

Imperial College London
Department of Physics

**Reconstruction of neutrino
induced neutral current neutral
pion events with the T2K ND280
Tracker and ECAL**

Pawel Guzowski

September 2011

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Physics of Imperial College London
and the Diploma of Imperial College London

Declaration

I herewith certify that all material in this dissertation which is not my own work has been properly acknowledged.

Pawel Guzowski

Abstract

T2K is a long baseline neutrino oscillation experiment in Japan. One of the goals is to search for electron neutrino appearance at the far detector (Super Kamiokande) in a muon neutrino beam produced 295 km away, to make a measurement of the unknown neutrino oscillation angle θ_{13} . A major background to this appearance search is neutral pions (π^0) produced in neutral current (NC) neutrino interactions. The π^0 s decay to photons, which can in cases mimic the electron appearance signal.

The near detector of the experiment (ND280) has the capability of measuring the rates of NC π^0 production processes. NC π^0 interactions in the Tracker of the detector (plastic scintillator targets, and time projection chambers for tracking charged particles) can be detected through photon conversions in the electromagnetic calorimeter that surrounds the Tracker. For these types of events, the photon reconstruction in the calorimeter has to have good energy and angular resolutions.

This thesis describes the T2K experiment, the ND280 detector and its offline software, and gives details on the ECal reconstruction of photons. The performance of the reconstruction in relation to the π^0 decay photons, their energy and angular resolutions, is studied.

An analysis is presented for the reconstruction of NC π^0 s produced in the trackers, with both photons converting in the downstream calorimeter,

developed purely with Monte Carlo (with GENIE as the interaction generator), and then applied to the first-year data from ND280. The Monte Carlo expectation is $17.7 \pm 1.1(\text{stat}) \pm 4.9(\text{syst})$ events, with $37 \pm 4\%$ $\text{NC}\pi^0$ purity; 21 events are seen in the first year data. The MC is consistent with data at the 0.5σ level.

As a crosscheck, the same analysis is applied to a different Monte Carlo production (using NEUT as the generator). The expectation from this production is $16.8 \pm 1.0(\text{stat}) \pm 4.7(\text{syst})$ events, with $30 \pm 4\%$ purity. This is consistent with data at the 0.7σ level.

Dedicated to our hosts in Tokai village, for their
hospitality in providing us with shelter, warmth, food and
water, in the aftermath of the Great Tohoku Earthquake
of March 2011

Contents

1. Introduction	15
2. Background to the T2K experiment	17
2.1. A history of neutrino oscillation physics	17
2.2. Neutrino oscillation	23
2.3. Neutrino-nucleon interactions	29
2.3.1. Neutral pion production	31
3. The T2K Experiment	34
3.1. Accelerator and neutrino beam line	35
3.2. Near detectors	39
3.2.1. INGRID	41
3.2.2. Magnet and SMRD	44
3.2.3. Tracker	46
3.2.4. Tracker Electromagnetic Calorimeters	49
3.2.5. P0D	52
3.2.6. Electronics	54
3.3. Far detector	58
4. ND280 offline software	63
4.1. Overall structure	63

4.2.	Monte Carlo simulation	65
4.3.	Calibration	67
4.3.1.	TFB calibration	67
4.3.2.	TPC calibration	69
4.3.3.	FGD calibration	70
4.4.	Reconstruction	70
4.4.1.	Tracker reconstruction	70
4.4.2.	Global reconstruction	72
4.5.	Analysis-ready summarisation	73
4.5.1.	Beam summary data	74
5.	Reconstruction performance of the Downstream Electro-	
	magnetic Calorimeter	75
5.1.	ECal reconstruction	76
5.1.1.	Hit bunching	76
5.1.2.	Clustering	76
5.1.3.	3D matching	78
5.1.4.	Angular reconstruction	80
5.1.5.	EM energy reconstruction	82
5.1.6.	Particle identification	84
5.1.7.	Michel tagging	86
5.2.	ECal Performance	87
5.2.1.	Monte Carlo	88
5.2.2.	Test beam	94
5.2.3.	Summary	97

6. Reconstruction of neutral pions in the Downstream Calorimeter	98
6.1. Data set	99
6.1.1. Data quality	101
6.1.2. Monte Carlo	101
6.2. Bunch timings	102
6.3. Preselection	106
6.3.1. Geometrical acceptance and reconstruction efficiency .	106
6.3.2. Preselection purity	110
6.3.3. Data excess	112
6.4. Selection cuts	117
6.5. Vertexing	126
6.5.1. Global vertex sample	127
6.5.2. FGD vertex sample	129
6.5.3. Cluster vertex sample	129
6.5.4. Resolution	136
6.6. Systematic errors	137
6.7. Selection summary	137
6.8. Crosscheck with NEUT	141
6.9. Summary and remarks	145
7. Conclusions	147
A. Distributions of likelihood input variables in data	148
Bibliography	158

List of Tables

2.1. Neutrino oscillation parameter measurements	26
6.1. Bunch timings	103
6.2. Event composition after preselection	111
6.3. Event composition after cluster-level likelihood selection . . .	123
6.4. Reconstructed π^0 mass and momentum resolutions and biases.	134
6.5. Final selection efficiencies and purities	138
6.6. Event composition in the final selection	139
6.7. Selection efficiencies and purities in the NEUT production. .	143
6.8. Event composition after preselection, after the likelihood cut, and in the final selection, for the NEUT production.	144

List of Figures

2.1. Feynman diagrams of coherent forward scattering processes.	28
2.2. Charged current cross section measurements	29
2.3. Feynman diagrams for various neutrino interaction processes.	30
2.4. $\text{NC}\pi^0$ production cross-section measurements	32
3.1. Baseline for the T2K experiment	35
3.2. Aerial view of J-PARC	36
3.3. Neutrino beamline at J-PARC	36
3.4. Predicted neutrino fluxes at Super Kamiokande	38
3.5. ν_μ flux uncertainty at ND280	39
3.6. Exploded view of the ND280 subdetectors	40
3.7. Exploded view of an INGRID module	42
3.8. Front view of the INGRID assembly	43
3.9. Cross section of a magnet yoke C-section	44
3.10. Photograph of an SMRD paddle	46
3.11. Drawing of the TPC design	47
3.12. Photograph of the MPPC photon detection area	54
3.13. Photoelectron spectrum of a MPPC device	55
3.14. Timing structure of MPPC noise digits with TFB readout.	56
3.15. Diagram of Super Kamiokande	59

3.16. Displays of Super Kamiokande events	60
3.17. NC1 π^0 reconstruction efficiency for K2K	61
3.18. Reconstructed neutrino energy spectrum for a Monte Carlo study of a ν_e appearance measurement	62
4.1. TPC truncated mean charge against momentum	72
5.1. Thrust for Monte Carlo particles	83
5.2. Track/shower discriminant for Monte Carlo particles	86
5.3. Monte Carlo π^0 decay photon energies	88
5.4. Reconstruction efficiency of Monte Carlo photons.	89
5.5. Reconstructed fractional energy bias of Monte Carlo photons in the DsECal.	90
5.6. Example energy slices from figure 5.5.	91
5.7. Energy resolution of Monte Carlo photons	92
5.8. Reconstructed angle deviations of Monte Carlo photons.	93
5.9. Example energy slices from figure 5.8	93
5.10. Angular resolution of Monte Carlo photons	94
5.11. Fractional energy resolution of test beam and MC electrons	95
5.12. Reconstructed energy bias of test beam electrons	96
6.1. Monte Carlo π^0 decay photon energies, for those π^0 s where both photons convert in the Downstream ECal.	99
6.2. Accumulated protons and protons per spill	100
6.3. DsECal cluster times	103
6.4. Daily variation of the first bunch cluster time and width	105
6.5. Raw data distributions for the preselected clusters	107
6.6. DsECal acceptance for π^0 s produced in the FGDs	108
6.7. π^0 reconstruction efficiency	109

6.8. True vertex distributions	113
6.9. Data excess distributions	115
6.10. Data excess distributions	116
6.11. Input variables for the likelihood	120
6.12. Linear correlation coefficients for the likelihood input variables	122
6.13. Likelihood output	122
6.14. Likelihood output for data and Monte Carlo.	123
6.15. Data excess distributions after likelihood cut	124
6.16. True vertex distributions after likelihood cut	125
6.17. Vertex resolution and angular deviation for the Global vertex sample	127
6.18. TPC PID	128
6.19. Global-based π^0 mass	129
6.20. Vertex resolution and angular deviation for the FGD vertex sample	130
6.21. FGD-based π^0 mass	130
6.22. Unconstrained vertex resolution for the Cluster vertex sample	131
6.23. Constrained vertex resolution and angular deviation for the Cluster vertex sample	132
6.25. Cluster-based π^0 mass	133
6.26. Mass resolution cross-check	134
6.27. Reconstructed π^0 mass and momentum biases	135
6.28. Mass spectrum with perfect vertex resolution	136
6.29. Reconstructed π^0 mass	139
6.30. Reconstructed π^0 momentum	140
6.31. Reconstructed π^0 angle	140
6.32. Likelihood output for the NEUT production	142

6.33. Reconstructed π^0 mass, NEUT	142
6.34. Reconstructed π^0 momentum, NEUT	144
6.35. Reconstructed π^0 angle, NEUT	145
A.1. Legend for the plots in the appendix	148
A.2. High track/shower discriminant	149
A.3. Low track/shower discriminant	149
A.4. High charge skew	149
A.5. Low charge skew	150
A.6. High EM likelihood	150
A.7. Low EM likelihood	150
A.8. High pointing	151
A.9. High first layer	151
A.10.Low first layer	151
A.11.High number of Michel tags	152
A.12.High incidence angle	152
A.13.Low incidence angle	152
A.14.High cone angle	153
A.15.Low cone angle	153
A.16.High mean position	153
A.17.Low mean position	154
A.18.High number of hits	154
A.19.Low number of hits	154
A.20.High fiduciality	155
A.21.Low fiduciality	155
A.22.High thrust	156
A.23.Low thrust	156

A.24. Energy asymmetry	156
A.25. Angle between the two clusters	157
A.26. Invariant mass	157
A.27. Distance between the two clusters	157

1. Introduction

This thesis presents an analysis of first year data taking with the near detector (ND280) of T2K, attempting to measure the rate of neutral current induced neutral pion production. Knowledge of this process is important for T2K as a whole, as it is one of the main backgrounds to an electron neutrino appearance search. The ND280 detector has a dedicated sub-detector to perform this measurement: the π^0 detector, or P0D. The other main subdetector of ND280 is the Tracker, which has capabilities of tracking and making momentum measurements of charged particles produced in neutrino interactions within it, and it is optimised for measurements of charged current quasi-elastic interactions. The Tracker is surrounded by electromagnetic calorimeter modules (the ECal). These are capable of detecting photons produced in the Tracker, and on that basis, π^0 reconstruction can be achieved. There is an order of magnitude less statistics for these interactions in the Tracker compared to the P0D; however, with the tracking capabilities, the events are cleaner, and exclusive measurements can be made. As the subdetectors use different technologies, the systematic effects are different, and the measurements in the two subdetectors can be crosschecked with one another.

Chapter 2 presents a historical perspective on the T2K experiment, the theoretical background to neutrino oscillations, and an introduction to the

important neutrino-nucleon interactions, including neutral pion production.

Chapter 3 is an overview of the experimental set up.

Chapter 4 describes the software used to perform the calibration and reconstruction of ND280 data. The author of this thesis has been active in development of the software, however this chapter is mostly a description of work that other members of the T2K collaboration have done, necessary for understanding the analysis performed in chapter 6.

Chapter 5 gives specifics on the reconstruction of photons in the calorimeters. The author has been involved with development of some of the reconstruction algorithms (reducing noise hits in clustering, improving the minimum hit threshold of low energy clusters, and implementing the thrust direction algorithm). The remainder of the algorithms have been implemented by various members of the collaboration, and are described here as their outputs are used in the analysis of chapter 6. The Monte Carlo study of the reconstruction performance is the author's work. The test beam analysis has been performed by other members of the collaboration, although the author was involved in test beam commissioning and data taking at CERN.

Chapter 6 describes reconstruction of neutral pions, produced in the Tracker, with the downstream calorimeter, and the analysis to improve the purity of the reconstructed sample. This analysis is the author's work.

Chapter 7 is the conclusion to this thesis.

2. Background to the T2K experiment

2.1. A history of neutrino oscillation physics

The existence of the neutrino was first postulated by Pauli in 1930 as a solution to the observed [1, 2] continuous β decay electron energy spectrum; if it is a two-body decay ($N(Z, A) \rightarrow e^- N'(Z + 1, A)$), then the decay electron should be monoenergetic. Pauli suggested [3] that the process was a three-body decay, and that a neutral particle of small mass (with respect to the proton) was carrying away momentum, and that this particle had to be of half-integer spin, for the conservation of spin in free neutron decay to remain valid. The neutrino was incorporated into Fermi's theory of β decay [4] in 1934.

Direct observation of the neutrino took place in the experiment of Reines and Cowan [5, 6] in 1956. Neutrinos were detected through the inverse β decay $\bar{\nu}_e p \rightarrow e^+ n$. A nuclear reactor provided the antineutrinos, and the detector target was a 7.5 cm thick layer of water doped with cadmium chloride, with two layers of liquid scintillator on either side of the water layer, to detect photons. The experiment relied on a coincidence between the positron annihilation photons, both 511 keV and back-to-back, and neutron capture on cadmium releasing photons totalling 9 MeV. The neutron capture

photons were delayed by around $5\ \mu\text{s}$ in relation to the positron annihilation photons, due to the neutron diffusion timescale. Reines and Cowan observed the expected rate of these events, which correlated with the reactor power.

Concurrently with Reines and Cowan, Davis was searching [7] for neutrinos through the inverse β decay $\nu_e\ ^{37}\text{Cl} \rightarrow e\ ^{37}\text{Ar}$, which had a threshold neutrino energy of 5.1 MeV. The experimental set up was two vessels containing 200 and 39001 of carbon tetrachloride. These were irradiated with antineutrinos from reactors at Brookhaven, and any argon gas produced was extracted. ^{37}Ar decays through electron capture with a half-life of 35 days, the ion releasing an X-ray photon in the process, which the experiment was designed to detect. Davis found no evidence of ^{37}Ar production through irradiation of ^{37}Cl with antineutrinos, and concluded that the neutrino was not identical to the antineutrino. However, the same experimental set up was also putting an upper limit on the rate of electron neutrinos produced in the Sun.

The identity of muon neutrinos as separate from electron neutrinos was observed in 1963 by Lederman, Shwartz and Steinberger [8]. They produced the first (deliberate) accelerator neutrino beam; 15 GeV protons colliding on a fixed beryllium target produced charged pions, which decayed to muons and muon neutrinos.¹ A spark chamber located 21 m from the target and behind 13.5 m of iron shielding detected the appearance of muons in coincidence with the accelerator protons. They observed the expected amount of muons, given the hypothesis that the neutrinos produced in the pion decay were only coupled to muons, and hence had a flavour identity different to neutrinos that coupled with electrons, like those produced in reactors.

In 1989 the SLC experiment at SLAC and the LEP experiments at CERN

¹There was some background from kaon decays, and decays to electron neutrinos.

deduced [9, 10, 11, 12, 13], from measurements of the width of the Z boson, that the number (N_ν) of light neutrino families ($M_\nu < M_Z/2$) was consistent with 3, while ruling out $N_\nu \leq 2$ and $N_\nu \geq 4$. The direct observation of the third light neutrino, the tau neutrino, was announced in 2000 by the DONUT collaboration [14].

Meanwhile, Davis was continuing his chlorine-based experiments to measure solar neutrinos, and by 1968 had expanded his experiment. It was now conducted in a 390,000 l vessel, 1.4 km under ground (to reduce cosmic backgrounds) at the Homestake mine. The upper limit on the solar electron neutrino flux was found [15] to be less than a seventh of the predicted flux. This discrepancy was termed the ‘solar neutrino anomaly’. The final result [16] of the Homestake measurement was a flux around a third of the predicted flux. Later experiments such as GALLEX [17] and SAGE [18] using gallium instead of chlorine (with a lower threshold neutrino energy of 0.2 MeV, hence sensitive to neutrinos produced by more theoretically precise reactions in the Sun). These experiments still measured an anomaly, although not as great as the chlorine ones.

In 1986 the IMB experiment, set up to search for proton decay, reported its findings [19] on the rate of atmospheric (cosmic-ray induced) muon neutrino interactions. The IMB detector was a $17 \times 17 \times 23 \text{ m}^3$ volume of purified water. Mounted on the walls of the volume were photomultiplier tubes (PMTs), facing inwards, which had wavelength shifting plates attached to increase light collection. One of the backgrounds to proton decay was atmospheric neutrinos. The neutrinos interacted in the water, creating charged leptons, and the PMTs picked up the Čerenkov light given off by these particles travelling through the water. IMB found that the flux of muon neutrinos was around 60% of the expected value, which was dubbed

the ‘atmospheric neutrino anomaly’.

At the same time, the similar Kamiokande experiment was running. This too was designed to detect proton decay, with a 16 m diameter, 20 m tall cylindrical tank, and surrounded by larger PMTs than the IMB’s. These allowed the experiment to utilise Čerenkov ring shape discrimination between muons and electrons. In 1988 it reported [20] that it too had measured only 60% of the expected muon neutrino flux. By 1994 the collaboration had analysed enough statistics for an angular measurement, showing [21] the anomaly was greater for upward-going neutrinos, those that had travelled through the earth over a longer distance. This provided evidence for neutrino oscillation being a cause of the anomaly, which as described in section 2.2 have a length dependence. The ratio of ν_μ to ν_e was also smaller than the expected 2:1. Kamiokande was also sensitive to solar neutrinos, and angular analysis allowed it to show [22] that the electron neutrinos were indeed coming from the Sun, but with a deficit from the expected amount of 0.46.

