






Electron identification in and performance of the ND280 Electromagnetic Calorimeter

by

Antony Carver

Thesis

Submitted to The University of Warwick

for the degree of

Doctor of Philosophy

Physics

March 2010

THE UNIVERSITY OF
WARWICK

Contents

Acknowledgments	vii
Declarations	viii
Abstract	ix
Abbreviations	x
List of Figures	i
List of Tables	xiv
Chapter 1 Introduction	1
1.1 Introduction	1
Chapter 2 Neutrino Physics	3
2.1 Neutrino Phenomenology	3
2.1.1 Neutrino Mass	3
2.1.2 Oscillation Probability	5
2.1.3 Neutrino Oscillations in Matter	8

2.1.4	Current 3 flavour neutrino oscillation model	9
2.2	Neutrino Interaction Physics	12
2.2.1	Charged Current Interactions	13
2.2.2	Neutral Current Interactions	16
2.3	A review of Neutrino Oscillations	17
2.3.1	The Solar Neutrino Problem	18
2.3.2	Neutrino Oscillation Experiments	25
2.4	T2K	31
Chapter 3 T2K and the ND280 Detector		33
3.1	Introduction to T2K	33
3.2	J-PARC Neutrino Beamline	36
3.3	INGRID on-axis detector	37
3.4	Super-Kamiokande	38
3.4.1	Super-Kamiokande Reconstruction	40
3.5	ND280 off axis near detector	42
3.5.1	P0D (π^0 Detector)	44
3.5.2	Fine Grained Detector	46
3.5.3	Time Projection Chamber (TPC)	48
3.5.4	Electromagnetic Calorimeter	49
3.5.5	Side Muon Range Detector	52
3.5.6	Scintillator Detectors	54
3.5.7	Multi Pixel Photon Counters	55

3.5.8	ND280 Electronics	57
3.6	Data Acquisition (DAQ)	61
3.7	ND280 Software Suite	61
3.7.1	oaEvent, oaRawEvent and oaUnpack	62
3.7.2	Monte Carlo simulation	63
3.7.3	Reconstruction	65
Chapter 4	Particle Identification in the ECal	70
4.1	Introduction	70
4.2	Monte Carlo simulation and Particle event types in the ECal .	71
4.2.1	Monte Carlo simulation of ECal	71
4.2.2	Tracks	73
4.2.3	Electromagnetic Showers	75
4.2.4	Hadronic Showers	76
4.3	Identification Techniques	79
4.3.1	Likelihood	80
4.3.2	Artificial Neural Networks	82
4.4	Reconstruction	88
4.4.1	Track Reconstruction	88
4.4.2	Shower Reconstruction	88
4.4.3	Angle Reconstruction	89
4.5	ECal Particle Identification	89
4.5.1	Particle Identification Variables	89

4.6	Particle Identification Technique	104
4.6.1	PID algorithm description	104
4.6.2	Network Training	106
4.6.3	Network Optimisation	106
4.7	Neural Network Performance	109
4.7.1	Neural Network Validation	109
4.7.2	Predicted Efficiency	111
4.8	Conclusions	116
Chapter 5	T9 Testbeam	118
5.1	ECal Testbeam Introduction	118
5.2	CERN T9 Beamline and ECal experimental layout	119
5.2.1	T9 beamline	119
5.3	Testbeam trigger and particle identification	121
5.3.1	Time of Flight	122
5.3.2	Čerenkov counters	130
5.3.3	Determination of beam composition	135
5.3.4	Sample selection	137
5.4	Analysis of testbeam data	138
5.4.1	Data Calibration	139
5.4.2	Cosmic Muon Calibration	141
5.4.3	Comparison of Cosmic Data to Monte Carlo	152
5.4.4	Comparison of electron data with simulation	156

5.5	Conclusions	167
Chapter 6	Electron Neutrino Analysis	170
6.1	Introduction	170
6.2	Neutrino Interactions in ND280	170
6.2.1	Neutrino interactions in the FGD	174
6.3	Electron Neutrino Analysis - Event Selection	176
6.3.1	Lepton Selection	177
6.3.2	Particle Identification	178
6.4	Analysis Performance	186
6.4.1	Systematic Errors	192
6.5	Results	194
6.6	Conclusions	195
Chapter 7	Conclusions	196

Acknowledgments

Thank you to Steve Boyd and Gary Barker for all their help over the past $3\frac{1}{2}$ years. Your wisdom and inspiration has been invaluable during my time on T2K. I would also like to thank everyone in T2K for the help and advice I have received over the years. In particular from the ECal software group and the electron neutrino analysis group.

I would also like to thank those I have worked with at Warwick; Martin, Phil, Leigh and Andy.

Thank you to my Mum, Dad and Grandma for all your support over the years. Finally, thank you to my wonderful fiancée, Alice, for all your understanding over the past few months.

Declarations

This work has been carried out as part of the T2K neutrino oscillation experiment. The first chapter is a review of the theory and current status of neutrino oscillation experiments and the second chapter describes the T2K experiment. The following three chapters describe the author's contribution to the experiment. The third chapter describes techniques used to separate classes of event and then the algorithm designed and implemented by the author. The development of the trigger and identification algorithms used in the testbeam analysis were also original work, as was the data to simulation comparison and energy resolution measurement carried out. The electron neutrino analysis presented in the final chapter was also implemented by the author as part of the T2K electron neutrino analysis group.

Abstract

T2K is an off axis neutrino beam experiment with a baseline of 295 km to the far detector, Super-Kamiokande. The near detector, ND280, measures the flux and energy spectra of electron and muon neutrinos in the direction of Super-Kamiokande. An electromagnetic calorimeter constructed from lead and scintillator surrounds the inner detector. Three time projection chambers and two fine grained scintillator detectors sit inside the calorimeter. This thesis describes the development of a particle identification algorithm for the calorimeter and studies how it can enhance a simple electron neutrino analysis.

A particle identification algorithm was written for the electromagnetic calorimeter to separate minimally ionising particles, electromagnetic and hadronic showers. A Monte Carlo study suggested that the algorithm produced an electron sample with a relative muon contamination of 10^{-2} whilst maintaining an electron efficiency of 80%. Data collected at CERN was then used to make comparisons between the Monte Carlo simulation used to train the particle identification, and experimental data. A reasonable agreement was found between the electron data and the Monte Carlo simulation, given that the available calibration framework was still preliminary. Cosmic data agreed well with simulation. The energy resolution of the DsECal for electromagnetic showers was estimated at $9\%/\sqrt{E}$. An electron neutrino analysis was developed that could be performed on T2K data from the first day of data taking. This analysis anticipated finding $33 \pm 10(sys) \pm 6(stat)$ CCQE electron neutrino events and $92 \pm 28(sys) \pm 10(stat)$ CCnQE electron neutrino events in the FGD after 12 months of nominal running.

Abbreviations

DsECal - Downstream Electromagnetic Calorimeter

FGD - Fine Grained Detector

PID - Particle Identification

SMRD - Side Muon Range Detector

P0D - Pi-0 Detector

CERN - Organisation Européenne pour la Recherche Nucléaire

TOF - Time Of Flight

T2K - Tokai to Kamioka

ND280 - Near Detector 280m

C/N C - Charged/Neutral Current

MLP - Multi Layered Perceptron

CC(n)QE - Charged Current (non) Quasi-Elastic

TPC - Time Projection Chamber

MIP - Minimally Ionising Particle

AMR - Axis Max Ratio

MPPC - Multi Pixel Photon Counter

List of Figures

2.1	Showing charged and neutral current coherent forward neutrino scattering.	8
2.2	The normal hierarchy has two light masses close together and a third much larger mass, following the pattern in the standard model where the particles in the thrid generation are typically much heavier than in the first two. In the inverted hierarchy the two mass states of similar magnitude are much more massive than the third. Figure modified from [13].	10
2.3	Figure from [14] showing the charged current neutrino cross section. Contributions from quasi-elastic scattering, single pion production and deep inelastic scattering are shown separately.	15
2.4	A resonant scattering interaction. A Δ is created which decays back to a pion and a nucleon.	15
2.5	Showing the flux of neutrinos from different reactions within the sun from [13].	20

2.6	Figure from [49] showing the tau and muon neutrino fluxes versus the electron neutrino flux. The total flux predicted by the SSM is shown as a dashed line. The SK result from elastic scattering is also shown.	24
2.7	The ratio of expected to observed neutrino events in the KamLAND experiment as a function of energy, Figure from [54]. Background due to geo-neutrinos has been subtracted.	27
2.8	A combined fit of the solar neutrino mass splitting and mixing angle, Figure from [11]. The SNO data comes from data in [48].	28
2.9	Global limits on solar and atmospheric neutrino parameters from [61]. MINOS and KamLAND provide the best measurements of the atmospheric and solar mass differences respectively.	30
2.10	The reconstructed neutrino spectrum (left) assuming $\Delta m_{atm}^2 = 2.7 \times 10^{-3}$ and $\sin^2 2\theta_{atm} = 1.0$, the hatched area shows non-quasi-elastic component. The Figure (right) shows the ratio of events between the oscillated and unoscillated neutrino cases as a function of energy [18].	32
3.1	The solid black line shows the energy spectrum of on-axis neutrinos in T2K. By changing the angle to 2° (dashed) or 3° (dotted) the peak energy is lowered [18].	35
3.2	Predicted neutrino flux at ND280 (top) and Super-Kamiokande (middle) and the near to far ratio (bottom) for both ν_μ (left) and ν_e (right) [18].	37

3.3	The structure of an individual module with lead and scintillator (left) and the cross shaped arrangement they are assembled in (right) [18].	38
3.4	A diagram of the Super-Kamiokande Detector. Figure adapted from [66].	39
3.5	Event displays from Super-Kamiokande. An electron event with a ‘fuzzy’ ring is shown (top left) and a sharper muon ring (top right). A two ring π^0 event is shown (bottom left) as well as a more complicated event containing three rings (bottom right). The relative timing of PMT hits is indicated by the colour of the square and the magnitude by its size [68].	42
3.6	A schematic view of the ND280 is shown where the P0D is at the upstream side of the near detector and the tracker section is at the downstream end of the detector.	44
3.7	A diagram showing a cross-section of a FGD scintillator bar. The titanium dioxide coating helps to contain light within a given bar. Optical fibre sits in the hole running through the centre of the bar to collect scintillation light [18].	47
3.8	Images from the ND280 technical review [18] showing the design of a TPC gas cage (left) and a simulated field map (right). A good field uniformity of close to 180 V/cm is predicted. Some fluctuations are seen at the edges due to simulation precision [18].	49
3.9	A view of one of the ND280 magnet yokes with air gaps visible (left) and an SMRD ‘slab’ (right) [18].	53

3.10	An end on view of a scintillator bar being illuminated with a UV LED. A hole in the middle to accommodate a fibre can be seen [69].	54
3.11	A diagram showing how a MPPC is connected to a TFB and Trip-T chip. Gain controlling capacitors can also be seen. . . .	59
4.1	Showing the energy spectra used to produce Monte Carlo samples to design and train PID variables.	72
4.2	Stopping power for muons for the momentum range 0.1 MeV/c up to 100 TeV/c from [19]. The region from 0.1 GeV/c to 1 GeV/c is of greatest interest to T2K and is around the MIP region for muons.	74
4.3	A typical MIP leaves a long and narrow track as shown by this simulated event. The charge deposited in each layer is similar as can be seen from the colour scale.	74
4.4	Showing a medium energy electron shower with track information (left) and without (right).	78
4.5	Showing two types of hadronic shower. The shower on the left is converting soon after entering the detector like an EM shower. The shower on the right traverses much of the detector as a MIP before converting. These events are called ‘lollipop’ events as the track and shower pattern bears a resemblance to a lollipop.	78
4.6	Demonstration of linear (left) and non-linear (right) separability. A linear boundary can only separate two classes of event if a straight line can be drawn between the classes.	83

4.7	A neural network connects an input data vector to an output via hidden nodes connected by weights. Indices are the same as those in Equation (4.6).	83
4.8	Showing a tanh activation function. In general, neural networks use sigmoidal activation functions such as $\tanh(x)$	84
4.9	The AMR is calculated by finding the average ratio of the principal component to each minor component.	90
4.10	AMR variable distributions for tracks and showers (left) and EM and hadronic showers (right).	91
4.11	AMR variable as a function of angle for shower hypothesis (left) and track hypothesis (right).	92
4.12	Maximum Charge Ratio variable distributions for tracks and showers (left) and EM and hadronic showers (right).	94
4.13	Maximum Charge Ratio variable as a function of angle for shower hypothesis (left) and track hypothesis (right).	94
4.14	EM fit likelihood variable distributions for tracks and showers (left) and EM and hadronic showers (right).	95
4.15	EM fit likelihood variable as a function of angle for shower hypothesis (left) and track hypothesis (right).	96
4.16	Charge Skew variable distributions for tracks and showers (left) and EM and hadronic showers (right).	97
4.17	Charge Skew variable as a function of angle for shower hypothesis (left) and track hypothesis (right).	97

4.18	Shower angle variable distributions for tracks and showers (left) and EM and hadronic showers (right).	98
4.19	Shower Angle variable as a function of angle for shower hypoth- esis (left) and track hypothesis (right).	100
4.20	Shower Width variable distributions for tracks and showers (left) and EM and hadronic showers (right).	100
4.21	Shower Width variable as a function of angle for shower hypoth- esis (left) and track hypothesis (right).	102
4.22	Shower asymmetry variable distributions for tracks and showers (left) and EM and hadronic showers (right).	102
4.23	Shower asymmetry variable as a function of angle for shower hypothesis (left) and track hypothesis (right).	103
4.24	Mean shower position variable distributions for tracks and show- ers (left) and EM and hadronic showers (right).	103
4.25	Mean shower position variable as a function of angle for shower hypothesis (left) and track hypothesis (right).	104
4.26	Training and testing error estimators showing the same be- haviour over training epochs indicates no over training occurs. The x-axis shows signal inefficiency (1-efficiency).	107
4.27	Contamination is shown as a function of efficiency for a vari- ety of different size training samples. The performance of the network can be seen to improve with increasing training sample size up to a few thousand events. No evidence of overtraining is seen.	110

4.28	Showing a comparison between neural network output and likelihood output. The figure on the left compares a single input neural network with a 1D likelihood. The figure on the right shows a comparison between a multidimensional PDF and a neural network of the same dimensionality.	110
4.29	The MLP output for muon sample is shown, left. On the right is the energy as a function of track vs. shower output. The effect of stopping muons is shown to create a small cluster of incorrectly identified events.	111
4.30	Comparing the MLP outputs for an on-axis pion and electron sample using energy distributions described in Section 4.2. Hadronic showers and electromagnetic showers are more similar than tracks and showers, making discrimination more difficult.	112
4.31	Showing the muon efficiency as a function of electron efficiency for different energy range, $E > 400\text{MeV}$ (left) and $E > 1000\text{MeV}$ (right).	114
4.32	The MLP output for the test electron sample is shown (left) and the test muon sample output is shown on the right.	114
4.33	Performance of competing discrimination techniques from Monte Carlo data.	115
5.1	The number of hits for each channel is shown at a given time. The time shown is relative to the trigger and measured in TDC counts (2.5 ns). The two distinct blocks represent the two RMMs in the DsECal.	122

5.2	Diagram of the NIM modules and logic used to trigger and tag events in the testbeam.	123
5.3	Measured TOF signal in high gain (right) and low gain (left) ADC channels. In the low gain channel a sharp peak can be seen due to electrons and pions with a broader peak due to protons. In the high gain channel electron/pion and proton events are also visible.	125
5.4	Showing the relationship between the high and low gain ADC channels. Two clusters due to electron/pion events and protons can be seen. A third cluster is more clearly separated on the two dimensional plot showing deuterons.	126
5.5	Showing the difference in time generated by the cable delay generator, as a function of difference in ADC count.	128
5.6	Showing the spread of the mean ADC values for the electron/pion TOF peak, for on axis data.	129
5.7	Reconstructed mass of the second largest peak seen in the TOF spectrum. The value is consistent with that of the proton. . .	130
5.8	Variation of the peak in the TOF ADC spectrum due to electrons and pions as a function of time. 24 hour intervals are shown with vertical lines.	131
5.9	Showing correlation plots of high and low gain ADC channels in integration periods 18 and 19.	132
5.10	Low gain ADC values for integration periods 18 and 19 where the position in the integration period is measured in tdc counts (1 tdc count = 2.5ns).	133

5.11	Cerenkov efficiencies found for positive and negative momentum data sets. The upstream Čerenkov counter was found to be systematically more efficient than the downstream Čerenkov counter.	134
5.12	Measured fractions of different particle species in the CERN T9 testbeam. Error bars include systematic and statistical error estimates.	137
5.13	Showing the measured temperature of TFB 1 channels, measured from the outside of the bulkhead, as a function of Unix time.	140
5.14	Showing the predicted distribution of number of photons deposited in a given bar for MIP like particles. The distribution has been fitted to a convoluted Landau-Gaussian function. . .	142
5.15	Showing the variation in $\chi^2/N.D.F.$ between simulated and measured number of hits distributions as a function of the threshold used in the simulation.	144
5.16	Showing the maximum value of the convoluted Landau-Gaussian function against the TFB temperature.	145
5.17	Showing the bar-to-bar variation in the maximum value of the number of photons generated at the face of the photosensor by a MIP.	147
5.18	Showing the bar-to-bar variation in the maximum value of the number of photons generated at the face of the photosensor by a MIP.	147

5.19	Showing the variation in the convoluted Landau-Gaussian fit maximum for a low light yield bar and two higher yielding bars. The bottom right plot shows the residuals from fitting a plot of maximum value against run number to a constant for all bars. For some runs a result is missing as not enough events were collected to fit the data.	149
5.20	Showing a comparison of the predicted and measured total hit charge in a cluster using a bar-to-bar calibration constant. . .	150
5.21	Showing a comparison of the predicted and measured hit charge spectrum and total hit charge, using a single calibration constant.	150
5.22	Showing the variation in calculated $\chi^2/N.D.F.$ as a function of smearing Gaussian parameters.	151
5.23	Comparing measured data with simulated predictions for the hit charge spectrum and the total hit charge using a single calibration constant smeared with a Gaussian to produce bar-to-bar variations.	152
5.24	Monte Carlo to data comparison for basic reconstructed quantities of cosmic muons. Shown is the reconstructed angle (a), the number of hits in a cluster (b) and the half length of the track (c).	153
5.25	Cosmic muon simulation to data comparison for PID variables.	154
5.26	Showing the correlation between AMR and shower width and shower angle	155

5.27	Monte Carlo to data comparison of basic quantities for 600 MeV electrons. While the hit charge spectrum and maximum hit charge agreed well, a discrepancy was seen in the total hit charge and number of hits distributions.	157
5.28	Comparison of testbeam data and simulated data after applying a 9% EM scaling factor.	160
5.29	Monte Carlo to data comparison for PID related quantities for 600 MeV electrons.	161
5.30	Comparing the agreement between data and simulation at 400 MeV and 1.4 GeV. Saturation effects can be seen in the hit charge spectrum in the data which are not modelled by the simulation. Distributions broadly agree although a discrepancy is seen for the maximum ratio.	162
5.31	Monte Carlo to data comparison for 800 MeV electrons orientated at 30° to the face of the DsECal.	164
5.32	Monte Carlo to data comparison for 800 MeV electrons orientated at 60° to the face of the DsECal.	165
5.33	Showing a preliminary measurement of the DsECal energy resolution.	167
6.1	Showing the distribution of neutrino interactions in the ND280 detector.	172
6.2	An event display of a CCQE electron event with recoil proton (left) and a CCQE muon event (right). Electrons are coloured red, protons are coloured blue and muons are coloured green. .	174

6.3	Showing an event where a pair of leptons were seen in the final state due a π^0 decay.	175
6.4	Genie predicted event rates in the FGD for electron neutrinos (left) and muon neutrinos (right).	176
6.5	The ratio of electron to muon neutrino events as a function of particle momentum.	177
6.6	Showing the stopping power for electrons, muons and charged pions over the momentum range in T2K. A powerful electron vs. muon separation can be made above 200MeV. Separating muons and pions is difficult in general. Figure modified from [99].	179
6.7	Histogram showing the calculated distance between the selected TPC track and ECal clusters. A large peak at approximately 20 cm indicates well reconstructed clusters with a gap due to structural components and electronics between the tracker and ECal.	182
6.8	Numbers of electron and muon neutrino events left in samples after lepton selection and particle identification.	183
6.9	Using cuts targeted at remaining CCnQE muon and NC events and π^0 events bring the signal to noise ratio to 1.84:1.	186
6.10	Showing a high multiplicity event that was mis-reconstructed by the reconstruction software.	187
6.11	Showing the efficiency and purity of the inclusive and exclusive analyses after each cut is applied.	190
6.12	Showing the efficiency and purity of the inclusive and exclusive analyses as a function of the neutrino momentum.	191

7.1	Showing one of the first candidate neutrino events in ND280. . .	198
-----	--	-----

List of Tables

2.1	Limits set on the absolute mass of the neutrino via different experimental techniques, data used from [5].	4
2.2	Global best fit values for three neutrino model mixing parameters as determined by [61].	30
3.1	Numbers of bars in each ECal module [18].	51
4.1	List of variables used in the PID algorithm and which discriminators they are used in.	105
5.1	Measurements of generated time delay from TOF calibration run.	127
5.2	Cut parameters used for testbeam PID.	135
5.3	Approximate number of triggers achieved in the T9 testbeam in six weeks with a total of approximately 2×10^6	138
5.4	Showing simulated track vs. shower discrimination efficiency for simulation, $(\epsilon_{\text{(MC)}})$, and data, $\epsilon_{\text{(Data)}}$, as well as the compatibility between them in terms of χ^2/ndf for a range of angles and energies.	166

6.1	The different categories of neutrino event expected in ND280 are shown along with their relative abundances [18]. An absolute number of events per 5×10^{21} POT per tonne is also shown. The largest fraction of events are CCQE events.	173
6.2	Predicted number of events after application of successive cuts.	185
6.3	Efficiency and purity of electron neutrino samples following application of successive cuts.	188
6.4	Number of signal and background events with statistical errors for inclusive and exclusive samples.	188
6.5	Systematic errors associated with the inclusive CC ν_e analysis. Systematic uncertainties are dominated by lack of knowledge of NC cross-sections.	193
6.6	Systematic errors associated with the exclusive CC ν_e analysis.	193
6.7	Predicted electron neutrino analysis event rates for one and five years of nominal running. The first error quoted is systematic and the second is statistical.	194

Chapter 1

Introduction

1.1 Introduction

The field of neutrino physics dates back to 1930 [1, 2], when quantum mechanics was still an emerging field attempting to explain the beta decay process. For many years progress in neutrino physics was held back by low event rates and much of the progress discussed in Chapter 2 has only crystallised in the last decade. Following the discovery in 1995 of the top quark [3], it has become one of the most prolific areas of experimental particle physics. In the past decade it has been established that neutrinos have mass and that they can oscillate between the three flavours that are currently known to exist [4]. A number of experiments are currently running, or in planning, in order to make precision studies of neutrino properties [4]. One of these is T2K.

The T2K (Tokai to Kamioka) experiment, discussed in Chapter 3, is one of a new generation of neutrino experiments looking to study the phenomenon of neutrino oscillation. T2K will compare the flux of different neutrino flavours and their energy spectra to measure neutrino oscillation parameters using two

detectors separated by 295 km.

This thesis presents a neural network based particle identification algorithm for the electromagnetic calorimeter in the near detector. In Chapter 5 data collected in the T9 beamline at CERN is compared to the Monte Carlo simulation used to train the particle identification algorithm. Finally, the particle identification algorithm is used in an electron neutrino analysis that could be used from the first day of data taking.

Chapter 2

Neutrino Physics

2.1 Neutrino Phenomenology

2.1.1 Neutrino Mass

The mass of the neutrino has been a subject of much debate since its existence was first postulated, and underpins many of the questions that make neutrinos interesting today. Pauli assumed neutrinos had a mass of the same order as the electron. In the standard model the neutrino is a massless particle [4]. However the existence of flavour change amongst neutrinos, be it through oscillation or another mechanism, precludes the neutrino from having no mass [4]. Neutrino oscillation experiments can only measure the mass difference between different neutrinos [4]. There exist a number of methods for determining the absolute mass of the neutrino, including limits set by cosmological models and terrestrial experiments such as tritium decay. These limits are summarised in [5] and below in Table 2.1.

Why the neutrino mass is so small ($< 1\text{eV}$) compared to the other standard

Measurement	Set Limit
Tritium Decay	$< 2.0eV$
π^+ decay	$< 0.19MeV$
τ decay	$18.2MeV$
SN1987A Time of flight	$< 5.7eV$
Cosmology	$< 0.7eV$

Table 2.1: Limits set on the absolute mass of the neutrino via different experimental techniques, data used from [5].

model particles is a mystery and one of the most important questions in particle physics. It is a possibility that the very small neutrino mass is a sign of ‘beyond the standard model’ physics [4].

If one assumes that the neutrino obeys the Dirac equation, then it’s possible that neutrino mass has the same origin as the other fermions, via the Higgs mechanism [4]. If the neutrino is a Dirac particle then one must introduce right handed neutrino fields. While this is the simplest possible addition to the standard model, known as the minimally extended standard model, the right handed neutrino would have the interesting property of being an electroweak singlet with hypercharge zero [4]. As any right handed neutrino would not feel the weak interaction, only gravitation, they are usually called sterile neutrinos [4].

For a massless fermion, only a two component spinor is required to describe the field [4]. Majorana found that this is also true for a massive field if the left and right handed fields are not independent. One consequence of a field having a two component (Weyl) spinor is that particle and antiparticle are the same particle [4], so a Majorana neutrino would be its own antiparticle. This allows a number of processes to occur that would not normally be allowed. One such process is neutrino-less double beta decay [4]. The discovery of neutrino-less double beta decay through the process ${}^A_ZX \rightarrow {}^A_{Z+2}X' + 2e^-$ would show not only

that the neutrino is a Majorana particle, but also allow a direct measurement of the mass of the electron neutrino [6].

2.1.2 Oscillation Probability

The phenomenon of neutrino oscillations is driven by the small mass difference between neutrinos [4]. A neutrino is created, in the standard model, through charged current interactions such as the beta decay process $n \rightarrow p + e^- + \bar{\nu}_e$, or by the decay of a Z^0 boson. At the moment of the interaction, the neutrino is therefore in a flavour eigenstate $|\nu_e\rangle$, $|\nu_\mu\rangle$ or $|\nu_\tau\rangle$. A neutrino will propagate in a mass eigenstate which is not necessarily equivalent to its flavour state [4]. A measurement of the neutrino flavour through another interaction, after propagating for some distance, need not yield the same flavour as the original neutrino. Neutrino oscillations are therefore a purely quantum mechanical effect. Section 2.3 discusses the historical context of the theory of neutrino oscillations. This section summarises the current theoretical basis of oscillations.

Neutrino oscillations are most often considered by treating the neutrino as a plane wave. A plane wave has an infinite extent in space-time and as such it cannot explain a local particle. To correctly express a superposition of plane waves one must use a wave packet analysis. A full wave packet analysis is given in [4]. A plane wave approximation, which gives the same result, is used below [4]. A neutrino in a flavour state is expressed as superposition of mass states

$$|\nu_\alpha\rangle = \sum_{\kappa} U_{\alpha\kappa}^* |\nu_\kappa\rangle \quad (\alpha = e, \mu, \tau), \quad (2.1)$$

where $|\nu_\alpha\rangle$ is the flavour basis and $|\nu_\kappa\rangle$ is the mass basis. The mixing matrix U , referred to as the PMNS matrix (Pontecorvo, Maki, Nakagawa and

Sakata) [4], is unitary because one should be able to represent any neutrino in either basis. In Equation (2.1) the number of mass states was not fixed, there could be more than three mass states corresponding to sterile neutrinos that do not contribute to weak interactions.

An important assumption of the plane wave formalism is the constant momentum assumption, which assumes that all the mass states have the same momentum. In the plane wave approximation the wave has a known energy and momentum. This cannot be the case for neutrinos, otherwise the mass of the neutrino would be known. In the more complete wave packet formalism, an integral over the momentum states in the packet circumvents the constant momentum assumption, restoring the observers ignorance of the neutrino mass. However, and perhaps rather curiously, the plane wave derivation yields the same result for the oscillation probability. The evolution of the neutrino in time and in space can be written as

$$| \nu(L, T) \rangle = e^{-i(E_\kappa T - \mathbf{p} \cdot \mathbf{L})} | \nu_\alpha \rangle. \quad (2.2)$$

In this case, when the phase is written

$$E_\kappa T - \mathbf{p} \cdot \mathbf{L} \approx |\mathbf{p}| \left(1 + \frac{m_\kappa^2}{2|\mathbf{p}|^2} \right) L - |\mathbf{p}| L = -\frac{m_\kappa^2}{2E} L, \quad (2.3)$$

where $E_\kappa = \sqrt{m_\kappa^2 + |\mathbf{p}|^2}$ and $T = L$, in natural units, as the neutrino is travelling close to the speed of light. Taking the difference between two different energy states, inserting the result as the phase in (2.1) yields the probability of oscillation

$$P_{\nu_\alpha \rightarrow \nu_\beta}(t) = |A_{\nu_\alpha \rightarrow \nu_\beta}(t)|^2 = \sum_{\kappa j} U_{\alpha\kappa}^* U_{\beta\kappa} U_{\alpha j} U_{\beta j}^* e^{-i \frac{\Delta m_{\kappa j}^2 L}{2E}}. \quad (2.4)$$

In (2.4) it should be noted that $t = L$ as neutrinos travel at the speed of light, to a good approximation. This is called the transition probability as it gives the probability that a neutrino will change from one flavour to another.

