends.

8.3.2 Attenuation Correction

The attenuation correction is performed to correct hits at different points along strips to their value at the centre of that strip. Since NearDet strips are offset from the centre of the detector by different amounts, this correction removes systematic responses that could arise from this, and




Figure 8.11: Distribution of PMT gains for strip-ends in the sample from the database. The nominal value, when none was present in the database, was 80 ADC/p.e. (hence the spike).




Figure 8.12: The ADC distributions of deposited light before and after the gain correction. Beforehand (black), the 200 ADC tracking cut can be clearly seen in the sample. Afterwards (red), the distribution is much narrower and this cut is smeared. The gain-corrected ADC unit is a scaled version of p.e..

from different lengths of strips, especially between partial and full planes. The variation of strip responses across the detector is shown in figure 8.13. If the attenuation correction does not come out to a sensible value, the hit is cut from the sample.




Figure 8.13: The attenuation corrections applied to the sample against position along the strip in both the U and V views from the MC sample. The values are not all 1 at the centre because of partial and full plane offsets. The tracker sometimes reconstructs hits as being outside the physical volume of the detector: these are cut from the sample.

Figure 8.14 shows the variation in strip response as a function of distance from the centre of the detector for MC. Before an attenuation correction is applied, the response at the near readout end of the strip can be up to double that at the far end. After the correction, there is no significant difference between responses across the length of strips in either the U or V views.

The mapper used to correct real data is not quite as accurate: two versions of this are shown in figure 8.15, highlighting the problems that can be caused to the strip-to-strip calibration by poor calibrations elsewhere in the chain. The data on the left comes from the original mapper, and on the right is the improved mapper fit, which was used for the data samples in this study. It can be seen that problems remain in the attenuation calibration at the 5% level, but this should only cause a small systematic effect on the strip-to-strip calibration, since most strips have hits along most of their length, and so the effect does not differ between different strips. It could, however, cause an offset in calibration constants between partially and fully instrumented strips, since these are systematically of different lengths.




Figure 8.14: The distribution of hits as a function of distance along strip (MC). When there is no attenuation correction applied, the light levels are higher towards the read-out end of the strips. The drop-off at either end is due to the light lost at the end of the scintillator strips (this is cut out of the sample used). After the attenuation correction, the response across the length of the strips is flat (bottom right), even on a zoomed scale there is no residual dependency in this MC sample (top right).




Figure 8.15: The distribution of hits as a function of distance along strip for two different mappers used on data, zoomed and unzoomed, after attenuation correction. The newer mapper, used for the data in this analysis, is shown on the right. There are still residual effects after attenuation correction at the few percent level in both cases.

8.3.3 Path-length Correction

The path-length correction takes account of the effect from muons entering the strips at different positions along the cross section of the strip and at different angles. Figure 8.16 shows the variation of deposited light with track angle, where the amount of light seen from a muon hit can double over the range of 1.1–2.5 cm, the range of most tracks seen and of those tracks used in the sample. Corner clipping means the path-length is less than $\frac{ds}{dz}$, the path-length of an infinitely long and wide strip (figure 8.17).




Figure 8.16: The variation of energy deposited with track angle highlights the need for a path-length correction, so that all strips see the same distribution of track hits

The distribution of path-lengths at NearDet is shown in figure 8.18. The minimum path-length is 1.1 cm through the strips for cosmic muons, due to the beamline cut. Figure 8.19 shows how good the correction is in correcting the deposited light as a function of average path-length through the strip. Note that the x -axis of figure 8.19 differs from that of figure 8.16, since the latter shows track angle rather than average distance through the strip (in the former).

There is still a small, residual effect in the path-length correction, and this is why the maximum track path-length used in this analysis is 2.5 cm.

8.3.4 Zero Correction

The zero correction takes account of muons travelling through the scintillator where the light produced by the energy deposit does not produce any photoelectrons at the PMT photocathode. The several possible methods of applying this correction were discussed in section 7.2.1. However, only the iterative method is appropriate for NearDet since the detector only has single-ended read-out.

The size of the zero correction is typically around 3%, and can be seen for data in figure 8.20. The shape of the distribution is asymmetric since strips at high light levels have a lower probability of getting a zero (e.g. 1%). Low light level strip-ends have a much higher zero fraction, up to 20%, and the difference has exponential terms (equations 7.7 and 7.8).




Figure 8.17: The average path-length through a strip (ds) for a given set of $\frac{dx}{dz}$ and $\frac{dy}{dz}$ against track angle $\frac{ds}{dz}$ shown for MC. The maximum ds is along the $x = y$ line.




Figure 8.18: The path-length corrections applied to the sample. Any tracks with angles that would give an average path-length through the strip of more than 2.5 cm have been cut. There is also a cut on horizontal muons, so the minimum path-length is 1.1 cm.




Figure 8.19: Strip hit distribution as a function of average path-length through the strip, before and after path-length correction, for U and V planes. The profile on the left-hand plots appears flatter than reality as it excludes hits that are outside the axis range. The path-length correction noticeably flattens the distribution (bottom right), but there is still a residual 10% effect over the range used (top right).




Figure 8.20: The distribution of fractional zero correction weights applied to hits. High light level strip-ends have few zeros, as the probability has exponential light level terms (equation 7.7).




Figure 8.21: Strip-end means (average hit), with and without zero correction. When the zero correction is applied, the strip-end means are reduced, here by around 5% on average.