In 1996, Super-Kamiokande started taking data. This was essentially an upgrade of Kamiokande, with 22 times the fiducial mass. The tank was now 39 m diameter and 41 m tall, and separated into inner and outer detectors. The outer detector could be used to veto external interactions and cosmic muons. In 1998 the collaboration presented strong evidence [23] for atmospheric neutrino oscillation, with a zenith angle dependence of the ratio of ν_μ to ν_e to the expected 2:1 ratio. The data was in good agreement with a two-flavour $\nu_\mu \rightarrow \nu_\tau$ oscillation model with maximal mixing.

The SNO collaboration in 2002 announced [24] direct evidence of solar neutrino oscillation. The SNO detector was a 12 m diameter sphere of heavy water, surrounded by PMTs fixed to a 18 m diameter support struc-

ture. SNO was sensitive to charged current (CC) interactions ($\nu_e d \rightarrow e^- pp$), neutral current (NC) interactions ($\nu_l d \rightarrow \nu_l np$) and elastic scattering (ES) ($\nu_l e^- \rightarrow \nu_l e^-$) of the electron neutrinos. The charged current and elastic scattering interactions were detected through the Čerenkov rings produced by the electron, whereas the neutral current interactions were detected through the Čerenkov rings produced by photons released in neutron capture. The experiment found that the flux sum of all three neutrino types (available through the NC and ES channels) was well predicted by solar models, whereas the exclusive electron neutrino flux (from the CC channel) was significantly below expectation. This measurement provided evidence that electron neutrinos produced in the sun were transforming into other flavours.

Oscillations were also being searched for using electron antineutrinos produced at nuclear reactors. One example is the CHOOZ experiment [25] which ran from 1997 to 1998. CHOOZ was a 5 t target of paraffin-based liquid scintillator doped with gadolinium, surrounded by PMTs, and positioned 1 km from two reactors at the Chooz B power plant in France. Similarly to the experiment of Reines and Cowan, the signal was two 511 keV positron annihilation photons, and a delayed neutron capture on gadolinium releasing photons totalling 8 GeV. The experiment found no evidence of $\bar{\nu}_e$ disappearance [25] within their sensitivity, and concluded that the atmospheric neutrino anomaly was due to a $\nu_\mu \rightarrow \nu_\tau$ oscillation.

KamLAND was also a reactor oscillation experiment. A 1 kt liquid scintillator target in the Kamioka mine searched for $\bar{\nu}_e$ disappearance from many reactors of order 100 km away. Their results published in 2002 [26] showed a deficit in observed interactions, and more recent analyses [27] make a precise measurement of the solar δm^2 oscillation parameter. KamLAND was

also able to fit an sinusoid-like function to a measured L/E distribution, providing more evidence of oscillation.

Particle accelerators have also been used in oscillation experiments, as they provide an independent atmospheric-like neutrino source. They have the advantage over natural atmospheric neutrino experiments, that the baseline is well known, the neutrino energies have a narrow band, can be tuned, have a relatively well-modelled flux, and beam timing can be used to greatly reduce uncorrelated backgrounds.

K2K was the first long-baseline accelerator neutrino experiment. The beam was produced by the 12 GeV proton synchrotron at KEK (with average neutrino energy of 1.3 GeV), and the detector was Super Kamiokande, 250 km away. The experiment ran from 1999 to 2004. As the neutrino beams are tertiary (the protons are primary, the produced pions are secondary), accelerator experiments require a near detector to accurately characterise the beam. K2K's was approximately 300 m from the production target, and consisted of a water Čerenkov detector (11 m diameter, 11 m tall), a water target with scintillator layers for tracking ($2.6 \times 2.6 \times 2 \text{ m}^3$), a lead glass calorimeter (which was later replaced with a tracker, made of scintillator bars, of dimension $3 \times 3 \times 1.7 \text{ m}^3$) and a muon chamber with iron of total length 2 m. K2K's final result [28], measurements of the atmospheric oscillation parameters, was consistent with results from Super Kamiokande.

MINOS is another long-baseline accelerator experiment, running since 2006 and still running as of 2011. The neutrino beam (with a tunable average energy between 3–10 GeV) is produced by the Main Injector at Fermilab, and the far detector is in the Soudan mine, 700 km away. Both the near and far detectors are similar: layers of plastic scintillator with steel plates between them, and the whole detectors are magnetised. The

near detector is 14 m long with a 980 t mass, and the far detector is 29 m long with a 5.4 kt mass. In 2008 the collaboration reported [29] improved limits on the atmospheric mixing parameters over K2K's, although MINOS doesn't provide the best constraint on the mixing angle.

T2K is a second generation long-baseline experiment, essentially the successor of K2K. The far detector is again Super Kamiokande, but the beam is produced at J-PARC, approximately 50 km from KEK. More details about T2K are given in chapter 3.

2.2. Neutrino oscillation

Neutrino oscillations can arise if there is a mixing of neutrino flavour ($\nu_{e,\mu,\tau,l}$) and mass ($\nu_{1,2,3,i}$) eigenstates, which is described by the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix² as given by:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = U \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}, \quad \text{or} \quad |\nu_l\rangle = \sum_i U_{li} |\nu_i\rangle, \quad (2.1)$$

where U is the PMNS matrix [30, 31]. U is a unitary matrix, as it is a rotation from one basis to another. U is commonly factorised into three terms: a term on which atmospheric ($\nu_\mu \rightarrow \nu_\tau$) oscillation strongly depends on, a term on which solar ($\nu_e \rightarrow \nu_{\mu,\tau}$) oscillation strongly depends on, and

²Analogous to the CKM matrix in the quark sector.

a third, ‘cross-mixing’ term:

$$\begin{aligned}
U &= \overbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & C_{23} & S_{23} \\ 0 & -S_{23} & C_{23} \end{pmatrix}}^{\text{‘atmospheric’ term}} \overbrace{\begin{pmatrix} C_{13} & 0 & S_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -S_{13}e^{i\delta} & 0 & C_{13} \end{pmatrix}}^{\text{‘cross-mixing’ term}} \overbrace{\begin{pmatrix} C_{12} & S_{12} & 0 \\ -S_{12} & C_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}}^{\text{‘solar’ term}} \\
&= \begin{pmatrix} C_{12}C_{13} & S_{12}C_{13} & S_{13}e^{-i\delta} \\ -S_{12}C_{23} - C_{12}S_{23}S_{13}e^{i\delta} & C_{12}C_{23} - S_{12}S_{23}S_{13}e^{i\delta} & S_{23}C_{13} \\ S_{12}S_{23} - C_{12}C_{23}S_{13}e^{i\delta} & -C_{12}S_{23} - S_{12}C_{23}S_{13}e^{i\delta} & C_{23}C_{13} \end{pmatrix}
\end{aligned} \tag{2.2}$$

where $C_{ij} \equiv \cos(\theta_{ij})$, $S_{ij} \equiv \sin(\theta_{ij})$, θ_{ij} are mixing angles, and δ is a CP-violating phase.³ Atmospheric oscillation depends strongly on the angle θ_{23} , and solar oscillation on the angle θ_{12} .

The evolution of the state vector of a neutrino mass eigenstate (ν_i) with four-position $\mathbf{x} = (t, \vec{x})$ is given by:

$$|\nu_i(\mathbf{x})\rangle = e^{-i\mathbf{p}_i \cdot \mathbf{x}} |\nu_i\rangle, \tag{2.3}$$

where $\mathbf{p}_i = (E_i, \vec{p}_i)$ is its four-momentum, with $|\nu_i\rangle$ the state at the origin ($\mathbf{x}_0 = (t_0, \vec{x}_0)$; for simplicity, the neutrino creation point), and working in natural units ($\hbar = c = 1$). The probability of measuring the neutrino flavour

³A further term involving two more phases, unobservable in oscillation experiments, is not shown.

eigenstate l' at \mathbf{x} after creating the eigenstate l is then

$$\begin{aligned}
P(\nu_l \rightarrow \nu_{l'}, \mathbf{x}) &= |\langle \nu_{l'}(\mathbf{x}) | \nu_l(\mathbf{x}_0) \rangle|^2 \\
&= \left| \sum_{ij} \langle \nu_i | U_{l'i}^* e^{i\mathbf{p}_i \cdot \mathbf{x}} U_{lj} | \nu_j \rangle \right|^2 \\
&= \left| \sum_i e^{i\mathbf{p}_i \cdot \mathbf{x}} U_{l'i}^* U_{li} \right|^2 \\
&= \sum_{ij} e^{i(\mathbf{p}_i - \mathbf{p}_j) \cdot \mathbf{x}} U_{l'i}^* U_{li} U_{l'j} U_{lj}^* \\
&= \sum_i U_{l'i}^* U_{li} U_{l'i} U_{li}^* + 2 \sum_{j < i} |U_{l'i}^* U_{li} U_{l'j} U_{lj}^*| \cos((\mathbf{p}_i - \mathbf{p}_j) \cdot \mathbf{x} \\
&\quad + \arg(U_{l'i}^* U_{li} U_{l'j} U_{lj}^*)) \\
&= \delta_{l'l} - 4 \sum_{i,j < i} \sin^2((\mathbf{p}_i - \mathbf{p}_j) \cdot \mathbf{x}/2) \Re(U_{l'i}^* U_{li} U_{l'j} U_{lj}^*) \\
&\quad + 2 \sum_{i,j < i} \sin((\mathbf{p}_i - \mathbf{p}_j) \cdot \mathbf{x}) \Im(U_{l'i}^* U_{li} U_{l'j} U_{lj}^*), \tag{2.4}
\end{aligned}$$

Working in the lab frame, with the neutrino momentum p along the direction $(\vec{x} - \vec{x}_0)$, with $|\vec{x} - \vec{x}_0| = L$ and $(t - t_0) = T$, then

$$\begin{aligned}
(\mathbf{p}_i - \mathbf{p}_j) \cdot \mathbf{x} &= (E_i - E_j)T - (p_i - p_j)L \\
&= (E_i - E_j)T - \frac{p_i^2 - p_j^2}{p_i + p_j} L \\
&= (E_i - E_j) \left(T - \frac{E_i + E_j}{p_i + p_j} L \right) + \frac{m_i^2 - m_j^2}{p_i + p_j} L, \tag{2.5}
\end{aligned}$$

and under a relativistic neutrino approximation, where $E_i \approx E_j \approx p_i \approx p_j \approx E$, equation 2.5 becomes

$$(\mathbf{p}_i - \mathbf{p}_j) \cdot \mathbf{x} = \frac{\Delta m_{ij}^2 L}{2E}, \tag{2.6}$$

Parameter	Best value	Experiment(s)
$\sin^2(2\theta_{12})$	$0.86^{+0.03}_{-0.02}$	solar expts., KamLAND
$\sin^2(2\theta_{23})$	> 0.92 (90% CL)	K2K, MINOS
$\sin^2(2\theta_{13})$	< 0.15 (90% CL)	CHOOZ
Δm_{21}^2	$(7.59 \pm 0.21) \times 10^{-5}$	solar expts., KamLAND
$ \Delta m_{32}^2 $	$(2.43 \pm 0.13) \times 10^{-3}$	MINOS

Table 2.1.: Neutrino oscillation parameter measurements [32].

with $\Delta m_{ij}^2 \equiv (m_i^2 - m_j^2)$.

The only terms in equation 2.4 that depend on position and time are the ones involving Δm_{ij}^2 , and are multiplied by a factor involving the mixing angles θ_{ij} . This implies that oscillation is only possible if i) there is mixing between mass and flavour eigenstates of the neutrino and ii) the masses of the mass eigenstates are not all degenerate, implying that at least one of the neutrinos is massive. This parametrisation also implies oscillation is at a maximum if L and E are tuned for optimising measurements of the Δm^2 parameters, explaining why experiments such as CHOOZ did not observe oscillation whereas KamLAND did.

The currently known values of the three mixing angles and two mass-squared differences (the third one is not independent of the other two) is given in table 2.1. The CP-violating phase δ is still unknown.

In the case of T2K, the relevant probabilities are $P(\nu_\mu \rightarrow \nu_\mu)$ for the muon disappearance measurement, and $P(\nu_\mu \rightarrow \nu_e)$ for the electron appearance measurement. Using the experimental measurements of the oscillation parameters allows for some approximations to be made when expanding equation 2.4, such as $|\Delta m_{32}^2| \gg |\Delta m_{21}^2|$ and $\sin^2(2\theta_{13}) \approx 0$. For the disap-

pearance measurement,

$$P(\nu_\mu \rightarrow \nu_\mu) = 1 - \sin^2(2\theta_{23}) \sin^2\left(\frac{\Delta m_{32}^2 L}{4E}\right). \quad (2.7)$$

For the appearance measurement,

$$\begin{aligned} P(\nu_\mu \rightarrow \nu_e) = & \sin^2(\theta_{23}) \sin^2(2\theta_{13}) \sin^2\left(\frac{\Delta m_{32}^2 L}{4E}\right) \\ & + [\cos(\theta_{13}) \sin(2\theta_{12}) \sin(2\theta_{13}) \sin(2\theta_{23}) \sin\left(\frac{\Delta m_{32}^2 L}{4E}\right) \sin\left(\frac{\Delta m_{21}^2 L}{4E}\right)] \times \\ & [\cos(\delta) \cos\left(\frac{\Delta m_{32}^2 L}{4E}\right) - \sin(\delta) \sin\left(\frac{\Delta m_{32}^2 L}{4E}\right)]. \quad (2.8) \end{aligned}$$

To improve the sensitivity of a measurement of θ_{13} , the other parameters θ_{23} and Δm_{32}^2 have to be known with good precision, as they appear alongside it in the first term of the expression in equation 2.8. δ can only be measured if $\sin(2\theta_{13})$ is non-zero, as it appears in the second term of the equation.

The oscillation phenomenology described so far applies only to neutrinos propagating through vacuum. Although the effect is small in the case of T2K, the phenomenology of propagation through matter is briefly introduced here for completeness. Essentially, the neutrino Hamiltonian is modified by a potential due to coherent forward scattering. Two processes contribute to this potential, charged current elastic scattering, and neutral current elastic scattering (figure 2.1). The neutral current scattering is mediated by a Z boson, and affects all flavours of neutrino equally. The charged current scattering is mediated by a W boson, and only affects electron neutrinos. This implies an asymmetry in the Hamiltonian between electron neutrinos and the other flavours. In a two-neutrino oscillation approximation (represented by the vacuum parameters θ and Δm^2), with neutrinos of energy E travelling through matter of constant electron number density N_e ,




Figure 2.1.: Feynman diagrams of coherent forward scattering processes.

this modifies the mixing angle and mass-squared difference in the following way:

$$\sin^2(2\theta_{\text{eff}}) = \frac{\sin^2(2\theta)}{\sin^2(2\theta) + (\cos(2\theta) - x(E, N_e))^2}, \quad (2.9)$$

$$\Delta m_{\text{eff}}^2 = \Delta m^2 \sqrt{\sin^2(2\theta) + (\cos(2\theta) - x(E, N_e))^2}, \quad (2.10)$$

$$x(E, N_e) = \frac{2\sqrt{2}G_F N_e E}{\Delta m^2}, \quad (2.11)$$

where θ_{eff} is the effective mixing angle due to this effect, Δm_{eff}^2 is the effective mass-squared splitting, G_F is the Fermi coupling constant. These effective parameters exhibit resonance-like behaviour for a critical energy,

$$E = \frac{\cos(2\theta)\Delta m^2}{2\sqrt{2}G_F N_e}, \quad (2.12)$$

when the mixing angle becomes maximal, and hence the oscillation is enhanced. This effect, known as the Mikheyev-Smirnov-Wolfenstein effect, is described further in [33, 34].




Figure 2.2.: Charged current cross section measurements, from [35].

2.3. Neutrino-nucleon interactions

The neutrino detection principles that T2K uses are to search for the charged lepton produced in a charged current interaction of the neutrino with a nucleon. The dominant charged current process for sub-GeV neutrino interactions, as shown in figure 2.2, is the quasi-elastic mode (CCQE) interaction. The Feynman diagram for this process is shown in figure 2.3(a). The next dominant interactions are single pion productions. In resonant production (figure 2.3(b)), the W boson excites a Δ resonance of the nucleon, which subsequently decays to a nucleon and a pion. The Rein-Sehgal model [36] is commonly used to calculate resonant cross-sections. The coherent pion production process (figure 2.3(c)) also produces a single pion, leaving the whole nucleus in the ground state. This process produces a pion that is strongly peaked in the forward direction. Another Rein-Sehgal model [37] is used for calculations of the coherent cross sections. Finally there are deep inelastic scattering (DIS) processes (figure 2.3(d)) which happen for neutrinos of higher energies (around 3 GeV or more). DIS processes produce

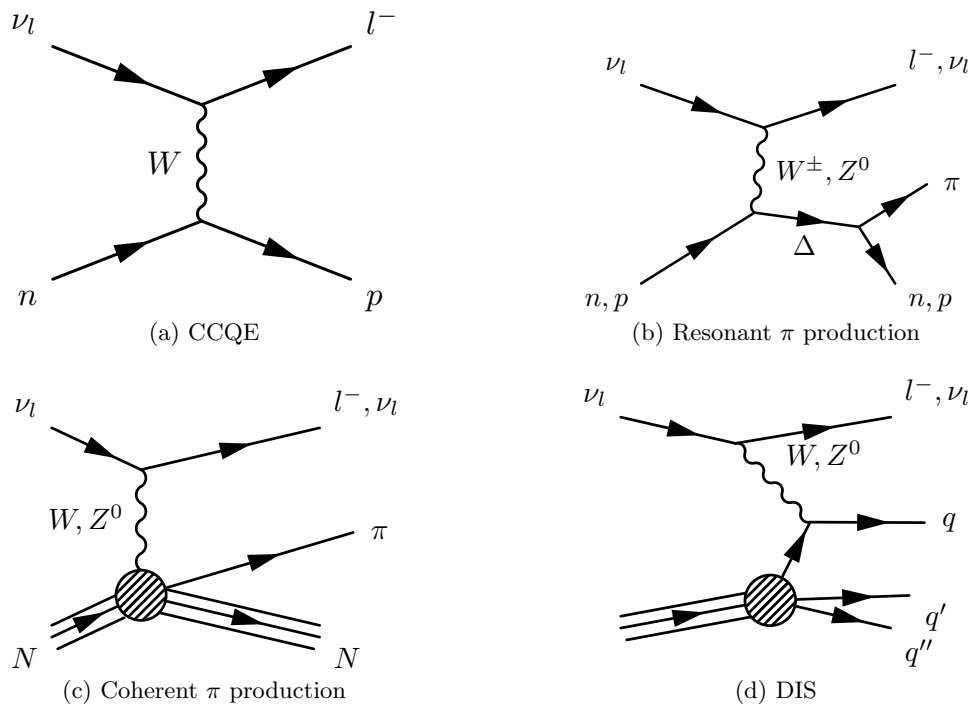


Figure 2.3.: Feynman diagrams for various neutrino interaction processes.

multiple pions in the final state.