The mixing matrix U obeys the unitarity relation

$$UU^\dagger = 1 \quad (2.5)$$

$$\sum_{\kappa} U_{a\kappa} U_{\beta\kappa}^* = \delta_{\alpha\beta}. \quad (2.6)$$

From Equation (2.6) we can write

$$\sum_{\kappa} |U_{\alpha\kappa}|^2 |U_{\beta\kappa}|^2 = \delta_{\alpha\beta} - 2 \sum_{\kappa > j} \Re[U_{\alpha\kappa}^* U_{\beta\kappa} U_{\alpha j} U_{\beta j}^*]. \quad (2.7)$$

The constant term $P_{\nu_\alpha \rightarrow \nu_\beta}(L = 0, E) = \delta_{\alpha\beta}$ is implied from the unitarity relation. The expression in (2.7) can be used to split Equation (2.4) into real and imaginary components [4], yielding

$$\begin{aligned} P_{\nu_\alpha \rightarrow \nu_\beta}(L, E) &= \delta_{\alpha\beta} - 4 \sum_{\kappa > j} \Re[U_{\alpha\kappa}^* U_{\beta\kappa} U_{\alpha j} U_{\beta j}^*] \sin^2\left(\frac{\Delta m_{\kappa j}^2 L}{4E}\right) \\ &+ 2 \sum_{\kappa > j} \Im[U_{\alpha\kappa}^* U_{\beta\kappa} U_{\alpha j} U_{\beta j}^*] \sin\left(\frac{\Delta m_{\kappa j}^2 L}{2E}\right). \end{aligned} \quad (2.8)$$

The first term in Equation (2.8) is, from the unitarity of the mixing matrix, zero if $\alpha \neq \beta$ and one if $\alpha = \beta$; and gives the observation probability if there are no oscillations. The second real term is the oscillation term which gives the probability that a neutrino of energy E will oscillate from flavour state α to flavour state β over a baseline L . This is true only if the flavour state is not also a mass eigenstate. The final term is zero if the mixing matrix is real, which is necessarily true in the case that there are only two neutrinos. The final term also vanishes if $\alpha = \beta$. The oscillation probability for the case $\alpha = \beta$, also called the survival probability, can be written as [4]

$$P_{\nu_\alpha \rightarrow \nu_\beta}(L, E) = 1 - 4 \sum_{\kappa > j} |U_{\alpha\kappa}^*|^2 |U_{\alpha j}^*|^2 \sin^2\left(-i \frac{\Delta m_{\kappa j}^2 L}{2E}\right). \quad (2.9)$$

2.1.3 Neutrino Oscillations in Matter

All three flavours of neutrino can interact with electrons or nuclei, via the neutral current. Interactions with matter can alter the oscillation probability. This phenomenon, analogous to a refractive index, was first studied by Wolfenstein [7]. Mikheev and Smirnov found that when a neutrino travels through a dense medium there can be a region where the effective mixing angle becomes maximal [8]. This effect is known as the MSW (Mikheev - Smirnov - Wolfenstein) effect and was used to explain the flavour conversion of solar neutrinos [4].




Figure 2.1: Showing charged and neutral current coherent forward neutrino scattering.

From Figure 2.1 we observe that there is a charged current component of forward scattering, accessible only to electron neutrinos interacting with electrons in matter. The scattering is forward as the interaction is coherent, negligible momentum is transferred. This results in a potential proportional to the electron density of the matter, $V_{CC} = +\sqrt{2}G_F N_e$, where G_F is the Fermi coupling constant and N_e is the electron density of the matter. The second potential is felt by all neutrino flavours via the neutral current and allows scattering from electrons and nucleons. If it is assumed that the matter is neutral so that there are equal quantities of protons and electrons, the proton and electron

contributions to scattering will cancel out leaving only a term due to neutrons,
 $V_{NC} = -\frac{\sqrt{2}}{2}G_F N_n$.

2.1.4 Current 3 flavour neutrino oscillation model

With the notable exception of the LSND experiment [9], all data collected so far can be described in a framework with only two mass differences. The atmospheric mass difference was confirmed by K2K [10] and the solar mass difference was confirmed by KamLAND [11]. One might infer that there are only three neutrino flavours from the invisible width of the Z-boson [12], although there could be any number of sterile neutrinos. Currently the simplest case is that there are three left handed neutrino fields, each with a mass, leading to the three mass squared differences

$$\Delta m_{21}^2 \equiv m_2^2 - m_1^2 \quad (2.10)$$

$$\Delta m_{31}^2 \equiv m_3^2 - m_1^2 \quad (2.11)$$

$$\Delta m_{32}^2 \equiv m_3^2 - m_2^2. \quad (2.12)$$

From the experimental measurements of the solar and atmospheric mass differences it is known that $\Delta m_{sol}^2 \ll \Delta m_{atm}^2$, where $\Delta m_{sol}^2 = \Delta m_{21}^2$ and $\Delta m_{atm}^2 = |\Delta m_{31}^2|$. The sign of the solar mass difference, Δm_{sol}^2 , is known from the MSW effect. The sign of the atmospheric mass difference is unknown. Experiments measuring the atmospheric mass difference such as K2K and MINOS can only observe an average of Δm_{13}^2 and Δm_{23}^2 [4]. In Figure 2.2 the two possible ways of arranging the known mass differences are shown along with the composition of each mass state in terms of the flavour states.

As shown in Figure 2.2 and Equation (2.1), each mass state is described as a superposition of flavour states. The mixing matrix for the three neutrino




Figure 2.2: The normal hierarchy has two light masses close together and a third much larger mass, following the pattern in the standard model where the particles in the third generation are typically much heavier than in the first two. In the inverted hierarchy the two mass states of similar magnitude are much more massive than the third. Figure modified from [13].

case is a 3x3 unitary matrix as described previously. The exact number of parameters depends on the nature of neutrino mass. If neutrinos are Dirac particles then the mixing matrix has three angles and a single complex phase. If the neutrino is a Majorana particle then an extra pair of complex phases must be added. The full mixing matrix has the form

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{13}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{-i\delta_{13}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{-i\delta_{13}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{-i\delta_{13}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{-i\delta_{13}} & c_{23}c_{13} \end{pmatrix}, \quad (2.13)$$

assuming that the neutrino is a Dirac particle. A common convention is to express the Dirac part of the mixing matrix as [4]

$$U = \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix}}_{\text{atmospheric}} \underbrace{\begin{pmatrix} c_{13} & 0 & s_{13} \\ 0 & 1 & 0 \\ -s_{13} & 0 & c_{13} \end{pmatrix}}_{\text{reactor/longbaseline}} \underbrace{\begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{\text{solar}} \quad (2.14)$$

to emphasise the de-coupling of atmospheric and solar neutrino sectors, due to the small value of θ_{13} . Inserting elements from Equation (2.13) into Equation (2.8) returns the precise probabilities for different oscillations. In the two neutrino case where the mixing matrix is simply

$$U = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}, \quad (2.15)$$

inserting into Equation (2.8) yields

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 L}{4E}\right). \quad (2.16)$$

This procedure is far simpler than solving Equation (2.8) and, in situations where only one mixing angle is significant, provides an excellent approximation to the oscillation probability. A primary physics goal for many next generation experiments is the determination of the mixing angle θ_{13} , by observing electron neutrino appearance. This must be analysed using the three neutrino formalism as the other mixing angles are large compared to θ_{13} . A full expression of $P(\nu_\mu \rightarrow \nu_e)$ is shown in [4] and shows the oscillation probability to be

$$\begin{aligned} P(\nu_\mu \rightarrow \nu_e) &= \sin^2 \theta_{23} \sin^2 2\theta_{13} \frac{\sin^2[(1-A)\Delta]}{(1-A)^2} \\ &\pm J\alpha \sin \delta_{CP} \sin \Delta \frac{\sin(A\Delta)}{A} \frac{\sin[(1-A)\Delta]}{(1-A)} \\ &+ J\alpha \cos \delta_{CP} \cos \Delta \frac{\sin(A\Delta)}{A} \frac{\sin[(1-A)\Delta]}{(1-A)} \\ &+ \alpha^2 \cos^2 \theta_{23} \sin^2 2\theta_{12} \frac{\sin^2(A\Delta)}{A^2}. \end{aligned} \quad (2.17)$$

J is the Jarlskog invariant, where $J = \cos \theta_{13} \sin 2\theta_{12} \sin 2\theta_{13} \sin 2\theta_{23}$, α is the ratio of the mass differences $\alpha = \frac{\Delta m_{12}^2}{\Delta m_{23}^2}$, $\Delta = \frac{\Delta m_{23}^2 L}{4E}$, $A = \frac{B}{\Delta m_{23}^2}$ and $B = \pm 2E\sqrt{2}G_F N_e$. For the current generation of long baseline experiments, the path length of the neutrinos through the Earth is not sufficient for matter

effects to be observed, so the term A will vanish as the matter term B vanishes. A common simplification for terrestrial neutrino experiments is that the term α also vanishes, as the atmospheric mass difference is much larger than the solar mass difference. In situations where the mass term, A , vanishes, the oscillation probability becomes

$$P(\nu_\mu \rightarrow \nu_e) = \sin^2 \theta_{23} \sin^2 2\theta_{13} \sin^2 \frac{(\Delta m_{23}^2)L}{4E}. \quad (2.18)$$

A long baseline experiment will therefore attempt to select a $\frac{L}{E}$ ratio that gives a value of $\sin^2 \frac{(\Delta m_{23}^2)L}{4E}$ as close to one as possible, in natural units, to resolve the effects of neutrino oscillations. This can be expressed in SI units by [4]

$$P(\nu_\mu \rightarrow \nu_e) = \sin^2 \theta_{23} \sin^2 2\theta_{13} \sin^2 \frac{1.267(\Delta m_{23}^2)[\text{eV}^2]L[\text{km}]}{4E[\text{GeV}]}. \quad (2.19)$$

Current experimental limits for the oscillation parameters are discussed in Section 2.3.2.

2.2 Neutrino Interaction Physics

As neutrinos only interact via the weak interaction the cross sections for all processes are very low. Assuming a cross section of $1 \times 10^{-38} \text{ cm}^2$ then the mean free path through a typical metal will be of the order of 10 light years. As the neutrino has no electrical charge it cannot be seen directly. Its kinematic properties are therefore inferred from the properties of the final state products of neutrino interactions. Weak interactions are normally divided into two types, charged current, where the interaction is mediated by $W^\pm$ bosons, and neutral current interactions mediated by the Z boson.

2.2.1 Charged Current Interactions

Charged current interactions are generally more useful for studying neutrinos at medium and high energies as a charged lepton allows the flavour of the neutrino to be determined. The nature of a charged current interaction is determined by the momentum transferred to the nucleus. Three different regimes are observed, as shown in Figure 2.3. At low momenta, below approximately 1GeV, quasi-elastic interactions dominate. A quasi-elastic interaction with a neutrino looks like

$$\nu_l + n \rightarrow l^- + p \quad (2.20)$$

for a neutrino and similarly for an anti-neutrino

$$\bar{\nu}_l + p \rightarrow l^+ + n. \quad (2.21)$$

The energy of the incoming neutrino can be approximated by a measurement of the energy and direction of the lepton and the energy of the nucleon using

$$E_\nu^{QE} \simeq \frac{M_N E_l - \frac{1}{2} m_l^2}{M_N - E_l + \sqrt{E_l^2 - m_l^2} \cos \theta_l}. \quad (2.22)$$

The above equation assumes that only the nucleon and lepton are released in the interaction, if any other particles are released then the neutrino energy will not be correctly reconstructed. The Fermi momentum of the nucleon that the neutrino scatters off is also not considered in the above equation. In the region of a few GeV a significant number of events are so called single pion events, as shown in Figure 2.3. Cross sections for specific interaction channels are poorly measured with most data coming from bubble chamber experiments, in particular Gargamelle [15], the 12 ft bubble chamber at ANL (Argonne National Laboratory)[16] and the BNL (Brookhaven National Laboratory) bubble chamber [17]. Single pion final states can be created by mechanisms such

as resonant pion production or by coherent pion production. Coherent pion production occurs when a neutrino scatters off of the nucleus as a whole. In coherent scattering the nucleus is left intact despite the neutrino scattering although a pion can be emitted. Coherent production can produce heavier mesons than pions although at the energies present in T2K such events will be rare [18]. A small fraction of ρ mesons are expected, accounting for approximately 0.5% of events some of which will be produced via resonant interactions [18]. Resonant pion production is where a proton interacts with a positively charged W-boson to create a resonant state such as a $\Delta^{++}(1232)$ which in turn decays back to a proton and charged pion. Figure 2.4 shows how resonant scatters can occur by the nucleon becoming excited by the weak boson and then decaying back into a pion and the original nucleon. A neutrino exchanging a positive W-boson with the nucleus can interact resonantly with a proton

$$p + W^+ \rightarrow \Delta^{++} \rightarrow \pi^+ + p, \quad (2.23)$$

or with a neutron

$$n + W^+ \rightarrow \Delta^+ \rightarrow \pi^+ + n. \quad (2.24)$$

As for coherent production, if the mediating boson transfers sufficient momentum then a more energetic resonance can be created, decaying into mesons heavier than pions. Similar interactions can be determined for anti-neutrinos by swapping for a W^- particle and adjusting the charge of the resultant Δ accordingly. Final states from any charged current interactions will also include a lepton. If a higher mass resonance is produced then the resultant decay can produce a multi-pion event where the final state consists of more than one pion.

At higher energies, Figure 2.3 shows that Deep Inelastic Scattering (DIS) interactions dominate. In DIS interactions the momentum transferred to the




Figure 2.3: Figure from [14] showing the charged current neutrino cross section. Contributions from quasi-elastic scattering, single pion production and deep inelastic scattering are shown separately.




Figure 2.4: A resonant scattering interaction. A Δ is created which decays back to a pion and a nucleon.

nucleus is sufficient to break the nucleus apart resulting in a hadronic shower. Such interactions result in a wide variety of final states where the exact number depends on the momentum transferred to the nucleus. At the energies used in the T2K experiment DIS interactions will only occur with neutrinos found in the high energy tail of the energy distribution, above a few GeV.

2.2.2 Neutral Current Interactions

A neutral current interaction is harder to reconstruct because a charged lepton will not be found in the final state. A neutral current interaction will therefore only be detected when the nucleus breaks apart as a result of the interaction. As no charged lepton is produced via the neutral current then the flavour of the neutrino present in the interaction will remain unknown. As such it is not useful for an oscillation study although can still be used in cross section analyses. At low momenta, where a nucleon is less likely to be emitted from the nucleus, neutral current events are usually undetected.

At higher momenta, resonant single pion production can occur through the same mechanism as for the charged current, except that the nucleon interacts with a Z-boson to create a Δ resonance. Neutral current π^0 are a major source of background at Super-Kamiokande because they can look like electron events. The π^0 decays in 8.7×10^{-16} s [19], producing two back to back photons in the rest frame of the π^0 . If, when boosted to the lab frame, one photon has a low energy and is missed, or both photons are colinear, then only a single ring will be observed. If the ring is sufficiently like an electron then the event can be mis-reconstructed as an electron neutrino event. Currently NC cross sections are poorly measured and are used to estimate systematic errors in Chapter 6.

2.3 A review of Neutrino Oscillations

The story of the neutrino began in 1930 when its existence was hypothesised by Wolfgang Pauli [1]. His proposal was conceived out of the apparent violation of angular momentum and energy conservation in beta decay. To conserve angular momentum and energy Pauli suggested a new, undetectable, spin 1/2 particle which would carry away the missing energy and momentum. Predicting a particle that could not be detected concerned Pauli and he never published his idea. In Enrico Fermi's theory of β -decay the term 'neutrino' was first used. Although the neutrino was not discovered until 1956 the success of Fermi's explanation of β -decay gave the hypothesis widespread acceptance.

The first experimental observation of the neutrino was made by Reines and Cowan [20]. They confirmed the existence of the neutrino by looking for the positron and neutron produced in the inverse β -decay process;

$$\bar{\nu}_e + p \rightarrow e^+ + n. \quad (2.25)$$

Their experiment studied the anti-electron neutrinos from the Hanford nuclear pile. Reines *et.al.* continued to study anti-electron neutrinos from the nuclear reactor at Savannah river. The cross section of the interaction between the anti-electron neutrino and the proton was determined to be $(11 \pm 2.6) \times 10^{-44}$ cm² [21]. The cadmium chloride based detector used at Savannah River is described in [22].

Flavour oscillation was first postulated by Pais, Gell-Mann and Piccioni who postulated oscillations between K^0 and $\bar{K}^0$ mesons [23, 24]. This can be said to be the genesis of the idea of flavour oscillations in particle physics. Bruno Pontecorvo then suggested that neutrino to anti-neutrino oscillations should be possible [25] after considering other systems such as $e^\pm \mu^\pm$ bound systems

[26]. In 1967 Pontecorvo went further and suggested that oscillations between electron and muon neutrinos could also be possible after confirmation that the electron and muon neutrino were distinct particles [27]. The following year Pontecorvo also suggested that this oscillation could be used to explain a discrepancy between the observed and expected number of neutrinos from the sun [28]. In 1962, before the confirmation that muon neutrinos and electron neutrinos were distinct particles [29], Maki, Nakagawa and Sakata formulated a theory that incorporated the idea of mixed neutrino states [30]. This was a year earlier than Cabbibo postulated the idea that non-conservation of strangeness was due to mixing between strangeness conserving and strangeness violating currents.

It was generally assumed that there was a third light neutrino to partner the τ lepton following its discovery at SLAC [31]. This motivated the development of a three neutrino oscillation model. Hints that the tau neutrino existed were provided by experiments at LEP measuring the invisible width of the Z-boson. The number of light neutrino flavours was found to be 2.9841 ± 0.0083 at LEP [12]. A direct observation of the tau neutrino remained elusive until one was finally seen by the DONUT collaboration in 2000 [32].

2.3.1 The Solar Neutrino Problem

The flux of neutrinos produced by the sun was predicted by John Bahcall [33]. Each neutrino producing fusion process within the sun emits neutrinos with different energy spectra. For example, PP (Proton-Proton) fusion reactions, the dominant hydrogen burning reactions within the sun, produce huge amounts of neutrinos with low energies as shown in Figure 2.5. Experiments attempting to confirm the solar model uncovered a mystery that was not solved

until 2002. The first experiments to try and experimentally measure the solar neutrino flux were radio-chemical experiments. These used the inverse β -decay process to detect low energy neutrinos from thermonuclear reactions within the sun. The first experiment of this type, the Homestake experiment, used 100,000 gallons of tetrachloroethylene. Tetrachloroethylene reactions have a threshold of 0.813 MeV. As can be seen from Figure 2.5, contributions from 8B processes dominate in this region. Argon atoms from the reaction,

$$\nu_e + {}^{37}\text{Cl} \rightarrow e^- + {}^{37}\text{Ar}, \quad (2.26)$$

were measured to determine how many interactions had occurred in the tank. The experiment reported results that were inconsistent with the SSM (Standard Solar Model) predictions from Bahcall [34]. The Homestake experiment continued to run until 1995, when a final result for the ratio of expected to observed solar neutrinos was reported of $2.56 \pm 0.16(\text{stat}) \pm 0.16(\text{sys})$ [13].

Other radio-chemical experiments searched for the process $\nu_e + {}^{71}\text{Ga} \rightarrow e^- + {}^{71}\text{Ge}$, in either metallic gallium or gallium dissolved in chlorine. Two experiments, carried out in the 1990s, using the latter approach were GALLEX and its successor GNO. Based in the Gran Sasso laboratory, Italy, GALLEX [35] was made from 30 tons of gallium allowed to dissolve in 100 tons of chlorine. GNO [36], containing 101 tonnes of GaCl_3 in water and HCL, measured a solar flux of $77.5 \pm 6.2^{+4.3}_{-4.7}$ SNU and GALLEX found a flux of $62.9^{+5.5}_{-5.3} \pm 2.5$ SNU, approximately half the 128 SNU expected by the SSM ($1 \text{ SNU} = 10^{-36}$ neutrino captures/(atom second)). The SAGE experiment, located in the Baksan mine, used metallic gallium as its target material. SAGE also found about half the expected number of solar neutrinos [37]. The Gallium based experiments were able to achieve higher numbers of neutrino events because they had a lower threshold energy of 0.233 MeV compared to 0.814 MeV for chlorine. The lower




Figure 2.5: Showing the flux of neutrinos from different reactions within the sun from [13].

threshold of the Gallium experiments made them sensitive to the very large flux of neutrinos coming from the PP processes inside the sun, as shown by Figure 2.5.

The first independent verification of the Homestake experiment came, slightly earlier than the gallium experiments, in 1989, from the Kamiokande experiment. Originally designed to search for the decay of the proton, the Kamiokande observatory turned out to be very effective at looking for neutrinos [38]. Not only did Kamiokande independently confirm the Homestake radio-chemical experiment; it confirmed that solar neutrinos were originating from the sun [38]. In Figure 2.5 we see that the water Čerenkov detectors have a high threshold compared to the radio-chemical experiments meaning they are only sensitive to neutrinos from 8B , therefore such detectors need to be much larger and/or more efficient to collect the same number of events.

Super-Kamiokande is a water Čerenkov detector and also the successor to Kamiokande. Water Čerenkov detectors observe neutrinos by detecting the Čerenkov light from charged leptons produced in charged current neutrino interactions. Large water Čerenkov detectors uncovered a new discrepancy between data and Monte Carlo simulations. These experiments compared the double ratio

$$\frac{(\nu_\mu/\nu_e)_{data}}{(\nu_\mu/\nu_e)_{MC}}, \quad (2.27)$$

of neutrino fluxes from neutrinos generated in the atmosphere, which was expected to be one in the absence of oscillations. Atmospheric neutrinos are created from the interaction of cosmic rays and gas nuclei. When a cosmic ray hits a nuclei particles such as muons and light mesons decay producing neutrinos. Super-Kamiokande [39], found a value for a ratio of $R = 0.57^{+0.08}_{-0.06} \pm 0.035$ for high energy events (above 1 GeV). This result was confirmed by IMB (Irvine-Michigan-Brookhaven) [40], another water Čerenkov experiment. Two smaller experiments Frejus [41] and NUSEX [42] failed to verify the results although their sample size was too small to be convincing. This second discrepancy

became known as the atmospheric neutrino problem.

In 1998 the Super-Kamiokande collaboration claimed to have solved the atmospheric neutrino problem by comparing the numbers of neutrinos passing through the Earth to the number coming from above [39]. Super-Kamiokande was instrumented with just over 11,000 PMTs (Photo Multiplier Tubes), reduced to around 5000 following an accident in 2001 and since rebuilt and returned to 11,000 PMTs [43]. Super-Kamiokande also boasts a total volume of 50 kt and a fiducial volume of 22.5 kt, 30 times larger than the original Kamiokande detector. Super-Kamiokande provided strong evidence for the existence of neutrino oscillations in 1998, before the SNO collaboration solved the solar neutrino problem [39].

Neutrinos created in the atmosphere above Super-Kamiokande travel only 30 km before reaching the detector, while a neutrino passing through the Earth will travel around 13,000 km to reach the detector. The angle with respect to vertical at Super-Kamiokande gives the distance travelled for the neutrino: $\cos \phi = 1$ indicates the neutrino has come from above while $\cos \phi = -1$ indicates the neutrino has travelled through the Earth [39].

Super-Kamiokande found a lower number of muon neutrinos passing through the Earth than coming from above. No difference was observed for electron neutrinos. Two further experiments MACRO [44] a liquid scintillator experiment and Soudan-2 [45] an iron tracking calorimeter experiment confirmed the Super-Kamiokande result. The confirmation of an apparent deficit in two separate systems lent credence to the idea that some new flavour changing mechanism was causing the deficit.

Neutrino oscillations had two main drawbacks as an explanation for the neutrino deficit. Firstly, if the mixing matrix was to be considered as an analogy

to the CKM (Cabibbo, Kobayashi, Maskawa) matrix in the quark sector, one might expect the mixing matrices to have similar values. In practise the neutrino mixing matrix (PMNS matrix) would need to have significantly larger mixing (off diagonal) terms to explain the deficit. The second problem lay in the mass of the neutrino. The standard model described the neutrino as massless which would preclude neutrino oscillations from occurring. If neutrino oscillations were the solution then the standard model would have to be revised to give the neutrino mass.

The solar neutrino problem was solved in 2002 by the SNO (Sudbury Neutrino Observatory) experiment. Located in a mine 2 km underground, SNO was a water Čerenkov detector like Super-Kamiokande. SNO however, was filled with heavy water, D_2O [46]. SNO could observe neutrino interactions through three separate channels;

$$CC : \quad \nu_e + d \rightarrow p + p + e^-, \quad (2.28)$$

$$NC : \quad \nu_\alpha + d \rightarrow p + n + \nu_\alpha, \quad (2.29)$$

$$ES : \quad \nu_\alpha + e^- \rightarrow \nu_\alpha + e^-, \quad (2.30)$$

unlike Super-Kamiokande which was sensitive to ES (Elastic Scattering) only. Elastic scattering and NC (Neutral Current) channels were sensitive to all flavours of neutrino. The third channel that SNO was sensitive to, the CC (Charged Current) channel, was only accessible to electron neutrinos. Comparing the rates of the different interactions allowed SNO to make a direct observation of the fraction of neutrinos from the sun that were electron neutrinos.

Initially SNO ran with just heavy water. In this mode the neutron from the neutral current interactions was detected via its interaction with a deuterium atom [47]. In its second phase, salt was dissolved into the heavy water so that




Figure 2.6: Figure from [49] showing the tau and muon neutrino fluxes versus the electron neutrino flux. The total flux predicted by the SSM is shown as a dashed line. The SK result from elastic scattering is also shown.

neutrons were detected through their interaction with the chlorine. Sensitivity to the neutral current channel was increased by the introduction of the chlorine as it gave an increased neutron cross section. The interaction also released more photons than the interaction with deuterium [48]. During the ‘salt phase’ the SNO collaboration collected 2176 ± 78 CC events, 2010 ± 85 NC events and 279 ± 26 ES events due to solar neutrinos. If the only flavour present in solar neutrinos was electron neutrinos then the ratio of the CC and NC fluxes would have been unity; however the ratio was found to be [49]

$$\frac{\Phi_{CC}^{SNO}}{\Phi_{NC}^{SNO}} = 0.340 \pm 0.023^{+0.029}_{-0.031}, \quad (2.31)$$

strongly suggesting that about $\frac{2}{3}$ of the neutrino flux had been converted to muon and tau neutrinos. Figure 2.6 shows the fluxes SNO found in the different channels. Results from SNO in Figure 2.6 show that approximately two thirds of the solar neutrinos had changed flavour. The complete mixing of the

neutrino flavours is now explained by a resonant mixing due to the density of the sun that gives the same effect as maximal mixing [7]. SNO favoured the LMA (Large Mixing Angle) model [4], which implied a “large” mass splitting and large mixing angle.

2.3.2 Neutrino Oscillation Experiments

Following the confirmation by the SNO collaboration that flavour change of neutrinos was the cause of the observed solar neutrino deficit, a new generation of neutrino experiments were built to study neutrino flavour change. Most experiments studying neutrino oscillations fell into one of two categories, determined by the mass splitting. Solar neutrino experiments such as SNO studied oscillations due to a small mass difference ($\approx 8 \times 10^{-5} \text{ eV}^2$). Atmospheric neutrino experiments examined oscillations driven by a larger mass difference ($\approx 2.5 \times 10^{-3}$). Super-Kamiokande and the long baseline experiments studied this regime. A third class also existed, often called the LSND regime. This explored possible oscillations due to a much larger mass difference, although only the LSND experiment claims to have observed a signal of this large mass difference.

Neutrino experiments looking at anti-electron neutrinos had continued since Reines and Cowan’s discovery of the neutrino. Experiments with a target mass in the region of 100 kg, such as Bugey [50], failed to observe any neutrino oscillations because they were simply too close to their respective reactors to be sensitive to $\bar{\nu}_e$ disappearance. Longer baseline reactor experiments, in particular CHOOZ and Palo Verde, were able to make useful contributions in the atmospheric neutrino sector. Looking for the disappearance of $\bar{\nu}_e$ neutrinos allowed limits to be placed on the value of θ_{13} , the Palo Verde experiment

found that $\sin^2(2\theta_{13}) \leq 0.17$ [51] while the CHOOZ experiment found a value of $\sin^2(2\theta_{13}) \leq 0.13$ [52].

Unique amongst the reactor experiments due to its comparatively large baseline was the KamLAND (Kamioka Liquid scintillator Anti-Neutrino Detector) experiment. KamLAND aimed to confirm that the LMA model for solar neutrino mixing was correct. Assuming a two neutrino scenario the probability of oscillation is given by,

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sin^2 2\theta \sin^2 \left(\frac{1.267(\Delta m^2/eV^2)(L/km)}{E/GeV} \right). \quad (2.32)$$

To observe oscillations directly therefore the ratio of the energy in GeV to the baseline in km should be of the same order as the square of the mass difference between the two neutrino flavours. The average distance from a nuclear reactor to the detector is 180 km and the energy of a $\bar{\nu}_e$ from nuclear fission is of order 2 MeV. Neutrino spectra for various nuclear processes found in nuclear reactors are shown in [4]. KamLAND was able to observe oscillations due to the same mass difference as the solar experiments because the splitting between the mass states was of order $10^{-4}eV^2$.

KamLAND consisted of a roughly spherical vessel containing 3000 tonnes of liquid scintillator, with photomultipliers around the side of the vessel [53]. The vessel was itself suspended in a non scintillating oil. Anti-electron neutrinos were detected via the reaction,

$$\bar{\nu}_e + p \rightarrow n + e^+. \quad (2.33)$$

In Figure 2.7 the ratio of expected to observed neutrino events is shown as a function of energy. The oscillation in neutrino deficit provides a direct observation of neutrino oscillation. KamLAND's ability to not just measure a deficit in the expected number of neutrinos but also to observe the oscilla-




Figure 2.7: The ratio of expected to observed neutrino events in the KamLAND experiment as a function of energy, Figure from [54]. Background due to geo-neutrinos has been subtracted.

tions in the energy spectra allowed it to make a sensible measurement of the mass splitting responsible for neutrino oscillations in the solar neutrino sector. Figure 2.8 shows the constraints on the solar parameters due to SNO and KamLAND. KamLAND accurately measured the energy spectra allowing a good measurement of the mass difference.

While SNO and KamLAND, with contributions from radio-chemical experiments, studied the solar neutrino sector, Super-Kamiokande was observing atmospheric neutrinos. These are neutrinos generated by interactions in the atmosphere from cosmic particles. The results from Super-Kamiokande are consistent with other experiments such as MACRO [44] and Soudan-2 [45], which used an iron detector as opposed to a water Čerenkov detector. As was demonstrated by KamLAND, terrestrial neutrino experiments are potentially very powerful because the baseline and energy of the neutrino can be chosen so as to lie on the oscillation maximum and give excellent measurements of the




Figure 2.8: A combined fit of the solar neutrino mass splitting and mixing angle, Figure from [11]. The SNO data comes from data in [48].

mass splitting, Δm_{atm}^2 . The favoured mass splitting for the atmospheric neutrino sector is approximately $2 \times 10^{-3} \text{eV}^2$ [55]. To resolve neutrino oscillations in the energy spectra the ratio $L/E \simeq 1000 \text{ km/GeV}$.

Long baseline experiments create an intense beam of neutrinos by colliding protons with a solid target and then guiding the resulting pions with magnets. A beam of neutrinos is created from the decay of the pions. The K2K experiment used Super-Kamiokande as its far detector. The K2K (KEK to Kamioka) experiment had a baseline of 250 km [10], while the MINOS (Main Injector Neutrino Oscillation Search) experiment has a baseline of 810 km [56]. MINOS had iron and scintillator tracking detectors as both near and far detectors. Having identical detector designs helped to reduce the complexity of a near-to-far extrapolation. The Fermilab Main Injector provided very high energy protons with a momentum of 120 GeV/c. Changing the current and position of the magnetic horns used to channel the pions allowed the neutrino momentum

to be varied. Typically the neutrino momentum was between 3 GeV and 9 GeV. At this momentum interactions were predominantly DIS (Deep Inelastic Scattering). Without performing a global fit combining data from more than one experiment, MINOS has the best measurement of the atmospheric mass splitting. MINOS measured Δm_{atm}^2 to be $|\Delta_{atm}^2| = (2.43 \pm 0.13) \times 10^{-3} eV^2$ at a 68% confidence level and the atmospheric mixing angle to be consistent with maximal.