The difference between zero-corrected and uncorrected distributions in terms of average strip hit light level in MC is shown by figure 8.21.

8.3.5 Sparsification Correction

This correction accounts for the part of the one photo-electron peak not seen due to sparsification [79] and its application is also discussed in more detail in section 7.4.5, for the FarDet case. The effect of this correction is almost negligible for most strips at NearDet, because the strip-end light levels are higher than at FarDet since the readout fibres are shorter. Figure 8.22 shows that the average effect of this correction on the calibration constants (in p.e.) is just 0.1%.




Figure 8.22: The distribution of sparsification correction weights applied to hits. It is much smaller than the zero correction and negligible for most strip-ends.

8.4 Results

The results of the calibration describe the response of the detector and estimate the statistical and systematic accuracy of the calibration.

8.4.1 Photoelectron Response

The only strip-to-strip variation that does not show up by looking at response differences in p.e. is PMT gain. Since the gain variation (figure 8.11) is a major strip-to-strip difference, and can be measured by light injection “gain curves”, it can be useful to deconvolve its effects on the strip-to-strip calibration constants in order to highlight other response patterns.

The strip-plane response maps of the results in terms of p.e. for MC and data are shown in figures 8.23 and 8.24, showing the calorimeter section and the first few planes of the spectrometer.

They show no definite patterns in strip response across the detector, although some consecutive sets of low or high output strips could be symptomatic of poor scintillator or connector efficiencies in those modules.




Figure 8.23: The response map of the calorimeter and the first few planes of the spectrometer in terms of p.e. (MC). This deconvolves gain differences from the results.

The spread of strip responses in p.e. is shown in figure 8.25 for the calorimeter section and the whole detector and the distribution of the statistical uncertainties on those is shown in figure 8.26. These show that the average response of strips is about 5.9 p.e. at NearDet, compared to about 3.9 p.e. at FarDet. The statistical uncertainty on the constants is much lower than at FarDet due to the higher data rate, and this relative uncertainty is around 0.8% on average for calorimeter strips in the data sample.

Some distributions of light seen by typical strips in different parts of NearDet for data and MC are shown in figures 8.27 and 8.28.

8.4.2 Validation

The results of the strip-to-strip calibration can be verified using MC to look for systematic errors and checked with data to show that the constants work in practice. There are some residual effects in the calibration, and fixes for these are described.

MC Tests

For the MC sample, the truth is known, so it can be compared to the calculated calibration constant and the accuracy of this method can be quantified. Figure 8.29 shows the distribution of true strip variations in the MC, the distribution of calibration constants and the difference between the two




Figure 8.24: The response map of the calorimeter and the first few planes of the spectrometer in terms of p.e. (data).




Figure 8.25: The distribution of calibration constants in terms of p.e. (data). The calorimeter section and whole detector are shown separately. Longer strips tend to have a lower response (more attenuation), so the spectrometer section increases the low light level tail. Also, MINOS scintillator is efficiency was measured before installation, and the lower light yield modules used for the spectrometer.




Figure 8.26: The relative statistical uncertainty in the p.e. calibration constants (data). The uncertainties for the whole detector and just the calorimeter section are shown separately. Lower statistics in the spectrometer section leads to poorer precision.




Figure 8.27: Some typical strip-end histograms from different parts of the detector from the MC sample, with an average of 240 entries per strip.




Figure 8.28: Some typical strip-end histograms from different parts of the detector from the data sample. The higher number of entries compared to figure 8.27 means that the distinctive histogram shape is discernable.

on a strip-by-strip basis. The distribution in strip responses is around 25%, as is the distribution in calibration constants, which is necessary if the calibration constants fully correct for all the strip-to-strip differences. On average the truth and calculated values agree in the MC sample (the mean is 1.002 in figure 8.29) and the width of the distribution is (4.5%) reflects the available statistics of the MC sample (240 entries per strip).




Figure 8.29: The spread of true strip-end response differences and calculated calibration constants are shown from the MC sample. The difference between truth and calculated constant is shown in red and its width (4.5%) is statistically dominated, showing good agreement.

Using the iteration convergence value of 1.29, figure 8.30 shows the relative discrepancy ($\frac{A}{B}$) between truth and calibration constant as a function of true strip light-level. This shows that the systematic error is $<1.0\%$ over the range of responses, with calibration constants in the range 0.4–2.0 times the detector average.

Figure 8.31 shows the normalised distributions of hits in p.e. for different light level strips, ranging from 2–9 p.e.. The strips with at the lowest light-level, around 2 p.e., corresponding to about a third of the detector average, have a narrower distribution of hits than strips at higher light levels. This accounts for the systematic mis-calibrations of around 7% at the low end of figure 8.30, although this affects few channels (figure 8.25).

Iteration Convergence Value & Truncation

In the FarDet case (section 7.6), an investigation was made into the ratio between MPV and mean for the strip-end histograms in order to improve the systematic error on the calibration constants. In the NearDet case, the discrepancy between truth and calculated constant is not greater than the expected statistical scatter over almost the full range of constants. There is no need to tune the convergence value away from 1.29 (i.e. tuning the scale of the zero reconstruction), since this value is appropriate for all but the lowest light level strips.




Figure 8.30: The gradient of the profile of calculated calibration constant as a function of true light level has a gradient of 1. This scatter plot shows the residual between the calculated constants and the true light level. The profile shows that all but the lowest light level strip-ends (< 0.5) have a systematic uncertainty of $< 1\%$ with this method.