Neutral current analogues of these interactions also exist. These are in general harder to detect, as there are no charged leptons in the final state, unless the interaction is neutral current elastic scattering off an atomic electron (figure 2.1(a)). The possible recoil protons are below the Čerenkov threshold at Super Kamiokande. They are also not desired interactions for an oscillation measurement, as no flavour information about the neutrino is observable through these interactions.

The pions produced in neutral current interactions can be a background to the CCQE signals. If the pions are charged, they can appear to be muon-like, especially if they stop without interacting hadronically, and their decay products are not detected. Neutral pions decay to two photons, which can look electron-like, if one of the photons is not reconstructed, or the opening angle of the decay is small.

There are added complications, that the state at the interaction vertex is not the final state of the event. As the interaction products move through the nucleus, they can reinteract with it, leading to different particles being seen at the detector level. For example, a charged pion produced in a resonant interaction can subsequently undergo charge-exchange with other nucleons, and the final state will contain a neutral pion.

2.3.1. Neutral pion production

A π^0 decays electromagnetically with a lifetime of (84 ± 5) as [32], producing two back-to-back photons of 67.5 MeV (in the rest frame), with a branching ratio of 0.99. Because of the boost due to the pion momentum, they will not be back-to-back in the lab frame, and will have higher energies on average. As already mentioned, these photons are a background to electron

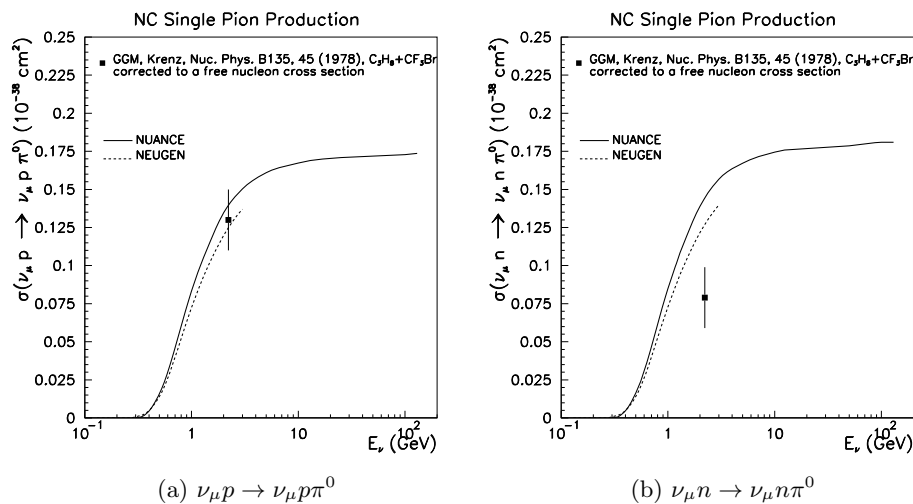


Figure 2.4.: NC π^0 production cross-section measurements, showing predictions from two interaction generators, from [38].

neutrino appearance searches, if one of the photons is not reconstructed, or the opening angle of the decay is small.

There are very few measurements of absolute NC π^0 production cross-sections. Figure 2.4 shows the data for this measurement pre-2009, and contains only one point per channel (made by the Gargamelle experiment [39]). A recent cross-section measurement from the MiniBooNE collaboration, published in 2010, is $\sigma = (4.76 \pm 0.05(\text{stat}) \pm 0.76(\text{sys})) \times 10^{-40} \text{ cm}^2/\text{nucleon}$ [40], where the mean neutrino energy was 808 MeV, and the nuclear target was carbon.

The near detector of T2K will be performing measurements of π^0 production cross sections. The Tracker subdetector is able to perform exclusive measurements differentiating CC and NC interactions, and also multipion production processes. This coupled with the relatively high statistics of the intense beam means that T2K is in a good position to add to this global

data set.

3. The T2K Experiment

The T2K experiment is a long-baseline neutrino oscillation experiment located in Japan. It consists of a neutrino production beamline, based at the Japanese Proton Accelerator Research Complex (J-PARC) in the village of Tokai, Ibaraki prefecture, with a complex of near detectors 280 m from the production target, and a far detector, Super Kamiokande (SK) 295 km to the west in Mount Ikenoyama, Gifu prefecture (figure 3.1). SK is located at an angle of 2.5° from the axis of the beamline. The experiment is designed to search for ν_e appearance in the ν_μ beam, and also perform a ν_μ disappearance measurement. It aims to make a precision measurement of the θ_{23} and Δm_{23}^2 neutrino oscillation parameters, and determine whether θ_{13} is non-zero, with a 20 times improvement of sensitivity over the CHOOZ limit. The far detector makes a measurement of the oscillation signal, while the near detectors are used to characterise the beam and reduce systematic uncertainties, and also to make exclusive neutrino cross section measurements.




Figure 3.1.: Baseline for the T2K experiment [41].

3.1. Accelerator and neutrino beam line

The accelerator complex at J-PARC (figure 3.2) primarily consists of a linear accelerator (Linac), a rapid cycling synchrotron (RCS) and a proton synchrotron (PS), which is also called the Main Ring. The Linac is designed to accelerate H^- ions from rest to a kinetic energy of up to 400 MeV, after which the electrons are stripped and the protons injected into the RCS. This bunches up the protons with a chopper, and boosts them to an energy of 3 GeV. The RCS can hold two proton bunches, with a cycling frequency of 25 Hz. The RCS feeds both the PS and the Materials and Life Sciences facility at J-PARC. The bunches that are fed into the PS are accelerated to 30 GeV. The PS has a circumference of 1,567 m, and can hold up to 9 bunches, with a bunch separation of 582 ns. The two experimental facilities that use protons extracted from the PS are the neutrino beamline and a hadron beamline.

Using kicker magnets, protons are extracted from the PS at a frequency of approximately 0.3 Hz and steered into the neutrino beamline (figure 3.3). Each of these spills (extractions) consists of eight¹ bunches of protons. The ninth bunch space of the PS is empty for the kicker magnets to turn on. The

¹Bfore Summer 2010 there were six bunches per spill.




Figure 3.2.: Aerial view of J-PARC [42].




Figure 3.3.: Neutrino beamline at J-PARC [41].

protons are steered with magnets towards a graphite target. This target is a 92 cm long cylinder of diameter 3 cm, its length parallel to the proton beam. The target length is approximately two interaction lengths. The protons enter it in the centre of the circular face, with various monitors measuring this targeting precision and the direction of the proton beam.

Inside the target, the protons interact with carbon nuclei, producing charged pions and kaons, amongst other products. A system of three coaxial magnetic horns [43] focus (or defocus, depending on charge and polarity) the charged mesons. The aluminium horns produce a toroidal magnetic field with a strength proportional to $1/r$, where r is the radial distance from the axis. The first horn has a radius of 40 cm and a length of 1.2 m, and is placed around the target (the target occupying the most upstream 92 cm of its length). The second horn has a length of 2 m and a diameter of 1 m, and is 2 m downstream of the first. The third horn has a diameter of 1.4 m and length of 2.5 m, and is 7.5 m downstream of the second. All the horns are supplied with a pulsed current of 250 kA to provide the magnetic field, and this current is monitored as it affects the flux of neutrinos.

The unit of statistical measure is the Proton On Target (POT). Every proton that enters the target has an equal chance² of producing a neutrino of a certain energy and in a certain direction, and so the integrated number of neutrinos at both the near detector and far detector is proportional to the number of protons on target.

The charged mesons are allowed to decay in a 96 m long decay volume filled with helium, giving rise to a mostly muon neutrino beam (with positive pions as their parent). This muon neutrino beam is contaminated with antineutrinos and electron neutrinos, which come from other pion, kaon,

²Disregarding small systematic uncertainties such as target decay, proton beam momentum, and horn currents.




Figure 3.4.: Predicted neutrino fluxes at Super Kamiokande, without oscillation [44].

and muon decays. The main contributors to the beam are $\pi^+ \rightarrow \mu^+ \nu_\mu$, $K^+ \rightarrow \mu^+ \nu_\mu$, $K^+ \rightarrow \pi^0 e^+ \nu_e$, and $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$. Figure 3.4 shows the unoscillated flux at Super Kamiokande of the four types of neutrinos and their sources. The ν_μ flux peaks at around 700 MeV

At the end of the decay volume, approximately 110 m from the target, is the beam dump. This is a structure made of graphite, with 3.2 m depth, and iron, with 2.4 m depth, that absorbs surviving pions and muons in the beam, and only muons above 5 GeV should penetrate it. Just beyond the beam dump lies the muon monitor. This can monitor the profile of the beam by measuring the muons that penetrate the dump. The monitor is built of

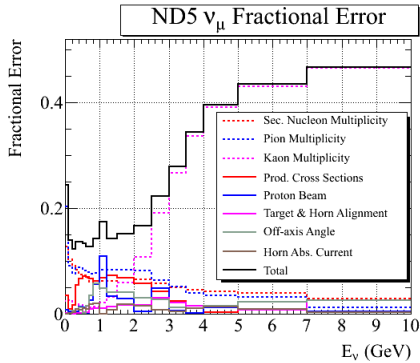


Figure 3.5.: ν_μ flux uncertainty at ND280 [44].

two layers separated by 1.2 m, with ionisation chambers in the first layer, and silicon photodiodes in the second layer. Each layer is a 7×7 square array, with 25 cm periodicity, of the detector elements.

The uncertainties on the neutrino flux depend on the beam direction, the horn currents, and uncertainties in hadron production rates from proton-carbon collisions. The former two uncertainties are estimated from the monitors (beamline and muon), and measuring the horn currents. The uncertainties of the latter are estimated with help from the NA61/SHINE [45] experiment at CERN, which has a data-sharing and collaboration agreement with T2K. Figure 3.5 shows the uncertainty in the ν_μ flux prediction at the off-axis near detector position. At the oscillation maximum, the uncertainty is around 15%, whereas in the high-energy tail the uncertainty is around 45%, and is dominated by uncertainties in the kaon production processes.

3.2. Near detectors

Approximately 280 m downstream of the production target is the complex of near detectors: the INGRID on-axis detector and the ND280 off-axis




Figure 3.6.: Exploded view of the ND280 subdetectors [41].

detector. These are situated in a cylindrical cavern of radius 8.75 m, with the centre of ND280 1 m downstream of the centre of the cavern.

INGRID is a system of modules arranged in a cross, with the centre of the cross aligned with the designed centre of the beam. ND280 is a system of subdetectors (the Tracker, the π^0 detector (P0D), the electromagnetic calorimeters (ECal), and the side muon range detectors (SMRD)) placed inside an iron electromagnet, with the centre of ND280 at an angle of 2° from the beam. The P0D, Tracker and downstream ECal module are placed in a steel frame (the ‘basket’) that is 6.5 m long, 2.6 m wide and 2.5 m high. The other ECal modules are fixed to the inner sides of the magnet, and the SMRD sits within the magnet yoke itself. Figure 3.6 shows this arrangement. The figure also shows the right-handed coordinate system used for ND280, in this chapter and also the remainder of the thesis, where the z axis is pointing downstream parallel to the beam direction, the y axis is vertical, and the x axis is horizontal (with ‘left’ in the positive x direction,

and ‘right’ in the negative direction).

3.2.1. INGRID

The primary purpose of INGRID is to measure the direction of the beam, along with the beam stability. It does this by having 14 identical modules arranged in a cross structure, 7 stacked horizontally and 7 stacked vertically, with the centre of the cross being the nominal centre of the beam.

Each of these modules is made of layers of scintillator and iron. Each scintillator layer is formed of 24 bars of thickness 1 cm, width 5 cm and length 1.2 m. The bars are extruded polystyrene, doped with 1% PPO (2,5-diphenyloxazole) and 0.03% POPOP (1,4-bis(5-phenyloxazol-2-yl) benzene) as fluors, and coated with titanium oxide. A 1 mm diameter wavelength shifting (WLS) fibre (Kuraray double-clad Y11) runs along the length of the centre of the bar, inside a hole of diameter 3 mm. Photosensors (described in more detail in section 3.2.6) are attached to the ends of the bars, connecting with the fibres. The layers have area 1.2 m × 1.2 m, and 22 layers make up the module. Pairs of layers, one with horizontal bars and one vertical, are kept as one unit, a tracking plane. There are ten 6.5 cm gaps between the 11 tracking planes, and in the most upstream 9 gaps there is an iron plate of thickness 6.5 cm, with area 1.24 m × 1.24 m. The iron serves as the neutrino target mass.

Around the top, bottom and sides of the modules are three or four³ veto planes: these are single layers, parallel to the module faces, composed of 22 bars. The bottom veto planes use 1.1 m long bars, the side and top veto planes use 1.3 m long bars. The length of these bars runs parallel to the beam direction. The distance between the faces of the main module and

³Veto planes between two neighbouring modules are shared, so only the module on the outside needs an additional veto plane.




Figure 3.7.: Exploded view of an INGRID module [41]. (left) Tracking planes are shown partially out, and the steel plates are in blue. (right) Four veto planes around the module.

the veto planes is 5 cm. These veto planes serve to veto interaction from outside the modules.

Overall these modules are 0.9 m long along z , 1.3 m high along y and 1.3 m wide along x . The mass of a module is 7.4 t. Figure 3.7 is an exploded view of one of the modules.

7 such modules make up the horizontal stack, with the central one being in the centre of the beam. Their separations are 1.5 m in x from the centre of one module to the centre of its direct neighbours. Another 7 make up the vertical stack. The central module of this stack is 4 m upstream of the central horizontal module. The modules have 1.5 m in y centre-to-centre separation from their neighbours. By measuring the rate of neutrino interactions in these modules, which should be symmetrical around the beam centre, the direction of the beam can be determined to within 0.4 mrad. This arrangement is shown in figure 3.8

There are two further modules like these positioned off-cross, which are




Figure 3.8.: Front view of the INGRID assembly, showing the 7 vertical modules, 7 horizontal modules, and 2 off-cross modules, and support structures, looking downstream [41].

used to measure the axial symmetry of the beam. Their centres are 1.5 m upstream, 2.2 m in y , and 3.1 m on either side along x , to the centre of the central horizontal module.

There is a 17th module built without the iron absorber, capable of detecting recoil protons in the interactions, and making an exclusive measurement of CCQE interactions. The layers have 32 1.2 m long bars, with the central 16 of dimensions 2.5 cm wide and 1.3 cm thick to provide more granularity. The remaining bars are the same type as in the rest of INGRID. The spacing between each layer is 2.3 cm. The module is positioned 1.2 m upstream of the central horizontal module (which is also used to detect muons that exit this proton module). The mass of this module is 0.6 t.

In relation to the ND280, the centre of the INGRID central horizontal module is 1.5 m upstream, 3.2 m along $+x$, and 9.5 m along $-y$ from the centre of ND280.

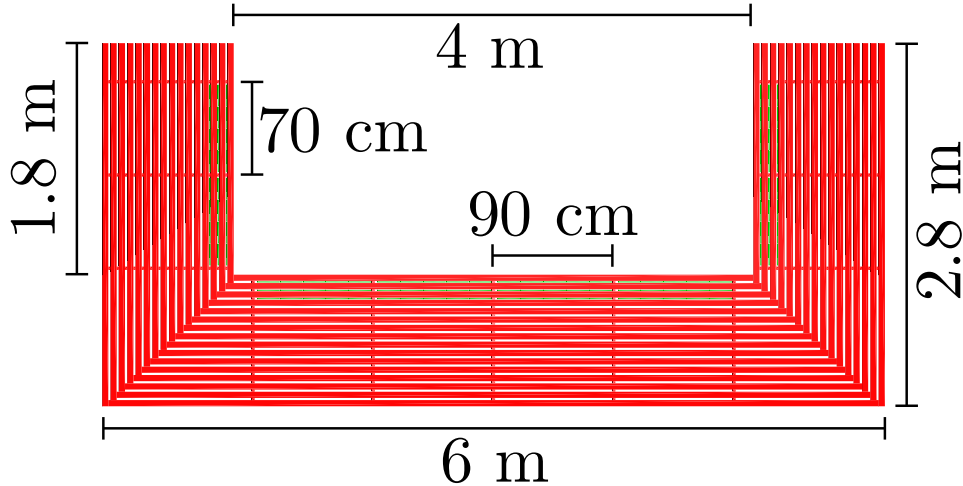


Figure 3.9.: Cross section of a magnet yoke C-section, lying on its side, with SMRD paddles (in green) occupying the three innermost gaps between iron layers (red).

3.2.2. Magnet and SMRD

The bulk of the ND280 detector is formed of the magnet, which was previously used in the UA1 [46] and NOMAD [47] experiments, and refurbished for use in T2K. It is formed of two halves (the ‘left’ clam or ‘right’ clam), split longitudinally along z . Each clam can be closed or opened up in the horizontal direction, to allow access to the basket. The clams are 7 m long, 6 m high and 2.8 m wide on the outside edges, and 4 m high and 1.8 m wide on the inside edges. The length of the clam is split into eight C-shaped sections, each made of 16 layers of iron, with 1.7 cm air gaps between the layers. The layers are 4.8 cm thick and 88 cm wide along z . The length of the layer depends on which layer it is, such that a concentric C-shape with square corners is formed (and so the inner layers are shorter than the outer layers). There are 2.5 cm iron spacers between the layers, such that the air gaps are separated into four sections of 90 cm on the sides, and two sections of 70 cm on the top and bottom. A diagram of the cross section of

a C-section is shown in figure 3.9. The 8 C-sections of a clam are separated by gaps of 10 cm along z .