The third mass splitting regime, the LSND (Liquid Scintillator Neutrino Detector) regime, so called because of the LSND experiment which claimed to have seen a signal [9]. The LSND experiment saw an excess of $\bar{\nu}_e$ events, corresponding to a $\Delta m_{LSND}^2 \approx 1 eV^2$, depending on the mixing angle. This however is larger than the sum of the atmospheric and solar mass differences, so it is not consistent with a three neutrino model. If the LSND result is confirmed then it would be confirmation of a fourth neutrino flavour. A number of other neutrino experiments (Nomad [57], Karmen2 [58] and more recently MiniBooNE [59]) have been sensitive to the same parameter space as the LSND experiment and have succeeded in refuting its results. The MiniBooNE experiment has excluded the LSND parameter space by searching for electron neutrino appearance [59]. Currently MiniBooNE is repeating the experiment by looking for anti-electron neutrino appearance and has currently excluded some of the LSND parameter space [60]. Data is still being collected for the anti-neutrino appearance analysis.

The discussion so far has concentrated on ‘solar’ and ‘atmospheric’ neutrino sectors. In practise it is a quirk of the oscillation parameters, the large difference between the mass splittings and the smallness of θ_{13} , that we can think of these two decoupled regimes. The latest three flavour oscillation results for

atmospheric and solar neutrino parameters as of 2008 are shown in Figure 2.9. A number of experiments contribute to the experimental limit on the value of




Figure 2.9: Global limits on solar and atmospheric neutrino parameters from [61]. MINOS and KamLAND provide the best measurements of the atmospheric and solar mass differences respectively.

θ_{13} , however the current limit is driven by the CHOOZ experiment [61]. In [61] it has been noted that the latest results from SNO suggest that θ_{13} maybe non-zero, $\theta_{13} = 0$ is now disfavoured at a level of 1.5σ . Work by [61] combines data from neutrino oscillation experiments to find best global values for the three flavour neutrino parameters. These are shown in Table 2.2.

Parameter	Best Fit	2σ	3σ
$\Delta m_{21}^2 [10^{-5} \text{ eV}^2]$	$7.65^{+0.23}_{-0.20}$	$7.25 - 8.11$	$7.05 - 8.34$
$\Delta m_{31}^2 [10^{-3} \text{ eV}^2]$	$2.40^{+0.12}_{-0.11}$	$2.18 - 2.64$	$2.07 - 2.75$
$\sin^2 \theta_{12}$	$0.304^{+0.022}_{-0.016}$	$0.27 - 0.35$	$0.25 - 0.37$
$\sin^2 \theta_{23}$	$0.50^{+0.07}_{-0.06}$	$0.39 - 0.63$	$0.36 - 0.63$
$\sin^2 \theta_{13}$	$0.01^{+0.016}_{-0.011}$	≤ 0.040	≤ 0.056

Table 2.2: Global best fit values for three neutrino model mixing parameters as determined by [61].

2.4 T2K

T2K has two main physics goals, in common with other long baseline experiments [18]. The first of which is a muon disappearance measurement. With this T2K will be able to make a more precise measurement of the atmospheric neutrino parameters, Δm_{atm}^2 and the mixing angle $\sin^2 2\theta_{atm}$, than has been achieved so far [18]. The second analysis is an electron neutrino appearance analysis. T2K will be sensitive to the value of the mixing angle $\sin^2 2\theta_{13}$ below the CHOOZ limit. After 5 years of running the T2K collaboration aims to have achieved close to $750 \text{ kW} \times 10^7 \text{ s}$.

Muon Neutrino Disappearance

In five years, without oscillations, 11,000 events are expected at Super-Kamiokande, with 4,805 events after cuts [18]. In the case that oscillations occur then one expects to see fewer neutrino events at Super-Kamiokande. The exact number depends on the oscillation parameters. Under the assumption that the atmospheric mass difference is $\Delta m_{atm}^2 = 2.3 \times 10^{-3}$ at maximal mixing 1,344 events would be expected at Super-Kamiokande after cuts [18]. Figure 2.10 shows the predicted reconstructed neutrino spectrum at Super-Kamiokande for fully contained muon-like events. From this the oscillation parameters can be determined. The ratio of expected events between the oscillating and non-oscillating neutrino hypotheses shows a significant deviation from unity around the T2K peak energy. During the first year of data taking the power of the beam will be limited to 43 kW permitting approximately 2×10^{20} POT. After this first data run, limits on the atmospheric neutrino parameters are predicted to be very competitive to similar experiments such as MINOS. After 2×10^{20} POT, T2K aims to measure the value of $\sin^2 2\theta_{23}$ with an uncertainty




Figure 2.10: The reconstructed neutrino spectrum (left) assuming $\Delta m_{atm}^2 = 2.7 \times 10^{-3}$ and $\sin^2 2\theta_{atm} = 1.0$, the hatched area shows non-quasi-elastic component. The Figure (right) shows the ratio of events between the oscillated and unoscillated neutrino cases as a function of energy [18].

of 0.03 and Δm_{23}^2 with an uncertainty of 1.6×10^{-4} [18].

Electron Neutrino Appearance

The value of θ_{13} is measured or constrained by the electron neutrino appearance analysis. This is a very important analysis at T2K. Background reduction is a significant challenge as a positive signal will be small with significant background channels. In particular neutral current processes with a π^0 in the final state are an important background, as one of the photons from the resultant decay can appear as an electron. Electron neutrinos in the beam constitute the other major source of background. After five years, assuming that $\Delta m^2 = 2.5 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta_{13} = 0.1$, 103 signal events are expected at Super-Kamiokande with a further 23 background events [18].

Chapter 3

T2K and the ND280 Detector

3.1 Introduction to T2K

T2K is a long baseline neutrino oscillation experiment studying the disappearance of muon neutrinos between two detectors [18]. Neutrinos are generated from the 30 GeV J-PARC (Japan Proton Accelerator Research Complex) proton beam located at Tokai-mura on the East coast of Japan. The experiment uses two detectors to measure the change in electron and muon neutrino flux. The near detector (ND280), so called because it is 280 m from the proton target, is located at the J-PARC site. The far detector is the Super-Kamiokande neutrino observatory in Kamioka, 295 km from the neutrino source. This creates a large value of L/E , hence the term ‘long baseline experiment’. Neutrinos from the T2K beamline have a mean energy of about 700 MeV. T2K is therefore sensitive to oscillations between neutrino flavours with a mass splitting in the region of $10^{-2} - 10^{-3} eV^2$. This corresponds to the ‘atmospheric sector’ discussed in Chapter 2.

The J-PARC beam produces mostly muon neutrinos which one expects to os-

cillate predominantly to tau neutrinos. Below energies of about 5 GeV the tau neutrino cannot interact via a charged current interaction as the production of the tau is kinematically forbidden. At Super-Kamiokande muon disappearance will be used to study the atmospheric neutrino parameters. Super-Kamiokande will also search for the appearance of electron neutrinos to test whether θ_{13} is non zero. This search is the primary goal for T2K. CCQE interactions are the most interesting to T2K as they are the dominant interaction in the region of hundreds of MeV. They are also favourable from an analysis perspective as the energy of the neutrino can be reconstructed from the lepton and proton as discussed in Section 2.2.

Its proximity to the neutrino source gives the near detector a very high event rate. This makes it ideal for determining key cross sections for signal processes such as CCQE electron and muon neutrino interactions. Also background cross sections for the far detector can be determined. Of particular importance is the cross section for single π^0 production as this is a large source of background at Super-Kamiokande. From the near detector the spectrum of electron and muon neutrinos can be determined. These are important because they provide unoscillated spectra which can be extrapolated to the far detector, producing an estimate of the expected spectra at the far detector in the absence of neutrino oscillations. Comparing this to what is observed at the far detector provides the basis of the oscillation analysis.

T2K is an ‘off-axis’ experiment, where the term ‘off-axis’ means that the neutrino detectors sit at an angle to the direction of the neutrino beam. The near and far detectors in T2K are offset by 2.3° with respect to the direction of the decay pipe. The energy spectra shown in Figure 3.1 can be seen to peak at a lower energy and become narrower, in particular the high energy




Figure 3.1: The solid black line shows the energy spectrum of on-axis neutrinos in T2K. By changing the angle to 2° (dashed) or 3° (dotted) the peak energy is lowered [18].

tail shortens. This combination of effects is beneficial for T2K. Critical to the oscillation analysis is to construct the experiment so that $L/E \approx 1000$ km/GeV, as can be seen in Equation 2.16. The distance is fixed at 295 km so, by having the beam off-axis, the peak neutrino energy can be moved so as to produce an oscillation maximum. Shortening the high energy tail of the neutrino beam brings reconstruction advantages. At energies much above 1 GeV DIS (Deep Inelastic Scattering) becomes significant. These types of event are harder to reconstruct for T2K. By reducing the energy more of the events are comparatively simple quasi-elastic interactions. DIS events are also irrelevant in oscillation studies as they do not oscillate significantly over the T2K baseline.

3.2 J-PARC Neutrino Beamline

The proton beam for T2K is accelerated in three steps [62] [63]. Firstly a linac accelerates negatively charged H^- ions to 400 MeV in bunches with a frequency of 50 Hz. Electrons are removed from the ionised Hydrogen by a copper stripper foil, creating a beam of protons. The second stage accelerates the protons to 3 GeV. Finally the beam is passed into the main PS (Proton Synchrotron) which accelerates the protons to 30 GeV.

Protons from the main ring beamline are extracted for T2K through a series of kicker magnets. Protons impinge upon a carbon target producing an intense source of pions and a smaller number of kaons. Secondary particles are collected, focused and sign selected by magnetic horns and directed down a decay pipe. The horns also purify the beam by deflecting away pions of a given polarity, allowing a pure neutrino or anti-neutrino beam to be produced.

In the decay pipe kaons and muons decay into muon neutrinos with some fraction decaying into electron neutrinos, anti-electron neutrinos and anti-muon neutrinos. The length of the decay pipe is 110 m with the dimensions of the cross section increasing from 2.2 m wide x 2.8 m high to 3 m wide x 4 m high furthest from the horns. The length is tuned to maximise the number of pion decays while minimising the electron neutrino contamination from the muon decay process, $\nu_\mu \rightarrow \bar{\nu}_\mu + \nu_e + e$. Ideally, muons from the pion decay lose their energy in a beam dump at the end of the decay pipe. Muons losing their energy in the beam dump produce low energy neutrinos isotropically, so are a less significant background than higher energy muons [64].

At the end of the beam pipe, past the beam dump, is a muon flux monitor which will operate in the early days of the experiment. The harsh radiation

environment of the beam pipe will severely limit the lifetime of the monitor but it will provide valuable information on the exact profile of the beam in the decay pipe which can be used to validate the Monte Carlo simulation of the beam. The projected beam flux at ND280 and Super-Kamiokande as a function of energy are shown in Figure 3.2.




Figure 3.2: Predicted neutrino flux at ND280 (top) and Super-Kamiokande (middle) and the near to far ratio (bottom) for both ν_μ (left) and ν_e (right) [18].

3.3 INGRID on-axis detector

INGRID (Interactive Neutrino GRID) [18] sits in the near detector pit, 280 m from the beam target, situated on axis with respect to the beam line. INGRID is an iron-scintillator detector with alternating layers of iron and active scintillator detector and is laid out with a cross-shaped profile, as shown in Figure

3.3. The long arms allow a broad measurement of the beam profile. A set of 16 modules make up INGRID. Each module contains 11 scintillator layers measuring $1 \times 1 \times 0.03 \text{ m}^3$ and 10 iron layers measuring $1 \times 1 \times 0.1 \text{ m}^3$. The purpose of the INGRID detector is to provide to provide daily characterisation of the neutrino beam. INGRID will also monitor the direction of the neutrino beam and characterise the profile of the beam.




Figure 3.3: The structure of an individual module with lead and scintillator (left) and the cross shaped arrangement they are assembled in (right) [18].

3.4 Super-Kamiokande

On the west coast of Japan, the T2K far detector is located in a former mine under Mt. Ikenoyama. Super-Kamiokande reported its first results operating as a neutrino observatory in 1998, studying atmospheric and solar neutrinos. To shield it from cosmic ray muons Super-Kamiokande is 2,700 m.w.e (Meter Water Equivalent) or 1000 m of rock underground. Filled with 50 kt of water,




Figure 3.4: A diagram of the Super-Kamiokande Detector. Figure adapted from [66].

Super-Kamiokande is the largest neutrino detector in the world [65]. Super-Kamiokande I completed data taking in 2001, when the detector was upgraded. During a refill of the detector, a damaged photosensor imploded, the shock of which destroyed many of the remaining photosensors. Run II of the Super-Kamiokande experiment was carried out with 47% of the photosensors used in the full detector, and stopped in 2006 for new photosensors to be installed [43]. Returned to full photosensor coverage, Super-Kamiokande III took data until it was upgraded to Super-Kamiokande IV. Super-Kamiokande IV has been completed for T2K data taking, and consists of new electronics with updated reconstruction software.

The Super-Kamiokande detector is 42 m in height and 39 m in diameter, split into two regions - the inner detector (ID) and outer detector (OD). The inner detector consists of 11,146 photomultipliers of 50 cm diameter and is 36.2 m

high and 16.9 m in radius, and is used for observing signal events. The outer detector is used as a veto for cosmic ray muons that penetrate the rock, as well as other particle interactions in the material surrounding the detector. The outer detector uses 1,885 photomultipliers attached to a metal bulkhead separating the inner and outer detectors.

Events in Super-Kamiokande are classed as either FC (Fully Contained) or PC (Partially Contained). An event is FC if a signal is measured on the inner detector. If however, a signal is measured in the outer detector consistent with an outgoing particle, then the event is tagged PC. In a PC event it is only possible to put a lower limit on the energy of the neutrino, as at least some energy has escaped. Stopping muon events tagged by the outer detector as having come from the atmosphere are useful for calibration, and occur with a frequency of approximately 2 Hz.

3.4.1 Super-Kamiokande Reconstruction

A water Čerenkov detector, such as Super-Kamiokande, uses high gain Photo Multiplier Tubes (PMTs) to detect the Čerenkov light given off by a medium when a particle travelling faster than the phase velocity of light in that medium passes through. When a charged particle passes through a non-conductive material temporary polarisation occurs. If the particle is travelling faster than the phase velocity of light, then at a certain angle with respect to the direction of motion of the charged particle light will be emitted [67]. The angle is given by the expression,

$$\cos\theta = \frac{c}{vn} = \frac{1}{\beta n}. \quad (3.1)$$

Projected into the detector, the cone of Čerenkov light creates a ring of hits in

the detector, centred on the direction of motion of the particle. Particle identification of electrons and muons considers how well defined the ring of PMT hits are. A massive particle such as a muon will not radiate bremsstrahlung or scatter significantly in the water, creating a well defined ring as shown in Figure 3.5. Electrons however, lose energy radiatively and scatter off atoms; showering to create a number of lower momentum secondary electrons and positrons, which can also give off Čerenkov light. A ring due to an electron therefore has a poorly defined edge and is known as a ‘fuzzy ring’. An example of such an event can be seen in Figure 3.5. In the case of a muon interaction, a second signal can be observed later in time, due to the Michel electron from the decay of the muon.

A large background in the electron neutrino analysis, due to π^0 events, motivates the P0D detector in ND280. The π^0 background is caused by co-linear photons, or highly asymmetric events where one photon takes most of the π^0 energy and the second photon is not reconstructed. If both gammas are close to co-linear, then the pair of rings will lie on top of each other and the ring finding algorithm will see a ‘fuzzy’ ring and incorrectly identify the particle as an electron.

In Chapter 6 an electron neutrino analysis is designed using the ND280 detector. This analysis measures the electron neutrino flux at the near detector. Extrapolating this measurement to the far detector provides an estimate of what would be seen by Super-Kamiokande in the absence of $\nu_\mu \rightarrow \nu_e$ oscillations. An excess of electron neutrinos at Super-Kamiokande would represent a non-zero value of θ_{13} .




Figure 3.5: Event displays from Super-Kamiokande. An electron event with a ‘fuzzy’ ring is shown (top left) and a sharper muon ring (top right). A two ring π^0 event is shown (bottom left) as well as a more complicated event containing three rings (bottom right). The relative timing of PMT hits is indicated by the colour of the square and the magnitude by its size [68].

3.5 ND280 off axis near detector

The ND280 near detector, located 280 m away from the proton target, will accurately characterise the T2K neutrino beam and measure key cross sections. The near detector consists of two principal sections: the tracker which measures the flux of electron and muon neutrinos, and the P0D (π^0 Detector) which measures the inclusive neutral current single π^0 production cross section, a key background at the far detector. A TPC (Time Projection Chamber) and a FGD (Fine Grained Detector) lie at the heart of the tracker part of ND280. Precision momentum measurements and efficient PID are provided by

the TPC. However, as a TPC is filled with gas it has a very low density, so very few neutrino events occur. The FGD, made of plastic scintillator, provides a denser medium where most tracker neutrino events will occur. Vertex finding is helped by using fine plastic scintillator bars. The tracker is surrounded by a lead-scintillator calorimeter which contains events, determines the energy of particles, and provides additional particle identification.

The ND280 coordinate system is defined such that the Z axis runs through the centre of the detector close to the direction of the neutrino beam. The Y axis points vertically out of the detector and the X axis is orientated to give a right handed coordinate system. The origin of the coordinate system is in the centre of the detector.

Plastic scintillator is a recurring theme in the ND280 detector. Water is held within water bags in the P0D and the downstream FGD allowing measurements of neutrino cross sections on water. As is discussed in more detail below in the discussion of Super-Kamiokande the neutral current π^0 cross section is particularly important as it is a dominant background.

The core of the detector is encased within a magnet to provide a magnetic field for momentum measurements. When energised the nominal magnetic field is 0.2 T however due to power supply constraints the initial field may be closer to 0.17 T. Interleaved between the yokes in the magnet is a muon range detector, called the SMRD (Side Muon Range Detector). The structure of the near detector is shown in Figure 3.6.




Figure 3.6: A schematic view of the ND280 is shown where the P0D is at the upstream side of the near detector and the tracker section is at the downstream end of the detector.

3.5.1 P0D (π^0 Detector)

Super-Kamiokande detects a π^0 by looking for Čerenkov radiation from the two photons that come from the decay. If one of the photons is missed, either because it is low in energy or because it overlaps with the other photon, then the event may be misidentified as an electron event. The primary objective of the P0D is to measure the inclusive π^0 cross section on water. This allows a Monte Carlo estimation of the number of π^0 events expected at Super-Kamiokande. Cross sections for processes such as,

$$\nu_\mu n \rightarrow \mu^- p \pi^0 \quad (3.2)$$

or

$$\nu_\mu N \rightarrow N \pi^0 \quad (3.3)$$

currently have poorly known production rates. The P0D will provide the data necessary to quantify backgrounds [18].

The P0D, like much of the near detector, is made from plastic scintillator, with a wavelength shifting (WLS) fibre coupled to a photosensor to read out signals. To aid with reconstruction, scintillator bars are triangular so a track is less likely to be contained within a single column of bars. If a particle hits two neighbouring bars, then the ratio of the charge can be used to improve the resolution of the hit position. Neighbouring layers of scintillator are arranged perpendicular to each other to aid with 3D reconstruction.

To optimise the conversion of gammas from π^0 decay, thin lead sheets are used to provide an interaction mass. The lead scintillator layers are located on either side of the detector to contain decay photons. Upstream of the water filled section of the P0D the lead sheets are 4 mm thick and downstream of the water filled section the lead is 4.4 mm thick. Water bladders are located on the inside of the detector to provide a water target for the neutrinos. Each of the water bladders holds approximately 100 kg of water with dimensions of 3 cm x 1.8 m x 2.1 m. Lead radiators are used on the outer parts of the P0D to contain photons as lead has a very short radiation length. However, this reduces the energy resolution. Around the water bladders brass is used as the radiator material as it has a shorter radiation length, giving a superior energy resolution.

3.5.2 Fine Grained Detector

Combined with the TPC and the tracker ECal (downstream and barrel), the FGD (Fine Grained Detector) forms part of the tracker section of ND280. In the first instance the tracker will measure the flux of ν_e and ν_μ neutrinos through charged current interactions. The Charged Current Quasi-Elastic (CCQE) channel is the most important channel as the single lepton and proton channel makes reconstruction more simple. Energy reconstruction is also easier as the energy of the neutrino can be reconstructed from the kinematics of the lepton and nucleon, up to assumptions about the Fermi momentum, as discussed in Section 2.2. If the proton travels far enough to be reconstructed then the neutrino can be completely reconstructed. These measurements provide the un-oscillated muon and electron neutrino event rates for the oscillation analysis. In the longer term the high neutrino flux will allow cross section studies of different neutrino interactions such as neutral and charged pion production cross sections.

As a TPC is filled with gas, the vast majority of reconstructable events in the tracker occur in the FGD. The absence of any lead in the FGD allows a lepton from a CCQE interaction to propagate into the TPC. In the TPC accurate momentum measurements can be made by measuring the curvature of a track caused by a 0.2 T magnetic field. Measuring the energy loss of the track as it travels forms the basis of PID in the TPC. A driving factor in the spatial resolution of the FGD is the need to reconstruct the recoil proton. As protons from CCQE interactions have very low momenta they will not generally leave the FGD. Scintillator bars in the FGD must therefore be narrow enough to allow the proton to travel far enough to be reconstructed. A narrow bar is also preferable as the resolution of the position of the neutrino vertex is partly




Figure 3.7: A diagram showing a cross-section of a FGD scintillator bar. The titanium dioxide coating helps to contain light within a given bar. Optical fibre sits in the hole running through the centre of the bar to collect scintillation light [18].

determined by the size of the scintillator bar.

The FGD is an extruded scintillator detector which uses wavelength shifting fibre to transport light to a photosensor as shown in Figure 3.7. Each scintillator bar is coated in a TiO_2 layer to help contain light within that scintillator bar, and to reduce optical cross-talk with neighbouring bars. Narrow scintillator bars are used in the FGD, measuring $0.96 \times 0.96 \times 184.3$ cm. Each fibre has a mirrored end to increase the signal at the other end which is coupled to a MPPC. Scintillator bars are arranged into layers in the XY plane with the direction of the scintillator bars alternating between the X and Y direction.

Two FGDs have been constructed for use in the tracker with identical dimensions, $230 \times 240 \times 36.5$ cm. The upstream FGD, closest to the P0D, is made entirely of scintillator bars, 5,760 in total. Fewer scintillator bars are present in the downstream FGD as six 2.5 cm thick water modules are used to create a fraction of neutrino events on water. By comparing the cross sections measured in the upstream and downstream FGDs the cross sections on water can

be determined.

3.5.3 Time Projection Chamber (TPC)

The tracker consists of three TPCs, each measuring 2.5 m x 2.5 m x 1 m, with the TPC being 1 m deep in the beam direction to achieve the necessary momentum resolution [18]. TPC 0 is the most upstream TPC, located between the P0D and first FGD. The second TPC is positioned between the two FGDs, measuring the momenta of forward going particle from TPC 1, or backwards going particles from FGD 2. The final TPC measures the momentum of the high momentum forward going particles. Only the TPC does not use a variant of scintillator detector.

The active volume of each TPC is filled with a mixture of Argon, CF_4 and C_4H_{10} with approximately 95% Argon with the remaining 5% made up of 3% CF_4 and 2% C_4H_{10} . This ratio of gases is chosen so as to maximise the drift velocity of the electrons. Each TPC is structured with an anode at either side of the outer box with a central cathode held at a potential of -25 kV. A field gradient of 200 V/cm is maintained within the sensitive regions of TPC modules. MICROMEGAS sensors are fixed to the opposite side of the gas box from the cathode, see Figure 3.8. Field uniformity is affected by the presence of the MICROMEGAS sensors. Their effect on the potential can be seen in Figure 3.8. Each MICROMEGAS module measures $36 \times 34 \text{ cm}^3$ with an individual pad size of $6.9 \times 9.7 \text{ cm}^2$. MICROMEGAS modules sit just 100 μm above the surface of the anode, detecting free electrons from the gas drifting into the anode. Only information in the YZ plane is provided by the MICROMEGAS modules. This was chosen because the high momentum CCQE leptons that the tracker is most interested in will be mostly forward




Figure 3.8: Images from the ND280 technical review [18] showing the design of a TPC gas cage (left) and a simulated field map (right). A good field uniformity of close to 180 V/cm is predicted. Some fluctuations are seen at the edges due to simulation precision [18].

going, along the Z axis, while the magnetic field will bend particles in the Y axis. A high spatial resolution is therefore most useful in the YZ plane to make accurate momentum measurements. The TPC recovers information in the X axis by looking at the time taken for the drift electrons to reach the MICROMEGAS modules. This requires a good initial hit position in the X axis. This must be provided either by the FGD or the ECal.

Physics goals for the tracker were discussed previously in the discussion of the FGD. A few thousand events every year will occur in the gas of the TPC. For these events the high spatial resolution of the tracker will be able to reconstruct all charged final state particles, allowing detailed studies of neutrino-nucleus/nucleon interactions.

3.5.4 Electromagnetic Calorimeter

Surrounding the P0D, TPC and FGD is an Electromagnetic Calorimeter (ECal). In keeping with other sub-detectors it is a lead/scintillator detector. Each scintillator bar carries a wavelength shifting fibre down its centre to capture

photons and transport them to ends of the bars. It differs from the P0D in that it has comparatively thick lead to contain any showers propagating from the tracker or P0D. The lead in the P0D ECal (the part of the ECal surrounding the P0D), at 4 mm, is thicker than the 1.75 mm lead in the tracker ECal (the part of the ECal surrounding the TPCs and FGDs). Compared to the tracker the ECal is a very coarsely grained detector. The scintillator bars vary in length from 1.6 m to 4 m; however they all have a cross section of 1 cm x 4 cm. Each layer of scintillator is orientated at 90° to the neighbouring layer, so as to give useful positional information in three dimensions. For example, in the barrel ECal one layer will give information in the XZ (top view) plane and the next layer will give useful information in the YZ (side view) plane.

A sampling calorimeter like the ECal uses a dense material like lead to induce showers in charged particles allowing their energy to be measured. Importantly, the lead in the ECal causes photons, which are completely invisible to the TPC, to shower. The ECal can catch bremsstrahlung photons from interactions in the tracker and help reconstruct the full energy of the original electron. The ECal is designed to have a resolution of approximately $7.5\%/\sqrt{E}$ for electromagnetic showers [18], providing a complementary measurement to any momentum measurement from the TPC and aiding in particle identification. As the ECal encloses the entire tracker section of ND280, apart from a few small gaps due to mechanical structures and electronics, very few charged particles will not be visible in the ECal.

The downstream (DsECal) sits at the downstream end of the tracker in the XY plane of the detector. All of the DsECal scintillator bars are 2 m long, with a photosensor present at each end of the bar. The attenuation length of the optical fibre is such that longer bars need sensors at both ends. Single

ended bars in the P0D ECal and barrel ECal are mirrored with evaporated aluminium at the un-instrumented end to recover light travelling away from the photosensor. Containing 34 layers of scintillator bars, the DsECal measures $200 \times 200 \times 50 \text{ cm}^3$ and is readout from the top and side. Six modules make up the barrel ECal, one on each side of the tracker and two on the top and bottom. The top and bottom modules are split into two as they need to come apart when the magnet is opened. Top and bottom barrel ECal modules have dimensions, $150 \times 50 \times 420 \text{ cm}^3$ and the side barrel ECal modules measure $50 \times 230 \times 420 \text{ cm}^3$. At the upstream end of the detector, the P0D ECal has similar dimensions to the tracker ECal, except it is slightly shorter as the P0D ECal is not as large as the tracker. Top and bottom P0D modules measure $140 \times 50 \times 230 \text{ cm}^3$ and the side modules measure $50 \times 230 \times 260 \text{ cm}^3$ [18]. The P0D ECal reads out data in a single view only. ECal modules are arranged so

ECal module	Layers	Bars in Layer	Total Channels
Downstream	34	50	1700
Top/Bottom Barrel	31	105/37	3255/1147
Side Barrel	31	105/37	3255/1147
Top/Bottom P0D	6	35	210
Side P0D	6	58	348

Table 3.1: Numbers of bars in each ECal module [18].

that every bar provides positional information for particle leaving the tracker, this means the ‘long’ axis of a bar is always parallel to the FGDs and TPCs. For example every layer in the DsECal always gives useful positional information in the Z axis.

3.5.5 Side Muon Range Detector

A characteristic of neutrino detectors is that interactions are not guaranteed to occur within the centre of the most sensitive part of the detector, as is the case with a collider detector. Unlike the far detector, ND280 is not deep underground, so there will be a significant flux of cosmic muons passing through the detector, although this will be small compared to the duty cycle of the beam. The SMRD (Side Muon Range Detector) is another extruded scintillator detector that lies within the steel yokes of the magnet enclosing the inner detectors. From its position around the periphery of the detector, it can be used to tag cosmic rays coming from above and to detect neutrino events in the surrounding sand. In turn these events can either be used for calibration in the case of cosmic rays or vetoed or reconstructed if it was a neutrino interaction in the sand. A further purpose for the SMRD is as a range detector, measuring the energy and direction of muons that leave the core of ND280.

The size and number of SMRD modules is determined by the structure of the magnet. The magnet yoke is split into 8 C-like segments on either side of the detector. This can be seen in Figure 3.6. Each yoke consists of 18 layers made from 48 mm thick steel, each layer separated by a 17 mm thick air gap, as shown by Figure 3.9. In turn, each air gap contains a series of 12 spacers which hold the SMRD modules. In each air gap only 8 of the 12 sections can be instrumented with SMRD modules due to the corners in magnet. The number of SMRD modules in different regions of the magnet varies with the dominant use of the SMRD in that region. For example, the top and bottom modules of the SMRD are to be used primarily to tag cosmic muon events. This demands an even coverage of modules so that there is no cosmic trigger bias. Each air gap in the top and bottom modules has four layers instrumented, with each air

gap containing three SMRD ‘slabs’. An example of an SMRD slab is shown in Figure 3.9. Modules in the side of the magnet are intended to observe muons ranging out from charged current interactions in the tracker. As these will be forward going particles, the downstream end of the magnet is more heavily instrumented than the upstream end. The three most upstream ‘C’s of the magnet have two air gaps instrumented, whereas the two most downstream sections have six air gaps instrumented and the three in between have four air gap instruments. Each of the modules in the side of the magnet is made up of five slabs.