Figure 8.31: The distributions of hits in p.e. on different light level strip-ends. The shape of the histogram changes slightly between different light level strip-ends.

The truncation of the histograms applied at FarDet does not need to be applied here because there is significantly more data at NearDet. It is not necessary to use such a method to reduce the statistical error and doing so may even be detrimental, since truncation may introduce small systematic biases.

Patterns of Miscalibration

Looking at the discrepancy between truth and constant as a function of position in the detector is another powerful way to look for systematic errors in this method, and this is shown in a strip-plane map for the calorimeter in figure 8.32. No discernable patterns (other than those of broken electronics) are discernable in the range shown, implying that the discrepancy is dominated by statistical scatter.




Figure 8.32: The discrepancy between truth and calibration constant is plotted for the calorimeter. This gives another good way for looking for potential systematic errors in the calibration by looking for patterns of discrepancy on the map in terms of position in the detector. No such patterns can be seen except for a region with low statistics between strips 10–20 in plane 19.

Data Validation

Confirmation that the calibration works on real data comes from analysing the energy deposition of stopping muons: a properly working strip-to-strip calibration reduces the response difference between different regions of the detector.

Figure 8.33 shows the cross-section of detector response at different points in the calibration chain. The raw values have large differences between strips, but after the strip-to-strip correction

is applied, these effects appear reduced. Once the attenuation correction is applied, which corrects for differences along strips, the response looks flat across the detector (especially in the U-view). Some residual differences remain, since the mapper used was the old mapper (mentioned earlier), and patterns can be seen running along the strips in the V-view.

Pathological Strip-Ends

Strip-ends that have too few entries or will not converge to a stable value in the iteration can be “fixed” to produce a best guess of true light level. The strip-end can be fixed using adjacent planes. The number of muons passing through the strip can be estimated by taking the average of the number of muons passing through the strip with the same number (0–96) on the planes either side of the plane with the bad strip. The correct number of zero hits are then added to the histogram and a new light level can be calculated. If the strip-ends are dead and no other information is available, setting the light-level to the detector-average value gives the best estimate of light level: the strip-ends may be subsequently fixed (in the case of faulty electronics, for example).

Figure 8.34 shows the average raw hit as a function of the calculated calibration constant. Strips outside of the main swathe have been fixed using this routine. The profile to the scatter also shows systematic variations with light level, which, together with the spread, illustrate the improvement of using the full, iterative method over just using the average seen muon response.




Figure 8.33: The cross-sectional variation in detector response at various stages of calibration. The plots show muon energy units and are made with stopping muons. The top plots show raw response and the middle plots show the response after strip-to-strip calibration. The bottom plots should show no systematic variation if both the strip-to-strip calibration and mapper-based attenuation corrections are working, but there are patterns indicative of poor attenuation correction, especially in the bottom right plot (V-view) [84].




Figure 8.34: When the raw hit average is compared to the fully calculated calibration constant, there is an almost linear relationship, but with some scatter. This scatter and non-linearity proves the necessity for a method more complicated than just using the average raw hit, such as the one used here.

Chapter 9

SUMMARY

This chapter summarises this thesis: neutrino physics, the role of MINOS in it and the part the work described in this thesis plays in MINOS. Some of this section is taken from [85].

9.1 Contemporary Neutrino Physics

Neutrino oscillations have become accepted natural phenomena from the results of experiments that have taken place over the last decade, most notable of which are Super-Kamiokande [19], SNO [17] and KamLAND [22]. In the “atmospheric sector” (oscillations of ν_μ) the data taken by Super-Kamiokande and K2K [27] has shown maximal mixing between the mass eigenstates 2 and 3. In the “reactor sector”, experiments such as Chooz [28] have shown minimal mixing between the mass eigenstates 1 and 2.

Experiments are needed that can address the hot topics in neutrino physics, summarised as “unanswered questions” in section 2.6.1. Maximal 2–3 mixing, minimal 1–3 mixing and the K2K proof of the oscillation hypothesis must all be confirmed. A measurement must also be made of whether there is CP violation in the neutrino sector. MINOS is an experiment that can address all of these topics.

9.2 The MINOS Project

The Main Injector Neutrino Oscillation Search is a long baseline experiment using muon neutrinos from Fermilab’s NuMI beam [86]. The experiment consists of two identical, magnetised, tracking calorimeter detectors: a near detector, 1 km from the target, at Fermilab (IL, USA); and a far detector at Soudan (MN, USA), 735 km from the target. The experiment has already measured atmospheric neutrinos since 2003 at the far detector, and the beam–phase of the experiment started in February 2005.

9.3 Beam, Beamline and Protons

120 GeV protons from Fermilab’s Main Injector are fired at a graphite target, producing pions and kaons. These intermediate particles are focused by magnetic horns and decay into (mainly) ν_μ

along the 675 m decay tunnel. The total distance from the target to the near detector is 1.0 km, an absorber covers much of the intervening distance to decrease the muon flux at the near detector. The magnetic horns can be moved to low, medium and high energy configurations, so the neutrino energy spectrum can be tuned to physics requirements.

The first neutrinos to MINOS from the NuMI beam were sent in January 2005, and the first beam neutrino events were seen at the near detector in January and then at the far detector in March. Apart from 3 weeks of beam shut-off in March and April due to a cooling water leak in the target, the beam has been running reliably since. The intensity of the protons fired at the target had reached 2.1×10^{13} per pulse by July, with a beam pulse every 3 s. The beam-line has been tested with up to 2.5×10^{13} protons per pulse in February, and this is the intensity goal with pulses every 2 s by the end of 2005. With this proton intensity, MINOS expects to see 1.1×10^4 neutrino interactions at the near detector every day.