There are four coils, two for each clam, that are fitted into the inner sides of the clam. The coils are 78 cm wide along x , 7 m along z and a 3.5 m high along y . They are made up of aluminium bars of cross section $5.4 \times 5.4 \text{ cm}^2$, arranged into ‘pancakes’. A pancake is a single piece of conductor, coiled up to make a layer of 2×4 . Six pancakes connected in series make the coil. In total there are 52 turns in a coil. The coils, when electrified with a current of 3.3 kA, create a magnetic field of 0.2 T along the x direction.

Most of the neutrino interactions in ND280 happen inside the magnet, as it has a mass of around 850 t (the total mass of the ND280 itself is around 1 kt).

The Side Muon Range Detector (SMRD) is a system of 440 scintillator paddles, placed into some of the air gaps between the layers of the magnet yoke. Five (four) paddles fill an individual air gap between spacers on the sides (top and bottom) of the yoke. As the flux of particles is in general greater for lower angles with respect to the beam, there are more of these paddles at the downstream end of the magnet than the upstream. All the top and bottom gaps of the yoke have their three innermost gaps instrumented. For the side gaps, the five most upstream C-sections have their three innermost gaps instrumented; the next section has four; and the two most downstream sections each have six.

Each paddle is a scintillator bar of 0.7 cm thickness, 16.7 cm (for the side sections) or 17.5 cm (for the top and bottom sections) width, and 87.5 cm length along z . The scintillator is extruded polystyrene and dimethylacetamide with admixtures of POPOP and para-terphenyl, coated with titanium oxide. A 2.2 m long, 2.5 mm deep S-shaped groove with bending radius




Figure 3.10.: Photograph of an SMRD paddle, showing the WLS fibre in the S-shaped groove [41].

of 3 cm is carved along the length of the innermost face of the paddles, and 1 mm diameter Y11 WLS fibre is placed into it. Figure 3.10 is a photograph of an SMRD paddle. At both ends of the paddle, the fibre connects with photosensors, described in section 3.2.6.

The SMRD is used to make range estimates for muons that exit the inner detectors, which can improve the momentum measurements. It can also serve as a veto for interactions that take place in the magnet, as they are backgrounds to the measurements of the inner detectors. Coincidences in the SMRD paddles can also act as a cosmic ray trigger.

3.2.3. Tracker

The Tracker region of the detector is formed of two types of subdetector: two ‘fine-grained’ detectors (FGDs) and three time projection chambers (TPCs).

The FGDs are made of scintillator bars arranged in layers. The bars are 0.96 cm wide, 0.96 cm deep and have a length of 1.8 m. The scintillator material is extruded polystyrene, doped with 1% PPO and 0.03% POPOP. They are coated with titanium oxide, to reflect light inside the bars. A 1 mm diameter Y11 WLS fibre runs down the length of the bar, in a hole of diameter 1.8 mm. One end of the fibre is connected to a photosensor, described in section 3.2.6. The end of the bar that the photosensor is on




Figure 3.11.: Drawing of the TPC design [41].

alternates from bar to bar. The other end of the fibre is mirrored, so that light heading along the fibre towards the mirrored end is reflected back to the instrumented end. Every layer is composed of 192 bars, giving layer dimensions $1.8\text{ m} \times 1.8\text{ m} \times 9.6\text{ mm}$ in the xy plane. Alternating layers have the bars in different orientations; if one layer has the bar lengths oriented horizontally, the neighbouring layer will have the bars vertical.

The most upstream FGD (FGD1) is built of 30 such layers, 15 horizontal and 15 vertical, with total dimension of 30 cm along z . The other FGD (FGD2) has 14 layers, 7 horizontal and 7 vertical. Between each vertical and the next horizontal layer there is a gap of 2.5 cm along z , and this gap is filled with a water bag. The total length is 30 cm along z . The centre of FGD2 is 1.4 m downstream of FGD1's centre. They are in-line horizontally and vertically. Both subdetectors have a mass of 1.1 t.

The TPCs (figure 3.11) are contained in aluminium boxes 2.3 m along x , 2.4 m along y and 1 m along z . An inner box made of copper-clad G10 of dimensions $1.8 \times 2.2 \times 0.8 \text{ m}^3$ holds a mixture of the gases argon, tetrafluoromethane, and isobutane (in the ratios 95:3:2 respectively). The box is horizontally divided in the centre by a cathode panel also made of copper-clad G10, perpendicular to the x direction. Given a nominal voltage of 25 kV, the cathode provides an electric field of 250 V/cm parallel to the magnetic field. Charged particles travelling through the TPCs ionise the gas, and the ions drift to the edges. The drift velocity is around $75 \text{ mm}/\mu\text{s}$, depending on the electric field and the gas pressure.

At both sides of a TPC, the ions are detected with Micromegas ('micro mesh gaseous structures' [48]). A Micromegas module is 36 cm wide along z and 34 cm high along y , and oriented parallel to the central cathode. It is segmented into 1728 rectangular pads of area $9.8 \times 7.0 \text{ mm}^2$, which connect to the readout electronics, described further in section 3.2.6. Two layers of six Micromegas modules form a readout plane, one on each side of the TPC, so in total across the three subdetectors there are 72 modules.

There are three TPC modules, TPC1 upstream of FGD1 (68 cm centre-to-centre displacement), TPC2 in between FGD1 and FGD2 (its centre 68 cm equidistant from the centres of the two FGDs), and TPC3 downstream of FGD2 (68 cm centre-to-centre displacement), and this combination of five subdetectors makes up the ND280 Tracker.

The purpose of the Tracker is to make exclusive measurements of neutrino interactions, with a focus on CCQE. The Tracker is capable of measuring the muon neutrino flux, and the electron neutrino contamination of the beam. The FGDs provide the target mass for the interactions, and they are thin so that particles can escape into the TPCs. The FGDs also provide

timing information for the TPCs, as the TPC tracks' position coordinate parallel to the drift direction is dependent on the track time. The curvature of tracks in the TPCs due to the magnetic field allows a measurement of particle momentum, with the design resolution of better than 10% at 1 GeV. The TPCs provide particle identification through measurements of dE/dx and the momentum, and the track curvature allows for determination of particle charge.

3.2.4. Tracker Electromagnetic Calorimeters

The Tracker region is surrounded by electromagnetic calorimeter (ECal) modules. There are seven in total, in three configurations: the Downstream ECal (DsECal), two Barrel Side modules, and four Barrel Top/Bottom modules. Each module is a segmented sampling calorimeter, formed of layers of lead and plastic. The lead acts as the absorber medium, the plastic as the sampling medium.

The scintillator unit is a bar of extruded polystyrene plastic doped with 1% PPO and 0.03% POPOP. Each bar is 1 cm thick and 4 cm wide. The bars are coated with titanium oxide, which acts to reflect inwards any light produced in the bars. A $1\text{ mm} \times 2\text{ mm}$ elliptical hole runs along the length of the bar, through the centre. A Y11 WLS fibre, of diameter 1 mm, runs down this hole, with the fibre being a few cm longer than the bar. The fibre picks up the scintillation light and directs it to the ends of the bar. A photosensor is placed at one or both ends of the fibre, described further in section 3.2.6. If the readout is single-ended, the uninstrumented end of the fibre is mirrored, so that light can be detected at the photosensor. Each layer is built of these bars arranged side-by-side, to make a layer of depth 1 cm and the required width (a multiple of 4 cm).

A module is made up of alternating layers of lead sheets and scintillator. The lead sheets are 1.75 mm thick (0.31 radiation lengths). Every other layer of scintillator is in a perpendicular orientation: if one has the length of the bars oriented horizontally, the next layer will have the bars vertically, and so on. Thus the module is said to have readout ‘views’, with every other layer producing hits in a different view; in the example given, these views will be the horizontal one and vertical one. There is one less lead sheet than scintillator layers; the first scintillator layer has no lead sheet before it, and the last scintillator layer has no lead sheet after it. There is however a 2.5 cm thick carbon fibre cover on both faces for light tightness.

The Downstream module is made of 34 scintillator layers, and the bars are 2 m long in each view, with each layer having 50 bars. All the bars have double ended readout. The most upstream layer is a horizontal layer (the length of a bar runs horizontally). The module is mounted within the basket, directly downstream of TPC3 (75 cm centre-to-centre). The total radiation depth (including scintillator effects) is 11 radiation lengths. The mass of the DsECal is 9 t.

The Barrel Side modules are made of 31 scintillator layers. The bars are 3.84 m long in the ‘long’ view, and 57 of these make up the ‘long’ layer. These bars have double-ended readout, and they are all oriented along z . In the other, ‘short’ view, the bars are 2.28 m long with single ended readout. These bars are oriented vertically, with the photosensor end being the top end. 96 of these bars make up the layer. The innermost layer of the module (the one closest to the basket) is the ‘short’ layer. These modules are mounted onto the inside of magnet. The offset of the modules’ centre from the centre of FGD1 is 10 cm upwards and 30 cm in the downstream direction, with the horizontal offsets nominally 160 cm, when the magnet is

fully closed. The Side modules have a fiducial mass of 17 t.

The Barrel Top/Bottom modules are also made of 31 scintillator layers. 38 double-ended bars of 3.84 m length make up the ‘long’ layers, with the bar lengths parallel to the beam. The ‘short’ layers have 96 single-ended 1.52 m long bars oriented horizontally, but perpendicular to the beam. The photosensors of these bars are on the ends closest to the basket centre. The innermost layer (bottom layer of the Top modules, top layer of the Bottom modules) is a ‘short’ layer. These modules too are mounted directly onto the inside of the magnet. The modules’ centres are offset 150 cm in $+y$ and 30 cm in $+z$ from the centre of FGD1, with the horizontal offset being 80 cm for the Top modules and 90 cm for the Bottom modules. The Top and Bottom modules have a mass of 12 t.

One purpose of the Tracker ECal is to make a momentum measurement, complementary to the TPCs, of particles produced in the FGDs. The energy resolution of the ECal improves with higher momentum, whereas for the TPCs it becomes worse for higher momenta. Another purpose is to provide some level of particle identification, based on whether particles shower or produce tracks in the modules. The ECal also acts to convert any photons produced in the Tracker. This is an important requirement for π^0 analyses using the FGDs and the ECals. The ECal not only provides energy reconstruction, but also angular reconstruction; for many neutral current π^0 events, the two decay photons in the ECal modules are the only indication of the interaction, and it is necessary to use the angular reconstruction to check if the two photons are consistent with a single vertex, and also to determine the location of that vertex. Discussion on how well the Downstream ECal module performs for these purposes is in chapter 5. A secondary purpose of the ECal modules is to make neutrino interaction measurements.

The Barrel modules have the largest length along the beam direction of any instrumented subdetector, and muons travelling in that direction can be contained.

3.2.5. P0D

The π^0 detector is composed of layers of scintillator, brass, lead and water. The scintillator bars are extruded polystyrene doped with 1% PPO and 0.03% POPOP and coated with titanium oxide. The shape of the bars is triangular, with a base of 33 mm and a height of 17 mm. A hole of diameter 1.5 mm runs through the centre of the bars, down which a 1 mm diameter Y11 WLS fibre is placed. One end of the fibre is attached to photosensors (described in section 3.2.6), and the other end is mirrored.

The bars are arranged into layers in an alternating fashion, with the sloped sides of neighbouring bars touching each other, such that two bars together form a parallelogram. The layers are in two orientations: horizontal and vertical, from the orientation of the length of the bar. 134 2.2 m long bars make up a vertical layer, and 126 2.32 m long bars make a horizontal layer. A horizontal and vertical layer together make up a ‘P0Dule’, and there are 40 of these P0Dules in the whole subdetector. The most upstream 7 P0Dules have 4 mm thick lead sheets in between the P0Dules. The next 26 P0Dules have 1.5 mm thick brass layers and 28 mm thick water bags between them. The water bags have an area of $1 \times 2 \text{ m}^2$, and there are two of these side-by-side per layer. The water bags can be filled or empty during running. The final 7 P0Dules are identical to the first 7. The total length of the P0D is 2.4 m, and it is 2.3 m wide and 2.2 m high. It has a fiducial mass of 16 t when filled with water, and 13 t when empty.

The P0D lies in the basket directly upstream of TPC1, with a centre-to-

centre distance of 1.7 m.

The P0D is designed to make a high statistics⁴ measurement of π^0 production, one of the main backgrounds to the ν_e appearance measurement at Super Kamiokande. Cross section measurements on an oxygen target are possible by comparing rates with water in and out of the P0D. The scintillator bars have enough resolution to reconstruct charged particle tracks, and the brass and lead allow for photon conversion and containment in the subdetector.

P0D Electromagnetic Calorimeter

The P0D is surrounded on four sides by six P0D-ECal modules. These are similar to the Tracker-ECal modules in overall design (lead/plastic layers), with some differences. The lead sheets used are 4 mm thick. There are only six scintillator layers, and they all have only one view, so independent 3D reconstruction is not possible. The 2.34 m long bars are all oriented along z , with single-ended readout at the upstream end. The layers of the side modules are made of 69 bars, while the layers of the top and bottom modules have 38 bars each. The centre-to-centre longitudinal displacement of each module from the centre of the P0D is 23 cm in the downstream direction.

The purpose of these modules is improve measurements that utilise the P0D. The orientation of the P0D, with all layers perpendicular to the beam, allows for an inefficiency if particles travel laterally from the interaction. There is also some inefficiency if the interactions are close to the edge of the P0D, with particles escaping completely. With the P0D-ECal in place, these escaping particles can be detected and tagged. Escaping photons from π^0 decays will shower, and escaping muons from charged current interactions

⁴From mass considerations, the statistics in the P0D will be around 15 times higher than in an FGD.




Figure 3.12.: Photograph of the MPPC photon detection area [41].

will produce tracks.

3.2.6. Electronics

The INGRID, P0D, FGDs, ECal and SMRD all use the same photosensors. These are Geiger-mode avalanche photodiodes (described in [49]) custom-made by Hamamatsu, with the trade name Multi-Pixel Photon Counters (MPPCs). Each MPPC is a matrix of square cells (also known as ‘pixels’) measuring $50\ \mu\text{m}$ across. 667 of these cells make up the MPPC, arranged in a 26×26 square, with a 3×3 square missing from one corner. A photograph is shown in figure 3.12. The light-collection coverage is $1.3 \times 1.3\ \text{mm}^2$, which is slightly larger than the 1 mm diameter WLS fibres used, although light somewhat spreads out on exiting the fibres.

Each cell is a reverse-biased diode, biased slightly beyond the breakdown voltage. The overvoltage (ΔV) is $V_{\text{bias}} - V_{\text{bd}}$, where V_{bias} is the bias voltage and V_{bd} the breakdown voltage. A photon incident on the cell may produce a photoelectron (p.e.), which creates an avalanche of electron-hole pairs. This avalanche would be self-sustaining, so a resistor is used to quench it. The voltage drops to V_{bd} , after which the avalanche stops, and the voltage

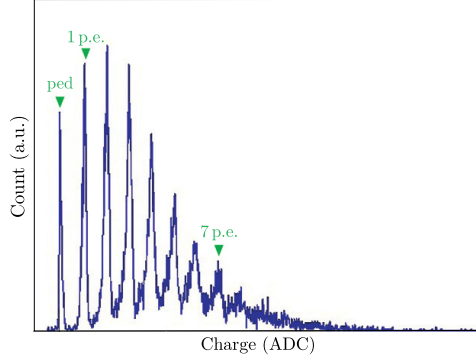


Figure 3.13.: Photoelectron spectrum of a MPPC device, from [50].

risks back to V_{bias} . The timescales of these voltage drop and rise times for the MPPCs are of order 100 ps and 10 ns respectively.

The cell can be thought of conceptually as a charged capacitor, at a voltage of ΔV , that discharges when triggered by a photon. The amount of electrons released per triggering photon (the gain) is typically 0.75×10^6 , depending on the overvoltage, and this number is independent of the number of electron-hole pairs produced in the avalanche. The cells of an MPPC are all connected in parallel, and so by measuring the charge released when a discharge happens, the amount of cells that were triggered can be counted. An example charge spectrum for a MPPC, which shows the photoelectron peaks, is in figure 3.13.

Noise from a MPPC comes in three forms: dark noise, cross-talk and afterpulse. Dark noise is when an avalanche is triggered not by a photoelectron but with thermal noise. The dark noise rate in the sensors used is approximately 500 kHz, depending on the overvoltage and temperature. Cross-talk is when an avalanche in one cell triggers an avalanche in a neighbouring cell. Afterpulse is when an electron or hole from a developing avalanche is trapped and subsequently released, which triggers another avalanche in the

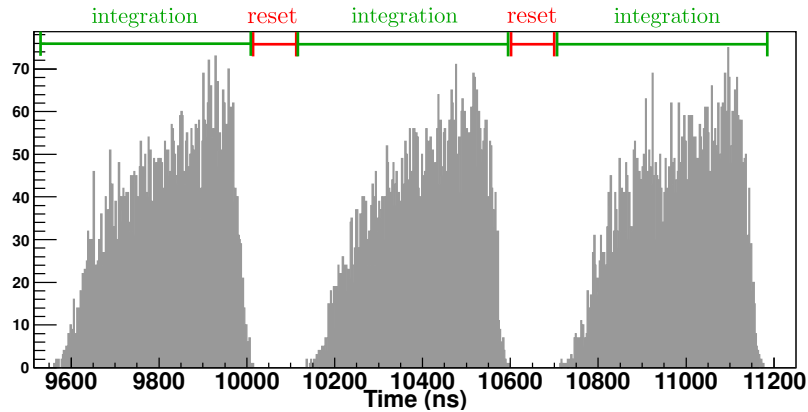


Figure 3.14.: Timing structure of MPPC noise digits with TFB readout.

same cell at a delayed time. Characterisation of the MPPC noise is reported in [50].