Figure 3.9: A view of one of the ND280 magnet yokes with air gaps visible (left) and an SMRD ‘slab’ (right) [18].

Each slab has a fibre in a groove in the surface of the bar; unlike the other scintillator bars which have a fibre running down the centre of the bar. Fibres follow an S-shaped path through the scintillator bar to improve light collection. This feature is needed only on the SMRD modules because they are comparatively wide compared to the other scintillator bar designs in ND280. Horizontal bars in the top and bottom of yokes measure $87 \times 16.7 \times 0.7 \text{ cm}^3$ and vertical bars in the sides of the yokes $87 \times 17.5 \times 0.7 \text{ cm}^3$.

3.5.6 Scintillator Detectors

With the exception of the TPCs, all active ND280 subdetectors use extruded plastic (polystyrene) scintillator as the main detection mechanism. Polystyrene itself is not a scintillating material, so scintillating dopants are added to the polystyrene during the extrusion process. Two scintillating materials (fluors) are used in the ND280 scintillator. The first, PPO (poly(2-5-diphenyloxazole)), emits light when a charged particle scatters off it. This emitted light is quickly attenuated by the polystyrene. A second fluor, POPOP (1,4-bis(5-phenyloxazol-2-yl)), absorbs the light from the PPO and emits light in a region of the UV (Ultra-Violet) spectrum less attenuated by the polystyrene. This two step process is described in [19]. PPO is present in a concentration of 1%, whereas the secondary fluor, POPOP, has a concentration of 0.03% [18].

Each bar is coated in a layer of TiO_2 (Titanium Dioxide), which is used as a reflective layer because it has a high reflection coefficient and a uniform reflectivity throughout the visible spectrum. A reflective layer also helps to reduce optical cross talk by preventing light leaking from one bar into a neighbouring bar.




Figure 3.10: An end on view of a scintillator bar being illuminated with a UV LED. A hole in the middle to accommodate a fibre can be seen [69].

The attenuation length for near UV light in the scintillator bar is of the order of 40 cm. Almost no light therefore would reach the end of a 2 m bar. Optical

fibres are used to transport light to the photosensor. Scintillation light from the polystyrene bars is transported to the MPPCs using WLS (Wavelength Shifting) fibres. Photosensors are most sensitive in the ‘green’ part of the visible spectrum whilst the secondary scintillator emits light in the UV. To convert the UV light to green light, Kuraray Y11 wavelength shifting fibre is used. Y11 is a three layer fibre, with an outer cladding, an inner cladding and a core which forms the ‘light pipe’ through which most of the light is guided. The fibres have a total diameter of 1 mm. To guide the light via total internal reflection the refractive index of layers increases outward from the core, from 1.42 in the core to 1.59 in the outer cladding.

The attenuation length of the fibre is of order 3-4 m, driving the decision to have photosensors at both ends of some of the longer bars found in the ECals. With the exception of the SMRD where the fibre is mounted in a groove on one surface of the scintillator bar; all the subdetectors thread a fibre into a hole running through the centre of the bar, as shown in Figure 3.10.

3.5.7 Multi Pixel Photon Counters

MPPCs (Multi-Pixel Photon Counters) are a comparatively new photosensor technology used in ND280 to detect the light from scintillator detectors. Sensors supplied by Hamamatsu are used in ND280. MPPC is a term used by Hamamatsu [70], other devices operating on the same principle exist under different names such as MRS-APD (Metal Resistive Semiconductor - Avalanche Photo-Diodes). Multi-pixel avalanche photodiodes were chosen for the ND280 because they can operate inside a magnetic field without any effect on their operation. Also, they are physically small, which is an advantage in the tightly packed ND280 detector.

MPPCs are Geiger mode semiconductor photosensors. When a photoelectron is created inside the p-n junction of the diode it is accelerated by the high electric field, initiating an avalanche (Geiger discharge) [71][72]. The diode is connected to a quenching resistor. The avalanche causes a voltage drop over the resistor. This reduces the electric field density and quenches the discharge. The bias voltage then causes the high electric field gradient to recover and the pixel is reset. The total charge released during the avalanche is given by

$$Q_{pixel} = C_{pixel}(V_{op} - V_b), \quad (3.4)$$

where Q is the charge released by a pixel, C is the pixel capacitance, V_{op} is the operating voltage and V_b is the bias voltage. Pixels are created on the MPPC by etching grooves into the surface of the device so that one section can avalanche whilst leaving the field intact elsewhere. A typical pixel capacitance is 50 fF. A typical value for the over voltage, $(V_{operating} - V_{bias})$, is 3 V, so the charge released in an avalanche is 150 fC, or approximately one million electrons. As one photoelectron initiates the avalanche this is equivalent to a gain of around 10^6 [18]. The recovery time of the pixel is determined by its RC constant which is on the order of 10 ns for the Hamamatsu MPPCs.

At low light levels individual photoelectron peaks can be resolved. For higher light levels a continuum is observed. Also, the devices become increasingly non-linear due to photons incident on pixels that have been hit and not recovered. Typically, a MIP-like energy deposit will give a signal of approximately 30 pe. The unit pe, or photon equivalents, refers to the charge released by the device due to a single electron. A threshold is needed in the electronics and in the reconstruction software to remove dark noise signals. The electronics threshold is set at 2.5 pe, removing single PE dark noise signals. The dark noise rate where two pixels are triggered is greater than 10 kHz, equating to

approximately 10 hits in the DsECal during a 550 ns integration cycle.

The avalanche from one pixel can trigger an avalanche in a neighbouring pixel; an effect known as crosstalk. This occurs because charge leaks into a neighbouring pixel. Another source of noise comes from after-pulsing, this is where a second avalanche occurs in one pixel due to the same photon. During an avalanche a mobile charge can become trapped on an imperfection in the semiconductor. When the charge is released some time later it can trigger a second avalanche. Dark noise is a third source of noise found in MPPCs. Thermal excitation of electrons in the semiconductor can initiate an avalanche without an incident photon, hence the term dark noise. In Hamamatsu MPPCs the single photoelectron dark noise frequency is in the region of 100-1000 kHz, although significant variation is seen between devices from different wafers.

3.5.8 ND280 Electronics

Individual subdetectors use either Trip-T (Trigger Pipeline and Timing Chip) or AFTER (ASIC For TPC Electronics Readout) electronics to read out data. AFTER is used in the TPC and the FGD, Trip-T electronics in all other subdetectors. Data collected by Trip-T electronics is sent to a RMM (Readout Merger Module). An RMM controls up to 48 TFBs (Trip-T Frontend Board), which hold Trip-T chips [18], sending run specific parameters and controlling the startup and stopping of TFBs. TFB firmware can also be updated via a RMM. Once a trigger has been received by a RMM it will collect the data from each TFB and pass it to a FPN (Front-end Processing Node). Data are sent by the FPN for storage and offline analysis. Timing and trigger information is sent to the RMM via a SCM (Slave Clock Module) of which there is one for each subdetector using the Trip-T electronics. During full ND280 running the

SCM passes through information from the MCM (Master Clock Module). It can generate its own timing signals, allowing each subdetector to run its own testing and/or calibration. The MCM manages the timing and triggering of the ensemble of ND280 subdetectors during full running. It receives a signal from the GPS system to synchronise beam triggers with the global T2K experiment as well as manage other triggers from within ND280, such as the cosmic trigger [18].

Trip-T Electronics

Trip-T is used in all subdetectors except for the tracker, which demands a very high time resolution. The ASIC (Application Specific Integrated Circuit) was originally designed for the D0 experiment [73]. A Trip-T chip has 32 channels and a 48 channel pipeline to store analogue signals from the MPPCs. The charge is stored in a capacitor before being digitised. The time over which charge is allowed to accumulate before being read out is called the integration period, typically of order 500 ns. A total of 23 integration periods worth of data can be stored per channel. This means that when a trigger is received data is read out covering a few microseconds depending on the selected integration time. Reading out data from the MPPCs requires the electronics to have a dynamic range from 1 to 500 photoelectrons. The Trip-T ASIC does not have the dynamic range required to be sensitive to 500 pe pulses while being able to discriminate 1.5 pe noise signals. To increase the dynamic range, the MPPC signals are split into high and low gain channels. This requires two ADC channels per sensor so each chip can accommodate 16 MPPC channels. Trip-T chips are mounted on a TFB. Each board holds four chips and so can read out 64 MPPC channels. Figure 3.11 shows how an MPPC is connected

to a Trip-T chip and the capacitive dividing of input into high and low gain channels.




Figure 3.11: A diagram showing how a MPPC is connected to a TFB and Trip-T chip. Gain controlling capacitors can also be seen.

The ratio of the gain between the high and low gain ADC channels is determined by the ratio of the capacitance of the coupling capacitors, C_{Hi} and C_{Lo} in Figure 3.11, currently set at 10 to 1. A TFB is also capable of supplying an individual high voltage trim between 0 V and 5 V to each MPPC. This allows gain variations between MPPCs to be controlled by changing the high voltage trim. A charge injection capacitor is also shown in Figure 3.11. This allows charge from a known capacitance to be digitised by the ADC. As ambient conditions can cause a drift in the behaviour of the ADC, taking periodic charge injection runs allows any drift to be tracked and corrected for.

As well as recording the magnitude of any charge deposit, an associated time stamp is also required. A Trip-T chip can assign a time stamp to each recorded MPPC hit with a time resolution of 2.5 ns. The signal from the high gain channel is passed to the TDC (Time to Digital Converter) which applies the time stamp.

AFTER Electronics in the FGD and TPC

Time stamping does not provide a sufficiently high time resolution for the FGD and TPC. In the tracker the AFTER ASIC [18] is used to read out data. An AFTER chip has 72 channels, each reading a single MICROMEGAS pad in the case of the TPC. In the FGD a pair of ASICs are used to read out 64 photosensors providing high and low gain outputs. Unlike Trip-T, the AFTER continuously samples data at a rate of up to 50 MHz into 511 memory bins. Charge collected in the memory creates a waveform of charge against time which is then fitted to produce a hit position in (X,Y,T) space. AFTER ASICs are mounted onto FECs (Front-End Cards).

Each FEC reads out 4 AFTER chips or 288 MICROMEGAS pads, and contains circuitry to protect the MICROMEGAS pads from damage due to high voltage sparks. A four-input ADC is also included on each FEC. Each AFTER chip is connected to one of the four ADC channels. The ADC, supplied by Analog Devices, is capable of 50 MSPS (Mega Samples Per Second). In much the same way that the Trip-T chip provides a charge injection for the ADC calibration, FECs hold a DAC which generates square wave pulses to calibrate the AFTER chips.

AFTER is also used in the FGD, where the improved timing information is useful for finding neutrino interaction vertices. Good time resolution is needed for track matching between the FGD and TPC. The TPC will be reliant on the FGD for a time value to seed its reconstruction of the particle position in the X plane. Processes such as charged pion decay, where the pion decays into a near stationary muon, which in turn decays into a Michel electron, will be more accurately measured using the AFTER electronics than they would with a Trip-T style time stamp.

3.6 Data Acquisition (DAQ)

The DAQ is responsible for collecting the data from the electronics systems of each subdetector and format the data as a complete event. The data is then be sent to be stored for later analysis. As a subdetector can be added or removed during the lifetime of the experiment the DAQ is modular so that it can handle the addition or removal of a subdetector. FPNs (Front-end Processing Nodes) are responsible for collecting the data from the subdetector electronics [18]. At this level trigger information and a GPS time stamp is attached to the data. Data is then passed to a backend network which merges the data from the different subdetectors. The interface to control the experiment and online monitoring is also provided by the DAW. The DAQ system in ND280 is based on the MIDAS software [74].

As well as controlling the movement of data from individual subdetectors to long term storage the DAQ is also responsible for monitoring and controlling the environment of the experiment. Global parameters such as the temperature and humidity are monitored and stored by the DAQ. Subdetector specific peroperties such as gain values, supply voltages and temperatures can be stored. These paramters are used for calbration and determination of data quality [18].

3.7 ND280 Software Suite

As of July 2009, an extensive suite of applications is available for the ND280 detector, ranging from external packages such as ROOT [75] and GEANT [76] to very specific calibration packages. A Monte Carlo package, nd280mc, simulates the ND280 near detector using the GEANT 4 particle simulation frame-

work. Support for two neutrino event generators, GENIE [77] and NEUT [78], seeds the Monte Carlo with neutrino interactions. Reconstruction algorithms are provided by each subdetector, with a global algorithm reconstructing complete events.

All work discussed in this thesis uses v6r1 of the ND280 software with GENIE used to generate neutrino interactions. Labelled a physics release, v6r1 was used for a Monte Carlo data challenge to verify the performance of the Monte Carlo and analysis chain, in preparation for the start of the experiment.

3.7.1 oaEvent, oaRawEvent and oaUnpack

Providing low level libraries for the ND280 software, oaEvent is heavily dependent on the ROOT framework. OaEvent defines the structure of the ROOT files used in ND280 and provides methods for opening, reading and writing files. A ND280 ROOT file contains a representation of the detector geometry, generated by the ROOT geometry tools, as well as the data. Classes are provided to store different types of data; such as raw data, simulated data or reconstructed data.

Each file stores a series of events. An event does not necessarily refer to a single neutrino interaction but a collection of recorded hits within the detector. This could be a single particle generated with a particle gun Monte Carlo, or a full spill simulation with many neutrino interactions contained within the same event. All hits are stored in ‘hit selections’. There is a hit selection for each subdetector. Various reconstruction algorithms look at these hit selections and attempt to reconstruct what the ‘event’ was. Raw data are stored in a format defined by the DAQ software, MIDAS. As this is not compatible with ROOT it cannot be read with oaEvent. The oaRawEvent package provides a series of

interfaces to read the raw data files with the offline software suite.

In v6r1 there is no completely generalised package to convert raw data files to the ND280 ROOT files defined in oaEvent. A new package, oaUnpack, has been developed for the subdetectors which use the Trip-T electronics, such as the ECals and INGRID, and is being generalised to convert all data to the ND280 format. While the interfaces for oaUnpack are still being finalised, development versions were successfully used in the ECal test beam to analyse the data collected.

3.7.2 Monte Carlo simulation

Detector Simulation

The simulation of ND280 is carried out in two steps, a detector simulation and an electronics simulation. GEANT 4 is used as the framework for the detector simulation package, nd280mc. It provides mechanisms to represent the geometry of complex detector structures with active and passive components. It also models the way a particle interacts with matter as it traverses a detector, not the physics of neutrino interactions.

Simple particle gun Monte Carlo can be started with GEANT 4 alone. For example, a sample of electrons can be simulated with a distribution of energies. To model the position and physics of a neutrino interaction, neutrino generators are used. These provide a Monte Carlo simulation of the different neutrino interactions, and output the kinematic properties of final state particles which can then be passed to GEANT to model in the detector. Currently nd280mc can be used with two neutrino interaction generators, GENIE and NEUT. GENIE has comprehensive support for different detector geometries

through the ROOT geometry libraries. This makes it very adept at calculating the number of neutrino interactions in different materials. In nd280mc, GENIE calculates the cross section for neutrino interactions for all the different types of nuclei present in the detector. It then calculates the number of interactions based on the density of the material and the beam profile.

NEUT, by contrast, was designed as the neutrino generator for the Super-Kamiokande experiment[78]. Therefore, NEUT is very well tuned for calculating neutrino cross sections on water, carbon and iron nuclei, but it is not designed to deal with detectors made of different materials. A substantial update of NEUT is underway to allow it cope with other nuclei. Currently the interactions are just located in the detector based on a density re-weighting, after being calculated on water.

Electronics simulation

The final stage of the Monte Carlo is the electronics simulation, contained within the elecSim package. Two types of distinct electronics are used in the ND280 detector, as discussed previously. Scintillator parts of ND280 (P0D, ECal and SMRD) use Trip-T electronics, while the FGD and TPC use AFTER electronics. The FGD is also a scintillator detector, so the response of the detector itself is more similar to the sub-detectors using Trip-T electronics. Scintillator bars are modelled by parametrising the number of photons generated for a given energy deposit provided by Monte Carlo. Photon transport in optical fibres and scintillator bars is also parametrised as an attenuation length, to provide an estimate of the number of photons incident at the face of the MPPC. The MPPCs are modelled as a grid of pixels, each of which may release charge if a photon is incident upon it. Second order effects such

as cross-talk and after-pulsing are also included in elecSim.

3.7.3 Reconstruction

Reconstruction software takes the calibrated hits and reconstructs the neutrino interaction. In ND280 the reconstruction is made harder than for a collider detector because the position of the event vertex is unknown; indeed, it need not even be in the detector. Past neutrino experiments such as MINOS and Super-Kamiokande used monolithic detectors, which made reconstruction simpler. ND280 contains a number of subdetectors with different detection methods making the global reconstruction challenging. In ND280 a package called RECPACK [79] carries out a global likelihood fit to associate clusters and tracks in separate clusters.

oaRecon

While global reconstruction algorithms are contained within the RECPACK software, oaRecon manages the order in which the data from each subdetector is processed. Each subdetector has a dedicated reconstruction package which can be run via oaRecon. Subdetector reconstruction algorithms return objects inheriting the TReconObject base class. Ultimately a subdetector will return a TReconPID class, containing a likelihood for different particle hypotheses which is used for the global likelihood fit. oaRecon then attempts to connect together clusters from the different subdetectors to reconstruct full events. At the time of writing this was implemented and tested for the TPC and FGD. The ECal is also implemented but untested.

TPC and FGD Reconstruction

The TPC and FGD detectors share a common reconstruction algorithm. It begins with the TPC as it possesses the highest spatial resolution of any part of ND280. Reconstruction looks for hits close in time, and on neighbouring MICROMEGAS pads in the y axis, giving a position in the XY plane. Reconstruction starts here because high momentum leptons travel along the z axis. A cellular automata algorithm [64] clusters segments of tracks that lie close to each other in time as well as in position. The longest chain of connected segments is then made into a track.

Tracker particle identification carried out by the TPC is discussed in detail in Chapter 6. Here it is sufficient to note that the TPC allows an accurate measurement of the energy loss of the particle to be made, as well as its momentum. By measuring the energy loss at a known momentum the species of particle can be identified.

Once tracks in the TPCs have been identified, a Kalman filter is used to extend the tracks into the FGDs. Tracks are propagated into a FGD one hit at a time, updating the position, direction, length and curvature of the track after each step. The Kalman filter is implemented in the RECPACK software. RECPACK is also capable of connecting together tracks that have passed through multiple TPC and FGD modules. The FGDs also have a standalone cellular automata based clustering algorithm. This is because some tracks will travel in or close to the XY plane and so not pass into a TPC.

ECal Reconstruction

Particle Identification in the ECal will be discussed extensively in Chapter 4. This section describes the structure of the package. The `ecalRecon` package runs a series of algorithms in a fixed order, each adding a layer of reconstruction before storing reconstructed clusters. A cluster is a set of hits that the reconstruction treats as being from the same particle. Attached to each cluster is an estimate of the energy and kinematic properties of the particle, as well as an identification of particle type. Before reconstruction begins, the hits from an event are ‘prepared’ by performing basic checks, such as a cut on charge deposit to remove detector noise. Hits passing this cut are clustered based on which ECal module the hit was found in and the time bucket the cluster was found in.

Charge released by the photosensors is the basic quantity measured by the ECal electronics. It is this charge that is digitised by the ADCs. Reconstruction however uses MEUs (MIP Equivalent Unit), a unit proportional to the energy of the particle. The MIP energy is defined as the energy deposited by a muon passing at normal incidence through 1 cm of scintillator. Currently the calibration to convert from a charge deposit to a number of MIPs is carried out in two stages. The first is to convert the measured charge in femtocoulombs to the number of photons on the face of the photosensors. The second reconstructs the position of the hit and uses that position to calculate how many photons were created by the original energy deposit.

A clustering algorithm is applied to associate hits due to the same particle. This step is split into four steps and is described fully in [64]. Basic clustering is the first level of reconstruction, using a nearest neighbour algorithm to separate hits into clusters representing individual particles. The second stage

of clustering, ‘combi clustering’, seeks to combine clusters split by the effects of layer efficiency or a gamma not converting straight away. A final stage of clustering expands the existing clusters by determining whether any of the un-clustered hits are on the periphery of a cluster. When the hits are grouped together, they are placed in separate lists, one for each view of the detector. For example, hits in the DsECal are arranged into two lists; one for hits in the XZ plane and another for hits in the YZ plane. After clustering, a set of clusters exists for each view of a given ECal module. The next step, matching, associates clusters in each view due to the same particle.

Once all the hits due to a common particle have been grouped together, then the energy of the particle can be reconstructed. Particles showering electromagnetically in the ECal, primarily electrons and gammas, currently have their energy measured by the ECal reconstruction software. The energy reconstruction algorithm is trained using Monte Carlo gamma rays passing into the ECal at a range of energies. Gammas are used instead of electrons to reduce the effects of pre-showering in the electronics and metalwork surrounding the ECal. Variables used in the training are the distributions of the mean, sigma and skew of the energy deposit, all of which are fit to a skewed Gaussian at each energy used in the training. A maximum likelihood fit is then used to find the most likely energy of the particle.

The energy of muons is measured by observing how far they travel before being ranged out either in the ECal or the SMRD. Low energy muons, up to about 250 MeV, will stop in the ECal. Higher energy muons will range out in the SMRD. The highest energy muons, of a few GeV or more, can escape the detector completely; although for T2K this is in the high energy tail of the neutrino energy spectrum. The visible energy of muons and other track

like particles passing through the ECal can be measured, providing a very useful measurement for calibrating the absolute energy scale in the detector. Hadronic showers also occur in the ECal, due mostly to pions. Currently there is no hadronic shower fitter in the ECal reconstruction, as hadronic showers are not as well understood as electromagnetic showers. On top of this, a hadronic shower loses a large fraction of its energy in neutrons which the ECal is not designed to detect.

Particle identification (PID) is the final stage of the ECal reconstruction. Development of the PID algorithm is the subject of Chapter 4. The ECal reconstruction uses a mixture of topological features and charged based quantities to decide what type of particle has caused a cluster. Each cluster is fit as both a track and a shower. Track fitting currently fits a track as a straight line using a chi-squared fit. Shower fitting uses PCA (Principal Component Analysis), as discussed in Section 4, to find a direction, centre and opening angle for the shower. Both track and shower fitters find start and ends points for a cluster. As with the clustering, the PID assumes that the starting point of a track or shower is the end on the inside of the detector. Outputs from the reconstruction are used as the discriminating variables in the PID and are discussed at length in the Chapter 4.

Chapter 4

Particle Identification in the ECal

4.1 Introduction

In the T2K beam the muon neutrino flux is approximately 200 times higher than the electron neutrino flux. The performance of the electron neutrino analysis is dependent on very efficient muon rejection. The background acceptance needs to be of order 10^{-3} or less to give a signal to noise ratio better than 5:1. The ECal is used for containing and reconstructing the energy of particles leaving the tracker. By looking at the topology and distribution of energy deposits in a cluster, the ECal identifies the species of particle that created the cluster. The ECal can aid in the separation of muons and pions by looking for pion showers, which is not done by the tracker. This chapter describes the development of a particle identification algorithm for the ECal using Monte Carlo simulation and examines its ability to separate electrons and muons.

4.2 Monte Carlo simulation and Particle event types in the ECal

The ND280 ECal is coarsely grained, which limits the number of different topologies that can be identified. In the case of the barrel and downstream ECals, event topologies can be categorised as electromagnetic shower, hadronic shower or track. Electromagnetic showers come from electron and photon interactions. Protons, charged pions, neutrons and kaons can shower through hadronic processes. Tracks are defined as MIP (Minimally Ionising Particle) particles: muons and non-interacting charged pions. A Monte Carlo simulation of the ECal is used to model events of different species of particle. The design of the simulation and the different types of event predicted are presented below.

4.2.1 Monte Carlo simulation of ECal

The passage of the particles through the ECal is simulated using the nd280mc GEANT 4 simulation. For each bar containing an energy deposit in the detector a hit is created. The hit stores the location and magnitude of the energy deposit. The response of the ECal is simulated by the elecSim package. ElecSim parametrises the number of photons created by an energy deposit in the scintillator - currently 25 photons per MeV. A Birk's law correction is applied and then an attenuation correction is used to estimate the number of photons found at the face of the photosensor, as discussed in Section 3.5.7. A photosensor simulation is also carried out using elecSim. This model is now quite advanced and includes the effects of crosstalk and afterpulsing. The output from elecSim is a prediction of the number of photoelectrons seen at the surface of each photosensor.

A particle gun was used to fire particles into the detector at 20° intervals from 0° to 80° with respect to the face of the ECal. The energy distributions of electron, muon and negative pion samples were based on the predicted spectra from neutrino interactions in ND280. These distributions are shown in Figure 4.1. A sample of photon events was produced by taking the energy spectrum of photons expected from neutral pion decays. This is also shown in Figure 4.1. A sample of 30,000 events were produced for each particle at each energy. Muons were used to create a track sample, electrons and photons were used to create the EM shower sample and showering negative pions were used to create the hadronic shower sample.




Figure 4.1: Showing the energy spectra used to produce Monte Carlo samples to design and train PID variables.

4.2.2 Tracks

As a relativistic charged particle passes through the ECal, it will lose some of its energy via the Coulomb interaction. For particle energies below 100 GeV, the familiar Bethe-Bloch equation [19]

$$-\left\langle \frac{dE}{dx} \right\rangle = \frac{4N_A Z \pi r_e^2 m_e c^2 z^2}{A \beta^2} \left[\ln \frac{2m_e c^2 \gamma^2 \beta^2 T_{max}}{I^2} - \beta^2 - \frac{\delta}{2} \right], \quad (4.1)$$

is used to calculate the mean energy loss of a charged particle. In Equation (4.1) E is the kinetic energy of the particle, T_{max} is the largest amount of energy that can be given to an atomic electron, r_e is the classical electron radius, Z is the atomic number of the media being traversed, β is the velocity of the particle, δ is a correction dependent on the density of the material, c is the speed of light in a vacuum, m_e is the mass of the electron, z is the particle charge, I is the mean excitation potential of the target and N_A is the electron density of the target. A muon with $\beta\gamma \simeq 2$ has the smallest ionising power. Such a muon is called a MIP. This can be seen in Figure 4.2. A particle much lighter than a muon will have a higher value of $\beta\gamma$ for the same energy, and so will lose its energy rapidly through the emission of bremsstrahlung radiation.

At energies of a few hundred MeV, muons and charged pions in the ECal lose only 1-2 MeV cm² g⁻¹; so only those with energies $\lesssim 250$ MeV might be expected to stop. MIP-like particles are typified by a very narrow track, with a charge distribution centred around the MIP charge. Figure 4.3 shows a narrow MIP-like track with a relatively small spread in the charge deposits. Variables effective at separating MIPs from showers look for the narrow, long shape, as well as the uniform charge deposition. Low energy, stopping MIPs, with a short length are the most likely to be misidentified as a shower because they can be confused with a narrow shower.




Figure 4.2: Stopping power for muons for the momentum range 0.1 MeV/c up to 100 TeV/c from [19]. The region from 0.1 GeV/c to 1 GeV/c is of greatest interest to T2K and is around the MIP region for muons.




Figure 4.3: A typical MIP leaves a long and narrow track as shown by this simulated event. The charge deposited in each layer is similar as can be seen from the colour scale.

4.2.3 Electromagnetic Showers

While muons lose most of their energy through ionisation, electrons lose energy primarily through the emission of bremsstrahlung photons as shown in Figure 4.2. Radiative processes are suppressed by a factor of $(m_\mu/m_e)^2 \approx 40000$, so bremsstrahlung is not a significant process for muons or charged pions until the TeV scale. As electrons are 200 times lighter, radiative effects dominate above a few tens of MeV, depending on the density of the medium.

Photons will interact with a medium in any one of four ways: Compton scattering, Rayleigh scattering, pair production and the photoelectric effect. For photons with energies at the MeV scale or higher, photoelectric effects quickly become negligible. In the range of hundreds of KeV to a few MeV, Compton scattering is the dominant effect. Rayleigh scattering, where the gamma is deflected by a nucleus with no loss of energy, is also significant at low energies. At high energies, greater than 10 MeV, pair production becomes dominant. The positron and electron created in pair production are then able to emit bremsstrahlung photons which can themselves pair produce. This process, of pair production to bremsstrahlung to pair production, initiates an electromagnetic shower. A shower from a high energy photon/electron will grow in size until the energy of the shower photons drops below the pair production threshold and the shower dies away. As a shower develops it will become wider, creating a cone shaped shower envelope with a characteristic transverse size described by the Molière radius, ρ_M [80]. An average of 90% of the energy from a shower will be deposited in a cylinder of radius ρ_M around the axis of the shower, where;

$$\rho_M = 21.2(\text{MeV}) \frac{X_0}{\epsilon_c}. \quad (4.2)$$

The critical energy, ϵ_c , is the energy where contributions from ionising and

radiative effects are equal. This parameter depends on the material, but is typically tens of MeV. The radiation length, X_0 , describes the length of an electromagnetic cascade and is parametrised in [19],

$$X_0 = \frac{716.4 \text{ g cm}^{-2} A}{Z(Z+1) \ln(287/\sqrt{Z})}. \quad (4.3)$$

For the materials in the ECal, carbon has a X_0 of 188 mm, an ϵ_c of 83 MeV and a ρ_M of 48 mm while lead has a X_0 of 5.6 mm, an ϵ_c of 7.4 MeV and a ρ_M of 16 mm [80, 19]. In total, the ECal has approximately three radiation lengths worth of lead and just under two radiation lengths of carbon.

Electromagnetic showers in the ECal are characterised by their width and charge distribution. Hits due to the primary particle are expected to be much larger than those created by secondary particles close to the threshold for a shower. Good resolution of the typical cone shape predicted for an electromagnetic shower is not expected due to the coarse granularity of the detector. Clusters are, however, expected to get wider as the shower propagates. Variables useful for separating electromagnetic showers will look for clusters with a large width compared to their length. Variables looking for a broad charge distribution centred in the inner part of the detector are also useful. Energy independence is a particular challenge when identifying showers as their charge distribution and spatial extent are strong functions of the energy of the particle. Figure 4.4 shows a typical medium energy EM shower of order 1GeV.

4.2.4 Hadronic Showers

Hadronic showers are vastly more complex than electromagnetic showers where most types of interaction have analytical descriptions. A neutral hadron can

pass through a medium leaving no trace until it collides with a nucleus. A charged hadron may travel through a medium causing ionisation via the same mechanisms as an electromagnetic shower, until it hits a nucleus and interacts hadronically. Neutrons released from the nucleus can leave the detector completely undetected, resulting in a component of missing energy. Particles such as neutral pions, commonly created in hadronic showers, decay rapidly into electrons or photons, which in turn initiate an electromagnetic shower at some angle to the track of the original hadron [80]. In T2K, most particles will have energies in the range of hundreds of MeV to a few GeV, which is not significantly above the rest mass of the π^0 . Hadronic showers in T2K therefore will be comparatively small as there is not enough energy to produce large cascades of hadrons. Neutral pions will be produced almost at rest and so will decay isotropically.