9.4 Detectors

The detectors consist of 6 cm composite planes of 1 in thick steel, 1 cm thick solid scintillator [87] and an air gap (for mechanical reasons). The scintillator is divided up into 4.1 cm wide strips and alternating planes have these strips oriented orthogonal to each other at 45° to vertical: this gives MINOS three-dimensional resolution. Light is collected from the scintillator using wavelength-shifting optical fibres running in a groove along the scintillator strips, and these are read out by multi-anode photomultiplier tubes [88].

The near detector has a total mass of 0.96 kton and is 282 planes long. The detector has a 1.5 T magnetic field centred away from the beam axis. Every 5th plane in the detector is fully instrumented (96 strips across). In the front 120 planes, the calorimeter section, the intervening planes are partially instrumented (64 strips across). This section is used for measuring shower energies. The back 160 planes constitute the spectrometer section, and there is no read-out between fully instrumented planes. This back section is used to measure high energy muon momenta by curvature in the magnetic field.

The far detector has 485 planes with 192 strips in each plane, it is 8 m across and 31 m long, with a fiducial mass of 5.4 kT. There is a 1.5 T magnetic field running through the centre of the detector which allows it to measure the charge-sign of charged-current atmospheric neutrino events. These can be distinguished from stopping cosmic muons using timing and a veto shield covering the top of the detector. The detector has an angular resolution of 1° and has already measured 10^7 through-going cosmic muon events since it commenced data-taking in 2003. It has detected the shadow of the moon in the direction spectrum of muons with energy >20 GeV, and the first oscillation results from atmospheric neutrinos will be announced soon.

9.5 Data Analysis at MINOS

MINOS is an unique experiment since it can distinguish the charge-sign of muons from neutrino events. By measuring the muon event rates from atmospheric neutrinos, MINOS can measure the difference between μ^+ and μ^- event rates, and hence whether there is any CP violating phase controlling neutrino oscillations.

By looking for ν_e appearance in the ν_μ NuMI beam, MINOS can make a measurement of the “reactor sector” (1–3) neutrino oscillation parameters. If no ν_e appearance is found in the beam, MINOS can halve the Chooz limit towards minimal mixing [28].

MINOS’ *raison d’être*, however, is to measure the “atmospheric sector” (2–3) parameters, by looking for a deficit of ν_μ events in the charged current (CC) channel. An analysis using simple topology cuts to distinguish neutral current (NC) events from CC events was presented in this thesis, making the first measurement of neutrino oscillations with MINOS using the NuMI beam.

An appropriately detailed analysis measures the difference between the NearDet and FarDet CC spectra as a function of event energy from 3 years of running, and from this MINOS can measure the “atmospheric sector” neutrino oscillation parameters to 10% precision. Such an analysis will also confirm that neutrinos oscillate, since part of the energy spectrum will have oscillated back, if current predictions of the value of these parameters are correct. Analysis of the CC energy spectrum requires good calibration, however.

9.6 Calibration

Calibration is performed in MINOS by a series of corrections for different effects. The first link in the calibration chain is on-board electronics calibration with charge injection. Then PMT gain and “gain drift” calibrations are performed with a light injection system [74]. Once the read out is calibrated, the response of each strip in the detector is normalised with cosmic ray muons. Then the inter-detector calibration is performed using muons that stop in each detector, to ensure the same part of the Bethe–Bloch curve is measured at each detector. Absolute calibration is carried out using Monte Carlo simulations and a special calibration detector to characterise the response to different particles at different energies.

The MINOS detectors are designed to be calibrated to 2% relative and 5% absolute uncertainty in order for systematic errors to be less than statistical errors after 3 years of NuMI beam running. This is an important aspect of MINOS, since MINOS will compare the energy scales at the near and far detectors in order to extract the atmospheric neutrino oscillation parameters to the 10% precision afforded by the full spectrum analysis.

9.6.1 Light Injection

The MINOS light injection system is used for measuring differences and monitoring changes of the photomultiplier tubes (PMTs) and readout electronics at all three MINOS detectors (calibration, far and near). There are three main requirements for the light injection system:

- “gain curves” to measure the non-linearity of PMT response with incident light. Relative amounts of light illuminate the PMT face via optical fibres from light emitting diodes (LEDs). Twenty points are taken over the dynamic range of the PMT and the relative measurement of light from the LEDs is measured by PIN photodiodes;
- “drift points” to monitor the change of gain over time between gain curves. Light is shone at a default light level point every 3 hours. This corrects mainly for diurnal temperature changes and drifts of gain throughout the month;

- occasional runs to make an absolute gain measurement.

The LI system is also a very effective debugging tool. A package was written during installation of the LI system at NearDet to find problems with the detector, and to analyse what these problems were, using a cascading pattern-finding technique.

A study was undertaken to ensure the linearity of the light input to LI system to ensure that the gain curves linearise the PMTs properly. Different LEDs, emitting light at different wavelengths, were compared until a particular model of UV LEDs was shown to work to the right specification.