Up to 64 MPPCs are connected to a Trip-T front-end board (TFB), for the INGRID, P0D, ECal and SMRD subdetectors. The Trip-T chip was originally developed for the D0 experiment [51], and a TFB contains four of these chips (each with a maximum capacity of 16 MPPCs). Charge from an MPPC is collected on one of 23 capacitors, dependent on the time: there are 23 integration windows, with a reset period between each window. The duration of the windows and reset periods can be controlled with 10 ns precision, and in nominal running they are 480 ns and 100 ns respectively, approximately adding up to the beam bunch period of 582 ns. The sampled duration is 13.2 μ s, with 11 μ s of that time active. The integration and reset periods are illustrated in figure 3.14.

The integrated charge in each window is converted into two 10-bit ADC channels. One is the ‘low-gain’ ADC, with a dynamic range of around 500 p.e.. The other is the ‘high-gain’ ADC (with around 10 times the gain of the low-gain channel), which has a dynamic range of around 50 p.e..

Individual photoelectron spectra (such as in figure 3.13) can only be resolved with the high-gain channel.

Each integration window may have an associated timestamp with 2.5 ns precision. The time stamping is triggered by a discriminator that fires when the integrated charge rises above a set threshold (nominally 3.5 p.e. for the ECal channels). MPPC dark noise that has triggered this threshold is shown in figure 3.14, showing that the incidence of noise-triggered timestamps has a sawtooth shape, which rises for times later in the integration window. This is because the probability of having four dark noise avalanches (the minimum to trigger the time stamping) in a given period is proportional to the period duration. If the discriminator did not fire, then no timestamp is associated with the window.

When readout of the electronics is triggered, the output for each MPPC on these detectors is a list of 23 ‘digits’, one for each integration window. A digit is a high-gain ADC, a low-gain ADC, and a possible timestamp.

The TPCs on the other hand use Micromegas. A voltage of 350 V is applied between a micromesh (woven $18\ \mu\text{m}$ diameter stainless steel wires with a pitch of $63\ \mu\text{m}$) and the collection pads, which are copper anodes. The distance between mesh and anodes is $128\ \mu\text{m}$. The drift ions from the TPCs passing through the mesh are accelerated by the high field, and trigger an avalanche of charge which is collected on the pads. The high electric field in the avalanche gap keeps the perpendicular spread of the avalanches low.

The TPC electronics uses custom-made ‘ASIC for TPC Electronics Red-out’ (AFTER) chips, with four such chips on a ‘front end card’ (FEC). 72 pads of a Micromegas module are connected to an AFTER chip, which pre-shapes the charges and then stores them on 511 capacitors, with a collection period of 40 ns. The capacitors are designed to have a dynamic range of 10

minimum ionising particles' charge deposits. When triggered for readout, the integrated charge on each capacitor is converted to a 12-bit ADC channel, and the final output digit is a 12-bit waveform sampled with a period of 40 ns, with duration $20.4 \mu\text{s}$, for each Micromegas channel.

The FGD MPPCs are similarly connected to AFTER chips. 32 MPPCs are connected to one chip, with the signal from an MPPC split into high- and low-attenuations as inputs to the chip. There is a factor of around 9 difference between high and low attenuations. The signals are preshaped and collected on 511 capacitors as with the TPCs, but with a collection period of 20 ns. The output digit is also a waveform of 511 12-bit ADC values, with sampling period 20 ns and duration $10.2 \mu\text{s}$.

3.3. Far detector

Super Kamiokande is located 1 km deep in the Ikenoyama mountain, with a water equivalent depth of 2.7 km (figure 3.15). It is formed of an Inner Detector (ID) and an outer detector (OD) and filled with purified water. Between the ID and OD is a wall of 50 cm thickness. This wall is a stainless steel scaffold covered with black panels to absorb photons and stop light contamination between inner and outer detectors. This is also the structure that the photosensors are mounted onto.

The ID is a cylinder of radius 16.9 m and height 36.2 m. The official fiducial volume is 2 m from any wall, with a fiducial mass of 22.5 kt. Around 11,000 photomultiplier tubes (PMTs) of 20 in diameter built by Hamamatsu are placed along the walls, facing inward, and placed at regular intervals. The coverage offered by the PMTs is approximately 40%. The PMTs are designed to pick up Čerenkov light given off by charged particles travelling





Figure 3.15.: Diagram of Super Kamiokande, showing the location in the mountain (inset), and a cut-out of the detector, from [52].

through the water, with the aim of classifying the neutrino interactions. Energy measurements are also possible, by summing the photoelectron measurements of the PMTs.

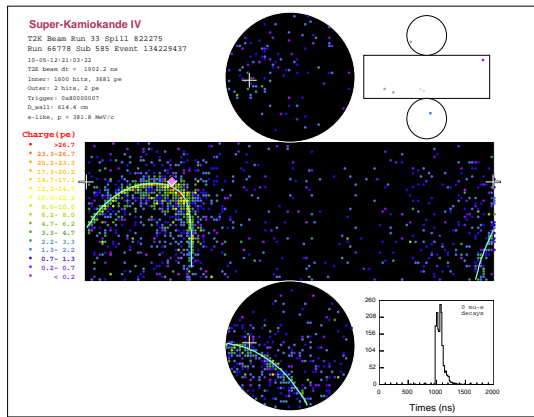
The OD is the region between the scaffold wall and a cylinder of radius 20 m and height 41 m. The OD uses around 2,000 8 in PMTs, also made by Hamamatsu, facing outward. The coverage is around 7%, so to improve this the outer wall of the OD is covered with a reflective coating, so photons can scatter and have more chance of being detected. The OD also detects Čerenkov photons, but is designed as a veto for interactions outside the detector, and cosmic rays.

Data is read out when triggered with a GPS system providing the beam timing information. All hit information in a 1 ms window around the expected beam time is read out, with a reduction of the data happening in offline analysis.

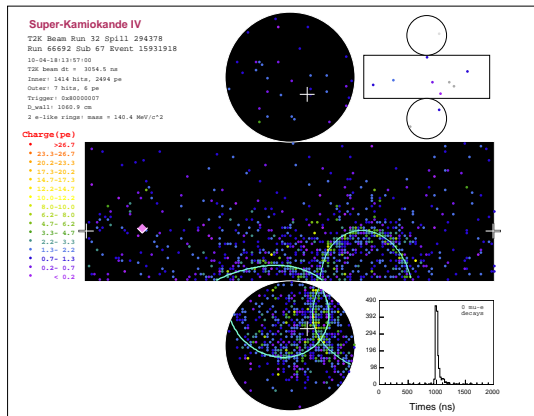
The ID is capable of particle identification by looking at the shape of the Čerenkov ring (figure 3.16). If the ring has a sharp edge, then it is



(a) Muon-like Čerenkov ring



(b) Electron-like Čerenkov ring



(c) π^0 candidate

Figure 3.16.: Displays of Super Kamiokande events [41]. Each circle represents a PMT, with the colour given by the charge on the PMT.

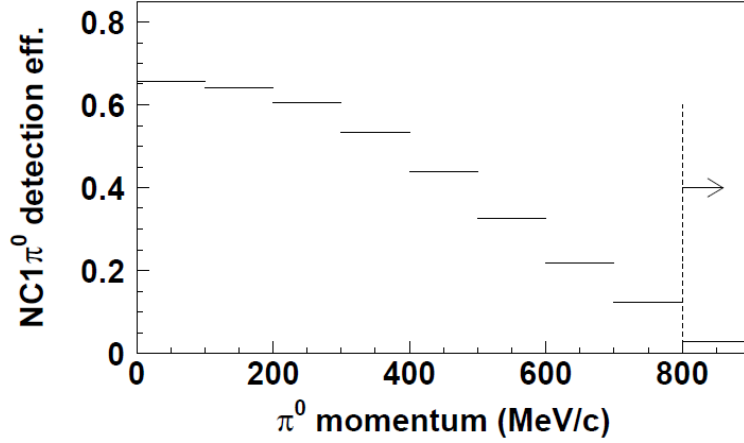


Figure 3.17.: NC1 π^0 reconstruction efficiency for K2K, from [53].

likely to be a muon, as it is unlikely to scatter. If it has a diffuse edge, then it is likely to be an electromagnetic shower, as the showering process produces daughter particles in multiple directions. A π^0 decay will look like two diffuse rings, one from each decay photon. However, under certain circumstances, one of the rings might be misreconstructed, and so the event will look like a single electromagnetic shower. This is also the signal for a ν_e appearance, hence why π^0 s are a background to this channel.

Figure 3.17 shows the reconstruction efficiency for single π^0 s produced at Super Kamiokande in the K2K neutrino beam. Although the situation is slightly different for T2K (with upgraded electronics and reconstruction), the general features still apply: for higher momentum π^0 s, the reconstruction efficiency falls off. This is because at higher momenta the decays are asymmetric, and the low energy photon has a chance of being too low in energy to be reconstructed. Also with higher momenta, the opening angle between the photons is smaller, and the two rings might overlap significantly enough to be misreconstructed as one ring.

The π^0 background to the ν_e appearance search is the same order of

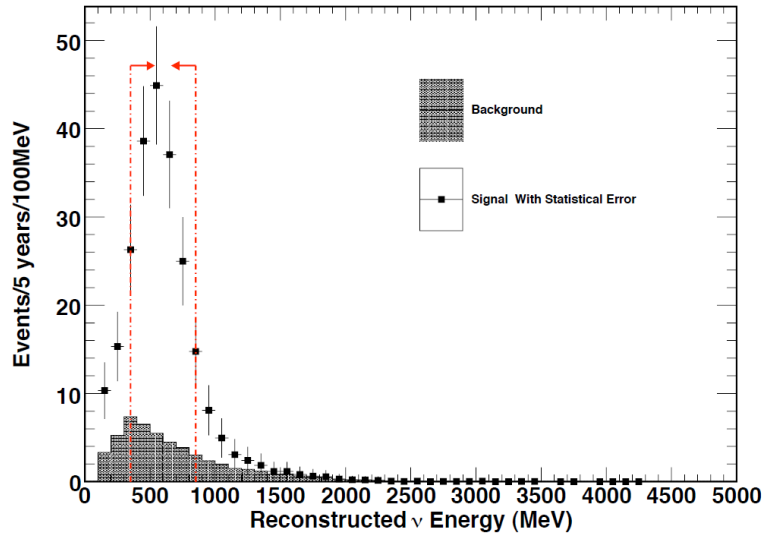


Figure 3.18.: Reconstructed neutrino energy spectrum for a Monte Carlo study of a ν_e appearance measurement, with $\sin^2 2\theta_{13}$ set to the CHOOZ limit of 0.1 [54].

magnitude as the background from the intrinsic ν_e component of the beam (figure 3.4(c)), both of which are much larger than any other backgrounds. Figure 3.18 shows the reconstructed energy spectrum in a ν_e appearance search study. In the selected region, 61% of the background is from intrinsic ν_e contamination, and 30% is from $\text{NC}1\pi^0$ production. Uncertainties in the production cross section are also a limiting factor of the sensitivity of the ν_e appearance search: for smaller values of $\sin^2 2\theta_{13}$, the signal peak might lie within the background uncertainty.

4. ND280 offline software

The ND280 offline software, based on the ROOT software library [55], encompasses the whole process of reading the raw data produced by ND280, calibrating it, reconstructing events, and producing analysis-ready summaries. Monte Carlo production is also handled by the software.

The software is mostly written in C++, with Python scripts used for overall runtime automation, and also as part of higher-level analysis.

4.1. Overall structure

The ND280 software has a modular structure of packages that perform specific tasks. This structure reflects the modular nature of the ND280 detector, with each subdetector group having control over packages specific to that subdetector. There are also several higher-level packages which control how these subdetector-specific packages interact with each other.

The `oaEvent` package is a basis package which provides the data format objects used by the rest of the packages, as well as methods of reading and writing these objects to ROOT-based files. Examples of some of these data objects are a `THit` object (which stores charge, position and time information for charge deposits on detector elements), a `TReconCluster` (which stores a collection of `THit` objects which have been clustered together, along with a position and time), or a `TReconPID` (which represents a recon-

structed particle, with associated goodness-of-fit for any particle identification performed on it). Any two packages that depend on oaEvent can consistently use these objects, and read (write) them from (to) a file that can be written (read) by the other package, and oaEvent ensures this consistency across the whole software.

The data processing chain is controlled by the nd280Control package, which provides Python scripts that run the executables of the other packages. The general processing chain feeds raw data into oaCalib, which feeds into oaRecon, which feeds into oaAnalysis. The oaCalib package performs the initial reading of the raw data files, and channel-by-channel calibration, described further in section 4.3. The oaRecon package runs the reconstruction algorithms on the calibrated data, described in section 4.4. oaAnalysis produces analysis summary trees, which translate the oaEvent-based files into smaller standalone files based only on ROOT objects and trees.

Running parallel to the data processing chain is a Monte Carlo production chain, described in section 4.2. The output of the Monte Carlo chain can be fed into oaCalib, and follows the same path as real data.

There are a few utility packages used at various stages. The oaGeomInfo package provides methods to look up geometrical information for positions in the detector geometry, for example the number of layer in a module at a certain position, or the layer's orientation in space. The oaChanInfo package provides a translation of electronic channels (a certain channel on a front end board, for example) into detector elements (a bar in a module), as well as methods to look up in a database if a certain channel is dead or faulty.

Software development by the author was focussed on the ECal reconstruction package (described in chapter 5, and smaller contributions were made to other packages, such as writing the general framework of the oaChanInfo

package, developing the digit-simulation for the ECal in `elecsim`, general bug solving in `oaEvent`, and also writing a module for `oaAnalysis` to read beam summary information from the T2K beamline group.

4.2. Monte Carlo simulation

The T2K beam group provides a simulation of the flux of neutrinos at the ND280 location. The FLUKA software library [56] simulates the hadronic interactions of the proton beam with the graphite target, and GEANT3 [57] is used to simulate the secondary particles, their transport through the magnetised horns, decay volume and beam dump, and their decays. The simulation produces a list of neutrino vectors, which are n-tuples containing the neutrino energy, the direction as it crosses a central plane of the detector ($z = 0$ in the ND280 coordinate system), the crossing position on that plane (the x, y coordinates at $z = 0$), and a weight for statistical normalisation.

External neutrino interaction generators simulate the interaction of the neutrinos with nuclei. The ND280 Monte Carlo is designed to be generator-agnostic, as long as the generator can handle the input and output data format requirements. The generator has to be able to read the neutrino beam simulation (the list of vectors), and a ROOT-based geometry simulation of the detector (i.e. the spatial distribution of the nuclei). The generators use their own internal mechanisms to simulate the neutrino-nucleon interactions, with all the appropriate weighting applied, and then they have to output a list of interactions. These are n-tuples containing an interaction vertex position, the input neutrino information, the nuclear target information, and a list of outgoing secondary particles along with their four-momenta. The generators can also provide more information, such as the

type of interaction, its cross-section, and any intermediate particle information, but these are not a requirement. The primary generators used on T2K are NEUT [58] and GENIE [59].

The simulation of the ND280 detector, using the nd280mc package, is based around the GEANT4 software library [60]. The detector geometry simulation is built up at first, then the output of the generators is used as an input. GEANT uses the list of interaction positions and the outgoing secondary particles, and simulates the passage (and possible decay) of these secondaries and their daughters through the detector. This is also the stage at which the accelerator bunch timing structure, and the appropriate number of interactions per beam spill, is simulated. GEANT produces a list of energy deposits in the sensitive units of the detector (scintillator bars or TPCs) for each beam spill.

This stage of Monte Carlo also has the capability of simulating a particle gun, which is used for example to test ECal reconstruction by firing a photon into a module. No neutrino beam or interaction simulation is needed, so this MC production just uses GEANT to produce a particle at a required position range, and with a required momentum range.

Electronics simulation, using the elecSim package, takes the deposition of energy in the sensitive detector units, and simulates the electronics response. This includes the attenuation of light in the bars and along the WLS fibres, the response of the photosensors, and effects of noise. Ion drift in the TPCs and the Micromegas responses are also simulated. This information is then digitised with a simulation of the TFB and the AFTER FEC responses, so that the MC output is in the same format as the raw data, a list of digits.

The Monte Carlo can then proceed along the same chain as the real data,

through calibration,¹ reconstruction, and analysis summarisation.

4.3. Calibration

The purpose of calibration is to change the electronics ‘digit’ output (described in section 3.2.6) into ‘hits’, which are calibrated energy deposits in a geometrical position of the detector, at a calibrated time. This is partially done through application of calibration constants on the digits (i.e. channel-by-channel), with some calibration (e.g. attenuation corrections) performed during reconstruction. This section describes the channel-by-channel calibration. The time-dependent constants (i.e. ones with possible spill-by-spill variation) are kept on a database accessed during run-time, whereas the time-independent constants are hard-coded.

4.3.1. TFB calibration

Calibration of the TFB-based subdetectors is identical for the ECal, the P0D and the SMRD.

For the first stage of calibration, pedestal² ADC values (which are continuously monitored as part of the data acquisition process) are subtracted from the ADC values of the digits. Following this, the ADC response is converted to a charge (linearised), using linearity constants generated in special ‘charge-injection’ (CI) runs, when the beam is off. The TFB boards have charge-injection circuitry, which inject known values of charge into the capacitors of the channels. The CI runs are used to build a linearity curve, by injecting various amounts of charge and measuring the ADC response. This step of calibration also takes care of whether to use the high- or low-gain

¹Separate MC-based calibration constants are used.

²The pedestal is the charge accumulated due just to noise in the electronics.

ADC value. The digits have now had the effects of the Trip-T electronics calibrated out.

Next the effects due to the MPPCs are calibrated out. The gain of the devices can be monitored by measuring the charge values of the photoelectron peaks, which are resolvable in the high-gain channel. Other effects accounted for are correlated noise, efficiency and saturation effects (characterised in test bench setups, described in [50]). Also the gain drift (time-dependent variations in the MPPC gain due to temperature and voltage variations) is accounted for, using data from interspill cosmic-triggered events and normalising the response due to muons.