For analyses of particles from the tracker the most significant source of hadronic showers in the ECal comes from charged pions. As they have a very similar mass to muons, charged pions are very hard to separate from muons ‘in flight’. An important task for the PID in the ECal is to look for charged pions converting in the detector. As pions pass through the detector as a MIP before converting, a strong signature for a pion is a ‘lollipop’ event as shown in Figure 4.5 (right). A ‘lollipop’ particle is track like at first, before converting in the back half of the ECal. Pions converting near the front of the detector can only be separated from electrons and gammas by looking for the more spherical topology of a hadronic shower. This can be seen in Figure 4.5 on the left.

Effective variables for separating hadronic showers from electromagnetic showers will look at where the charge is deposited. In many hadronic events this will be towards the outside as opposed to the inside of the detector for electro-




Figure 4.4: Showing a medium energy electron shower with track information (left) and without (right).




Figure 4.5: Showing two types of hadronic shower. The shower on the left is converting soon after entering the detector like an EM shower. The shower on the right traverses much of the detector as a MIP before converting. These events are called 'lollipop' events as the track and shower pattern bears a resemblance to a lollipop.

magnetic showers. Hadronic showers close to the inner part of the detector are hard to separate from electromagnetic showers. However, from Monte Carlo studies, the detector appears to have some sensitivity to the different shower topologies.

4.3 Identification Techniques

To classify a pattern as belonging to a given class, a set of feature variables is designed. The distribution of the feature variables should be different for the classes being separated. A number of techniques exist to use the feature variables to identify which class a given pattern belongs to. The simplest method to combine the information from a number of variables is to use cuts. A cut is formed by taking a point in each of the feature variables. Above this point one hypothesis is preferred, below it another is selected. As this technique can produce only straight line cuts in the feature variable space, it cannot take into account any correlation between variables; nor is it capable of producing a non-linear cut through the feature space. The diagram on the left side of Figure 4.6 shows two classes of event being separated with two feature variables which are perfectly separable with a linear cut. The variable on the y-axis has no discriminating power, but a linear cut on the variable along the x-axis can perfectly separate the two samples. The diagram on the right of Figure 4 shows another two classes of event, again separated with two feature variables. Neither variable can discriminate between the two classes with a linear cut; there is no way to draw a straight line through the plot to separate the two classes. Multi-variate techniques such as likelihoods or non-linear methods can separate the two samples. Such techniques are the subject of this section.

4.3.1 Likelihood

The likelihood method uses Bayesian inference to determine the most probable hypothesis from a set of hypotheses. One can obtain the probability of a certain hypothesis being true, given a variable value, by rearranging Bayes' theorem [81]

$$P(H_k | x) = \frac{P(x | H_k) P(H_k)}{P(x)}, \quad (4.4)$$

where H_k is the hypothesis k and x is the value of a given variable. The value $P(H_k)$ is known as the prior value. This is the probability that a hypothesis is true without any extra information. In the ND280 detector, one might naively expect this to be the fractions of different neutrino species traversing the detector; although kaon and pion production is not independent of neutrino flavour. The $P(x | H_k)$ term is the probability that the variable x has a given value, given a certain particle hypothesis. These values are taken from appropriate Probability Density Functions (PDFs). The denominator, $P(x)$, is the probability that the variable x has that value and acts as a normalising factor. This guarantees that the total posterior probability is equal to one. This unconditional probability density can be expressed as,

$$P(x) = \sum_{k=1}^n P(x | H_k) P(H_k). \quad (4.5)$$

The likelihood method provides the best possible discrimination between two variables. However, this assumes that the form of the PDFs is known to arbitrary precision. In practise, this is rarely the case. The distributions are usually made up from histograms which are limited not only by the binning of the histogram but also by knowledge of the variable distribution. Particles can originate from a variety of angles and positions; and the shapes of the distributions may depend on location within the detector. The most limiting

factor, however, is the number of events with which the histograms are filled. The number of entries required to fill a histogram increases with the power of the number of dimensions in the histogram. This makes filling a histogram with greater than three dimensions impractical as the time taken to fill it goes up by an order of magnitude for each dimension added. The PID algorithms considered for the ECal contain at least five variables. Populating a five dimensional histogram is not practical. Other techniques, such as neural networks, while not able to achieve ideal discrimination, are often able to provide better discrimination under more practical training conditions.

Errors due to misclassification in likelihoods

A simple way to estimate the error from the likelihood method is to find the probability of mis-identification. More formally this is expressed as $P(x \in H_1, C_2)$ where x is found to belong to hypothesis one (H_1) when it is actually a member of class two (C_2) [81]. In general the result of a likelihood will be used such that the probability of misclassification is minimised over all hypotheses. For a single feature variable separating two hypotheses, this would be the point where the two distributions intersected. In some situations it may be necessary to reject one hypothesis particularly strongly. A good example of this is a medical diagnosis, where a false positive is often far less harmful than a false negative. This can be achieved by demanding that the likelihood for one hypothesis is particularly low. This is done at the expense of rejecting a greater number of signal events.

4.3.2 Artificial Neural Networks

Single Layered Perceptron [81]

An artificial neural network models the human brain by simulating the behaviour of neurons and synapses. A neuron will fire if the sum of the inputs from connecting neurons reaches some threshold level. In software this is represented by nodes connected by a series of weights from each input. A typical network with a single hidden layer is shown in Figure 4.7.

The simplest artificial neural network has no hidden layers, as shown in Figure 4.7, and consists of a series of inputs connected to outputs by a series of weights. This case is described by the following expression

$$y = F \left(\sum_{i=0}^n (w_i b_i(x)) \right), \quad (4.6)$$

where w is the weight connecting node y to node b_i , b_i is the value of node i and F is the activation function. In Figure 4.7 each line has a weight (w in Equation 4.6) which is multiplied by the value of the node it connects to (b_i in Equation 4.6). In the expression above, inputs one to n represent the input variables. The term corresponding to $n = 0$ in Equation 4.6 is the bias or threshold term. It can be thought of as a weight associated with a constant input of one. Its exact interpretation depends on the training of the network [81]. The above expression can be written in vector form

$$y = F \left(\mathbf{w}^T \mathbf{x} + x_0 \right). \quad (4.7)$$

This describes a straight line in the feature space of the variables. The single layered perceptron, is therefore a type of linear discriminant with a threshold applied. To construct a neural network capable of creating a non-linear boundary in the space of the feature variables, a multi-layered neural network




Figure 4.6: Demonstration of linear (left) and non-linear (right) separability. A linear boundary can only separate two classes of event if a straight line can be drawn between the classes.

Input Nodes Hidden Layer Output Node(s)




Figure 4.7: A neural network connects an input data vector to an output via hidden nodes connected by weights. Indices are the same as those in Equation (4.6).




Figure 4.8: Showing a tanh activation function. In general, neural networks use sigmoidal activation functions such as $\tanh(x)$.

must be used.

The function, F , is called the activation function. This simulates the firing of a neuron in the brain. As a neuron either fires or not, the activation function should be close to a step function. The activation function is typically linear in the region $|x| < 1$ and saturates quickly to -1 for $x < -1$ and 1 for $x > 1$ [81]. Figure 4.8 shows a typical sigmoid activation function, $\tanh(x)$ in this case.

Multi-Layered Perceptron [81]

If the neural network has one or more hidden layers, then it is no longer restricted to having a linear decision boundary. A Multi Layered Perceptron (MLP) is a type of feed forward network. This means that there are no loops in the network. The first hidden layer can connect to the second, but not the other way around. This limits the number of possible connections, but guarantees that the output is an explicit function of the input variables. The

general form of a MLP with a single hidden layer is

$$y = F \left(\sum_{i=0}^n (w_i G \left(\sum_{j=0}^m (w_j b_j(x)) \right)) \right). \quad (4.8)$$

All terms are the same as in Equation 4.6, except this time we note that the output for one node is nested as the input for the next node, i.e. the output of the hidden node is an input to the output node, after passing through the activation function, G . It has been shown that by adding a third layer of weights, this type of network can describe any decision boundary; although there is no algorithm for finding such a solution [82][81]. Figure 4.6 demonstrates graphically the difference between the linear discriminant created by a single layer perceptron and the non-linear decision boundary described by a multi-layered perceptron. In effect, the MLP can create the curved boundaries that are required to separate the two classes on the right of Figure 4.6.

Neural Network Training

While the output of a neural network is deterministic and straightforward to analyse, for small feed-forward networks at least, their training is often opaque. Most training techniques involve minimising an error function to find the values for the optimum weights. For a simple perceptron with no hidden layers a sum of squares error function is suitable, as it is simple and works for feed-forward networks. A sum of squares error function has the form

$$E(w) = 1/2 \left(\sum_{n=1}^N \sum_{k=1}^c [y_k(X^n, W) - T_n^k]^2 \right) \quad (4.9)$$

where N is the number of training samples and c is the number of outputs. The vector X is the vector of input variables and T represents the target output values. W is the vector of weights to be calculated. If an activation function is used, as will generally be the case in a neural network, then the sum of squares

error does not have a closed form. This means an exact solution cannot be found by a computer using matrix methods.

The gradient descent method can be applied to an activation function if the function has a differentiable form. The gradient descent method begins with approximate values for the weights vectors, and then minimises the error function by moving in the direction of maximum decrease of the error function. This algorithm will converge towards the optimum values for the weights, where

$$w_{jk}^{n+1} = w_{jk}^n - \eta \frac{\partial E}{\partial w_{kj}}|_{w^n}, \quad (4.10)$$

E is the error as shown in Equation 4.9. The learning parameter, η , determines the rate at which the gradient descent method converges; too large and it can oscillate, too small and convergence will be unnecessarily slow.

Interpreting Neural Network Output

When using a comparatively opaque technique such as a neural network, it is important to be able to interpret what the network is modelling. The neural network used in the ECal PID has been designed to model the posterior probability that a given input vector belongs to a certain particle type. The addition of an activation function does not make the neural network a non-linear discriminant, as the activation function is still monotonic with respect to the node input. The non-linearity comes from the presence of hidden layers. A logistic activation function (shown in (4.12)) can be shown to allow the output to be interpreted as a posterior probability [81]. Assuming a single discriminator is represented by a Gaussian distribution such that

$$P(x|C_k) = \frac{1}{|\Sigma|^2 2\pi} \exp \left(-0.5(x - \mu_k)^T \Sigma^{-1} (x - \mu_k) \right), \quad (4.11)$$

where $P(x|C_k)$ is the probability of finding the data given membership of class k , μ_k is the mean and Σ is the covariance matrix. Also, assuming a logistic function is used as the activation function with the form [81],

$$F(a) = \frac{1}{1 + \exp(-a)}, \quad (4.12)$$

where a is the output from a given node. The expression for the posterior probability from Bayes' theorem is shown in (4.4). If the output of a node models the posterior probability then we set $F(a) = P(x|C_k)$, which gives [81]

$$a = \ln \frac{P(x|C_1)P(C_1)}{P(x|C_2)P(C_2)}. \quad (4.13)$$

Substituting equation (4.11) into the expression for a general linear discriminant, $a = w^T x + w_0$, the following expressions for w and w_0 are reached [81]

$$w = \Sigma^{-1}(\mu_1 - \mu_2), \quad (4.14)$$

and

$$w_0 = -1/2\mu_1^T \Sigma^{-1} \mu_1 + 1/2\mu_2^T \Sigma^{-1} \mu_2 + \ln \frac{P(C_1)}{P(C_2)} \quad (4.15)$$

This shows that in principle a non-linear discriminant with a sigmoidal activation function can model a posterior probability. The expressions for the weights above shows how they may be interpreted for Gaussian distributions.

To interpret the output as a posterior probability, then, the training must also be able to model the posterior probability. Training algorithms are discussed in detail in [81]. Here it is sufficient to consider the training as the minimisation of an error function. The sum of squares error function used to train the ECal PID was described above. It is shown in [83] that the sum of squares error function can be interpreted as a posterior probability, given that the outputs sum to unity.

4.4 Reconstruction

4.4.1 Track Reconstruction

Track reconstruction in the ECal takes all the hits in a given layer and creates a single point. Each hit position is referred to as a track node. A track node has a position and an associated direction and curvature. The start point for a track is taken to be on the inside of the ECal. This assumes that all particles come from the tracker. To find the direction and refine the position, MINUIT [84] is used to fit a node and its four nearest neighbours to a straight line. The node closest to the tracker section of the ECal is taken as the starting point for the fit. The curvature in the ECal is not defined, as the granularity is low and the high density materials make the effects of scattering much larger than any curvature. Ideally, the track fitting could be done with an algorithm such as a Kalman filter, which can include the effects of scattering. Again, the low granularity of the ECal means that a more complex algorithm may not enhance the performance significantly.

4.4.2 Shower Reconstruction

Shower reconstruction is a more poorly defined task than track reconstruction. Whereas track fitting involves fitting a line to some functional form, often a helix or a straight line, there is no single way to fit a shower. In the ECal, PCA (Principal Component Analysis), explained in Section 4.5.1, is used to create an ellipse representing the shape of the shower. The direction assigned to the shower is given by the direction of the principal component in the detector coordinate system. An opening angle is found by calculating the angle between the two largest eigenvalues. The starting point of the shower is taken to be

the position of the hit closest to the tracker section of the detector.

4.4.3 Angle Reconstruction

The angle of the particle is reconstructed by measuring the angle between the principle component from the shower reconstruction, and a line normal to the face of the detector. Many of the variables in the PID are a function of the angle of the particle with respect to the plane of the detector. To account for this dependence, the PID has a different training depending on the angle of incidence of the particle. The PID is trained at angles of 0° , 20° , 40° , 60° and 80° with respect to the face of the ECal module.

4.5 ECal Particle Identification

4.5.1 Particle Identification Variables

The positions and magnitudes of hits in the scintillator bars can be used to determine the species of cluster in the detector. As discussed previously, the ECal defined three particle classes, electromagnetic showers, hadronic showers and MIPs. The variables used to discriminate between different particle classes are discussed below. When selecting a variable, angular dependence and energy dependence are considered as well as discriminating power.

AMR

The AMR (Axis Max Ratio) uses PCA to calculate the principal direction of the track and to measure the track's ellipticity. PCA transforms data to a basis that has the largest variance in the data as the first coordinate. In the

case of the AMR variable the coordinates transformed are the positions of the hits in the detector. This means that the principal component lies parallel to the direction of the particle, as shown in Figure 4.9.




Figure 4.9: The AMR is calculated by finding the average ratio of the principal component to each minor component.

The AMR is the ratio of the principal eigenvalue to the orthogonal eigenvalue, weighted by the charge of each hit. A highly elliptical track is more likely to be a muon than an electromagnetic shower or hadronic shower. The AMR is also slightly sensitive to charged pions. The hard scatters in a hadronic cascade make them more spherical than a cone-shaped EM shower causing a lower AMR, as can be seen in Figure 4.10.




Figure 4.10: AMR variable distributions for tracks and showers (left) and EM and hadronic showers (right).

The AMR is a three dimensional algorithm. It splits the hits into two groups, one for each orientation of scintillator bar. A two dimensional AMR is then calculated for each of the two views. The final step is to find the mean of the two values. The AMR is limited to 300 to prevent anomalous results due to small denominators. In Figure 4.10 this can be seen to cause peaks in the distributions at a value of 300, due to very narrow tracks. Similar peaks are seen close to 150 where the AMR in one view has been set to 300 and the AMR in the other view was much less. As the AMR transforms the basis of the coordinate system to the direction of the particle, this variable is robust against the angle of the particle with respect to the ECal. Looking at

the difference between the two small components can be useful for separating electromagnetic and hadronic showers, and is discussed later. The AMR is used to separate tracks and showers.




Figure 4.11: AMR variable as a function of angle for shower hypothesis (left) and track hypothesis (right).

Maximum Charge Ratio

The maximum charge ratio is the ratio of the charge deposited in the layer with the highest deposited charge to the charge in the layer with the lowest deposited charge. The maximum charge ratio was intended to address the vulnerabilities that taking the ratio of largest and smallest charges had to varying angle and anomalous charges, by considering the layers with the highest and lowest (non-zero) charge deposits. By considering a whole layer it was expected that the effects of anomalous values would be reduced. The maximum ratio is not dependent on the location of the shower in the detector.

This variable is good for separating short tracks, or long showers, as it is very independent of the topology. Figure 4.12 shows that, as with the maximum charge ratio, it is very effective for discriminating tracks from showers. It

also has some power to discriminate between electromagnetic and hadronic showers. It can be seen from Figure 4.13 that the maximum charge ratio only has a dependence on the angle of incidence of the particle at very high angles. This is an advantage common to many of the charge based variables discussed below. This variable is used in the discrimination of tracks and showers.




Figure 4.12: Maximum Charge Ratio variable distributions for tracks and showers (left) and EM and hadronic showers (right).




Figure 4.13: Maximum Charge Ratio variable as a function of angle for shower hypothesis (left) and track hypothesis (right).

Electromagnetic Energy Fit Likelihood

Energy reconstruction in the ECal uses a maximum likelihood fit to determine the most likely energy of a cluster, assuming that it was an electromagnetic shower. The likelihood from the fit can be used as a PID feature variable. An electromagnetic shower will produce a good fit to an electromagnetic shower energy distribution, while hadronic showers and MIPs do not fit as well. As the likelihood is often a very small number, the variable used for discrimination purposes is, $-\log(\text{EMLikelihood})$, often a more manageable quantity.




Figure 4.14: EM fit likelihood variable distributions for tracks and showers (left) and EM and hadronic showers (right).

Again, the electromagnetic energy fit likelihood is very effective in track and shower discrimination, as shown by Figure 4.14. The peak at about 46 is caused by software cuts to protect against anomalously low likelihoods from failed fits. The distribution of hadronic shower events has a long tail, implying that they look less like EM showers in general, although fewer hadronic events fail to fit at all. As is expected for charge based variables the EM likelihood of showers is quite resilient against a change of angle, as demonstrated by Figure 4.15, although not as resilient as the AMR. The EM likelihood of tracks show

more variation with angle. The EM likelihood is used to separate tracks and showers.




Figure 4.15: EM fit likelihood variable as a function of angle for shower hypothesis (left) and track hypothesis (right).

Charge Skew

The charge skew variable looks at the third standardised moment of the charge distribution as measured by the energy reconstruction. It was expected that the skew would be larger for showers than for tracks.

Figure 4.16 shows that in practise the charge skew does not separate tracks and showers very strongly. However, it does separate hadronic and electromagnetic showers comparatively effectively, when compared to many of the other variables considered. Like the EM Likelihood variable, QSkew has no dependence on geometry. It is also not a strong function of angle, particularly for showers, as shown by Figure 4.17. The QSkew is currently used in the discrimination of electromagnetic and hadronic showers.




Figure 4.16: Charge Skew variable distributions for tracks and showers (left) and EM and hadronic showers (right).




Figure 4.17: Charge Skew variable as a function of angle for shower hypothesis (left) and track hypothesis (right).

Shower Angle

This variable uses the reconstructed opening angle of the shower as calculated by the reconstruction algorithm. The angle of the shower is calculated by using PCA to determine the direction of the shower. The shower angle is then defined as an angle between the principal component and the largest of the minor components.




Figure 4.18: Shower angle variable distributions for tracks and showers (left) and EM and hadronic showers (right).

The shower angle is very effective for separating MIPs and showers as shown in Figure 4.18. It is less effective at separating hadronic showers and electromagnetic showers. There is a long tail in the hadronic distribution which gives the shower angle some discriminating power. Shower angle is used to discriminate between tracks and showers. Figure 4.19 shows that the shower angle is resilient against changes in angle at low angles, although it is not as resilient as charged based variables.

Shower Width

Using a similar principle to the shower asymmetry, the shower width seeks to exploit the more spherical structure of a hadronic shower. The shower width calculates how wide the shower is. This variable would be effective in events where a particle scatters at right angles to the direction of the shower. In Figure 4.20 the variable is very effective in the separation of tracks and showers. It also has some power to separate electromagnetic and hadronic showers. The hadronic shower distribution is broader than the electromagnetic distribution. The tail at high widths is due to larger showers releasing protons and other pions. In Figure 4.21 the shower width is shown to be quite independent of angle except at 80° in the shower case. This is due to the use of PCA to calculate the width of the shower, instead of simply using the width measured in scintillator bars. Shower width is used in the discrimination of electromagnetic and hadronic showers.

Shower Asymmetry

The asymmetry of a shower is similar in principle to the AMR. Whereas the AMR measures how long a shower is compared to its width, shower asymmetry compares the width of a shower in two orthogonal directions perpendicular to the shower direction. The shower asymmetry is defined as the smaller component divided by the larger component to restrict the variable to between zero and one. An electromagnetic shower develops as a cone, so would be expected to be comparatively circular in the plane perpendicular to the principal component. A hadronic shower is expected to have scattered mesons and nucleons giving the shower a more chaotic shape. The more circular the cluster is, the smaller the asymmetry, a value of one indicates that both smaller principal




Figure 4.19: Shower Angle variable as a function of angle for shower hypothesis (left) and track hypothesis (right).




Figure 4.20: Shower Width variable distributions for tracks and showers (left) and EM and hadronic showers (right).

components are equal.

The very small numbers often generated by tracks mean that the output from narrow tracks causes a sharp peak at close to zero. The shower asymmetry does, however, have discriminating power when separating electromagnetic and hadronic showers, as hadronic showers are often more asymmetric than electromagnetic showers, as shown by Figure 4.22. This variable does have some angular dependence, particularly at very high angles, as shown in Figure 4.23. The shower asymmetry is currently used to discriminate between electromagnetic and hadronic showers.

Mean Shower Position

A type of hadronic event is the ‘lollipop’ event as shown in Figure 4.5. This is characterised by the particle depositing most of its energy in the outer half of the ECal. The mean shower position calculates the charge weighted mean hit position along the principal axis of the shower, found using PCA. This position is measured relative to the starting point for the shower. For a ‘lollipop’ type event this charge weighted mean will be strongly shifted towards the outer edge of the detector, making the mean shower position a good variable for finding certain hadronic showers. Mean shower position is effective for both track versus shower and electromagnetic versus hadronic shower discrimination. For MIP particles there is a peak at half the track length, as approximately the same amount of charge is deposited along the track. For both electromagnetic and hadronic showers the mean shower position depends on the size of the shower, but generally the mean position is further along the shower for a hadronic shower as it can propagate as a MIP for some distance before showering, as shown in Figure 4.24. Mean position has a slightly unusual




Figure 4.21: Shower Width variable as a function of angle for shower hypothesis (left) and track hypothesis (right).




Figure 4.22: Shower asymmetry variable distributions for tracks and showers (left) and EM and hadronic showers (right).




Figure 4.23: Shower asymmetry variable as a function of angle for shower hypothesis (left) and track hypothesis (right).




Figure 4.24: Mean shower position variable distributions for tracks and showers (left) and EM and hadronic showers (right).

response to angle. At very high angles the path length decreases because particles leave the side of the detector if they don't convert, as shown by Figure 4.25. Mean shower position is used to separate electromagnetic and hadronic showers.




Figure 4.25: Mean shower position variable as a function of angle for shower hypothesis (left) and track hypothesis (right).

4.6 Particle Identification Technique

4.6.1 PID algorithm description

The particle identification algorithm uses an artificial neural network to identify the species of a particle from the position and charge information of hits in a cluster. As a neural network is most effective when applied to a single task, two neural networks are used. The first determines whether the cluster is more like a track or a shower. The second determines whether it is more like an electromagnetic shower or a hadronic shower. If the neural network output can be interpreted as a probability, then it can be easily applied to a global reconstruction algorithm that multiplies likelihoods together. The running of

the neural network is very fast once the variables have been calculated. Training is, by comparison, slow and also critical to producing a stable network architecture.

Two likelihood techniques are also implemented. The first is a product of one dimensional likelihoods. The second uses a PDE range search to estimate the population of a PDF with the same dimensionality as the neural network input. The PDE range search method is explained fully in [85] and [86]. The product of one dimensional likelihoods is simple to interpret and computationally fast, but is far from optimal as it has no knowledge of the correlation of the input variables. The PDE range search method, on the other hand, is capable of finding non-linear correlations in the same way as a neural network, but is less desirable computationally. The PDE method trains quickly but is slow when classifying events.

Properties of the Monte Carlo simulation used to generate the training data were discussed previously in Section 4.2. Effective discriminating variables were described in Section 4.5.1. Table 4.1 summarises which variables were used for separating tracks and showers and which were used for separating electromagnetic and hadronic showers. To generate different weights files, Monte

Variable	Track vs. Shower	EM vs. Hadronic
AMR	✓	×
Max Ratio	✓	✓
EM Likelihood	✓	×
Shower Angle	✓	×
Shower Width	✓	✓
Asymmetry	×	✓
Mean Position	×	✓
QSkew	×	✓

Table 4.1: List of variables used in the PID algorithm and which discriminators they are used in.

Carlo data was generated using a particle gun orientated at 0° , 20° , 40° , 60° and 80° to the face of the ECal using energy distributions discussed in Section 4.2. The networks were then trained using each data set and the reconstruction software chose which weights file to open based on the reconstructed angle of the particle.

4.6.2 Network Training

Neural networks used for the PID were trained and read using the TMVA (Toolkit for Multivariate Analysis) package in ROOT [86]. The network architecture had $n+1$ hidden nodes in a single hidden layer, where n is the number of input variables. Each hidden node used a sigmoidal activation function. There was a single output node guaranteeing an ‘either or’ output with a linear activation function. The training of the network used a sum of squares error function. This configuration was intended to model a posterior probability. This property of the network was explored using two methods. The results are discussed in Section 4.7.1. The first method was to train the neural network with a single variable and compare this to a simple one dimensional likelihood. The second method was to use a range searching probability density estimation algorithm to create a PDF of the same dimensionality as the neural network input.

4.6.3 Network Optimisation

Node Optimisation

The MLP had many parameters that could be varied to optimise the performance. Increasing the number of hidden layers or the number of nodes

increased the number of degrees of freedom available to the network. The optimum number of nodes for a given network has been studied by Kolmogorov [81] assuming a training algorithm capable of achieving the best possible separation. In practise, with a real training algorithm, the optimum separation was unlikely to be found, but the performance was still expected to increase until the number of nodes were greater than the number of input nodes.




Figure 4.26: Training and testing error estimators showing the same behaviour over training epochs indicates no over training occurs. The x-axis shows signal inefficiency (1-efficiency).

To verify this was true for the network described, the ability of the network to remove muons from an electromagnetic shower sample was calculated for different numbers of nodes in the hidden layer. The network was trained using a sample of 5000 muons and 5000 electron and photon events, using all five variables. From Figure 4.26 it could be seen that, for signal efficiencies lower than 90%, the background contamination was significantly lower with six or seven hidden nodes, where five input nodes were used. The signal contamina-

tion is the fraction of the sample that consists of background events. At very high efficiencies there was much less difference in performance, although the six node hidden layer still performed slightly better.

Training Sample Optimisation

Overtraining in a neural network occurs when the training algorithm trains on very subtle features and statistical fluctuations of the training sample. This results in apparently very good performance when being tested on data very similar to the training sample, but poorer performance in general. The opposite of overtraining is where too few training samples are used, and the number of degrees of freedom of the network is of similar size to the training sample. This leads to similar symptoms to overtraining because the neural network has sufficient degrees of freedom to almost completely describe the training sample. In between these two extremes the optimum number of training events can be found.

Figure 4.27 shows the contamination of the electron sample as a function of signal efficiency for a single hidden layer network. In the plot, inefficiency (one minus the signal efficiency) is shown so that the most pertinent region of the plot can be highlighted with log-log axes. On the right hand side of the plot, at an efficiency of about 80%, the 5000 and 7000 event training samples have a contamination of around 1%, increasing to 2% and 3% for the 500 and 50 event training samples respectively. At higher efficiencies of around 95% all training samples exhibit similar performance before diverging again at very high efficiencies with the 5000 and 7000 event training samples again having similar and superior performance. No drop in performance is observed with the 7000 event sample indicating that overtraining is not occurring.

4.7 Neural Network Performance

The ECal PID has been tested with Monte Carlo, statistically independent of the training sample. The set of discriminating variables used for track versus shower discrimination is AMR, Maximum Charge Ratio, EM Likelihood, Charge Ratio and Shower Angle. For the discrimination of electromagnetic and hadronic showers, Asymmetry, Charge Ratio, Mean Position, Qskew and ShowerWidth variables are used.

4.7.1 Neural Network Validation

Consideration has been given to interpreting the output of the neural network as a posterior probability. The architecture and training of the network were chosen such that the neural network would model a posterior probability. This was tested with two different methods. Firstly, a separate network was trained using just the AMR variable. The output was compared to the posterior probability found using the likelihood method with a single variable, the AMR. Secondly, a probability density estimator was used to create a PDF of the same dimensionality as the neural network input. This was also done using the TMVA libraries. The results of these tests are shown in Figure 4.28. The neural network outputs were very similar to the likelihood outputs, although they were not perfectly matched. The neural network was minimising an error function, so was not bounded between zero and one, producing a tail below zero and above one.




Figure 4.27: Contamination is shown as a function of efficiency for a variety of different size training samples. The performance of the network can be seen to improve with increasing training sample size up to a few thousand events. No evidence of overtraining is seen.




Figure 4.28: Showing a comparison between neural network output and likelihood output. The figure on the left compares a single input neural network with a 1D likelihood. The figure on the right shows a comparison between a multidimensional PDF and a neural network of the same dimensionality.

4.7.2 Predicted Efficiency

The performance of the PID was examined for particles with a momentum above 400 MeV. The PID efficiency increased with energy because clusters were larger, making the topology easier to identify. Raising the energy threshold of a ν_e study has the effect of increasing the purity because PID efficiencies increase, but decreasing the ν_e efficiency as signal events are removed. The plot in Figure 4.29 shows the output of the EM shower vs. hadronic shower neural network, against the track vs. shower neural network output, for a sample of muon events. The sample is strongly populated in the track region of the graph. A small number of events have been identified as EM showers. These are stopping muons, occurring mostly at low energies as can be seen on the right hand side of Figure 4.29.




Figure 4.29: The MLP output for muon sample is shown, left. On the right is the energy as a function of track vs. shower output. The effect of stopping muons is shown to create a small cluster of incorrectly identified events.

Figure 4.30 shows how well electron and pion events were separated by the neural network. The pion sample is populated more towards the hadronic side, while the electron sample is largely in the EM shower region. The dis-

crimination is not as clear as the track vs. shower discrimination, since this is a far more challenging discrimination to make.




Figure 4.30: Comparing the MLP outputs for an on-axis pion and electron sample using energy distributions described in Section 4.2. Hadronic showers and electromagnetic showers are more similar than tracks and showers, making discrimination more difficult.

In the electron neutrino analysis the most important discrimination is removing muons from the electron sample. As the muon neutrino flux is approximately 200 times that of the electron neutrino flux, it is important to achieve an overall efficiency of about 1000:1 to give a signal to noise ratio of better than 1:1.