Drift points taken using the standard method (mean of 1,000 pulses) can give erroneous values due to hardware changes and problems with monitoring light levels, e.g. not reading out PIN photodiodes properly. The different effects that could cause problems to the LI system were investigated. In order to reduce the size of the offline database tables and overcome the problems discussed, different drift methods based on averaging were proposed. The drift method based on PMT-averaging of gains calculated by photo-electron statistics was shown to be optimal.

9.6.2 FarDet Strip-to-strip Calibration

The strip-to-strip correction is the link in the MINOS calibration chain that deals with normalising any strip-end response differences that can be observed with cosmic ray muons. Factors affecting strip-end response include clear readout fibre lengths, wavelength-shifting readout fibre “pigtailed”, gain, scintillator efficiency (light output), wavelength-shifting fibre collection efficiency and connector efficiency.

In order to properly normalise the detector, the data sample used for each strip-end must be the same. There are several factors that have to be corrected for in the data before this is the case. The main corrections that need to be applied are an attenuation correction (to account for wavelength-shifting fibre length differences along the strips) and a path-length correction so that tracks coming in at different angles with respect to the strip orientation give the same response.

Several methods of measuring the strip-to-strip response using cosmic ray muons are compared for use at FarDet. The iterative method described gives a consistent way of normalising the response of all three detectors. Improvements to the method at FarDet show that it is a viable calibration method to give constants accurate to within 2.0% from about 3 months’ worth of cosmic ray neutrino data. This method is suitable for a rolling calibration using the procedure described.

9.6.3 NearDet Strip-to-strip Calibration

The iterative method has been shown here to work accurately at NearDet. At NearDet, this method is applied to give a viable calibration producing constants accurate to within 1.0% statistical and systematic uncertainty from around 3 days’ worth of cosmic ray neutrino data.

9.6.4 Calibration Summary

Good calibration is vital for MINOS in order for the project to progress from simple $\mathcal{T}$ -test analyses to analyses of the energy dependence of the CC spectrum. The calibration systems described in this thesis are some of those in place to ensure that MINOS meets its calibration requirements, so that the atmospheric sector neutrino oscillation parameters can be measured by the experiment to the 10% specified precision.



Figure 9.1: Cartoon by S. Harris

Bibliography

- [1] W. Pauli, letter addressed to the Tübingen Conference on radioactivity (1930).
- [2] MINOS Collaboration, The MINOS Detectors Technical Design Report, NuMI-337 (1998).
- [3] W. Greiner & B. Mueller, Gauge Theory of Weak Interactions, 3rd ed., Springer-Verlag (2000).
- [4] P. Langacker, J. Erler & E. Peinado, Neutrino Physics, arXiv:hep-ph/0506257 (2005).
- [5] S. Eidelman et al., Phys. Lett. B 592, 1 (2004) and 2005 partial update for the 2006 edition available on the PDG WWW pages (URL: <http://pdg.lbl.gov/>)
- [6] R. N. Mohapatra et al., Theory of Neutrinos, arXiv:hep-ph/0412099 (2004).
- [7] F. Reines and C. L. Cowan, Detection of the free neutrino, Phys. Rev. 92, 830 (1953).
- [8] J. N. Bahcall and M. H. Pinsonneault, What do we (not) know theoretically about solar neutrino fluxes?, Phys. Rev. Lett. 92, 121301 (2004).
- [9] R. Davis & D. S. Harmer, Attempt to Observe the $Cl^{37}(\bar{\nu}, e^{-})Ar^{37}$ Reaction Induced by Reactor Antineutrinos, Bull. Am. Phys. Soc. 4, 217 (1959).
- [10] SAGE Collaboration, A. I. Abdurashitov et al., Phys. Rev. Lett. 77, 4708 (1996).
- [11] GNO Collaboration, M. Altmann et al., Phys. Lett. B 490, 16 (2000).
- [12] GALLEX Collaboration, W. Hampel et al., Phys. Lett. B 388, 384 (1996).
- [13] M. Kobayashi and T. Maskawa, CP violation in the renormalizable theory of weak interaction, Prog. Theor. Phys. 49, 652 (1973).
- [14] B. Pontecorvo, Inverse beta processes and nonconservation of lepton charge, Sov. Phys. JETP 7, 172 (1958).
- [15] M. A. Barker, Hadronic Response of the MINOS Detectors, DPhil thesis, University of Oxford, 2003.
- [16] SNO Collaboration: B. Aharmim et al., Electron energy spectra, fluxes, and day-night asymmetries of B-8 solar neutrinos from the 391-day salt phase SNO data set, (2005).
- [17] SNO Collaboration: Q. R. Ahmad et al., Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory, Phys. Rev. Lett. 89, 011301 (2002).