Finally there is some bar-by-bar variation, measured using cosmic runs, that is normalised out across each subdetector. The output is a normalised estimate of the number of photons incident on the MPPC.

The timing calibration involves normalising the TFB-by-TFB time offsets due to different cable lengths (these offsets are calculated from cosmic events), and also calibrating out time-walk effects (charge-dependent timestamps). There are two independent contributions to time-walk. One arises because large energy deposits create more photons. The arrival times of these photons at the MPPC surface have some distribution, for example a normal distribution with mean t_μ . If only a few photons are produced, they will likely all have time close to t_μ (hence the earliest photon time will be close to t_μ); however with more photons produced, some of them will have arrival times in the tails of the distribution, and the earliest photon time will be somewhat earlier than t_μ . The second contribution to time-walk is that a sharper voltage peak is created when more MPPC pixels trigger at the same time. The sharper rise time triggers the Trip-T timestamping closer to the true time that the photons arrived at the sensor. A time-walk

correction is applied by calibrating the times to be earlier if the charges are small, with the level of correction measured in cosmic runs.

4.3.2. TPC calibration

The TPC calibration accounts for the time-independent effects of the electronics linearity (measured in a once-off charge injection test bench run) and Micromegas pad-by-pad response (also measured with a test bench, exposing the pads with a known source of ionisation). The time-dependent effects are due to the gas density, the drift velocity, and Micromegas-by-Micromegas variation.

The gas density is continuously monitored, as it affects the ionisation energy loss of charged particles, the drift velocity, and the Micromegas gain.

The drift velocity of the TPC gas is monitored in two ways: through mini-TPCs, and through cosmic runs. There are two mini-TPCs set up connected to the same gas line as the TPCs. Each contains two ^{90}Sr sources a known distance apart, and measuring the time difference between the two ionisation timing peaks allows the drift velocity to be measured. The other method of monitoring the drift velocity is to use cosmic events where the muons have crossed both the central cathode and the readout plane (both planes a known distance apart), and the time difference between the earliest and latest hits is used to measure the drift velocity.

The Micromegas-by-Micromegas variation is accounted for through measurements made with cosmic events, with minimum ionising particles being the standard candle (once gas density has been accounted for), and comparing the responses of the Micromegas that the particles went through.

4.3.3. FGD calibration

FGD calibration is very similar to the TFB calibration, in that it is a scintillator-bar detector using MPPCs. The main difference is the timing calibration, due to the FGD timing being important for the TPCs. This is performed with the aid of timing markers that are injected into the front end boards in each trigger, and help reduce the timing jitter in the system.

4.4. Reconstruction

Once the data has been calibrated, it passes through a chain of reconstruction algorithms. Reconstruction is performed first on the hits from each individual subdetector. Each subdetector group is tasked with creating and maintaining their own reconstruction algorithms. The reconstruction performed on hits in the Tracker ECal modules is described in chapter 5. Following this, there is a Global reconstruction step, where all the reconstructed objects in the subdetectors are connected and combined (if possible), and interaction vertices are searched for.

Briefly summarised here are Tracker reconstruction and Global reconstruction, as some of the information is used in the analysis presented in chapter 6.

4.4.1. Tracker reconstruction

Reconstruction in the Tracker (TPCs and FGDs) starts with the TPCs. Hits in the TPCs have a three-dimensional position resolution (up to a linear translation along the drift axis), and do not rely on matching of tracks between views, as would be required for FGD reconstruction.

TPC hits close in space and time are clustered together, and tracks are

formed from these clusters. An absolute time for the TPC tracks (t_0) is calculated if the tracks cross the central cathode: as the ions drift in both directions from the central cathode, the maximum drift time is used to calculate t_0 . These tracks are then extrapolated into the FGDs, and using a Kalman filter [61], any FGD hits near the positions where the tracks should be are associated with the tracks. The FGD hits provide t_0 for those TPC tracks which do not cross the central cathode. Using t_0 , the position along the drift axis can be calculated.³

The FGD reconstruction also includes standalone algorithms, used in the case of tracks exiting the FGDs without passing through a TPC. This method relies on a cellular automaton (described in Appendix A of [62]) to connect small straight-line segments of hits in the FGD into longer tracks. Following this, there is matching between two views, making 3D tracks. If only one view has a track, that is left as an unmatched track.

TPC PID

Particle identification in the TPC works on the principle of measuring dE/dx for a track, and using the momentum measurement from the track curvature.

First of all a truncated mean charge is calculated for a TPC track. The truncated mean charge is the mean of the lowest 70% TPC cluster charges in the track, with correction factors due to track lengths and angles. Figure 4.1 shows the truncated mean charge against particle momentum for data and Monte Carlo, for tracks originating in an FGD, in neutrino beam spill data.

For each particle hypothesis (muon, electron, proton, pion, and kaon), a

³For TPC tracks which do not connect with the FGDs, their t_0 is calculated by extrapolation into the ECAL, POD or SMRD, but no connection is done to those subdetectors until later. If no t_0 information is available, then the default value is 5000 ns.

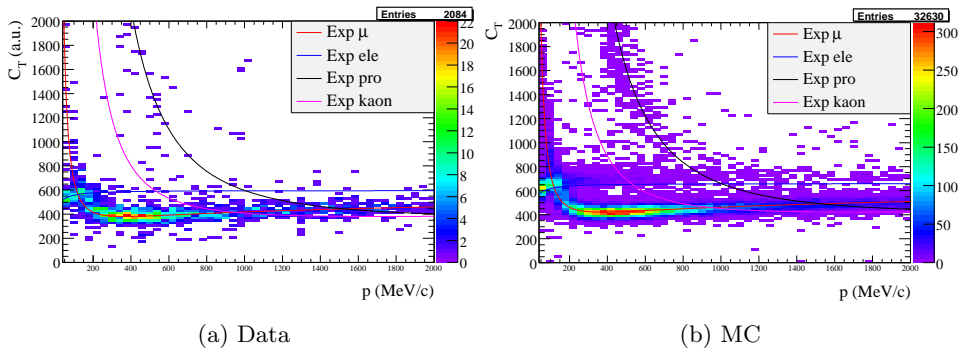


Figure 4.1.: TPC truncated mean charge (in arbitrary units) against momentum for tracks originating in an FGD [63]. The lines are the expected values for muons, electrons, protons and kaons.

pull is calculated. This is a measure of how many σ the truncated mean charge is away from the expected value for that particle and momentum, where σ is a combination of the error on the truncated mean charge, and the error on the momentum measurement.

As an example of the PID performance, if a selection is made for electrons below $1 \text{ GeV}/c$ by selecting the electron pull to be less than 1σ , the fake rate of muons in that sample is 0.19% [64]. If the pull is less than 2σ , the fake rate is 0.72%.

More detailed discussion on the TPC PID and its performance is in [65].

4.4.2. Global reconstruction

Global reconstruction takes all the reconstruction objects in the subdetectors and tries to connect them together with a Kalman filter, and then tries to find common vertices of the objects. The output objects of this algorithm are ‘Global tracks’ and ‘Global vertices’. Global reconstruction uses the RECPACK [66] software library .

The Global reconstruction starts off with Tracker objects, and tries to

match them with objects in the neighbouring subdetectors. It does this through an extrapolation of the Tracker track to the entrance plane of the subdetector, and searches for objects in the subdetector (within 300 ns of the track) near that entry point. A χ^2 fit is performed on the object position and direction, with their uncertainties, of the subdetector objects near the point, and used as a cut: if $\chi^2 < 100$, then the objects are matched together. A Kalman filter is used to refit the new larger object, to recalculate its position and direction. These larger objects are also used recursively until no more pairings are possible. The Global matching then moves onto the other subdetectors, for example trying to match POD objects with SMRD hits, and again this is done recursively until no more matches are possible.

Finally, the reconstruction tries to find vertices using a Kalman filter, based on an implementation used by the COMPASS experiment [67]. If only one Global track exists in a time bin 300 ns wide, then the most upstream end of that track is a Global vertex. Otherwise, Kalman filtering is used to find the best fit vertex position using the tracks. If there are two or more vertices found, then the one with the highest momentum track associated with it is selected as the primary vertex.

4.5. Analysis-ready summarisation

The final step of the software chain is to reduce the reconstructed information into summary trees. These summary trees contain standalone-ROOT readable information (i.e. anyone with access to the trees can read them with their local installation of ROOT, and do not need to install the whole ND280 software). The summary trees contain high-level information, such as the Global vertices, the Global tracks associated with them (and their

TPC PID information, if any), and various reconstructed quantities of tracks and showers in the ECal. Each of these trees is produced by a module in oaAnalysis, for example an ECal reconstruction module produces the ECal-related summary tree. Whether a module runs on a given input file can be independently controlled.

Each subdetector has its own module to output the related summary information. There are also Global-level modules, outputting the summary of the Global reconstruction.

If the summarisation is running on Monte Carlo, then truth information is also stored in these outputs. The truth stored is the interaction vertices, the primary and secondary GEANT trajectories, and the neutrino interaction generator information, if that is available.

Although this summary output is standalone, the ND280 software installation provides example macros (both ROOT '.C' macros and python scripts) that are able to read it, and that analysers can develop to perform their analyses.

4.5.1. Beam summary data

An example oaAnalysis module, initially written by the author, is a beam summary data module. This takes information from the T2K beamline group and attaches it to the oaAnalysis output file on an event-by-event basis. The information attached includes the number of protons per spill, and whether the spill is 'good' (i.e. has been approved for physics analysis), horn currents, and the outputs of the various beamline monitors. An analysis script that runs on the oaAnalysis output can use this information to decide if to use the event in the analysis, or to cut it.

5. Reconstruction performance of the Downstream Electromagnetic Calorimeter

The performance of the ECal reconstruction is important for a π^0 measurement, in particular the reconstruction efficiency of low energy photons, and their angular resolution. The author was involved in developing the reconstruction algorithms with the aim of addressing these issues, these primarily being the low-level clustering algorithms, especially the time-based noise filtering (section 5.1.2), the effort to reduce the minimum number of hits per cluster from 6 to 4 (section 5.1.3), and implementing the thrust-based angular reconstruction (section 5.1.4). The remainder of the algorithms are also described in this chapter, as they are used as part of the analysis in chapter 6.

To test the performance, a particle gun Monte Carlo study was performed by the author, to extract the reconstruction efficiency, energy and angular resolutions for photons. Also summarised in this chapter is a study (performed by others on the experiment) with real data, using a test beam, as a cross-check.

5.1. ECal reconstruction

5.1.1. Hit bunching

The input to the ECal reconstruction stage is a list of calibrated hits. Each hit corresponds to a digit of the electronics output.

The first stage of the reconstruction is to break up each event into ‘bunches’ of hits. Each bunch contains only hits in one view of one module, and the time separation between bunches (between the latest hit of one bunch and the earliest hit of the next) is at least 50 ns. After this is done, the hits in each bunch which occur on the same bar (at both ends) within 21.25 ns^1 of each other are calibrated into a single reconstructed hit. An estimate is made for the position along the length of the bar, based on the times at the two ends. If the hit is single-ended, then the position is estimated to be at the centre of the bar. The time of this reconstructed hit is calculated from the times at the ends and the effective speed of light in the bar. The charge of the reconstructed hit is calculated using the attenuation measurement of the bar and the estimated position of the hit. The calibrated charge unit of the hit at this stage is the MIP-equivalent unit (MEU). This is the most probable value of charge deposited by a minimum-ionising particle travelling directly incident into the bar (i.e. with a path length of 1 cm).

5.1.2. Clustering

The next stage of reconstruction is to cluster hits together. This is performed on each 2D view and bunch separately. Hits which are in next-to-nearest-neighbouring layers, and nearest-neighbouring bars, are clustered

¹This number is a combination of the time light takes to propagate along the 2 m bar, and the timing resolution of both hits. The time cut will be different for the Barrel modules, as the double-ended bars are 4 m long.

together if the time difference between them is less than 15 ns. The time cut is used to reduce the incidence of noise hits being clustered with real hits or even with other noise hits. There is a three-hit minimum threshold to produce a cluster in this step.

For each such cluster, a principal component analysis (PCA) method (see section 5.1.4) is used to determine the cluster's major and minor axis (\vec{a}_{\parallel} and \vec{a}_{\perp} respectively), and the eigenvalues of these axes (σ_{\parallel} and σ_{\perp}).

Next, the clustering algorithm looks for cases where two clusters can be combined. This can happen for example in showering, where clusters could become separated. For each pair of clusters, the distance d is given by $d = |(\vec{x}_2 - \vec{x}_1) \times \vec{a}_{\parallel,1}| \sec \theta_1$, where \vec{x} is the cluster's centre, the subscripts 1,2 denote the higher- and lower-charged clusters, and θ is the angle of incidence of \vec{a}_{\parallel} into the module. The pair of clusters will be combined together if d is less than 80 mm, and the charge-weighted mean times of the clusters are within 40 ns of each other. This algorithm is applied recursively (working with the newly formed clusters, recalculating the PCA axes and eigenvalues) until no more pairings pass the cut. There is a slight probability that a single showering particle will produce two or more clusters, even after this stage of clustering.

Following this, any remaining hits in the same bunch are checked to see if they can be clustered with a pre-existing cluster. For each remaining hit and cluster pair, if the time difference between the two is less than 40 ns, a weight w is given by $w^2 = [(\vec{p} - \vec{x}) \cdot \vec{a}_{\parallel}/\sigma_{\parallel}]^2 + [(\vec{p} - \vec{x}) \cdot \vec{a}_{\perp}/\sigma_{\perp}]^2$, where \vec{p} is the hit position and \vec{x} the cluster centre. For each hit, the cluster with the lowest weight, but not exceeding 80 mm, is chosen, and the hit is added to that cluster.

The performance of the clustering has been previously assessed in [68].

5.1.3. 3D matching

There are now 2D clusters in the two views of a module, and these are combined to make 3D clusters. Only one 2D cluster from each view makes up the 3D cluster, and they have to be in the same bunch.

The matching process constructs likelihoods for all combinations of 2D clusters. The likelihood is a product of one-dimensional pdfs, with no correlations taken into account. As inputs to the likelihood, these variables are calculated: the charge ratio of the two views (r_Q), the ratio of ‘axis max ratios’ in the two views (r_{AMR}), the difference in the innermost layers of the two views (d_-), and the difference in outermost layers of the two views (d_+). These are defined as:

$$r_Q = \frac{\sum_{i,1} q_{i,1}}{\sum_{i,2} q_{i,2}}, \quad (5.1)$$

$$r_{\text{AMR}} = \frac{\sigma_{\parallel,1}/\sigma_{\perp,1}}{\sigma_{\parallel,2}/\sigma_{\perp,2}}, \quad (5.2)$$

$$d_{\pm} = |l_{\pm,1} - l_{\pm,2}|, \quad (5.3)$$

where the subscripts 1 and 2 indicate the two views, q_i is the charge of the i th hit in the cluster and i runs over all the hits in that cluster, σ_{\parallel} and σ_{\perp} are the PCA eigenvalues of the cluster as defined in section 5.1.2, and $l_{-(+)}$ are the innermost (outermost) layers of the cluster.

There is also a ‘seeded’ likelihood constructed. Reconstructed Tracker tracks are extrapolated into the ECal modules, and used as seeds for the matching. For each 2D cluster, the hits in the innermost two layers of the module (if there are any) are checked to see how far they deviate from the Tracker tracks’ positions and times at the inner face of the module, and the likelihood is based on this deviation.

When the total likelihood is worked out, there is a likelihood cut, such that even if there are two 2D clusters that could be matched together, they might not have a large enough likelihood, and so remain as unmatched 2D clusters.

The algorithm combines the 2D cluster combinations with the largest likelihoods together into 3D clusters. There is a minimum of six hits in a 3D cluster produced in this manner.

There may also be leftover unmatched clusters in one of the views, for example if there is one cluster in a view and two clusters in the other view, after the 3D matching there will be a 3D cluster and an unmatched 2D cluster. In this situation, an attempt is made to match these with unclustered hits in the other view. The only requirement is that the hits in the other view be in time with them, and in overlapping layers. The timing requirement is that the unclustered hits be between 10 ns before the start time and 10 ns after the end time of the unmatched cluster. The layer requirement is that the unclustered hits must be between one layer before the first, and one layer after the last, layers of the unmatched cluster. There are no bar-distance requirements of the hits, meaning situations can arise where two hits are on opposite sides of the module. The motivation for this step is to reduce the minimum hit threshold for a reconstructed cluster from six hits to four hits, thereby increasing the efficiency for reconstructing low-energy photon showers.

Following the 3D clustering, the hits in each bar are reassigned a position along the bar, using information from the other view of the cluster. The reconstructed hit charge and times are recalibrated using this new information.

5.1.4. Angular reconstruction

Good angular reconstruction is a very important requirement for a π^0 analysis. In most neutral current π^0 events, the only indication of the interaction will be two showers in the ECal modules, and the interaction vertex will not be evident. The reconstruction has to be able to take the two clusters and point them back to an estimated vertex, and the location of this vertex is strongly dependent on the angles of the two clusters.

A robust method of angular reconstruction of a cluster is by using the aforementioned principal components analysis. This is generally used in data analysis to transform a multi-dimensional data set to fewer dimensions which should be uncorrelated, returning eigenvectors and eigenvalues of the principal components of the data set. In the case of PCA being used in angular reconstruction, the data set input is the positions of the hits in a cluster, weighted by charge. The eigenvectors returned are the directions of the major and minor axes of the ellipsoid surrounding the hits, and the eigenvalues are proportional to their lengths. The reconstructed angle of the cluster is then given by the major axis.

The PCA of a cluster is used in other algorithms in the reconstruction. If the algorithm runs on a 2D cluster, then the third component of the position of the hits (the position along the bar) is explicitly set to 0, and PCA returns two axes and two eigenvalues. If the algorithm runs on a 3D cluster, all components of the hit positions are used, and the method returns three axes and three eigenvalues. The PCA centre is sometimes needed as well; this is centre of the ellipsoid.