The predicted performance of the ECal PID in the particle energy regions $E > 400$ MeV and $E > 1000$ MeV can be seen in Figure 4.31. A background efficiency of about 10^{-2} was found for a signal efficiency of about 80 percent. In Figure 4.32 it can be seen that the distributions are very strongly peaked around one and zero. Cutting in these regions is difficult without an excep-

tionally good knowledge of the distribution. This is partly by construction. The neural network is designed to try to move the output value close to one or zero, but this does not suit making cuts on the regions of the plot with a large gradient. As a result, a cut in practise has to be closer to 0.5 in Figure 4.32.

A comparison between the two likelihood approaches discussed, and the artificial network, is presented in Figure 4.33. The one dimensional likelihood was significantly worse. The artificial neural network and the multidimensional PDF gave comparable results for the Monte Carlo sample considered. This suggested that the neural network and the range search algorithm are converging towards a common decision boundary.




Figure 4.31: Showing the muon efficiency as a function of electron efficiency for different energy range, $E > 400 \text{ MeV}$ (left) and $E > 1000 \text{ MeV}$ (right).




Figure 4.32: The MLP output for the test electron sample is shown (left) and the test muon sample output is shown on the right.




Figure 4.33: Performance of competing discrimination techniques from Monte Carlo data.

4.8 Conclusions

The ECal PID algorithm was able to achieve a 10^{-2} muon efficiency for an 80% electron efficiency at energies above 400 MeV. The further factor of ten required for the electron neutrino analysis can be easily achieved using the tracker section. A combined analysis using the FGD, TPC and tracker ECal is presented in Chapter 6. In the ECal the minimum energy that could be reconstructed was restricted by the number of hits a particle creates in the detector. If a particle fails to leave more than approximately five hits then reconstruction becomes very challenging. Another problem with identifying lower energy clusters was that more muons stop in the ECal. These were the most likely to be misidentified as electrons.

The neural network algorithm provides better performance than a one dimensional likelihood. It also processed significantly faster than a PDE range search algorithm, the exact speed of each algorithm depended on a number of factors; such as number of events, computer, number of training samples and number of variables, but the PDE method often took an order of magnitude longer in time to process than the neural network. This is because large amounts of data are stored in a binary tree for the PDE range search algorithm which are read at run time. The neural network by comparison is slow to train but a comparatively small list of constants are generated which can be multiplied and summed rapidly at run time. The size of tree is determined by the number of training events. Not all networks matched the distributions produced by likelihoods; however the classification power was comparable. The output distributions can be normalised and used as PDFs for likelihood algorithms in higher levels of reconstruction, even if the values themselves do not always represent posterior probabilities.

A neural network output will have an associated error. One approach to solve this problem is to perturb the inputs to the network slightly and determine the effect that this has on the output of the network. This approach can be time consuming because the effect of perturbing one node will be dependant on the values of the other nodes. A second approach is the use the second derivative of the error, defined in Equation (4.9). In [81] the propagation of Gaussian uncertainties on the input distributions and the weights values are discussed.

Future improvements to the ECal PID algorithm could include a Michel electron search. This would make the PID more effective at low energies by detecting the electron from decaying muons.

Chapter 5

T9 Testbeam

5.1 ECal Testbeam Introduction

To provide calibration and tuning data the ECal was taken to CERN. Data was acquired using the T9 beamline, connected to the PS (Proton Synchrotron). The PS was capable of providing a number of particle species in the energy region expected in the T2K near detector. Data was also collected from cosmic particles to provide a sample of muons.

The Monte Carlo simulation of the near detector used in the previous chapter required data for verification and tuning. The data taken from cosmic rays provided a source of muons with which to test the precision of the simulation in predicting the behaviour of muons. Electrons made up a large fraction of the particles in the testbeam. These particles were used to test the ability of the simulation software to predict the behaviour of electromagnetic showers. As muons lose energy in matter predominantly by ionisation it was expected that the simulation software would accurately predict the behaviour of muons. Electromagnetic showers deposit energy in a more complex manner which is

harder to model accurately so the simulation was expected to be less precise for electromagnetic showers.

5.2 CERN T9 Beamline and ECal experimental layout

5.2.1 T9 beamline

Protons from the PS were collided with a solid target in CERN's east experimental area. The resulting shower of particles provided a beam for a number of beamlines in the east area. T9 was a medium momentum beamline, supplying particles from as low as 300 MeV up to 15 GeV. The secondary particles from the target were guided along the beamline to the experimental area. The momentum of particles at the focal point in the beam area was determined by the electrical currents passed through the guiding magnets in the beamline, as specified in [87]. The intensity of the beam was controlled by the current through the focusing magnets.

Switching the direction of current through the bending magnets reversed the polarity of the particles guided down the beamline. Data collected in negative polarity running consisted primarily of electrons at low momenta and charged pions at high momenta. In positive mode, the beam consisted primarily of positrons at low momenta. At high momenta it was dominated by protons and pions. There was also a small contribution from kaons and deuterons (deuterium nuclei). The composition of the beam as predicted by CERN is given in [87].

A pair of Čerenkov counters were included in the beamline apparatus provided

by CERN and were useful for identifying particles, in particular electrons. Each detector was filled with carbon dioxide. The threshold could be varied by changing the pressure of the gas. At the end of each tube there was a thin layer of mylar. This was used instead of metal to contain the gas, so the beam was not significantly attenuated by the counters.

Time of Flight (TOF) Detector

To separate particles which are not highly relativistic a TOF detector is often effective. A TOF detector was built to separate heavier particle species below the Čerenkov threshold. If two particles are travelling at different velocities, then measuring the time taken for them to travel between two points can be used for identification, if the momentum of the particles is known. The time taken for a particle to travel between two points can be simply calculated using relativistic kinematics, by combining

$$p = \gamma m v \tag{5.1}$$

and

$$p^2 = E^2 - m^2, \tag{5.2}$$

giving the result

$$\gamma^2 m^2 v^2 = E^2 - m^2. \tag{5.3}$$

Re-writing γ in terms of the velocity of the particle gives

$$(1 - v^2)^{-1} m^2 \left(\frac{L}{t} \right)^2 = E^2 - m^2 \tag{5.4}$$

$$\frac{E^2}{m^2} - 1 = \frac{L^2}{t^2 \left(1 - \frac{L^2}{t^2} \right)}. \tag{5.5}$$

Finally, converting from natural units yields the result

$$t^2 = \frac{L^2}{c^2} \left\{ \left(\frac{E^2}{m^2 c^4} - 1 \right)^{-1} + 1 \right\}. \quad (5.6)$$

5.3 Testbeam trigger and particle identification

The particle identification algorithm designed for the T9 testbeam used a Čerenkov detector to detect electrons and the TOF to identify heavier species. The Čerenkov detector recorded a signal when an electron passed through the tube, but not if a heavier particle such as a proton or pion passed through at the momentum ranges studied. The TOF detector could not separate pions and electrons except at very low momenta (<400 MeV) as they were highly relativistic, although it could separate particles of higher mass. The TOF separated protons from pions below 1.8 GeV. Therefore, above 1.8 GeV, only an electron or hadron identification was possible.

A beam enable signal was provided by the accelerator system. While the beam enable was on, the ECal DAQ system read out the data from the TFBs when the MCM (Master Clock Module) received a trigger pulse. As all particles incident on the ECal were charged, the TOF was also used as the trigger. A coincidence trigger, demanding a hit in both TOF paddles, was implemented so that noise in the TOF was less likely to trigger the ECal when no particle was present. As the T2K beam trigger is based on a GPS signal sent from the T2K beamline, the beam trigger could not be used at the testbeam. A modified cosmic trigger was used for the testbeam. Until a trigger was received the TFBs cycled through the 23 windows that store the integrated charge for each time window. When a trigger was received each capacitor was read out.

The data corresponding to the triggering event was found in, or close to, integration period 18. In Figure 5.1, the distribution of hit times is shown for




Figure 5.1: The number of hits for each channel is shown at a given time. The time shown is relative to the trigger and measured in TDC counts (2.5 ns). The two distinct blocks represent the two RMMs in the DsECal.

each channel. The large number of hits around 5300 TDC counts shows the triggering events occurred in integration period 18. The two distinct blocks are from the two RMMs (Readout Merger Modules). The thin line close to the first RMM is the Čerenkov and TOF channels plugged into a separate TFB.

5.3.1 Time of Flight

Two scintillator paddles, approximately 8 cm x 12 cm and separated by 14 m for initial data runs, formed the basis of the TOF detector. When the detector was at normal incidence to the beam, the downstream TOF paddle was only a few centimetres from the ECal. When the detector was rotated, the paddle closest to the ECal had to be moved upstream, reducing the distance between the paddles. The scintillator paddles were a fishtail design, guiding

scintillation light to a Hamamatsu PMT. The processing of the signal was carried out using NIM electronics. Each scintillator signal was read into a constant fraction discriminator (CFD), which produced a NIM logic pulse if both inputs were 'high' at the same time. A CFD was used because the cable lengths were of order 20 m, so dispersion of the signal became significant. The discriminated signal was then passed to a coincidence unit, to ensure that a pulse was present in both scintillators before calculating the time between the two pulses. The output from the coincidence unit was also used to send a trigger to the DAQ. A TAC (Time to Amplitude Converter) was used to determine the time between the pulses.




Figure 5.2: Diagram of the NIM modules and logic used to trigger and tag events in the testbeam.

Figure 5.2 shows a schematic of the TOF electronics. The output from the pair of CFDs was 'fanned out'. Each CFD signal was sent to the trigger coincidence unit and to another coincidence unit. The output from the trigger coincidence unit was passed to a gate generator. This opened a gate that was sent to the TAC input coincidence units. As a gate was only opened if both TOF paddles saw a signal, the TAC only started timing if it was guaranteed to receive a stop

signal. The gate had to open before the direct signal from the CFDs arrived at the TAC coincidence unit, so that the time between the output signals from the coincidence units was the same as the time between the signals leaving the CFDs. Therefore, the direct signals from the CFDs to the TAC coincidence units were delayed by a period of time exceeding the sum of the time taken for the signal to pass through the coincidence units, plus the gate generator, plus the maximum expected travel time (that of a low momentum proton). The component latency times were typically measured to be 10-20 ns; however gate generators were significantly larger, typically 30-40 ns.

The output from the TAC was recorded via a modified TFB. The TAC output was a voltage pulse proportional to signal size, whereas the ADCs expected a current pulse of the type from an MPPC. A converter was built to output a current proportional to the voltage from the TAC, allowing the data to be output via the TFB. The output pulse from the TAC was of order $2 \mu\text{s}$ in length, so the output pulse covered three integration periods. As shown above with the modified cosmic trigger used for the beam, data typically was found in integration period 18. Latencies due to processing in the NIM electronics caused the TOF data to be delayed by around a microsecond, pushing the TOF output to integration window 20.

A rising pulse was observed in integration window 18 as the TAC measured the time difference between the TOF paddles. During the following three integration periods, two peaks were seen in the ADC spectrum, a large peak due to electrons and pions and a smaller peak due to protons. These were seen in both high and low gain ADC channels. Event selection in the TOF is based around checking how close a signal lies to the mean value of a peak. Figure 5.3 shows the ADC spectra for both high and low gain channels in integration

period 21. Electron/pion and proton peaks can be seen in both channels.




Figure 5.3: Measured TOF signal in high gain (right) and low gain (left) ADC channels. In the low gain channel a sharp peak can be seen due to electrons and pions with a broader peak due to protons. In the high gain channel electron/pion and proton events are also visible.




Figure 5.4: Showing the relationship between the high and low gain ADC channels. Two clusters due to electron/pion events and protons can be seen. A third cluster is more clearly separated on the two dimensional plot showing deuterons.

TOF Time Calibration

To measure the time taken for a particle to travel between the paddles, the number of ADC counts per nanosecond was required. To make this measurement, a NIM delay box was used to create a series of delayed signals into the two CFDs, simulating particles of different velocities. Time differences created by the delay box were in addition to an unknown baseline that was due to the NIM module latency. The difference between each measurement was used to find the number of ADC counts per nanosecond. Measurements were made at time differences of 1 ns, 3 ns, 7 ns, 15 ns and 31 ns. Each ADC distribution was fitted with a Gaussian function. The mean and the standard deviation were recorded and the results displayed in Table 5.1.

Time Difference/ns	Mean/ADC counts	Standard Deviation/ADC counts
1	259	4.45
3	303.9	4.29
7	359.8	4.5
15	477.1	5.23
31	706.6	7.86

Table 5.1: Measurements of generated time delay from TOF calibration run.

To remove the effect of the unknown baseline, the differences between times and ADC counts were plotted. The error on each measurement was assumed to be approximately equal to the standard deviation of the fit. Fitting the data in Figure 5.5 gave a value of $7.22 \pm 0.36 \frac{ADC}{ns}$. The intercept was compatible with zero, as would be expected if the effects of the baseline time difference had been removed.

A large uncertainty was associated with applying an absolute time calibration because a large run-to-run spread was observed in the electron/pion peak. The spread of mean values for the electron/pion peak is shown in Figure




Figure 5.5: Showing the difference in time generated by the cable delay generator, as a function of difference in ADC count.

5.6. Only data where the detector was at normal incidence were included, as the downstream TOF paddle was moved when the detector was rotated, introducing a systematic shift in the mean ADC value. When fitted to a Gaussian, the standard deviation was found to be 15.77, approximately four times larger than the standard deviation of a single run.

In practise, the electron/pion peak was trivial to find, as only electrons and pions were present in significant quantities in the beam. This was useful for confirming the identity of the proton peak and helping to identify minor peaks in the data. Figure 5.7 shows the reconstructed mass of the second largest peak in the positive beam momentum data. It was assumed that the largest peak was due to electrons and pions. The data was fitted to a Gaussian, which, with a mean of (0.92 ± 0.0211) GeV, was compatible with the well established proton mass of 0.93 GeV [5]. This suggested that the calibration had worked well, although it had to be applied on a run-by-run basis due to the drift in the electron/pion peak.




Figure 5.6: Showing the spread of the mean ADC values for the electron/pion TOF peak, for on axis data.

No cooling was applied to the TFB which read data from the trigger detectors, leaving it more vulnerable to changes in ambient temperature. Similarly, there was no cooling for the circuitry converting the TAC and PMT pulses to MPPC style pulses. Each run had a time stamp applied, giving the time the run started in Unix time. In Figure 5.8, the variation in the electron/pion mean ADC value is shown as a function of time. Vertical lines indicate successive 24 hour periods. A strong diurnal behaviour was seen in the data. The x-axis had been offset so zero was the start of the first beam run, sometime before midday. Peaks therefore, corresponded to sometime in the middle of the afternoon, when the ambient temperature was at its highest, while troughs corresponded to the early hours of the morning when the TFB would have been coolest.




Figure 5.7: Reconstructed mass of the second largest peak seen in the TOF spectrum. The value is consistent with that of the proton.

5.3.2 Čerenkov counters

A pair of carbon dioxide filled Čerenkov detectors were used in the testbeam to reduce efficiency losses. To verify that the Čerenkov detectors were functioning correctly, the beam momentum was set to 3.6 GeV. At this momentum, the pressure in the Čerenkov detectors could be set below both the muon and pion thresholds, or below the pion threshold and above the muon threshold, or above both muon and pion thresholds. The electron threshold was many orders of magnitude below the minimum pressure for the gas system available, so a signal was observed for electrons at all pressures. The electronics for the Čerenkov detector were considerably simpler than the TOF as no coincidence electronics was necessary, as shown by Figure 5.2. PMTs were used to detect light from the gas. The output of a PMT was closer to that of a MPPC than the TAC, so the interface to a TFB was simpler.

While taking testbeam data the electronics was configured to integrate the




Figure 5.8: Variation of the peak in the TOF ADC spectrum due to electrons and pions as a function of time. 24 hour intervals are shown with vertical lines.

charge over a 540 ns period, with a 120 ns reset period between each integration cycle. For a given trigger, the electronics provided a read out of data covering approximately 15 μ s; however any signal occurring during the reset time was lost. From Figure 5.1 it can be seen that the Čerenkov signals arrived close in time to the data, so few events would have lost Čerenkov information if the event itself was not also lost in the reset period.

Each Čerenkov channel was connected to a high and low gain ADC, in the same way as the TOF detector. Figure 5.9 shows high and low gain ADC spectra for both Čerenkov counters in integration period 18. The signal peak in the low gain channel can be seen above approximately 160 ADC counts. For the high gain channel most of the signal was found around the saturation region. A correlation existed between the high and low gain channels. This can be seen in Figure 5.9. The correlation between high and low gain was lost if a hit fell at the start of the integration window. In this region the charge measured could be unreliable. Vetoing events, where the correlation between high and




Figure 5.9: Showing correlation plots of high and low gain ADC channels in integration periods 18 and 19.

low gain was lost, allowed events occurring at the start of the time window to be removed. In Figure 5.10, the effect of hits early in the integration period on the measured charge can be seen. Events where the charge measurement has been affected by the start of the integration window could be vetoed, either by a cut on the low gain ADC channel, or by a cut on the time of the event. As the Čerenkov counters were only used to tag electrons, the quality of the charge measurement was not important. Events where the Čerenkov signal was close to the beginning of the integration period also had hits in the ECal with distorted hit values, so were removed.

Čerenkov Efficiency

Electron and π^- events (or positron and π^+ events) can only be separated by the Čerenkov counters, so it is important to know the efficiency of the




Figure 5.10: Low gain ADC values for integration periods 18 and 19 where the position in the integration period is measured in tdc counts (1 tdc count = 2.5ns).

Čerenkov counters to determine the electron (positron) contamination of the π^- (π^+) sample. The efficiency of a Čerenkov counter can be defined as

$$\epsilon = \frac{n}{N}, \quad (5.7)$$

where n is the measured number of events above threshold and N is the actual number of events above threshold. The total number of events in a sample can then be expressed as

$$N = \frac{n_{ud}}{\epsilon_u \epsilon_d}, \quad (5.8)$$

where n_{ud} is the number of events seen by both upstream and downstream Čerenkov counters and ϵ_u and ϵ_d are their respective efficiencies. Substituting into Equation (5.7) gives

$$\epsilon_u = \frac{n_u \epsilon_u \epsilon_d}{n_{ud}}. \quad (5.9)$$

So the efficiency of the counters can be expressed as

$$\epsilon_d = \frac{n_{ud}}{n_u} \quad (5.10)$$

for the downstream counter. Swapping the u and d subscripts gives an expression for the efficiency of the upstream Čerenkov counter.

Signals were observed in the Čerenkov counters over two integration periods. For the purposes of calculating efficiency, a signal was accepted if it was found in either time window. When selecting samples for analysis, an event was accepted if a signal was present in one integration window only. If more than one event was present there was no way to know which was associated to any TOF signal. A signal also had to be above 160 ADC counts as a pedestal peak was found at 150 ADC counts.




Figure 5.11: Čerenkov efficiencies found for positive and negative momentum data sets. The upstream Čerenkov counter was found to be systematically more efficient than the downstream Čerenkov counter.

Systematic errors were estimated from the spread of efficiencies measured in different runs, statistical errors were small by comparison. It can be seen from Figure 5.11 that the Čerenkov efficiency dropped sharply at high momentum. In many of the data sets taken at high momentum the signal recorded in the Čerenkov detector appeared to be smaller, resulting in a lower recorded ADC value. It was found that data had been shifted into the region of the ADC spectrum occupied by detector noise where it was removed by threshold cuts. As the signal was removed by a hardware level cut it was not possible to recover the events and use a different analysis method, as a result the efficiency for the high momentum data was significantly reduced.

5.3.3 Determination of beam composition

Before the testbeam data could be used to verify the ECal PID algorithm, TOF and Čerenkov data were needed to create high purity samples of electrons and pions. Electrons and pions were separated from all heavier species using the TOF detector. The Čerenkov detector separated electrons from all heavier species. Protons were only separable from pions up to a momentum of 2 GeV using the TOF, as protons became suitably relativistic that a distinct peak was no longer visible. Above 2 GeV, electrons were separable from hadrons using the Čerenkov counters, but no further separation was possible. Below 600 MeV, the time taken for the proton to pass between the TOF paddles was long enough that the TAC saturated and failed to give an output. Consequently, separate pion and proton fractions were only given between 600 MeV and 1.8 GeV. Table 5.2 summarises the cuts used to create the samples from TOF and Čerenkov data.

Particle Species	TOF Cut	Cerenkov Cut
Electron/Positron	within 3σ of e/π peak	both above 160 ADC
Charged Pion	within 3σ of e/π peak	no signal
Proton	within 3σ of proton peak	no signal

Table 5.2: Cut parameters used for testbeam PID.

As the Čerenkov counters were configured such that only electrons gave a signal, the purity of the sample was expected to be very high. A possible source of impurity for all samples was where two particles passed through the detector at almost the same time. This impurity could not be easily removed by the PID method described here, if a TOF signal still fell within a cut boundary. This was not expected to be a common pathology, as the beam spill was distributed evenly over approximately 400 ms; during which about

10^5 particles would reach the ECal, depending on momentum. Assuming an upper limit of 5×10^5 particles in 400 ms, one particle was expected every 2 μ s, while a low momentum proton only took approximately 100 ns to traverse the TOF paddles.

As shown previously the Čerenkov counter efficiency was found to be of order 0.95 for the upstream counter and 0.92 for the downstream counter. An electron contamination of approximately 0.4% was therefore expected in the pion sample. The electron efficiency was 87%. The proton sample was expected to be very pure as only a particle with the same rest mass as the proton should have caused an impurity.

Another source of impurity for pions came from the fraction of pions that decay to muons in flight. Particles from the proton target travelled approximately 30 m to the DsECal. The momentum of a forward going muon from pion decay approaches that of the parent pion as momentum increases. The momentum in the laboratory frame of a forward going muon from pion decay can be calculated by [88]

$$\rho_\mu^{fwd} = \frac{(\beta_\pi - \beta_\mu) \rho_\pi}{\beta_\pi (1 + \beta_\mu)}. \quad (5.11)$$

The term β_μ is the velocity of the muon due to the 25.9 MeV momentum it has in the pion rest frame. From this expression we see that in a pion decay with a forward going muon, the muon will have a very similar momentum to the parent pion. At 400MeV the difference in momentum between the pion and the daughter muon is approximately 1%, the same order as the uncertainty in the momentum of the parent pion. A forward going muon from pion decay would not, therefore, have been separated by the TOF. Although a full beamline simulation would be required to accurately estimate the muon contamination in the beam, work by [89] estimates the contamination to be in the region of

10% below 1 GeV, falling to less than 1% above 2 GeV.

A measurement was made of the beam composition as a function of momentum for positive and negative beam polarities. Where the beam polarity was negative, only electrons and charged pions were expected, as shown by Figure 5.12. For both beam polarities, the fraction of electrons fell with increasing momentum, while the fraction of hadrons was found to increase. The measurement of the beam composition has been corrected for Čerenkov inefficiency. Systematic errors were estimated from the efficiency difference between the two Čerenkov counters.




Figure 5.12: Measured fractions of different particle species in the CERN T9 testbeam. Error bars include systematic and statistical error estimates.

During the course of the testbeam over 2 million triggers were recorded between 300 MeV and 4 GeV. To characterise the trigger a sample of just over a million triggers was used. Table 5.3 shows how many of these were due to an identified particle and how they were distributed amongst different angles and momenta.

5.3.4 Sample selection

Once an event had been recorded and identified, it had to be associated to a cluster in the ECal. Events due to identified particles in the beam appeared in

Position	Angle /degrees	Triggers
central	0	620,000
off-centre	0	150,000
centre	15	50,000
centre	30	720,000
centre	60	510,000
centre	75	80,000

Table 5.3: Approximate number of triggers achieved in the T9 testbeam in six weeks with a total of approximately 2×10^6 .

one of two of the 23 integration cycles. Any event where a cluster was present in both possible integration cycles was rejected, as it could not be known which of the two events was due to the particle causing the trigger. No more precise timing cut, beyond identifying the integration window, could be applied to proton and pion events, as no signal was found in the Čerenkov counters. For events tagged as an electron a more precise timing cut was used, by comparing the time of the Čerenkov counter hit to the average time of a hit in the cluster. The time associated with each hit was measured relative to the time of the trigger. For an event identified as an electron to be associated with a cluster, the event in the ECal had to be less than 65 ns after the Čerenkov hit and not before.

5.4 Analysis of testbeam data

An experimental verification of the Monte Carlo simulation used to design the PID algorithm (described in Chapter 4) was carried out using the testbeam data. Events were tagged with the trigger and PID procedures described previously. The energy scale of the detector was calibrated using cosmic ray data. Cosmic ray data was also used to test the accuracy of the simulation

for muon events. Testbeam data was used to compare simulation to actual detector behaviour for electrons. Pions were not considered in this study.

5.4.1 Data Calibration

Calibration within the ECal can be separated into three stages. A TFB calibration is the first step in the calibration chain, followed by an MPPC calibration and finally a calibration of the bar itself. The calibration of the scintillator bars is implemented within the reconstruction software because some reconstruction information, such as the position of the hit, is required before the attenuating effects of the bar can be calibrated. The calibration of the electronics is implemented in the tfbCalib package. Over time, changes in ambient conditions, such as temperature, cause the number of ADC counts per MIP to drift. MppcCalib, the photosensor calibration package, calculates the gain of a given MPPC, from the bias voltage and device capacitance, on an event-by-event basis to provide the number of photons that caused an avalanche.

On the order of every hour, the pedestal ADC count was recorded to track its drift. A subtraction of the pedestal could then be applied. It was found that the drift in the pedestal over time was comparatively low, varying by about 0.5 ADC counts over a 10 hour period. After the pedestal was subtracted, a linearity calibration was applied. Both high and low gain channels had a non-linear response to an input charge. To map the response of each ADC, special CI (Charge Injection) runs were carried out. In these runs a series of known charges were injected into each ADC. The applied charge was plotted as a function of ADC count for each channel. These curves were then fitted with a third order polynomial. The fit parameters were stored and used to correct the data. TFB calibration constants were generated from a single CI

run. After this step, 99.7% of the channels had a charge calibrated to within 5% [90].

MPPC gain was parametrised as a function of bias voltage and temperature. Although the bias voltage was a constant throughout the data recorded, the temperature certainly was not. Figure 5.13 shows the measured temperature variation at the bulkhead, close to the MPPCs, as a function of the time the run was initiated. Over the data taking period, the temperature varied from 28°C to 16°C. A diurnal oscillation was seen due to day-night temperature variations. A large drop in the temperature was seen after 1.241×10^9 s, due to the water chiller used to cool the detector being set to a lower temperature.




Figure 5.13: Showing the measured temperature of TFB 1 channels, measured from the outside of the bulkhead, as a function of Unix time.

5.4.2 Cosmic Muon Calibration

To compare the data taken in the testbeam to simulation the energy scale of the simulation needed to be calibrated to the detector. This was done using events from cosmic particles. Specific data sets were acquired to collect cosmic particle events in the ECal. By comparing the observed and predicted number of photons at the face of each MPPC the behaviour of each simulated scintillator bar was calibrated to the physical bar in the detector.

Cosmic Muon Simulation

A sample of muon events was collected by looking for muons from cosmic particle events high in the atmosphere. The CoRSiKa (Cosmic Ray Simulations for Kascade) package [91, 92] was used to simulate particle interactions produced in the upper atmosphere. Using CoRSiKa, cosmic muon events were generated for a given time of year and position on the Earth's surface (including CERN) [93]. Kinematic properties of these muons were used to seed particle events in the ND280 Monte Carlo simulation as described in Section 3.7. Events were seeded into the back of the detector so that the first layer hit was layer 34. This configuration was used because the ECal was upside down for the muon data taking runs. Muons were generated from a 3 m x 3 m surface, positioned 10 cm above the detector. As the surface from which muons were emitted was larger than the detector, tracks up to an angle of approximately 60° , the limit in the simulation, were modelled.

To identify the MIP peak in the simulation, the distribution of number of photons per hit was fitted to a convoluted Landau-Gaussian distribution. The distribution of energy deposited by a charged particle passing through a thin layer of material was well modelled by a Landau distribution. The Gaussian

function accounted for signal induced noise. From Figure 5.14, the convoluted function can be seen to fit the distribution around the peak. The fit was also effective at modelling high numbers of photons, although struggled to model the shallowing gradient of the data at low charges as the Landau distribution has no tail [94]. An extra term would need to be added to the fit to include the low charge tail.




Figure 5.14: Showing the predicted distribution of number of photons deposited in a given bar for MIP like particles. The distribution has been fitted to a convoluted Landau-Gaussian function.

Cosmic Muon Data

After the completion of data taking in the CERN T9 testbeam, the detector was laid down so that it faced downwards (layer 34 was on top). Cosmic muon data was then collected. These data were used to calibrate the MIP peak in the simulation to the data, as well as to perform a data to Monte Carlo comparison of the detector response to muons.

Each data file contained approximately 10,000 events, before reconstruction failures removed events. After events had been calibrated and reconstructed, cuts were applied to remove events not due to cosmic particles. Most background events were removed by only looking at integration periods 18 and 19, as events caught by the cosmic trigger fell only in those integration windows. Cuts to remove events with an anomalously high or low charge deposit were also used to remove noise events. Reconstructed angles greater than 40° were cut out of the analysis sample, along with events entering into the side of the ECal. This prevented discrepancies due to differences in the angular distribution at high angles, as the Monte Carlo simulation did not simulate cosmic muons at angles higher than 60° .

Threshold Calibration

A threshold was set in the electronics of 190 ADC counts (approximately two photon equivalents) to remove noise hits. As the gains of the MPPCs were not identical, the number of photons required to pass the electronics threshold varied from channel to channel. This meant that after the MPPC calibration, the hard cutoff in the electronics looked more like an attenuation. A threshold of 15 photons was used in the data to remove any remaining low charge noise hits, where approximately 90 photons were equal to a MIP, depending on bar.

Simulated cosmic data was produced using a number of different threshold values, ranging from 10 to 20 photon equivalents, where a photon equivalent is equal to approximately one third of a photon. Simulated thresholds were in units of photon equivalents, not photons, as this was the unit output by the simulation software. For each simulated threshold value, the distribution of the number of hits in a cluster was compared to the data using the $\chi^2/N.D.F.$ statistic. The minimum $\chi^2/N.D.F.$ was found to be at a value of 18 p.e. as shown by Figure 5.15. This was used as the threshold value for the testbeam simulation.




Figure 5.15: Showing the variation in $\chi^2/N.D.F.$ between simulated and measured number of hits distributions as a function of the threshold used in the simulation.




Figure 5.16: Showing the maximum value of the convoluted Landau-Gaussian function against the TFB temperature.

Charge Calibration

To calibrate the energy deposited in a simulated hit to a hit in the data, MIP-like events were isolated in the data. The spectrum of number of photons in the data was fitted to the same function used to find the MIP peak in the simulated data.