- [18] Super-Kamiokande Collaboration: Y. Fukuda et al., Evidence for oscillation of atmospheric neutrinos, *Phys. Rev. Lett.* 81, 1562 (1998).
- [19] Super-Kamiokande Collaboration: Y. Ashie et al., *Phys. Rev. Lett.* 93, 101801 (2004) and earlier papers referenced therein.
- [20] L. Wolfenstein, Neutrino oscillations in matter, *Phys. Rev. D* 17, 2369 (1978).
- [21] S. P. Mikheev and A. Y. Smirnov, Resonant amplification of neutrino oscillations in matter and solar neutrino spectroscopy, *Nuovo Cim.* C9, 17 (1986).
- [22] KamLAND Collaboration: K. Eguchi et al., First results from KamLAND: Evidence for reactor anti-neutrino disappearance, *Phys. Rev. Lett.* 90, 021802 (2003).
- [23] KamLAND Collaboration: T. Araki et al., Measurement of neutrino oscillation with KamLAND: Evidence of spectral distortion, *Phys. Rev. Lett.* 94, 081801 (2005).
- [24] Soudan 2 Collaboration: M. C. Sanchez et al., Observation of atmospheric neutrino oscillations in Soudan 2, *Phys. Rev. D* 68, 113004 (2003).
- [25] MACRO Collaboration (M. Ambrosio et al.), Measurement of the Atmospheric Neutrino Induced Upgoing Muon Flux Using MACRO, *Phys. Lett. B* 434, 451 (1998).
- [26] C. L. F. Howcroft, Atmospheric Neutrinos In The MINOS Far Detector, PhD thesis, University of Cambridge (2004).
- [27] K2K Collaboration, E. Aliu et al., *Phys. Rev. Lett.* 94, 081802 (2005).
- [28] M. Apollonio et al., Search for neutrino oscillations on a long base-line at the Chooz nuclear power station, *Eur. Phys. J. C* 27, 331 (2003).
- [29] J. Jochum, The search for θ_{13} with the Double-Chooz experiment, IDM 2004 (2004).
- [30] Daya Bay Neutrino Experiment: <http://bes.ihep.ac.cn/dayawane/>
- [31] The Braidwood Experiment: <http://braidwood.uchicago.edu/>
- [32] MINOS Collaboration, First Observations of Separated Atmospheric Muon Neutrino and Muon Anti-Neutrino Events in the MINOS Detector, hep-ex/0512036 (2005).
- [33] H. Georgi, H. R. Quinn & S. Weinberg, Hierarchy of Interactions in Unified Gauge Theories, *Phys. Rev. Lett.* 33 (1974).
- [34] M. Winter, Determination of the number of neutrino species from the Z line shape parameters, Prepared for Workshop on Tau Lepton Physics, Orsay, France, 24-27 Sep 1990.
- [35] LSND Collaboration: A. Aguilar et al., Evidence for neutrino oscillations from the observation of anti- ν_e appearance in a anti- ν_μ beam, *Phys. Rev. D* 64, 112007 (2001).
- [36] KARMEN Collaboration: B. Armbruster et al., Upper limits for neutrino oscillations $\bar{\nu}_m \rightarrow \bar{\nu}_e$ from muon decay at rest, *Phys. Rev. D* 65, 112001 (2002).

- [37] MiniBooNE Collaboration: H. L. Ray, Current status of the MiniBooNE experiment (2004).
- [38] WMAP Collaboration (D. N. Spergel et al.), First Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters, *Astrophys. J. Suppl.* 148, 175 (2003).
- [39] SDSS Collaboration (C. Stoughton et al.), The Sloan Digital Sky Survey: Early Data Release, *Astron. J.* 123, 485 (2002).
- [40] C. Kraus et al., Latest results from the Mainz Neutrino Mass Experiment, *Nucl. Phys. A* 721, 533 (2003).
- [41] V. M. Lobashev et al., Study of the Tritium Beta-spectrum in Experiment Troitsk ν -mass, *Prog. Part. Nucl. Phys.* 48 (2002).
- [42] KATRIN Collaboration, <http://ik1au1.fzk.de/>
- [43] H.V. Klapdor-Kleingrothaus, et al., Evidence for Neutrinoless Double Beta Decay, *Mod. Phys. Lett. A* 16, 37 (2001).
- [44] T2K Collaboration: Y. Hayato, T2K at J-PARC, *Nucl. Phys. Proc. Suppl.* 143, 269 (2005).
- [45] NOvA Collaboration: D. S. Ayres et al., NOvA proposal to build a 30-kiloton off-axis detector to study neutrino oscillations in the Fermilab NuMI beamline, (2004).
- [46] MINERvA Collaboration: S. Boyd, MINERvA: A high statistics neutrino scattering experiment in the NuMI neutrino beam, *Nucl. Phys. Proc. Suppl.* 139, 311 (2005).
- [47] M. Thomson, Status of the MINOS Experiment, *Nucl. Phys. B (Proc. Suppl.)* (2005).
- [48] R. Brun & F. Rademakers, ROOT: An object oriented data analysis framework, *Nucl. Instrum. Meth. A* 389, 81 (1997).
- [49] The MINOS Off-line Software User's Manual (R1.18), http://www-numi.fnal.gov/offline_software/srt_public_context/WebDocs/WebDocs.html (2005).
- [50] D. Boehnlein et al., Steel Mass and Meltcode Distribution in the MINOS Far Detector, *NuMI-1061* (2004).
- [51] A. Belias et al., The MINOS data acquisition system, *IEEE Trans. Nucl. Sci.* 51, 451 (2004).
- [52] G. Drake, J. Dawson, and C. Nelson, Overview of the front end electronics for the MINOS Near Detector, *NuMI-628* (1999).
- [53] N. Tagg et al., Performance of Hamamatsu 64-anode photomultipliers for use with wavelength-shifting optical fibres, *Nucl. Instrum. Meth. A* 539, 668 (2005).
- [54] H. Gallagher, The NEUGEN neutrino event generator, *Nucl. Phys. Proc. Suppl.* 112, 188 (2002).