Another method of determining the angle is to use a ‘thrust’ algorithm, adapted from a similar method used in jet physics [69, 70]. This method is less robust because makes a prior assumption that a particle is travelling

from the inner detectors (from the FGDs or TPCs) outwards when it enters the calorimeter module at its innermost face. This assumption is valid for the signal photons used in the analysis presented in this thesis, but in general it will not be valid; the cases where the assumption fails include particles coming in from interactions in the magnet, interactions within a calorimeter module itself, and cosmic rays.

The first step is to find the ‘thrust origin’ ($\vec{\sigma}$), which is the estimated position where an electromagnetic shower starts. The algorithm finds the hits in the innermost layer of the cluster (for the DsEcal, the most upstream layer) and then takes the charge-weighted mean position of these hits. This is mathematically expressed as

$$\vec{\sigma} = \frac{\sum_j q_j \vec{x}_j}{\sum_j q_j}, \quad (5.4)$$

where q_j is the j th hit’s charge and \vec{x}_j its position, and j runs over the hits in the innermost layer. By placing the estimated shower start in the innermost layer of the cluster, this reinforces the assumption that the particles are travelling outward through the module.

Having found the thrust origin $\vec{\sigma}$, the quantity t , which depends on two angles θ and ϕ , is defined as

$$t(\theta, \phi) = \frac{\sum_i q_i |\vec{n}_{\theta, \phi} \cdot (\vec{x}_i - \vec{\sigma})|}{\sum_i q_i |\vec{x}_i - \vec{\sigma}|}, \quad (5.5)$$

where q_i is the i th hit’s charge and \vec{x}_i its position, i runs over all the hits in the cluster, and $\vec{n}_{\theta, \phi}$ is a unit vector with polar angles θ, ϕ . This quantity is essentially the charge-weighted mean longitudinal (along $\vec{n}_{\theta, \phi}$) displacement of all the hits from the thrust origin. The algorithm maximises this quantity t over all angles θ, ϕ . The ‘thrust axis’ (\vec{a}) is $\vec{n}_{\theta, \phi}$ for the angles θ, ϕ which

maximise t , and that value of t is the ‘thrust’. Using this method, the reconstructed angle of the cluster is the angle of the thrust axis.

The thrust can also be used to discriminate between tracks and showers. For hits lying in a perfect straight line, as could happen with a MIP-like track, the thrust of this cluster will be 1. This is because \vec{n} will also lie along that line, and the displacements of the hits from the origin will only have longitudinal components, and so the numerator and denominator of equation 5.5 will be equal. For any other cluster the thrust will be less than 1: at least one of the hits will lie off the straight line, and so the displacements of the hits will have a longitudinal and transverse component, and the numerator of equation 5.5 will be smaller than the denominator. Figure 5.1 shows the thrust for Monte Carlo photons and muons entering the DsEcal from upstream. The muons have a thrust strongly peaked at 1, and photons have a thrust less than (but close to) 1.

5.1.5. EM energy reconstruction

The EM energy reconstruction of a cluster relies on the distribution of charges in the cluster. Four variables are calculated for the cluster: the charge sum Q_{sum} , the charge mean Q_{mean} , the charge RMS Q_{rms} , and the charge skew Q_{skew} . These are calculated as follows:

$$Q_{\text{sum}} = \sum_i q_i \quad (5.6)$$

$$Q_{\text{mean}} = \frac{Q_{\text{sum}}}{n} \quad (5.7)$$

$$Q_{\text{rms}} = \frac{1}{Q_{\text{mean}}} \sqrt{\frac{\sum_i (q_i - Q_{\text{mean}})^2}{n}} \quad (5.8)$$

$$Q_{\text{skew}} = \frac{1}{Q_{\text{mean}}^3 Q_{\text{rms}}^3} \frac{\sum_i (q_i - Q_{\text{mean}})^3}{n} \quad (5.9)$$

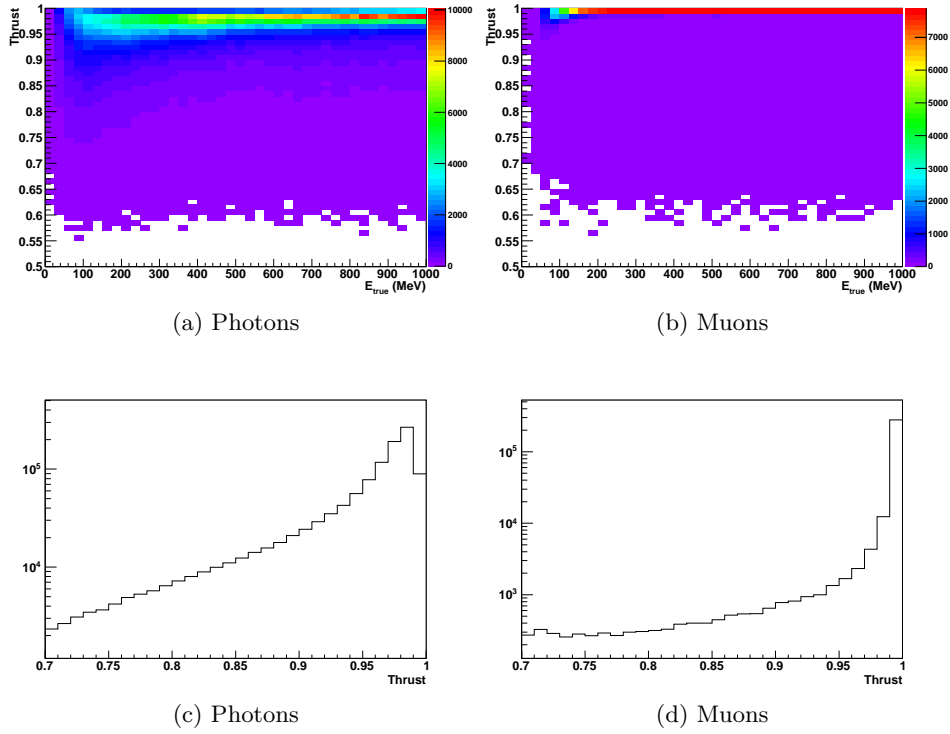


Figure 5.1.: Thrust for Monte Carlo particles of kinetic energies 0–1 GeV, with a uniform position distribution in a $1\text{ m}\times 1\text{ m}$ box parallel to and centred with the front face of the module, and at all angles.

where q_i is the charge of the i th hit, and i runs over all the hits in the cluster, with n being the total number of hits.

The energy reconstruction is first tuned with Monte Carlo by simulating photons being fired into the DsECal. 53 energy points in the range 0.075–25 GeV are used. For each energy, distributions of Q_{sum} , Q_{rms} and Q_{skew} are created, and then fitted with a Gaussian. Cubic splines are then produced for these means and widths as a function of energy. Splines are also produced for the correlations between pairs of the three distributions, to produce an error matrix.

To reconstruct the energy of a cluster, a log-likelihood fit is used. This uses the splines from the tuning process as one-dimensional pdfs, and takes into account the correlation between the variables. The likelihood depends on the deviations of the cluster’s values of Q_{sum} , Q_{rms} and Q_{skew} , from the expected values for a given energy hypothesis. The reconstructed EM energy is then the energy which optimises this likelihood, and that optimised value of the log-likelihood (L_{EM}) is also stored for use in other algorithms.

5.1.6. Particle identification

Particle identification (PID) in the ECal reconstruction relies on charge and shape distributions of the cluster. In addition to the Q_{skew} and L_{EM} variables described in section 5.1.5, the process also requires calculation of the ‘axis max ratio’ (AMR), the maximum/minimum layer charge ratio, the shower cone (or opening) angle, shower width, asymmetry, and the mean position.

The AMR is calculated by first performing a two-dimensional PCA on the hits from each view of the cluster. In each view, the ratio of the major to the minor axes is calculated ($\sigma_{\parallel}/\sigma_{\perp}$, using the labelling of section 5.1.2),

with a cut at 300 to stop the ratio becoming very large if σ_{\perp} is small, as can happen with straight tracks. This ratio is a measure of how long a cluster is to its width. The AMR for the whole three-dimensional cluster is then the average of the two view's ratios.

For each layer in the cluster, the layer charge is the sum of charges of hits in that layer. The layer charge ratio of the cluster is then (maximum layer charge)/(minimum layer charge).

A three-dimensional PCA is used for calculating the shower cone angle, shower width, asymmetry, and the mean position. The outputs of the PCA are three axes (\vec{a}_{\parallel} , \vec{a}_{\perp} and \vec{a}_3 , in order of significance), and their eigenvalues (σ_{\parallel} , σ_{\perp} and σ_3 respectively). The shower cone angle is given by $\arctan(\sigma_{\perp}/\sigma_{\parallel})$. The shower width is just σ_{\perp} . The asymmetry is given by σ_3/σ_{\perp} (a measure of how circularly symmetrical a cluster is in the plane perpendicular to its major axis). The mean position of a cluster is given by $|\min((\vec{x}_i - \vec{\sigma}) \cdot \vec{a}_{\parallel})|$, where $\vec{\sigma}$ is the PCA centre of the cluster, \vec{x}_i the position of the i th hit, and i runs over all the hits in the cluster. This finds the hit with the most negative component along the PCA axis, and the shower mean is then the absolute value of this component; this is a measure of how far the shower mean position (the PCA centre) is from the shower start.

The PID is designed to produce two discriminants: a track/shower discriminant, and a hadronic/EM shower discriminant. The inputs to the track/shower discriminant are the AMR, the EM likelihood, the AMR, the layer charge ratio, the shower cone angle, and the shower width. The inputs to the hadronic/EM shower discriminant are Q_{skew} , the layer charge ratio, the shower width, asymmetry and mean position.

The PID is performed using an artificial neural network (ANN) which has been trained with Monte Carlo particles (photons, electrons, muons,




Figure 5.2.: Track/shower discriminant for Monte Carlo particles of kinetic energies 0–1 GeV, with a uniform position distribution in a $1\text{ m}\times 1\text{ m}$ box parallel to and centred with the front face of the module, and at all angles.

and pions). More details on the ANN training process and its performance is given in [71].

Each discriminant is a value nominally between 0 and 1. For the track/shower discriminant, 0 indicates that a cluster is shower-like, and 1 indicates it is track-like. For the hadronic/EM discriminant, 0 indicates that a shower is hadronic-like, and 1 indicates it is EM-like. Figure 5.2 shows the track/shower discriminant distributions for Monte Carlo muons and photons.

5.1.7. Michel tagging

A Michel electron is the electron produced in a muon decay [72, 73, 74]. Muons that stop in an ECal module should produce one of these, and an algorithm is designed to look for their evidence, which will look like a low-energy delayed cluster near the track end point. If a cluster has its end point more than 8 cm from any side of the module, then a search is performed for Michel tag candidates with this cluster as a seed. The algorithm searches

for any other clusters and unclustered hits which satisfy the following cuts:

$$t - t_c > 30 \text{ ns}, \quad (5.10)$$

$$|\vec{x} - \vec{x}_c| < 10 \text{ cm}, \quad (5.11)$$

$$1 \text{ MEU} < Q_{\text{sum}} < 35 \text{ MEU}, \quad (5.12)$$

$$N_h < 10, \quad (5.13)$$

where t is the candidate hit or cluster time, t_c is the time of the seed cluster, \vec{x} is the candidate hit or cluster position, \vec{x}_c is the seed cluster end position, Q_{sum} is the candidate cluster charge sum, and N_h is the number of hits in the candidate cluster; the cuts of equations 5.12 and 5.13 only apply to candidate clusters, not hits. The hits from these candidates are split into ‘bunches’, with at least 30 ns between the hit times in two bunches, and each of these bunches becomes a Michel tag.

5.2. ECal Performance

Figure 5.3 shows the energies of the two π^0 decay photons, for all π^0 s produced in neutral current interactions in the FGDs. The low energy photon has a peak at around 75 MeV, while the high energy peaks around 150 MeV.

For the analysis presented in this thesis, it is important to know the performance of the calorimeters with photons of these relatively low energies. Two datasets are used to check this: a Monte Carlo simulation of photons being fired into the DsECal, and data from when the DsECal was placed in a test beam at CERN.

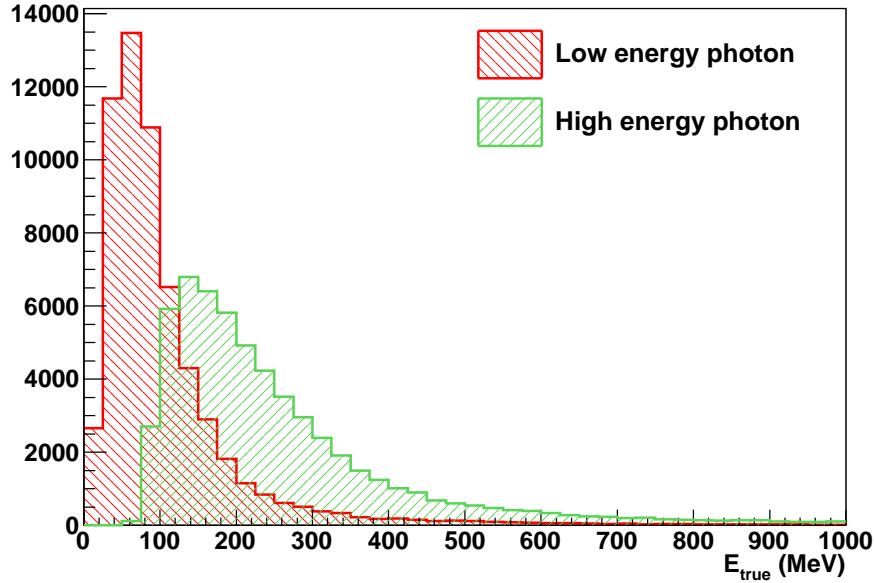



Figure 5.3.: Monte Carlo π^0 decay photon energies

5.2.1. Monte Carlo


To study the performance of the reconstruction with Monte Carlo, photons were fired into the DsECal. The simulated photons started 6.5 cm upstream of the module, in a $1\text{ m}\times 1\text{ m}$ box parallel to and centred with the front face of the module. They had energies of 0–1000 MeV and incidence angles of 0° – 90° .

Reconstruction efficiency


The reconstruction efficiency is defined as the fraction of reconstructed clusters for photons which convert in the DsECal module. Figure 5.4 shows the dependence on energy and angle of the reconstruction efficiency for the photons. Below around 100 MeV the efficiency falls quickly. This is because the clustering requires a minimum of four hits, and at lower energies this is



(a) Energy and angular dependence



(b) Energy dependence



(c) Angular dependence

Figure 5.4.: Reconstruction efficiency of Monte Carlo photons.




Figure 5.5.: Reconstructed fractional energy bias of Monte Carlo photons in the DsECal.

less likely to happen; an additional factor at low energies is that the hits might not be close together, so will fail to cluster as they are not nearest neighbours, even if the photon produces three or more hits in a view. The cut-off coincidentally happens to be around the mean low-energy π^0 -photon energy, meaning in general the π^0 reconstruction efficiency is badly affected by this. The angular dependence is flat until very oblique angles, where it falls off. For photons produced in the FGDs and entering the DsECal, there is a maximum possible incidence angle of 1.23 rad, which is below the point where the efficiency starts falling off.

Energy resolution and bias

Figure 5.5 shows the fractional energy bias, $(E_{\text{reco}} - E_{\text{true}})/E_{\text{true}}$, of the reconstructed EM energy. There is an overall bias of 6% for higher energy




Figure 5.6.: Example energy slices from figure 5.5.

photons, which becomes smaller as the energy decreases below 200 MeV. The bias arises because the EM energy fitter has not been tuned to this version of the reconstruction. In particular, the 3D matching has been substantially modified since the fit was last tuned.

An energy resolution can be extracted by taking each energy slice of figure 5.5 and fitting a Gaussian to it. The fractional energy resolution, σ_E/E , is given by the width of the fitted Gaussian for that energy slice. Figure 5.6 shows two such Gaussian fits at different energies. The fractional energy resolution at each energy point can then be fitted with the standard energy resolution function [32]:

$$\sigma_E/E = \sqrt{(a/\sqrt{E})^2 + (b/E)^2 + c^2}. \quad (5.14)$$

Figure 5.7 shows the fractional energy resolution, with the dotted line being the fit. The fitted terms of the resolution (with E in GeV) are: $a = (7.01 \pm 0.06) \%\sqrt{\text{GeV}}$, $b = (1.08 \pm 0.05) \%\text{GeV}$, and $c = (7.36 \pm 0.07) \%$. The stochastic term, 7.0%, is comparable to the design requirement of 10% or better.




Figure 5.7.: Energy resolution of Monte Carlo photons

Angular resolution

Figure 5.8 shows the angular ‘deviation’ of the clusters, which is the angle between the reconstructed direction and the true direction, for both types of angular reconstruction.

To extract the angular resolution, energy slices of figure 5.8 are fitted with a function of the following form:

$$\alpha \cdot \theta \cdot \frac{1}{\sqrt{2\pi}\sigma_\theta} e^{-\theta^2/(2\sigma_\theta^2)} + \beta \cdot e^{-\theta/\lambda},$$

where θ is the angular deviation, σ_θ is the angular resolution, α and β are normalisation factors, and λ a decay parameter. This is a Gaussian of width σ_θ centred on 0 and multiplied by the deviation; in addition there is an exponential decay that represents a background, which is included to




Figure 5.8.: Reconstructed angle deviations of Monte Carlo photons.