Figure 5.16 shows the variation in the MIP peak by showing the fitted Landau MPV (Most Probable Value) as a function of temperature. The data suggested a small residual temperature dependence of approximately 2% over 2.5°C. The MIP peak was taken to be the number of photons corresponding to the maximum in the convoluted Landau-Gaussian function, not the Landau MPV.

A bar-to-bar calibration was not present in the calibration software. As not all bars were equally well coupled to their MPPC, a bar-to-bar dependence on the calibration constant was expected. A sample of 93 cosmic data runs were combined, producing a high statistics sample. This allowed all bars to be individually fitted to the convoluted Landau-Gaussian. Only bars where more than 250 hits had occurred in that bar during the run were fitted, to prevent fitting to very few data points. From Figure 5.17 it was found that the maximum value of the convoluted Landau-Gaussian varied from around 65 photons to 95 photons, with a number of outliers. Outliers with a photon peak value of much greater than 100 were found to be due to unstable fits. The data points with a lower peak, between 40 and 50 photons, were checked and found to be real effects. Some of these were due to known dead MPPCs, so only half the photons were expected. In others, one MPPC appeared to be giving a lower output, resulting in a similar effect to a dead MPPC. The error on the fit, calculated using the MINOS error method, was found to be approximately $\pm 10\%$. The bar-to-bar variation in fit value was therefore larger than the estimated error on the fit. A visual inspection of Figure 5.17 also suggested a TFB dependent structure to the fit value, implying that the variation was partly a systematic effect due to the location of the channel, not a random fit effect. This suggested that the constant used to calibrate the MIP peak needed to be a function of bar.




Figure 5.17: Showing the bar-to-bar variation in the maximum value of the number of photons generated at the face of the photosensor by a MIP.




Figure 5.18: Showing the bar-to-bar variation in the maximum value of the number of photons generated at the face of the photosensor by a MIP.

The maximum photon values were plotted in a map showing physical location in the detector, as shown by Figure 5.18. The most obvious artifact from Figure 5.18 was the gap on one side of the detector due to the non-functional TFB. The effects that this had on the reconstruction efficiency on that side of the detector are clearly visible. A number of the poor fits were clustered in bars around the dead TFB. Similarly, a smaller effect can be seen on the other side of the detector where the influence of being near the side of the detector limits the number of clusters. There were cold channels in the centre of the detector not on the dead channel list, for example bars 23 and 24 in layer 22. To study the effects of run-to-run variations, twenty of the cosmic data sets were selected. For each data set, the distribution of the number of photons per hit was fitted to the convoluted Landau-Gaussian function. Figure 5.19 shows the results for three separate bars, a low light yielding bar and two higher light yielding bars. It can be seen that the low light yield bar, bar 24 in layer 22, was consistently low and not being influenced by anomalous fitting results. In three of the runs low statistics caused the fit to fail and yield an anomalously high result. The plot of maximum value against run number was then fitted to a constant and the fit residuals were used to quantify how much the light yield from a given bar varied between runs. The bottom right hand plot of Figure 5.19 shows the residuals from the constant fit for all bars in a given run. It could be seen that most bars fell within ± 10 photons of the best fit value, with very few bars having a maximum deviating by more than 20 photons from the best fit value in any run. The run-to-run variation was generally within the fit errors and less than the bar-to-bar errors.

Initially the charge was calibrated using data from the bar-to-bar comparison study. The maximum of the convoluted Landau-Gaussian function for a given




Figure 5.19: Showing the variation in the convoluted Landau-Gaussian fit maximum for a low light yield bar and two higher yielding bars. The bottom right plot shows the residuals from fitting a plot of maximum value against run number to a constant for all bars. For some runs a result is missing as not enough events were collected to fit the data.

bar was compared to the maximum point in the simulated distribution. The ratio of the two maxima was used to create a calibration constant for each scintillator bar in the DsECal. This constant was then applied to the simulated events. From Figure 5.20 it was found that this systematically over estimated the charge by approximately 10%. The total hit charge shown in Figure 5.20 is the sum of the charge deposited in all hits in a cluster, measured in MEUs (MIP Equivalent Units). One MEU is equal to the charge deposit corresponding to the maximum of the convoluted Landau-Gaussian function, approximately 90 photons. As shown previously, the variation in the bar-to-bar response of the detector was shown to be approximately the same as the error on the fit of the

convoluted Landau-Gaussian function.




Figure 5.20: Showing a comparison of the predicted and measured total hit charge in a cluster using a bar-to-bar calibration constant.




Figure 5.21: Showing a comparison of the predicted and measured hit charge spectrum and total hit charge, using a single calibration constant.

A single calibration constant was considered as an alternative method to calibrate charges. Using a calibration constant of 2.75 gave a good agreement in the total charge deposited in a cluster. Using a single calibration constant narrowed the width of the hit charge distribution, as shown in Figure 5.21, because the effect of the different scintillator bar responses was removed. The hit charge spectrum shown in Figure 5.21 is the spectrum of all individual hits,

measured in MIPs. To simulate the smearing effect of the bar-to-bar variation, a calibration constant was generated for each bar randomly, using a Gaussian function. To optimise the mean and standard deviation of the smearing function, simulated cosmic data was generated using a variety of parameters. The hit charge spectra for each simulated smearing function was compared to the measured hit charge spectra, using the $\chi^2/N.D.F.$ statistic. Figure 5.22 shows how the calculated $\chi^2/N.D.F.$ varied as a function of mean and standard deviation. The lowest $\chi^2/N.D.F.$ was found at a peak value of 2.75, with a mean of 0.6, suggesting that the true bar-to-bar variation was closer to 20% than the 10% implied from the MIP peak positions of individual scintillator bars. These parameters were used to calibrate the charge for all the testbeam data. After smearing a uniform calibration constant with a Gaussian function




Figure 5.22: Showing the variation in calculated $\chi^2/N.D.F.$ as a function of smearing Gaussian parameters.

with a standard deviation of 0.6, a good agreement was seen in the hit charge distribution, as shown in Figure 5.23. By construction, therefore, the total charge in the cluster was found to agree well between data and simulation.






Figure 5.23: Comparing measured data with simulated predictions for the hit charge spectrum and the total hit charge using a single calibration constant smeared with a Gaussian to produce bar-to-bar variations.

5.4.3 Comparison of Cosmic Data to Monte Carlo


After the parameters for the threshold and charge calibration had been decided upon, the chosen simulated data set was used to study the agreement between the data and simulation for a range of parameters. The reconstructed angle of the tracks predicted by the Monte Carlo agreed well with data as shown in Figure 5.24. This suggested that the PCA based angle reconstruction algorithm was functioning well. As angles above 60° were not generated by the Monte Carlo, the modelled and experimental angular distributions were not expected to agree well at higher angles. The predicted and measured number of hits agreed well, although as this parameter had been used in the threshold calibration a good agreement was expected. The majority of events were found to pass through the detector leaving a hit in each layer, as would be expected for a muon sample. A small fraction of events stopped in the detector, either because they lacked the momentum to pass through, or because they were incident at a higher angle. A good agreement was found in the mean position of the track. This was the distance from the start of a shower to the charge



(a) Reconstructed Angle




(b) Number of Hits




(c) Half length of track

Figure 5.24: Monte Carlo to data comparison for basic reconstructed quantities of cosmic muons. Shown is the reconstructed angle (a), the number of hits in a cluster (b) and the half length of the track (c).


weighted centre, as defined in Chapter 4. The majority of tracks had a half length of just greater than 20 cm, approximately half the depth of the DsECal.




(a) Shower Angle




(b) Shower Width




(c) AMR



(d) Maximum Charge Ratio



(e) EM Fit Likelihood



(f) Network Output

Figure 5.25: Cosmic muon simulation to data comparison for PID variables.

Figure 5.25 compares the response of the PID variables, discussed in Chapter 4, for data and Monte Carlo. Once again, a generally good agreement is found between distributions of data and simulation. An excess of data events was found at higher values of both shower angle and shower width. The AMR distribution in the data was similar in shape to the Monte Carlo prediction. Although, the distribution in data was shifted lower than in simulation. Shower width, shower angle and AMR all derived from the PCA fit to the shape of the shower, so events that caused the lower AMR peak in data were the same events causing the longer tail in the shower width and shower angle, as shown in Figure 5.26. This implied that events in the data looked more 'shower-like' than in the simulation. One possible reason for this was that the magnitude of small charges, from hits just clipping the edge of a bar, were underestimated in the simulation. Therefore, tuning the threshold by finding the best fit in the number of hits may have caused low charge hits in the simulation to be incorrectly cut.




Figure 5.26: Showing the correlation between AMR and shower width and shower angle

The maximum ratio variable was very sensitive to thresholds and saturation effects, as it divided the largest charge by the smallest charge. Despite this

potential instability, a good agreement was found between the data and the simulation. The EM likelihood, which measured the probability that the cluster was due an electromagnetic shower also broadly agreed between data and Monte Carlo.

After being passed through the neural network, the track vs. shower output variable was found to be very resilient against the differences in the input distributions. A small increase in shower-like events was observed in the data. These were correlated very strongly to the low AMR events due to a larger fraction of events with lower numbers of hits in the data that were misidentified as showers. The efficiency of the neural network output was 97.1% in simulation and 93.8% in data when a cut of 0.5 was applied

5.4.4 Comparison of electron data with simulation

Testbeam Simulation

A sample of pure electron events was generated to simulate electrons in the T9 testbeam. The electrons were generated 10 cm behind the back face of the ECal, the ECal sitting back-to-front in the testbeam, and fired towards the centre of the ECal. A detailed simulation of the beamline was not available, so the geometry of the beam was estimated from the position of the TOF trigger paddles. The energy of the electron was distributed uniformly, with a width of 10 MeV centred on the nominal beam momentum in order to simulate the momentum uncertainty. At all angles, the position the beam entered the detector was chosen so that the beam would pass through the centre of the detector, as had been attempted when taking data.

Testbeam Data

The PID described in Section 5.3 was used to create a sample of electron events. A signal was required in both Čerenkov counters no more than 65 ns before, and not after, the cluster appeared in the ECal for the event to be classed as an electron event. On top of this, only one cluster was allowed in the two integration periods containing triggering events, preventing any ambiguities over which cluster the Čerenkov signal was in response to. As for the cosmic sample, a threshold of 15 photons was applied to the data to remove low charge noise hits. Figure 5.27 compares testbeam data to Monte Carlo predictions




Figure 5.27: Monte Carlo to data comparison of basic quantities for 600 MeV electrons. While the hit charge spectrum and maximum hit charge agreed well, a discrepancy was seen in the total hit charge and number of hits distributions.

for 600 MeV, on-axis, electrons. The total hit charge in a cluster was lower by approximately 10% in the Monte Carlo than in data. This shows that the detector and simulation have different responses to an electromagnetic shower. The hit charge spectrum agreed well except for a small discrepancy at very low charges. The distribution of number of hits also peaked lower in the simulation by approximately three hits. Moving the charge threshold in the simulation by only four photon equivalents brought the number of hits distribution into much better agreement, raising the possibility that small differences between data and simulation at low charges could have been responsible.

An EM scale correction of 9% was applied to the simulated charges. Figure 5.28 shows that the total hit charge was brought into much better agreement by the scaling factor. The agreement in the hit charge spectrum was worse after adding the scaling factor suggesting that lost hits were responsible for at least some of the missing charge. The PID variables are shown not be significantly affected by the application of the scaling factor as can be seen by comparing to the unscaled data in Figure 5.29.


As was observed in the cosmic data, the shower angle and the shower width, as shown in Figure 5.29, appeared wider in the data than the Monte Carlo, although the agreement was generally good. The wider showers were reflected in the lower peak found in the AMR distribution, due to dividing by a larger denominator. Again, these results were compatible with the hypothesis that the simulation was underestimating the spatial extent of the shower. The EM likelihood showed a good agreement between simulation and data. This suggested that the energy fitter was robust when applied to showering particles as no specific tuning had been carried out. In the discussion of the muon data to Monte Carlo comparison it was noted that the maximum ratio variable

was very sensitive to low value charges and threshold values. For the electron data an offset was also seen. Once again the peak in the distributions broadly agreed although a shorter tail was seen in the simulation.


The output from the neural network had a very similar distribution in both data and Monte Carlo. There was no significant discrepancy in the number of events misidentified as tracks. Interestingly, the neural network output was closer to the ideal output in data rather than simulation. This demonstrated the assumption made in Chapter 4, that cutting on the rapidly changing region of the neural network output is not sensible without a thorough understanding of the input variables.

Electron energy comparisons


The agreement between data and simulation did not change significantly with increasing energy. At high energy, the agreement in the maximum ratio variable improved significantly for the same simulated threshold value. This was expected because the magnitude of the largest energy deposit increases with energy, while the smallest is approximately constant. Any uncertainty in the magnitude of the smallest charges was, therefore, less significant when finding the ratio. On top of this effect a saturation was seen in the data which was not modelled in the simulation, as shown in Figure 5.30. This caused the largest energy deposited in the simulation to be larger than in the data, compensating for the discrepancy at low charges. For all energies the neural network output is very consistent, with a similar offset observed between data and Monte Carlo. Once again, the measured distribution of neural network output values at 1.4 GeV was more tightly peaked around zero in the data than the simulation, as shown in Figure 5.30.




(a) Total Hit Charge




(b) Hit Charge Spectrum




(c) AMR



(d) Shower Width




(e) EM Likelihood




(f) Track/Shower Output


Figure 5.28: Comparison of testbeam data and simulated data after applying a 9% EM scaling factor.




(a) Shower Angle




(b) Shower Width




(c) AMR



(d) Maximum Charge Ratio




(e) EM Fit Likelihood




(f) Track/Shower Output


Figure 5.29: Monte Carlo to data comparison for PID related quantities for 600 MeV electrons.




(a) Hit Charge Spectrum 400 MeV




(b) Hit Charge Spectrum 1400 MeV




(c) Maximum Charge 400 MeV



(d) Maximum Charge 1400 MeV



(e) Neural Network Output 400 MeV




(f) Neural Network Output 1400 MeV


Figure 5.30: Comparing the agreement between data and simulation at 400 MeV and 1.4 GeV. Saturation effects can be seen in the hit charge spectrum in the data which are not modelled by the simulation. Distributions broadly agree although a discrepancy is seen for the maximum ratio.

Higher angle electrons


Data was also taken with the beam orientated at 30° and 60° with respect to the face of the ECal. Figure 5.31 compares the results for data and simulation with the ECal orientated at 30° to the beam, while Figure 5.32 shows the same data but with the detector at 60° . The agreement between data and simulation was similar to the on-axis data, although the agreement in the EM likelihood variable was poorer at higher angles. The good agreement between the half length of the track demonstrates the extent of shower was well modelled. Apparent saturation effects were seen in the charge spectrum in the 60° sample, with a similar cut off to that seen for the 1.4 GeV, on-axis data. As was seen for the on axis data, the shower width implied that lateral hits were not being modelled correctly. This behaviour was possibly due to uncertainties in the modelling of particles which just clipped the corner of a bar. This type of hit would be expected to be more abundant at higher angles. The maximum ratio showed the poorest agreement of the PID variables due to its strong dependence on the lowest charge. The EM likelihood demonstrated a poorer agreement at higher angle although the distributions were broadly similar. Once again the neural network outputs were similar, although the data was more ‘shower-like’ than the simulation.




(a) Track Half Length




(b) Hit Charge Spectrum




(c) Shower Width



(d) Maximum Charge Ratio




(e) EM Fit Likelihood




(f) Neural Network Output


Figure 5.31: Monte Carlo to data comparison for 800 MeV electrons orientated at 30° to the face of the DsECal.




(a) Track Half Length




(b) Hit Charge Spectrum




(c) Shower Width



(d) Maximum Charge Ratio



(e) EM Fit Likelihood



(f) Neural Network Output

Figure 5.32: Monte Carlo to data comparison for 800 MeV electrons orientated at 60° to the face of the DsECal.

Momentum/MeV/c	Angle/°	$\epsilon_{(\text{MC})}$	$\epsilon_{(\text{Data})}$	compatibility/ χ^2/ndf
400	0	96.8	98.1	3.46
600	0	96.8	98.3	4.16
800	0	98.9	99.2	3.66
1400	0	99.7	99.7	4.27
1600	0	99.7	99.4	3.43
400	30	93.2	97.2	9.6
600	30	95.4	98.2	9.37
800	30	96.6	98.1	5.91
1000	30	97.8	98.7	1.32
1400	30	98.7	99.7	0.91
400	60	91.5	93.4	7.49
600	60	95.8	97.5	6.51
800	60	97.3	98.8	3.19
1000	60	98.4	98.5	3.72
1200	60	98.5	99.3	4.31

Table 5.4: Showing simulated track vs. shower discrimination efficiency for simulation, ($\epsilon_{(\text{MC})}$), and data, $\epsilon_{(\text{Data})}$, as well as the compatibility between them in terms of χ^2/ndf for a range of angles and energies.

Table 5.4 shows the predicted and measured PID efficiency at a range of angles. In general the efficiency was found to increase slightly with increasing energy in both data and simulation. At low angles the data and simulation were found to agree well. At increasing angle the simulation was found underestimate the efficiency.

ECal Energy Resolution

To make a preliminary measurement of the energy resolution of the DsECal, the total hit charge for a given energy was fitted to a Gaussian function to estimate the width of the reconstructed energy distribution, in units of MEU (MIP Equivalent Units). A factor, α , to convert the width into units of GeV, was determined by comparing the known beam momentum to the mean MEU

value such that

$$\langle MEU \rangle = \alpha E(\text{GeV}). \quad (5.12)$$

From Figure 5.33 the value of α shows a small decrease with increasing energy, approximately 3% over 0.5 GeV, suggesting the energy response of the detector is not perfectly linear. The standard deviation in units of GeV was then fitted to the function

$$\frac{\sigma}{E} = \frac{A}{\sqrt{E}} \oplus B, \quad (5.13)$$

where A is a stochastic component to the resolution, B is a constant component and $\oplus$ indicates summing in quadrature. Figure 5.33 shows that the stochastic component to the resolution is approximately 9%, with a negligible constant component. This is close to the design resolution of $\frac{7\%}{\sqrt{E}}$.




Figure 5.33: Showing a preliminary measurement of the DsECal energy resolution.

5.5 Conclusions

A PID algorithm was successfully designed and implemented to separate data collected in the T9 beamline into different species. The efficiency of the

Čerenkov counters was found to be between 90% and 95% for almost all of the momentum points sampled. A measurement of the beam composition as a function of momentum was produced to test the algorithm and found to be similar to existing studies.

In general, a good agreement was found between the data and Monte Carlo. Some pathological types of event were identified and removed using cuts. Simple properties such as number of hits, total charge deposited and hit charge spectra agreed well for the cosmic data. Overall the simulation of cosmic muons agreed well with data. In the electron data an EM scale factor of approximately 9% was found to give a good agreement in the total hit charge distribution. A discrepancy was also seen between data and simulation in the number of hits distribution. This could be explained by the simulation underestimating the energy deposited by low energy hits. Consequently some of these hits were removed by the threshold cut, and so lost entirely. This was supported by the marginally narrower clusters produced by the simulation. An effect of approximately four photon equivalents was required to account for this effect.

Variables used in the PID had a similar structure in both data and Monte Carlo. Differences noted between measurements and simulation appeared to be due to effects in low charge hits. Most obviously, the maximum charge ratio was explicitly a function of the lowest charge measured and so was very sensitive to applied threshold values. The AMR, shower width and shower angle were also linked to threshold values, albeit less strongly. The PCA algorithm used to calculate those variables was charge weighted so the size of the lowest charge affected the shape of the PCA ellipse.

Further work is required to calibrate the data more accurately. A more refined

low level calibration of the MPPCs and electronics is required. In particular, the interplay between the thresholds and the charge calibration of small hits needs to be studied in detail.

Chapter 6

Electron Neutrino Analysis

6.1 Introduction

Determining the fraction of electron neutrinos in the beam is key to the ν_e appearance search and therefore the determination of the value of θ_{13} . ND280 will provide a measurement of the electron neutrino contamination in the beam. This can be extrapolated to Super-Kamiokande to predict the electron neutrino flux in the case where oscillations do not occur. In this chapter an electron neutrino analysis is presented using the tracker section of ND280, including the downstream and barrel ECals. The performance of the particle identification algorithm developed in Chapter 4 is explored, along with other methods for reducing backgrounds using the ECal.

6.2 Neutrino Interactions in ND280

The Genie [77] neutrino event generator was used to simulate neutrino interactions in the ND280. Genie produces neutrino events from neutrinos of a

given energy upon a given material. From the output, one gets the position and momenta of final state particles which can be used to seed a simulation of the event. The type of the interaction, for example CCQE (Charged Current Quasi-Elastic), is also given so that different types of event can be selected for simulation. Event types expected in ND280 were discussed in detail in Section 2.2.1. Genie was used to generate a sample of neutrino events equivalent to 5×10^{21} POT (Protons On Target), or five years nominal running, from four different detector configurations as part of the MDC 0 (Mock Data Challenge) exercise [95]. A startup configuration was defined in which the barrel and P0D ECals were not installed, a full detector configuration was also defined with all the ECals present. In total, 30% of the samples were produced in the startup configuration and 70% were produced in full mode. The P0D was simulated containing water for 50% of events in both full and startup configurations and containing air for the remaining 50%.

A key test of the PID algorithm is the electron neutrino analysis, due to the small expected signal. Electron neutrino contamination comes from the decay in flight of mesons and muons in the beamline decay pipe. At low momenta, contamination is primarily due to the decay of muons; and at high momenta, kaons and pions. CCQE events formed the basis of the electron neutrino analysis, both because reconstruction was assumed to be simpler and also CCQE interactions are the dominant process in the T2K energy regime.

Iron in the magnet makes up the majority of the mass in the ND280 detector, whilst the densest material is the lead in the ECal. In contrast, the tracker is made of plastic in the FGDs and low density gasses in the TPCs. Proportionately, very few interactions occur in the TPC gas. The events studied were from neutrinos interacting with the carbon and oxygen atoms in the FGDs,

as shown in Figure 6.1. The structure of the magnet can clearly be seen in




Figure 6.1: Showing the distribution of neutrino interactions in the ND280 detector.

this figure. An asymmetry in the number of interactions between the top and bottom parts of the magnet can also be seen. This is due to the direction of the neutrino beam. The lead in the ECal is also visible directly inside the magnet. A full ECal analysis should be able to identify the numerous neutrino events occurring within the ECal due to the dense lead.

If an interaction is neutral current, then it is not possible to identify the species of neutrino that has interacted, as there is no lepton from the primary vertex. A charged current interaction produces a lepton at the primary vertex. This can be identified and used to tag the species of the neutrino. The CCQE channel accounts for approximately one third of the charged current ν_e events. A CCQE measurement is defined as an exclusive CC (Charged Current) measurement because it includes only a subset of the entire CC sample. An inclusive CC measurement studies all CC neutrino interactions. The efficiency of the cuts applied to both types of analysis are discussed in Section 6.3. The different combinations of final state particles and their expected abundance is shown in [96].

Interaction mode	Fraction	No. of events / 10^{21} POT / tonne
CCQE	33.3%	65038
CC $p\pi^+$	11%	17846
CC $p\pi^0$	2.9%	4887
CC $n\pi^+$	3.0%	5107
CC Coherent	1.3%	2189
CC multi- π	7.0%	11943
CC p ρ^0	0.5%	835
CC DIS	7.7%	13057
NC n $p\pi^0$	1.7%	2837
NC p $n\pi^-$	1.1%	1931
NC Coherent	0.6%	1099
NC multi- π	2.1%	3639
NC n ρ^0	0.1%	150
NC p ρ^0	0.1%	120
NC ΛK^0	0.0%	31
NC DIS	2.4%	4022
NC Elastic p	8.0%	13581
NC Elastic n	9.2%	15671

Table 6.1: The different categories of neutrino event expected in ND280 are shown along with their relative abundances [18]. An absolute number of events per 5×10^{21} POT per tonne is also shown. The largest fraction of events are CCQE events.

Identifying the track belonging to the lepton and then identifying the species of the selected track was the first step of the analysis. In neutral current events where no lepton was present, the particle identification should have discarded the event, with the only background coming from mis-identified tracks. Possible sources of background arose from NC events that generated leptons. For example, the NC multi-pion channel was estimated to account for about 2% of events in ND280 as shown in Table 6.1. A π^0 will decay into two gamma rays or a gamma ray and an electron/positron pair almost immediately. If the lepton selection found a lepton from the pion decay (or from conversion of a gamma) then it could have passed selection cuts. Such events were revealed by the presence of two leptons close to each other, which

would not occur if a single lepton was created from the primary vertex of a CC interaction. An example of an event with a π^0 in the final state is given in Figure 6.3. CC pion events were predicted to constitute approximately 10% of the event sample. These were more easily rejected in the case of muon neutrino events as there was a muon which could be identified. The important CCQE channel made up approximately one third of events. An example of a CCQE electron and a CCQE muon event is shown in Figure 6.2. In the electron case a recoil proton can also be seen.




Figure 6.2: An event display of a CCQE electron event with recoil proton (left) and a CCQE muon event (right). Electrons are coloured red, protons are coloured blue and muons are coloured green.

6.2.1 Neutrino interactions in the FGD

In Figure 6.1 it can be seen that very few events occur in the FGD or TPC due to their low density. Conversely, the high spatial resolution of the TPC makes it the most sensitive part of the detector for particle reconstruction. The electron neutrino analysis presented concentrates on events with their vertex in the FGD, with final state particles passing through the TPC.

Figure 6.4 shows the neutrino flux as a function of momentum, for neutrino




Figure 6.3: Showing an event where a pair of leptons were seen in the final state due a π^0 decay.

interactions in the FGD. The ratio of the electron to muon neutrinos is shown in Figure 6.5 as a function of momentum. Around the T2K energy peak the electron neutrino fraction is less than 1%, although this increases at higher energies due to the electron neutrino contribution from charged and neutral kaon decays. To achieve a signal to noise ratio of better than 10:1 requires PID capable of a signal acceptance to background rejection of around 1:1000. This assumes the electron neutrino fraction is 1% and a high efficiency signal acceptance is achieved.

Leptons from CC interactions in the FGD, passing through the DsEcal, are high momentum particles from low Q^2 interactions where the lepton takes most of the neutrino's momentum, where Q^2 is the momentum transferred between the neutrino the the body it scatters off of. In the barrel, leptons from CC interactions will have been from higher Q^2 interactions where a larger fraction of the neutrino momentum has been imparted to the nucleus, scattering the lepton at a higher angle. One expects therefore to observe different lepton




Figure 6.4: Genie predicted event rates in the FGD for electron neutrinos (left) and muon neutrinos (right).

momenta in the barrel ECal compared to the DsECal.

6.3 Electron Neutrino Analysis - Event Selection

To quantify the amount of signal and background present in the selected sample after each cut, the efficiency and the purity of the sample was defined. The efficiency of the sample after a given cut, j , was defined as

$$E_j = \frac{S_j}{S_{j=0}}, \quad (6.1)$$

where S_j was the number of events passing the cut and $S_{j=0}$ was the total signal in the sample before any cuts were applied. Purity was used to measure the amount of background in a sample. The purity of the sample after a given cut was defined as

$$P_j = \frac{S_j}{S_j + B_j}. \quad (6.2)$$

All of the cuts in the study used reconstruction information available with real data. Truth information, from the Monte Carlo, was only used to determine




Figure 6.5: The ratio of electron to muon neutrino events as a function of particle momentum.

the true topology of the event, and to calculate the efficiency and purity of a given cut. At the time of the study the reconstruction was not capable of reconstructing the event fully. As a result, cuts used in the analysis presented were based on individual sub-detector PID algorithms and simple topological features of the event. The analysis used the ECal PID algorithm described in Section 4, the TPC PID algorithm developed by Giganti [97] and was developed from the analysis framework from George [98] as part of the ν_e analysis group.

6.3.1 Lepton Selection

To identify possible CC (Charged Current) events, a lepton selection identified the track most likely to be due to a lepton from the primary vertex. Particle identification then cut out events where the selected particle was not

an electron. The first step in the lepton selection was to check that at least one reconstructed track originated in the FGD. If more than one track originated from the same vertex then the highest momentum particle was made the electron candidate. A momentum threshold of 200 MeV was applied to the selected track to remove low momentum tracks.

The second step was to check whether the track originated within the fiducial volume of the detector. Currently, the reconstruction software does not fully reconstruct the times of hits in the tracker so, for an event leaving only a single track, the origin of the vertex is not known. The analysis described here defined the fiducial volume as a box 10cm inside the edge of each FGD on all sides. The only events removed were those which fell close to the edge of the FGD.

Particles with a curvature consistent with being positively charged were removed in the final stage of the lepton selection. This removed anti-neutrino events from the sample. For an anti-neutrino study this would be reversed to remove particles with negative curvature. This lepton selection algorithm assumed that the highest momentum particle was always the lepton. This assumption was not completely valid. High Q^2 events could have created a negative pion or kaon with a low momentum lepton.

6.3.2 Particle Identification

TPC Particle Identification

For momenta greater than approximately 250 MeV, the difference in the energy deposited per unit distance between an electron and a muon is around 30%, although this decreases as energy increases, as can be seen in Figure 6.6.

The TPC has been designed to have an energy resolution of 10%, separating electrons and muons by a measurement of the energy loss. Figure 6.6 shows the difficulty in using energy loss measurements to separate charged pions and muons. These have an almost identical mass and this results in very similar energy losses in a material at a given momentum. Below 200MeV the energy deposit per unit length of the muon increases, making it very difficult to make a 3σ separation of muons and electrons. The TPC PID algorithm described below was developed by Giganti and is described fully in [99]. To find the




Figure 6.6: Showing the stopping power for electrons, muons and charged pions over the momentum range in T2K. A powerful electron vs. muon separation can be made above 200MeV. Separating muons and pions is difficult in general. Figure modified from [99].

energy loss of a particle, the TPC PID does not use the mean energy loss as calculated by the Bèthe-Bloch equation, due to large statistical fluctuations in the data. Instead, a distribution of energy deposits per unit distance recorded

in the TPC, called a straggling function, is used. The straggling function has a long tail due to high energy deposits. The peak of the straggling function is the most probable value for the energy loss of the particle. To find the most probable value a truncated mean, defined as [99]

$$\bar{C}_T = \frac{1}{\alpha N d_0 f(N)} \sum_i^{\alpha N} g(d_i) C_C(i), \quad (6.3)$$

is used to remove the effects of the high energy deposit tail. In Equation 6.3, N is the number of entries, α is the truncation fraction controlling how far into the tail entries are included in the calculation and $C_C(i)$ is the energy deposited in a single TPC cluster. $f(N)$ and $G(d)$ are factors added to take into account, respectively, effects from the number of clusters and the gap between MICROMEGAS modules. Both factors are close to one. A width σ_T may also be associated with the truncated mean. This is defined as [99]

$$\sigma_T = \sigma_0 \alpha(N) \beta(d) \sqrt{\frac{C_E}{C_E(MIP)}}. \quad (6.4)$$

$C_E(MIP)$ is the expected energy deposit for a MIP track and σ_0 is its associated width. C_E is the expected energy deposit of the particle species for which the width is being calculated. Again, α and β are constants which depend on the number of clusters and the TPC gap size.