- [55] GEANT - Detector Description and Simulation Tool, CERN Program Library, Long Writeup, W5013.
- [56] J. Hylen et al., NuMI Technical Design Handbook, Internal NuMI report (2003).
- [57] D. A. Petyt, Low Dm^2 sensitivity of the T-test, NuMI-481 (1999).
- [58] D. A. Petyt, A Study of Parameter Measurement in a Long Baseline Neutrino Oscillation Experiment, DPhil thesis, University of Oxford (1998).
- [59] D. A. Harris, Systematic Uncertainties on m^2 from Neutrino Physics, using Calorimetric Energy Reconstruction, NuMI-948 (2003).
- [60] M. Diwan et al., A study of $\nu_\mu \rightarrow \nu_e$ in MINOS, NuMI-714 (2001).
- [61] B. J. Rebel, Neutrino-Induced Muons in the MINOS Far Detector, PhD thesis, Indiana University (2004).
- [62] MIPP Collaboration: H. Meyer, Physics of the MIPP Experiment, Nucl.Phys.Proc.Suppl. 142 453-458 (2005).
- [63] A. Cabrera Serra, Systematic Comparison of the MINOS Near and Far Detector Readout Systems, DPhil thesis, University of Oxford (2005).
- [64] R. Nichol, Calibration of the MINOS Detectors, PhD thesis University College London (2003).
- [65] P. Adamson, An LED Calibration System for the MINOS Long Baseline Neutrino Oscillation Experiment, DPhil thesis, University of Sussex (2001).
- [66] J. J. Hartnell, Measurement of the Calorimetric Energy Scale in MINOS, DPhil thesis, University of Oxford (2005).
- [67] P. L. V. Vahle, Electromagnetic Interactions in the MINOS Detectors, PhD thesis University of Texas at Austin (2004).
- [68] M. A. Kordosky, Hadronic Interactions in the MINOS Detectors, PhD thesis University of Texas at Austin (2004).
- [69] A. Blake, A Study of Atmospheric Neutrino Oscillations in the MINOS Far Detector, PhD thesis, University of Cambridge (2005).
- [70] P. Adamson, et al., The MINOS Light Injection Calibration System, Nucl. Instrum. Meth. A492-3,325 (2002) & NuMI-743 (2002).
- [71] R. Nichol, Light Injection Calibration at the Calibration Detector, NuMI-974 (2003).
- [72] E. Falk Harris, Changing the Light-Injection Flashing Frequency at the Far Detector from 50 Hz to 300 Hz, NuMI-1065 (2004).

- [73] P. Adamson et al., The MINOS light injection calibration system, Nucl. Instrum. Meth. A492, 325 (2002).
- [74] P. Adamson et al., On the linearity of the MINOS light-injection calibration system, Nucl. Instrum. Meth. A521, 361 (2004).
- [75] P. Adamson et al., The MINOS Calibration Detector, to be published in Nucl. Inst. Meth. A.
- [76] P. A. Symes, L. I. Hardware & L. I. Analysis at NearDet: Current Status, NuMI-1357 (2004-2006).
- [77] P. Adamson et al., Photoelectron Counting by Several Methods, NuMI-661 (2000).
- [78] R. Lee & S. Seun, Calibration of the MINOS Detectors Using Cosmic Ray Muons, NuMI-694 (2000).
- [79] C. B. Smith & R. Nichol, Pathlength Correction and Zero Reconstruction in the MINOS Detectors, NuMI-971 (2003).
- [80] R. Nichol and C. Smith, Strip-to-strip Calibration at the Calibration Detector, NuMI-0972 (2003).
- [81] Christopher Smith, Various presentations to the CalDet and Calibration Working Groups (since 2003).
- [82] P. A. Symes, Response Normalisation and Strip-to-Strip Correction at FarDet Using Cosmic Ray Muons, NuMI-1358 (2006).
- [83] C. B. Smith, Calibration of the MINOS Detectors and Extraction of Neutrino Oscillation Parameters, PhD thesis University College London, 2002.
- [84] Jeffrey Hartnell, Various presentations to the Calibration Working Group (2005).
- [85] P. A. Symes, MINOS: Design & Detectors, Proceedings of the 28th International Workshop on Neutrino Physics (2005).
- [86] A. G. Abramov, et al., Nucl. Inst. Meth. A485, 209-227 (2002).
- [87] P. Adamson, et al., IEEE Trans. Nucl. Sci. 49, 861-863 (2002).
- [88] J. Oliver, N. Felt, G. Feldman, A. Lebedev, and R. Lee, Design and performance of the readout system of the MINOS Far Detector, IEEE Trans. Nucl. Sci. 51, 2193 (2004).

Appendix A

The MINOS Collaboration

The MINOS Collaboration Main Injector Neutrino Oscillation Search November 2005

P. Adamson^{8,15}, C. Andreopoulos²¹, K.E. Arms¹⁶, R. Armstrong¹¹, D.J. Auty²⁴, S. Avvakumov²³,
D.S. Ayres¹, B. Baller⁸, B. Barish⁵, P.D. Barnes Jr.¹⁴, G. Barr¹⁸, W.L. Barrett³⁰, E. Beall^{1,16},
B.R. Becker¹⁶, A. Belias²¹, R. H. Bernstein⁸, D. Bhattacharya¹⁹, M. Bishai⁴, A. Blake⁶, B. Bock¹⁷,
G.J. Bock⁸, J. Boehm⁹, D.J. Boehnlein⁸, D. Bogert⁸, P.M. Border¹⁶, C. Bower¹¹,
E. Buckley-Geer⁸, C. Bungau²⁴, J.D. Chapman⁶, D. Cherdack²⁷, S.K. Chernichenko²⁰,
S. Childress⁸, B.C. Choudhary⁸, J.H. Cobb¹⁸, A.J. Culling⁶, J.K. de Jong¹⁰, M. Dierckxsens⁴,
M.V. Diwan⁴, M. Dorman^{15,21}, D. Drakoulakos², T. Durkin²¹, A.R. Erwin³², C.O. Escobar²⁸,
J. Evans¹⁸, E. Falk Harris²⁴, G.J. Feldman⁹, T.H. Fields¹, R. Ford⁸, M.V. Frohne³,
H.R. Gallagher²⁷, G.A. Giurgiu¹, A. Godley²², J. Gogos¹⁶, M.C. Goodman¹, P. Gouffon²⁹,
E. Grashorn¹⁷, N. Grossman⁸, K. Grzelak¹⁸, A. Habig¹⁷, D. Harris⁸, P.G. Harris²⁴, J. Hartnell²¹,
E.P. Hartouni¹⁴, R. Hatcher⁸, K. Heller¹⁶, C. Howcroft⁵, J. Hylen⁸, D. Indurthy²⁶, G.M. Irwin²³,
M. Ishitsuka¹¹, D. Jaffe⁴, C. James⁸, L. Jenner¹⁵, D. Jensen⁸, T. Kafka²⁷, H.J. Kang²³,
S.M.S. Kasahara¹⁶, M.S. Kim¹⁹, G. Koizumi⁸, S. Kopp²⁶, M. Kordosky¹⁵, D.J. Koskinen¹⁵,
S.K. Kotelnikov¹³, A. Kreymer⁸, S. Kumaratunga¹⁶, K. Lang²⁶, J. Ling²², P.J. Litchfield¹⁶,
R.P. Litchfield¹⁸, P. Lucas⁸, V. Makeev⁸, W.A. Mann²⁷, A. Marchionni⁸, A.D. Marino⁸,
M.L. Marshak¹⁶, J.S. Marshall⁶, A. McGowan^{1,16}, J.R. Meier¹⁶, G.I. Merzon¹³, M.D. Messier¹¹,
D.G. Michael^{5†}, W.H. Miller¹⁶, S.R. Mishra²², C. Moore⁸, J. Morfin⁸, L. Mueale¹⁶, S. Mufson¹¹,
S. Murgia²³, J. Musser¹¹, D. Naples¹⁹, J.K. Nelson³¹, H. Newman⁵, T.C. Nicholls²¹,
J.P. Ochoa-Ricoux⁵, W.P. Oliver²⁷, V.A. Onuchin²⁰, T. Osiecki²⁶, R. Ospanov²⁶, J. Paley¹¹,
V. Paolone¹⁹, A. Para⁸, T. Patzak⁷, Z. Pavlovich²⁶, G.F. Pearce²¹, C.W. Peck⁵, E.A. Peterson¹⁶,
D.A. Petyt¹⁶, H. Ping³², R. Piteira⁷, R. Pittam¹⁸, R.K. Plunkett⁸, D. Rahman¹⁶, R.A. Rameika⁸,
T.M. Rauber¹⁸, B. Rebel⁸, J. Reichenbacher¹, D.E. Reyna¹, C. Rosenfeld²², H.A. Rubin¹⁰,
K. Ruddick¹⁶, V.A. Ryabov¹³, R. Saakyan¹⁵, M.C. Sanchez⁹, N. Saoulidou⁸, J. Schneps²⁷,
P. Schreiner³, V.K. Semenov²⁰, P. Shanahan⁸, W. Smart⁸, V. Smirnitsky¹², C. Smith¹⁵, A. Sousa²⁷,
B. Speakman¹⁶, P. Stamoulis², P.A. Symes²⁴, N. Tagg¹⁸, R.L. Talaga¹, E. Tetteh-Lartey²⁵,
J. Thomas¹⁵, M.A. Thomson⁶, G. Tinti¹⁸, I. Trostin¹², V.A. Tsarev¹³, G. Tzanakos², J. Urheim¹¹,
P. Vahle¹⁵, C. Velissaris³², V. Verebrysov¹², B. Viren⁴, C.P. Ward⁶, D.R. Ward⁶, M. Watabe²⁵,
A. Weber^{18,21}, R.C. Webb²⁵, A. Wehmann⁸, N. West¹⁸, C. White¹⁰, S.G. Wojcicki^{23†},
D.M. Wright¹⁴, Q.K. Wu²², T. Yang²³, F.X. Yumiceva³¹, H. Zheng⁵, M. Zois², and R. Zwaska²⁶.

[†] Co-Spokespersons

Argonne¹ - Athens² - Benedictine³ - Brookhaven⁴ - Caltech⁵ - Cambridge⁶ -
Collège de France⁷ - Fermilab⁸ - Harvard⁹ - IIT¹⁰ - Indiana¹¹ - ITEP-Moscow¹² -
Lebedev¹³ - Livermore¹⁴ - UCL-London¹⁵ - Minnesota¹⁶ - Minnesota-Duluth¹⁷ -
Oxford¹⁸ - Pittsburgh¹⁹ - Protvino²⁰ - Rutherford²¹ - South Carolina²² - Stanford²³ -
Sussex²⁴ - Texas A&M²⁵ - Texas-Austin²⁶ - Tufts²⁷ - UNICAMP-Brazil²⁸ -
USP-Brazil²⁹ - Western Washington³⁰ - William & Mary³¹ - Wisconsin³²