To distinguish between different particle species a quantity referred to as the 'pull' is defined. The pull quantifies how far the measured energy loss is from what would be expected from theory. It is defined by [99]

$$\delta_E(i) = \frac{\bar{C}_T - C_E(i)}{\sigma_o(i)}, \quad (6.5)$$

where σ_o is the total width associated with the pull. This width is defined as

$$\sigma_o(i) = \sigma_T(i) \oplus (dC_E/dp)\sigma_p, \quad (6.6)$$

where i is the particle hypothesis. $\bar{C}_T$ is the measured truncated mean and $C_E(i)$ is the value expected from theory. The expected value of the most probable energy loss has been determined from a Monte Carlo study using mono energetic electrons, muons and protons, and is parametrised as a function of $\beta\gamma$. If the calculated pull for a given particle hypothesis is low, it means that the measured energy loss is consistent with that hypothesis. A large absolute value for the pull means that a significant difference between measured and expected energy loss has been measured, making that hypothesis less likely.

To pass the TPC cut recommended by the FGD group [97], the absolute value of the electron pull had to be less than two. On top of this, there was an explicit requirement that the muon pull had to be greater than two. This helped remove contamination due to the large muon background.

ECal Particle Identification

The global reconstruction software was not capable of connecting clusters in the ECal to tracks in the tracker. A simple connecting algorithm was included in the analysis to overcome this limitation. The distance between the end of the selected lepton track, defined as the last reconstructed hit in the FGD or TPC, was compared to the start point of all ECal clusters. If the closest cluster was within 40 cm of the end of the TPC track then the cluster was matched. Figure 6.7 shows the distance between the end of the lepton candidate track in the FGD, and the start of the nearest ECal cluster. The ECal PID was implemented as described in Chapter 4, by cutting on a two-dimensional plane of EM shower versus hadronic shower network output, against track versus shower network output. Neural network outputs for electrons, muons and charged pions are shown in Figures 4.29 and 4.30. Both neural networks were




Figure 6.7: Histogram showing the calculated distance between the selected TPC track and ECal clusters. A large peak at approximately 20 cm indicates well reconstructed clusters with a gap due to structural components and electronics between the tracker and ECal.

trained so that their output values lay between zero and one; an electron would occupy the region close to zero along the track vs. shower axis and the region close to one on the EM vs. hadronic shower axis. In the analysis presented here, to pass the ECal PID cut, a cluster had to have a track vs. shower value less than 0.5 and an EM vs. hadronic shower value of more than 0.5. Using the track vs. shower cut, the muon background was reduced further, removing events missed by the TPC PID algorithm. The EM vs. hadronic shower cut was primarily intended to remove events where a charged pion had been incorrectly identified as the muon.

After the TPC and ECal PID, the fraction of events due to electron neutrinos increased from about 1% to almost 30% for the inclusive sample. The number of muon neutrino events fell by a factor of 100, as can be seen in Figure 6.8. The cut from the TPC PID was the single most powerful cut in both analyses. Approximately 20% of CCQE electron events and 33% of CCnQE events were

lost due to the simple ECal connecting algorithm. From Figure 6.8 it was seen




Figure 6.8: Numbers of electron and muon neutrino events left in samples after lepton selection and particle identification.

that CCQE muon events were well suppressed by the TPC and ECal particle identification algorithms. After the ECal PID, only 23 out of approximately 60,000 CCQE ν_μ events remained in the sample. The background consisted mainly of CCnQE ν_μ interactions and NC interactions.

Specific Background Cuts

If the measurement of a track's curvature had a large associated error, then the pull calculated by the PID was also often affected. The FGD reconstruction group recommended

$$C - 2\sigma(C) < 0, \quad (6.7)$$

as a cut to remove events where the reconstruction had failed, where C is the curvature and $\sigma(C)$ is the error of the measurement. This was combined with cutting out events with a momentum error greater than 0.2. In Figure 6.9 the curvature error cut is shown to have a small effect on the performance of the analysis. Muon background was reduced by a factor of a third, on top of the lepton selection and PID cuts alone.

After the curvature error cut all remaining tracks were electron-like, having passed TPC PID cuts and ECal cuts. There were two possible sources of background from the PID cuts above. The first was due to ν_μ events where the lepton selection had identified an electron from particles ejected from the nucleus. These events were identified by searching all other clusters for a track consistent with being a muon. The second source of background came from π^0 neutral current events. A charged current event either had a muon and could be rejected by finding the muon, or it was a signal event. An example of a π^0 event is shown in Figure 6.3. To remove NC π^0 events, all TPC tracks were examined to check whether they had a start (or end) point within 4cm of the selected lepton track, as suggested by Giganti [97]. A more effective cut was found to be only to accept events with a single lepton.

The targeted background cuts reduced the number of signal events to 282 with the single lepton cut. Background events were reduced to 153 with the single lepton cut, increasing the signal to noise ratio to 1.84:1 as shown by Figure 6.9 and Table 6.2. After applying the targeted cuts, NC events constituted the largest source of background. Typically these were high multiplicity events where the reconstruction had failed to reconstruct the tracks correctly. A high multiplicity event with a large number of electrons which should not have passed cuts is shown in Figure 6.10.

Cut	CCQE ν_e	CCnQE ν_e	CCQE ν_μ	CCnQE ν_μ	NC event
Full sample	683	1922	75493	83138	22949
No. of tracks cut	683	1922	75943	83138	22949
Fiducial cut	617	1608	62014	64980	19808
Negative Lepton cut	525	1446	49317	57027	7869
TPC PID cut	384	793	158	1367	1101
Connected to ECal	276	605	76	994	748
ECal PID cut	264	547	23	660	666
Curvature cut	261	534	14	446	538
Single Lepton cut	151	131	10	31	112
CCQE cut	101	37	8	6	51

Table 6.2: Predicted number of events after application of successive cuts.

Particle Identification in the ECal has been discussed at length in Chapter 4. The ECal can employ other techniques to aid in PID. At low momenta, where both ECal and TPC PIDs suffer efficiency losses, the muon is more likely to convert in the ECal. The muon can then be detected by observing the Michel electron from the muon decay. A Michel electron search is implemented here, based on work by Murdoch [100]. A track stopping in the ECal is detected by checking whether or not it leaves any hits in the last two layers of the detector. If not, it is considered a stopping muon candidate. If a second cluster with fewer than seven hits is found within 20 cm of the outer-most layer, then it is tagged as a Michel electron event and rejected from the electron sample.

An additional cut, demanding that only a single cluster was present in the ECal, was applied in exclusive analysis calculations. This was implemented to remove gammas from pion decay in CC electron events. Although such a cut significantly improved the purity of a CCnQE sample the efficiency of the sample was degraded as a significant fraction of electrons were predicted to lose energy by Bremsstrahlung before converting.




Figure 6.9: Using cuts targeted at remaining CCnQE muon and NC events and π^0 events bring the signal to noise ratio to 1.84:1.

6.4 Analysis Performance

An inclusive electron neutrino signal event was defined as an event that was either CCQE or CCnQE. An exclusive signal event was defined as a CCQE electron neutrino event. The event selection described previously was tested as part of both inclusive and exclusive analyses. Figure 6.11 shows the efficiency and purity of the inclusive and exclusive selections as a function of cut, numerical values are shown in Table 6.3. Binomial errors were assumed when calculating efficiencies and were determined by

$$\delta\epsilon = \frac{1}{N} \sqrt{k(1 - (k/N))}, \quad (6.8)$$

where k is the number of signal events and N is the total number of events passing cuts. The purity of the exclusive sample was 51% while the inclusive




Figure 6.10: Showing a high multiplicity event that was mis-reconstructed by the reconstruction software.

sample had a purity of 69%. The tight cuts applied to reduce the CC muon contamination successfully reduced the background by approximately three orders of magnitude. A CCQE electron event had one of only two topologies: lepton only, or lepton and proton. These were found to be more resilient under tight cuts. The CCnQE inclusive analysis was more strongly affected by the tight cuts as the more complex final states were more likely to be rejected. The efficiency of the CCQE sample after the last cut was 14.8% while the efficiency of the inclusive analysis was lower at 11%. Approximately half of the total simulated neutrino events were not included in the analysis as they did not pass through the reconstruction algorithms either, because the reconstructed lepton momentum was below the threshold or the event failed to reconstruct. Lepton selection cuts had a relatively small effect on the efficiency and purity of either sample. TPC and ECal PID algorithms however, had a significant effect on the purity of the inclusive and exclusive selections. After the TPC PID algorithm was applied, the purity of the exclusive sample rose to 11.0% from approximately 1%. After applying the targeted cuts this rose further to

Cut	CCQE	CCQE	CCnQE	CCnQE
	efficiency	purity	efficiency	purity
Full sample	100 ± 0	0.38 ± 0.015	100 ± 0	1.44 ± 0.028
No. of tracks cut	100 ± 0	0.38 ± 0.015	100 ± 0	1.44 ± 0.028
Fiducial cut	90.29 ± 1.14	0.43 ± 0.017	85.36 ± 0.7	1.53 ± 0.033
Negative Lepton cut	76.76 ± 1.61	0.48 ± 0.021	75.36 ± 0.85	1.75 ± 0.04
TPC PID cut	56.03 ± 1.9	10.98 ± 0.53	45.42 ± 0.99	33.26 ± 0.8
Connected to ECal	40.14 ± 1.88	11.28 ± 0.64	33.81 ± 0.94	35.51 ± 0.97
ECal PID cut	38.38 ± 1.86	13.45 ± 0.77	31.17 ± 0.92	40.82 ± 1.11
Curvature cut	37.94 ± 1.86	15.89 ± 0.91	30.6 ± 0.91	47.9 ± 1.24
Single Lepton cut	22.21 ± 1.59	36.03 ± 2.35	10.9 ± 0.62	66.35 ± 2.31
CCQE cut	14.85 ± 1.36	50.75 ± 3.54	5.39 ± 0.44	68.8 ± 3.28

Table 6.3: Efficiency and purity of electron neutrino samples following application of successive cuts.

36%. A similar improvement was seen in the inclusive sample where the purity was 0.8% before the TPC PID cut and 33% after.

After the application of further cuts to remove remaining muons and electrons from NC π^0 decay, the final purity of the exclusive analysis was 51% with an efficiency of 15%. As expected, the inclusive selection had a higher purity than the exclusive analysis, although the efficiency was lower. After all cuts had been applied, the purity of the inclusive selection was 66% and the corresponding efficiency was 11%. Table 6.4 shows the total number of signal and background events that passed the final cut for the inclusive analysis and the exclusive analysis.

Event	Inclusive Analysis	Exclusive Analysis
Signal	282 ± 17	101 ± 10
Background	153 ± 12	102 ± 10

Table 6.4: Number of signal and background events with statistical errors for inclusive and exclusive samples.

Both ECal and TPC PID algorithms were less efficient at high energies. The

purity of the analysis was therefore expected to decrease with the momentum of the incoming neutrino. Most of the background was found not to be due to CC muon contamination, but to high multiplicity NC events coming from events with a large Q^2 . Figure 6.12 shows the efficiency and purity of both inclusive and exclusive samples as a function of momentum. In the inclusive sample the main source of contamination was due to NC muon neutrino events. In the exclusive sample, contamination due to NC muon interactions was a large source of background, but CCnQE electron events were the largest source of inefficiency.




Figure 6.11: Showing the efficiency and purity of the inclusive and exclusive analyses after each cut is applied.





Figure 6.12: Showing the efficiency and purity of the inclusive and exclusive analyses as a function of the neutrino momentum.

6.4.1 Systematic Errors

Neutrino energies used in long baseline neutrino experiments are chosen to complement the baseline in order to achieve an $\frac{L}{E}$ ratio close to the oscillation maximum. This means using neutrino energies of around 1 GeV, as baselines are usually hundreds of km. As was discussed previously in Section 2.2.1, this is a transition region for neutrino interactions. Quasi-elastic processes dominate at energies below 1 GeV and DIS processes dominate at energies above a few GeV. In between these two regimes, single pion interactions also contribute. Despite being essential for oscillation experiments, there have been relatively few experiments measuring cross sections in this region. Until recently, much of the data were from bubble chamber experiments, although K2K, MiniBooNE and SciBooNE have now made measurements [101]. Large theoretical uncertainties also exist due to the difficulties of QCD calculations at low energies. Uncertainties on the cross sections used by the neutrino generators are assumed to be the dominant source of systematic error in this study. Another potential source of error not considered here is the uncertainty in the beam flux.

A review of the current knowledge of neutrino interaction cross sections and their agreement with Monte Carlo models was carried out by Zeller [14]. In general, the charged current cross sections are better constrained than the NC cross sections. Figure 2.3 shows the state of knowledge of CC cross sections. The majority of the data in the Figure was from bubble chamber experiments such as ANL (Argonne National Laboratory) [16](CCQE) [102][103][104] (CC π^+) and BNL (Brookhaven National Laboratory) [17](CCQE) [105] (CC π^+). From this, a reasonable estimate of the uncertainty of the CC cross section is in the region of 30%. This uncertainty falls at higher energies where the DIS

regime has been well probed by experiments [14]. This regime is not applicable to T2K where neutrinos with energies of over 10 GeV do not form a significant fraction of the beam as shown in Figure 3.2.

In the neutral current case, the situation was found to be significantly worse. Again, almost all of the data acquired so far comes from bubble chamber experiments. Data pertaining to cross sections for neutral current pion production processes are very sparse [15]. In one channel shown in [15], Monte Carlo simulations fail to fit to the single experimental result. Due to this lack of knowledge, a conservative estimate of 100% was placed on NC interactions. To determine the systematic error on the modelled number of events, the signal and background samples were split into CC and NC components and the error of each group was found. CC and NC components were then combined by adding the errors in quadrature. Table 6.5 shows the systematic error for the inclusive analysis whilst Table 6.6 shows the systematic error for the exclusive analysis. The importance of removing NC events is shown by the large systematic errors attached.

Inclusive Analysis	CC Events	Syst Error	NC Events	Syst Error
Signal	282	85	0	0
Background	41	12	112	112

Table 6.5: Systematic errors associated with the inclusive CC ν_e analysis. Systematic uncertainties are dominated by lack of knowledge of NC cross-sections.

Exclusive Analysis	CC Events	Syst Error	NC Events	Syst Error
Signal	101	30	0	0
Background	51	15	51	51

Table 6.6: Systematic errors associated with the exclusive CC ν_e analysis.

1 Year Running		
	Inclusive	Exclusive
Signal	$92 \pm 28 \pm 10$	$33 \pm 10 \pm 6$
Background	$50 \pm 36 \pm 7$	$33 \pm 17 \pm 6$
5 Year Running		
	Inclusive	Exclusive
Signal	$459 \pm 138 \pm 21$	$165 \pm 49 \pm 13$
Background	$249 \pm 182 \pm 16$	$165 \pm 86 \pm 13$

Table 6.7: Predicted electron neutrino analysis event rates for one and five years of nominal running. The first error quoted is systematic and the second is statistical.

6.5 Results

An inclusive analysis is the more important analysis in the short term for T2K because it accumulates statistics more rapidly. The poorer energy resolution prevents a precise measurement of the energy spectrum, but a simple event counting experiment is possible. In the longer term, when more statistics have been accumulated, then a well reconstructed CCQE analysis can be used to measure the energy spectrum. In a year of nominal running, T2K aims to achieve 1×10^{21} POT, although the first year will be lower while the experiment is commissioned and the power of the beam is being incrementally increased. The data set used for the Monte Carlo study presented here represents 3.069×10^{21} POT. After combining errors, the inclusive analysis was predicted to find $282 \pm 85(sys) \pm 17(stat)$ signal events and $153 \pm 112(sys) \pm 12(stat)$ background events. In the exclusive analysis $101 \pm 30(sys) \pm 10(stat)$ signal events and $102 \pm 53(sys) \pm 10(stat)$ background events were expected. Expected events for one and five years of nominal running are shown in Table 6.7.

6.6 Conclusions

Combined with the TPC PID algorithm, the ECal reconstruction algorithms described in this thesis are capable of performing both inclusive and exclusive electron neutrino analyses. After one year of nominal running the signal for the inclusive analysis is $92 \pm 28 \pm 10$ events, almost a factor of two higher than the background. After five years of running both analyses would be possible, although the signal to noise ratio for the exclusive analysis remained worse than 1:1. The analysis could be improved if the performance of the PID algorithms could be improved below 200 MeV. In many cases this was because the reconstruction software had failed to reconstruct the event.

The presented analysis could be improved by using kinematic information. Only topological information and PID cuts were used here due to limitations of the reconstruction software. A global reconstruction algorithm, incorporating all the sub-detectors, would be able to more accurately veto the high momenta events responsible for a significant fraction of the predicted background. This could, to a limited extent, have been achieved with cuts on the allowed number of clusters; but in general more than one interaction is possible within a spill, so cutting on clusters not known to have come from the vertex within the FGD is undesirable.

Chapter 7

Conclusions

In this thesis a particle identification algorithm has been designed for the ND280 ECal. In particular, it has been developed to separate electrons and muons as part of an electron neutrino analysis. The Monte Carlo simulation used to design the ECal was then compared to data collected using the DsECal, with both cosmic muons and electrons from the T9 testbeam. Finally, the algorithm was tested as part of an electron neutrino analysis using the full tracker section of ND280.

A neural network based particle identification algorithm has been designed for the electromagnetic calorimeter of the ND280 detector. The neural network was compared to one dimensional likelihoods and multi-dimensional likelihoods. It was found to have far superior performance to a combination of one dimensional likelihoods. A comparable performance was found between the neural network and the multi dimensional likelihoods but the neural network was found to be a much faster algorithm in the reconstruction algorithm, although it was slower to train. The network was capable of an electron acceptance of 80% with a muon contamination of 10^{-2} .

Cosmic ray data was collected for the DsECal and used to calibrate the threshold and energy scale of the detector for the data collected in the T9 testbeam at CERN. Data from cosmic rays was also used to test the simulation of the DsECal simulation for muons. A good agreement was observed between data and simulation for MIP like particles. A high purity sample of electron events was created from data collected in the T9 testbeam. The data and simulation did not agree as closely for electron data as for cosmic data, although PID variable distributions were comparable. A deeper understanding of the EM energy scale and in particular the modelling of threshold is required to further improve the agreement. The electromagnetic shower scale of the detector was estimated at 9% larger than the MIP scale. The energy resolution of the detector was determined to be approximately $\frac{9\%}{\sqrt{E}}$, close to the design resolution of $\frac{7\%}{\sqrt{E}}$.

When applied to an electron neutrino analysis, the ECal PID in combination with the TPC PID achieved a signal to noise ratio of better than 1:1 in both inclusive and exclusive channels. To achieve this, strong cuts were placed on the data which resulted in a very low efficiency on the scale of a few percent. This analysis relied on the accuracy of the Monte Carlo simulation of ND280. Analysis of the data from the T9 testbeam indicated that muons were being quite accurately simulated. Electrons were not as well simulated. Although most of the discrepancies arose from small charges on the edge of showers and most distributions broadly agreed. The advantage of the analysis presented was that it used simple cuts and PID output so it can be used on data from the first day of data taking, before more complex reconstruction parameters are available. After one year of running the number of events in an inclusive electron neutrino analysis was predicted to be $92 \pm 28(sys) \pm 10(stat)$ with a background of $50 \pm 36(sys) \pm 7(stat)$. At the time of writing, the DsECal

collected its first data from neutrino events in the ND280, as shown in Figure 7.1.




Figure 7.1: Showing one of the first candidate neutrino events in ND280.

Bibliography

- [1] W. Pauli, Letter to Lise Meitner and Tübingen conference on radioactivity, <http://cdsweb.cern.ch/record/83282>, 1930.
- [2] A. Franklin, *Are there really neutrinos? An evidential history* (Perseus, 2000).
- [3] F. Abe *et al.*, Phys. Rev. Lett. **74**, 2626 (1995).
- [4] C. Giunti and C. W. Kim, *Fundamentals of Neutrino Physics and Astrophysics*, First ed. (Oxford, Oxford UK, 2008).
- [5] W. Yao *et al.*, Journal of Physics G **33**, 1 (2006).
- [6] K. Zuber, ACTAPHYS.POLON.B **37**, 1905 (2006).
- [7] L. Wolfenstein, Phys. Rev. D **17**, 2369 (1978).
- [8] S. Mikheev and A. Smirnov, Sov. J. Nucl. Phys. **42**, 913 (1985).
- [9] A. Aguilar *et al.*, Phys. Rev. D **64**, 112007 (2001).
- [10] K2K collaboration, Physical Review Letters **96**, 181801 (2006).
- [11] T. Araki *et al.*, Phys. Rev. Lett. **94**, 081801 (2005).

- [12] ALEPH, DELPHI, L3, OPAL and SLD Collaborations, Phys. Reports. **427**, 4257 (2006).
- [13] U. Dore and D. Orestano, REPT.PROG.PHYS **106201** (2008).
- [14] G. Zeller, Low energy neutrino cross sections: Comparison of various monte carlo predictions to experimental data, Preprint Archive, 2003, arXiv:hep-ex/0312.0611.
- [15] Krenz *et al.*, Nucl. Phys. **B135**, 45 (1978).
- [16] ANL Collaboration, Phys. Rev. **D16**, 3103 (1977).
- [17] BNL Collaboration, Phys. Rev. **D23**, 2499 (1981).
- [18] ND280 Collaboration, Document for ND280 review, Internal Document, 2006.
- [19] C. Amsler *et al.*, Review of Particle Physics **B667**, 1 (2008).
- [20] F. Reines and C. Cowan, Phys. Rev. **92**, 830 (1953).
- [21] F. Reines and C. Cowan, Phys. Rev. **113**, 273 (1959).
- [22] F. Reines and C. Cowan, Phys. Rev. **117**, 159 (1960).
- [23] M. Gell-Mann and A. Pais, Phys. Rev. **97**, 1387 (1955).
- [24] A. Pais and O. Piccioni, Phys. Rev. **100**, 1487 (1955).
- [25] B. Pontecorvo, J. Exptl. Theoret. Phys. **34**, 247 (1958).
- [26] B. Pontecorvo, J. Exptl. Theoret. Phys. **33**, 559 (1957).
- [27] B. Pontecorvo, Zh. Eksp. Teor. Fiz. **53**, 1717 (1967).

- [28] B. Pontecorvo, Phys. Lett. **28B**, 493 (1969).
- [29] G. Danby *et al.*, Phys. Rev. Lett. **9**, 36 (1962).
- [30] Z. Maki, M. Nakagawa, and S. Sakata, Prog. Theo. Phys. **28**, 870 (1962).
- [31] M. Perl *et al.*, Phys. Rev. Lett. **35**, 1489 (1975).
- [32] K. Kodama *et al.*, Phys. Lett. **B504**, 218 (2001).
- [33] J. Bahcall, Phys. Rev. Lett. **12**, 300 (1964).
- [34] R. Davis *et al.*, Phys. Rev. Lett. **20**, 1205 (1968).
- [35] W. Hempel *et al.*, Phys. Lett. B **447**, 127 (1999).
- [36] M. Altmann *et al.*, Phys. Lett. B **616**, 174 (2005).
- [37] J. Abdurashitov *et al.*, Phys. Rev. C **60**, 55081 (1999).
- [38] K. Hirata and others., Phys. Rev. Lett. **63**, 16 (1989).
- [39] Y. Fukuda *et al.*, Phys. Rev. Lett. **81**, 1562 (1998).
- [40] D. Casper *et al.*, Phys. Rev. Lett. **66**, 2561 (1991).
- [41] C. Berger *et al.*, Phys. Lett. **B227**, 489 (1989).
- [42] M. Aglietta *et al.*, Europhys. Lett. **8**, 611 (1989).
- [43] C. Yanagisawa, Report on the super-kamiokande accident II,
http://superk.physics.sunysb.edu/nngroup/publication_theses/SKaccident02.ps,
 2002.
- [44] MACRO Collaboration, Eur. Phys. J. **36**, 323 (2004).
- [45] W. W. M. Allison *et al.*, Physics Letters B **391**, 491 (1997).

- [46] J. Boger *et al.*, Nucl. Instrum. Meth. **A449** (2000).
- [47] Q. Ahmed *et al.*, Phys. Rev. Lett. **87** (2001).
- [48] S. Ahmed *et al.*, Phys. Rev. Lett. **92** (2004).
- [49] B. Aharmin *et al.*, Phys. Rev. **C72** (2005).
- [50] Y. Declais *et al.*, Nucl. Phys. **B434**, 503 (1995).
- [51] F. Boehm *et al.*, Phys. Rev. D. **64**, 112011 (2001).
- [52] M. Apollonio *et al.*, Europhys. Lett. J. **C27**, 331 (2003).
- [53] K. Eguchi *et al.*, Phys. Rev. Lett. **90** (2003).
- [54] S. Abe *et al.*, Phys. Rev. Lett. **100**, 221803 (2008).
- [55] Y. Ashie *et al.*, Phys. Rev. D. **71**, 112005 (2005).
- [56] M. Kordosky *et al.*, Physical Review D **77**, 072002 (2008).
- [57] NOMAD Collaboration: P. Astier *et al.*, Physics Letters B **570**, 19 (2003).
- [58] E. D. Church *et al.*, Physical Review D **66**, 013001 (2002).
- [59] The MiniBooNE Collaboration, Physical Review Letters **98**, 231801 (2007).
- [60] A. A. Aguilar-Arevalo *et al.*, A search for electron antineutrino appearance at the $\delta m^2 \simeq 1\text{eV}^2$ scale, Preprint Archive, 2009, arXiv:hep-ph/0904.1958.
- [61] T. Schwetz, M. Tortola, and J. Valle, NEWJ.PHYS **113011** (2008).

- [62] T. Ishia, Proceedings of the Particle Accelerator Conference PAC07 (June 2007).
- [63] Accelerator Group, JAERI/KEK Joint Project Team, Accelerator technical design report for j-parc, Internal document.
- [64] I. Taylor, *Development of T2K 280m Near Detector Software for Muon and Photon Reconstruction*, PhD thesis, Imperial College, London UK, 2008.
- [65] Y. Fukuda *et al.*, Nucl. Inst. Meth. **A501**, 418 (2003).
- [66] Super-Kamiokande collaboration, Super-Kamiokande diagram, <http://neutrino.kek.jp/Images/Figures>.
- [67] D. Green, *The Physics of Particle Detectors* (Cambridge, 2000).
- [68] T. Barszczak and the Super-Kamiokande collaboration, Images of super-kamiokande events, <http://www.ps.uci.edu/tomba/sk/tscan/pictures.html>.
- [69] G. Davies, Private Communication, 2009.
- [70] M. Yokoyama *et al.*, Mass production test of hamamatsu MPPC for T2K neutrino oscillation experiment, Preprint Archive, 2008, arXiv:physics/0312.0611.
- [71] B. Bondarenko *et al.*, Nucl. Inst. Meth. A **442**, 187 (2000).
- [72] O. Mineev *et al.*, Scintillator counters with multi-pixel avalanche photodiode readout for the ND280 detector of the T2K experiment, Preprint Archive, 2006, arxiv:physics/0606037.

- [73] T. Estrada *et al.*, Mcm II and the trip chip, 2002, D0 Note 4009.
- [74] S. Ritt *et al.*, Midas data acquisition, <http://midas.psi.ch/html/doc/index.html>, retrived 10/06/2010.
- [75] R. Brun and F. Rademakers, Nucl. Inst. Meth. in Phys. Res. A **389**, 81 (1997).
- [76] S. Agostinelli *et al.*, Nucl. Inst. Meth. A **506**, 250 (2003).
- [77] C. Andreopoulos *et al.*, The GENIE neutrino monte carlo generator, Preprint Archive, 2009, arXiv:hep-ph/0905.2517.
- [78] Y. Hayato, Nucl. Phys. Proc. Suppl. **112**, 183 (2002).
- [79] A. Cervera-Villanueva, J. Gomez-Cadenas, and J. Hernando, Nucl. Inst. Meth. A. **534**, 180 (2004).
- [80] R. Wigmans, *Calorimetry Energy Measurement in Particle Physics* (Oxford, 2000).
- [81] C. Bishop, *Neural Networks for Pattern Recognition* (Clarendon Press., 1997).
- [82] R. Lippmann, IEEE ASSP Magazine **April**, 4 (1987).
- [83] M. Richard and R. Lippmann, Neural Computation **3**, 461 (1991).
- [84] F. James and M. Roos, Comput. Phys. Commun. **10**, 343 (1975).
- [85] T. Carli and B. Koblitz, Nucl. Instrum. Meth. **A501**, 5756 (2002).
- [86] A. Hocker *et al.*, TMVA Toolkit for Mulitvariate Data Analysis with ROOT, <http://tmva.sf.net>, 2007.

- [87] D. Simon, Secondary beams for tests in the PS east experimental area, http://cern-accelerators-optics.web.cern.ch/cern-accelerators-optics/FT61_DIRAC/PS_PA_Note_93_21.pdf, 1993.
- [88] K. Nagamine, *Introductory Muon Science*, First ed. (Cambridge, Cambridge UK, 2003).
- [89] M. Kordosky, *Hadronic Interactions in the MINOS Detectors*, PhD thesis, University of Austin, Texas US, 2004.
- [90] A. Waldron, Private Communication, 2009.
- [91] D. Heck, CORSIKA an air shower simulation program, <http://www-ik.fzk.de/corsika/>, 2009.
- [92] D. Heck *et al.*, Forschungszentrum Karlsruhe Report FZKA 6019 (1998).
- [93] M. George, Cosmic flux simulation files for the DsECal cosmics tests, Internal T2K Document.
- [94] R. Fruehwirth, Convolutd landau and gaussian fitting function, <http://root.cern.ch/root/html/tutorials/fit/langaus.C.html>.
- [95] N. McCauley, Mock data challenge 0, Internal T2K Document.
- [96] T2K Collaboration, Nominal Event Rates In ND280, Internal T2K Document.
- [97] C. Giganti, Tracker nue analysis, Internal T2K Document.
- [98] M. George, Nue analysis framework, Internal T2K Document.
- [99] C. Giganti, Particle identification with the T2K TPC, 2009.
- [100] M. Murdoch, Private Communication, 2009.

- [101] G. Zeller, J. Phys. Conf. Ser. **136**, 022028 (2008).
- [102] ANL Collaboration, Phys. Rev. **D19**, 2521 (1979).
- [103] ANL Collaboration, Phys. Rev. **D25**, 1161 (1982).
- [104] Campbell *et al.*, Phys. Rev. Lett. **30**, 335 (1973).
- [105] BNL Collaboration, Phys. Rev. **D34**, 2554 (1986).