Figure 5.9.: Example energy slices from figure 5.8

improve the fit. The reason the function is of the form $x \times \text{Gaussian}(x)$ is that there is less phase space for reconstructing angles closer to the true angle. The angular resolution for the energy slice is the width of the Gaussian, σ_θ . Figure 5.9 shows examples of these fits. Figure 5.10 shows the angular resolution as a function of the true energy of a photon, for both types of angle reconstruction. The thrust method has a better resolution than the PCA method. The fits applied to these are the same form as equation 5.14:

$$\sigma_\theta = \sqrt{(a/\sqrt{E})^2 + (b/E)^2 + c^2},$$

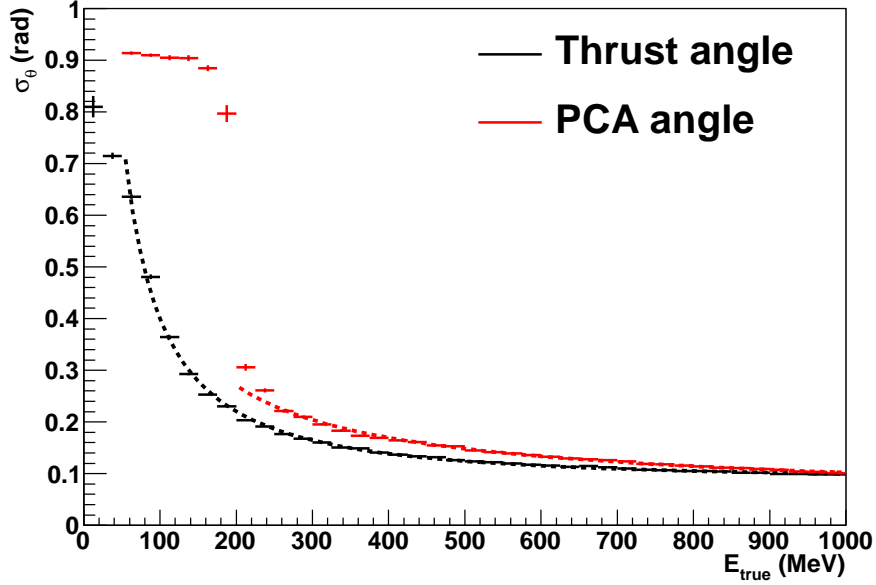


Figure 5.10.: Angular resolution of Monte Carlo photons

with E in GeV. For the PCA angle, the fit results are $a = (86 \pm 3) \text{ mrad}\sqrt{\text{GeV}}$, $b = (37 \pm 2) \text{ mradGeV}$ and $c = (40 \pm 5) \text{ mrad}$. For the thrust angle, the results are $a = (37 \pm 2) \text{ mrad}\sqrt{\text{GeV}}$, $b = (37 \pm 0) \text{ mradGeV}$ and $c = (85 \pm 1) \text{ mrad}$. The thrust method performs much better at energies below 300 MeV that are typical of a π^0 decay photon.

5.2.2. Test beam

In April–July 2009, the Downstream ECal module was placed in the T9 secondary beam line of CERN’s Proton Synchrotron, which was used to test the performance of the module, in particular the particle identification and EM energy resolution. The test beam provided electrons, charged pions and protons, with tunable momenta between 300 MeV/c and 15 GeV/c. The particles were incident onto the DsECal from the downstream face, due to




Figure 5.11.: Fractional energy resolution of test beam and MC electrons, as a function of beam momentum, with fitted resolution functions drawn as dashed lines [75].

triggering requirements. Various angles of incidence (0° , 15° , 30° and 75° horizontally) and entry positions (in the centre of the module face, and 25 cm from the side edge) were used. Two external detectors, a time-of-flight (TOF) counter with a baseline of 14 m, and two Čerenkov detectors, were used for beam characterisation. The TOF counter was used to distinguish protons from the lower-mass particles, and the Čerenkov detectors were set up to distinguish electrons (above threshold) from the higher-mass particles (below threshold), requiring that both Čerenkovs agreed.

The test beam electrons can be used to test the performance of the EM energy reconstruction. Using the TOF and Čerenkov information as a cut, and requiring only one reconstructed cluster per event, the electron sample is over 99% pure [41].

Figure 5.11 shows the measured fractional energy resolution of electrons, σ_E/E , which has been calculated from Gaussian fits to $(E_{\text{reco}} - p_{\text{beam}})/p_{\text{beam}}$,




Figure 5.12.: Reconstructed energy of test beam electrons as a function of beam momentum, with a linear fit in red [41].

and the energy of an electron is approximately p_{beam} . A resolution function, $\sigma_E/E = a/\sqrt{E[\text{GeV}]}$, is fitted for both test beam data and MC. In MC $a = (7.3 \pm 0.2)\%$, while in data $a = (9.8 \pm 0.2)\%$ [41]. The poorer resolution in data is probably due to thermal variations during the test beam (there was a 3% variation in the DsECal muon energy scale over the whole test beam run due to temperature fluctuations).

Figure 5.12 shows the reconstructed energy of the electrons against the beam momentum, along with a linear fit, constrained at the origin. The fit gradient ('p1' on the plot) shows that there is a bias of around 8% in the data. This bias is a different type to the one in section 5.2.1. That bias comes about due to the energy fit being detuned, and is assumed to affect data and Monte Carlo in the same way. This bias is a correction between the data and Monte Carlo, and was extracted using an older version of the reconstruction, where the energy fit was tuned. An assumption that this correction still applies should be checked, by running the test beam

analysis with the same version of the simulation and reconstruction used in this thesis.

Even though there are minor differences between electrons and photons, it is assumed that the bias applies to both, and so this 8.3% energy correction is applied to all the data clusters used in the analysis.

5.2.3. Summary

Using Monte Carlo particle gun simulations, the reconstruction efficiency for photons is found to be around 50% for the typical low-energy π^0 decay photon, and around 95% for the high-energy photon. The energy resolution has a stochastic term of $7\%/\sqrt{E/\text{GeV}}$. The angular resolution has a stochastic term of $86\text{ mrad}/\sqrt{E/\text{GeV}}$.

6. Reconstruction of neutral pions in the Downstream Calorimeter

This analysis aims to make a measurement of neutral current π^0 production in the FGDs using reconstructed clusters in the Downstream ECal. The definition of the signal is any neutral current neutrino interaction which produces a π^0 in the final state, and both π^0 decay photons convert in the DsECal.

The reason that only the DsECal reconstruction is used in the analysis is that, at the time of writing, the calibration and reconstruction for that module was more mature and better understood. The DsECal had been extensively tested, not only with the CERN test beam (section 5.2.2) but also with cosmic rays, at CERN and in ND280.

The remaining ECal modules were installed almost as soon as they were built, with very little opportunity for testing them beyond basic checks, e.g. if the electronics worked, and the photoelectron spectra of the channels were fine. Some aspects of their calibration, for example calculating their bars' attenuations *in situ*, had not been performed as of writing, and there were a few issues with the reconstruction that hadn't been ironed out.

This means that the phase space of π^0 s available to be analysed is higher momentum and very forward-going with respect to the beam direction (dis-




Figure 6.1.: Monte Carlo π^0 decay photon energies, for those π^0 s where both photons convert in the Downstream ECal.

cussed further in section 6.3.1). This also means that the decay photon energies are higher than normal (figure 6.1, in comparison with figure 5.3), and so the single photon reconstruction efficiency fall-off for photon energies below around 100 MeV (figure 5.4) should be less of an issue.

The analysis cuts presented in this chapter are all Monte Carlo driven, with no regard for the distributions in data.

6.1. Data set

The data used in this analysis was taken between March 2010 and March 2011. It is split into two physics runs, Run I (March–June) and Run II (November–March), with a summer maintenance shutdown in between. The spills in Run I contained six bunches of protons, with eight bunches per spill




Figure 6.2.: Accumulated protons and protons per spill (pulse) for MR runs 31–34 and 36–38 [54].

delivered in Run II. Of the ECal modules, in Run I only the Downstream module was working, with the Left Top Barrel (LTB) ECal module installed but, due to a short in the electronics, unable to read data out, so essentially acting as dead material with a large mass. In the summer shutdown, the remaining ECal modules were installed, and the LTB module was fixed.

Each physics run was further subdivided into Main Ring (MR) runs, numbered 31–34 in Run I and 36–38 in Run II. These were periods where the accelerator complex was delivering protons with a stable operation, and the ND280 hardware and DAQ system were kept in a continuous configuration. In between the MR runs, there was possibility for alterations in the accelerator and/or detector configurations, leading to some variations in the data on a run-by-run basis. This is especially evident with the bunch timings, and is discussed further in section 6.2. Figure 6.2 shows the accumulated protons and the protons per spill for the data set.¹

¹The accumulated number of protons in the figure includes MR runs before 31, however the number is very small compared to the total. This data is used in the SK and INGRID analyses, but not in any ND280 analyses, as the detector was not fully commissioned at the time.

6.1.1. Data quality

Over the course of the data taking, there were checks made to ensure the data was of a good quality. These were made at both a hardware level, e.g. checking that voltages were in a good range, and also at a data level, e.g. looking at the photoelectron spectra across all the channels, or seeing if the hit occupancy for interspill cosmic events was consistent. A data quality group provides spill-by-spill information on whether the data is of a good quality. Over the whole period of data taking, there was one significant deficiency in the data quality: in December TPC3 had to be turned off due to a failure in the front end electronics, which was not fixed until a shutdown between MR runs 36 and 37. Analyses that do not depend on the operation of TPC3 are still able to use the data in this period; however this analysis relies on TPC3 to be working. The statistics lost because of this was approximately 2×10^{19} POT.

A further data quality cut is to use the beam summary data (described in section 4.5.1) to select only those spills where the beamline group has said that the data is of a physics analysis quality. This means for example the horn currents and beamline monitors were within nominal operating parameters.

The total analysed data is 2.93×10^{19} POT in Run I and 7.71×10^{19} POT in Run II. The total POT sum, 10.64×10^{19} POT, is referred to as POT_{data} in the plots later on in the chapter.

6.1.2. Monte Carlo

The GENIE neutrino interaction Monte Carlo generator [59] was used to produce 1.71×10^{21} POT worth of statistics, 33% of the statistics being in the Run I configuration (in the data this ratio is 28%). The beam power

simulated for Run I was a constant 50 kW (equivalent to 37.8×10^{12} protons per spill), and 100 kW (66.7×10^{12} protons per spill) for Run II.

The full spill Monte Carlo sample contains multiple interactions per spill, a mean of 3.7 interactions per spill for Run I, and 6.9 for Run II. Most of these interactions occur in the Magnet, due to its large mass.

6.2. Bunch timings

As a first step in the analysis, the times of the individual bunches that make up a spill have to be determined, as these are used in the preselection, and later on for time offsets between subdetectors.

For each MR run, a histogram of DsEcal cluster times is filled. A fit is made to this histogram, consisting of six or eight Gaussian peaks. Each peak has the same width, and they are separated by 582 ns. Examples of these histograms, showing the fit result, are in figure 6.3. For both data and Monte Carlo, the bunch peaks are superimposed over a background noise spectrum. These are clusters produced purely from noise hits, which can happen if the hits are close together in time and space. The noise spectrum extends up to around 200 ns after the bunch peak, and has a sawtooth shape, with more noise clusters later than earlier. The reason for the sawtooth shape is discussed in section 3.2.6. The Monte Carlo models the noise well, with the only major difference between data and MC being the time of the bunch with respect to the noise background; this time is earlier in the Monte Carlo. The relative amount of noise is also lower in the MC: the peak noise level is around 200 times smaller than the peak bunch level, while in the data that ratio is around 50–100.

Table 6.1 shows the Gaussian mean times of the first bunch peak, and


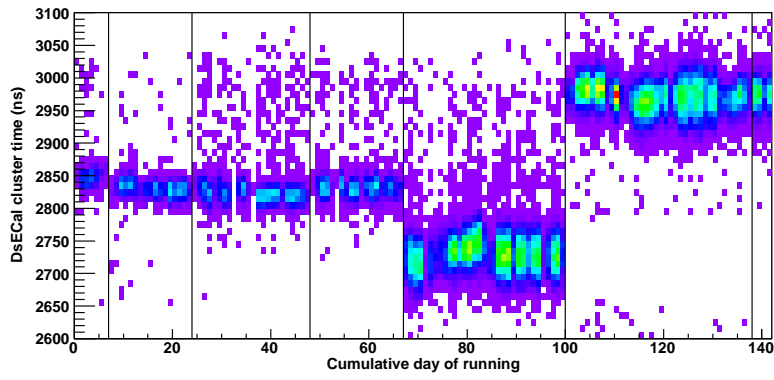


Figure 6.3.: Histograms of DsECal cluster times in green, with the bunch timing fit result drawn as a black line.

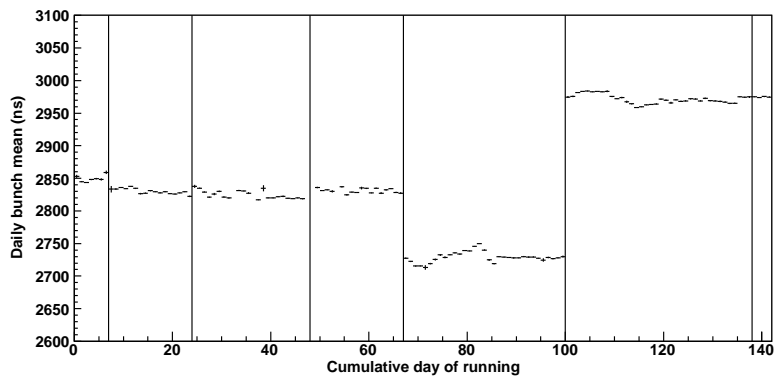
MR run number	μ_b (ns)	σ_b (ns)
31	2848.12	10.93
32	2830.25	11.47
33	2822.74	13.01
34	2832.36	12.86
36	2735.18	26.89
37	2982.97	21.82
38	2974.74	20.79
Monte Carlo	2744.46	9.07

Table 6.1.: First bunch mean times (μ_b) and widths (σ_b) for reconstructed clusters in the Downstream ECal for the Main Ring (MR) runs

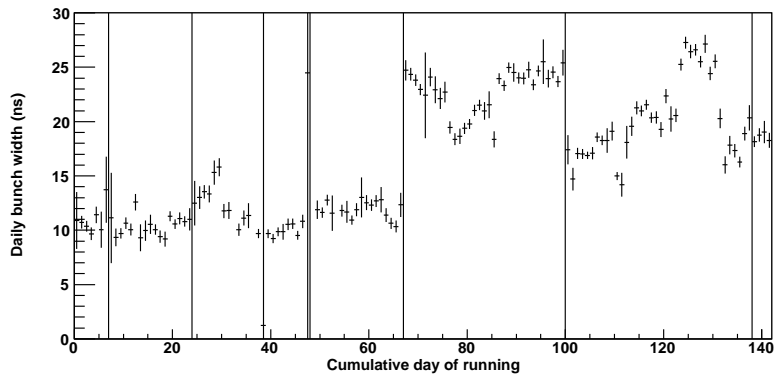
the Gaussian width of all the bunches, for each MR run and also for the Monte Carlo. There is some variation of the bunch times between MR runs, due to accelerator tuning and/or hardware changes in the periods between the MR runs. The bunch width also changes between Run I and Run II, approximately doubling. This is most probably due to the beam power increasing, with more protons per bunch being harder to keep tight in time. The variation of the bunch times on a daily basis is shown in figure 6.4. There is some drift in the day-to-day bunch mean time, which could account for some of the increase in the measured width. For example, in MR run 36, the overall bunch width is 26.89 ns, yet the daily widths do not exceed 25 ns, so the width is probably increased by the drift in the daily mean time. The daily drift in the mean time is probably due to accelerator variations.



(a) Cluster times



(b) Bunch mean



(c) Bunch width

Figure 6.4.: Variation of the first bunch cluster time, bunch mean time and bunch width on a daily basis. Vertical black lines separate the MR runs, which are (l-r) 31, 32, 33, 34, 36, 37 and 38.

6.3. Preselection

A preselection is performed on the output of the reconstruction, which is described in chapter 5. All DsECal objects in an event are split up into bunches, with a bunch defined as a period of 582 ns duration centred on the mean bunch time, as calculated in section 6.2. The window is chosen to be the full 582 ns bunch spacing to be as inclusive as possible for the preselection.


Within each bunch, for each pair of isolated DsECal clusters a π^0 candidate is preselected. A cluster is determined to be isolated if the global reconstruction has not connected it with any other track, cluster or hit in another part of the detector. This is especially true for TPC3 tracks, as most charged particles entering the DsECal from upstream should leave tracks in this TPC, and is the reason that this analysis relies on TPC3 to be operational.

If there are more than two isolated clusters in the same bunch, then each possible pairing produces a candidate, under the assumption that later stages of selection will reduce the candidates to only one (or zero) viable combination of clusters in the bunch.


Plots of the two clusters' reconstructed energies, angles of incidence, and the reconstructed mass assuming they are π^0 decay photons (described in section 6.5.3) are shown in figure 6.5.

6.3.1. Geometrical acceptance and reconstruction efficiency


The signal is a NC π^0 produced in one of the FGD modules, with it decaying to two photons which both have to convert in the DsECal. Due to the solid angle subtended by the DsECal, this leads to a geometrical acceptance of




(a) HE energy




(b) LE energy



(c) HE incidence angle



(d) LE incidence angle



(e) Reconstructed mass

Figure 6.5.: Raw data distributions for the preselected clusters, the higher energy (HE) cluster and the lower energy (LE) cluster.




Figure 6.6.: Acceptance for π^0 s produced in the FGDs. The momenta and angles of all the produced π^0 s (irrespective of the photon conversion position) is shown in (a) for FGD1 and (b) for FGD2, with arbitrary units on the colour axis. The accepted fraction of π^0 s, i.e. those where both photons convert in the DsECal, is shown in (c) for FGD1 and (d) for FGD2.

2.42% of all NC π^0 s produced in the FGDs.² The acceptance for FGD1 is much smaller than FGD2 (0.38% and 4.56% respectively), due to the smaller subtended angle of the DsECal, and the material between it and the DsECal (mainly FGD2). A plot of acceptance dependencies on π^0 momentum and angle is shown in figure 6.6.

A reconstruction inefficiency arises due to the kinematics of the two decay photons. As described in section 5.2.1 there is a dependence on photon

²For comparison, allowing the photons to convert in any of the Tracker ECal modules, as will be done in future for a Tracker-NC π^0 measurement, the acceptance is 19%.

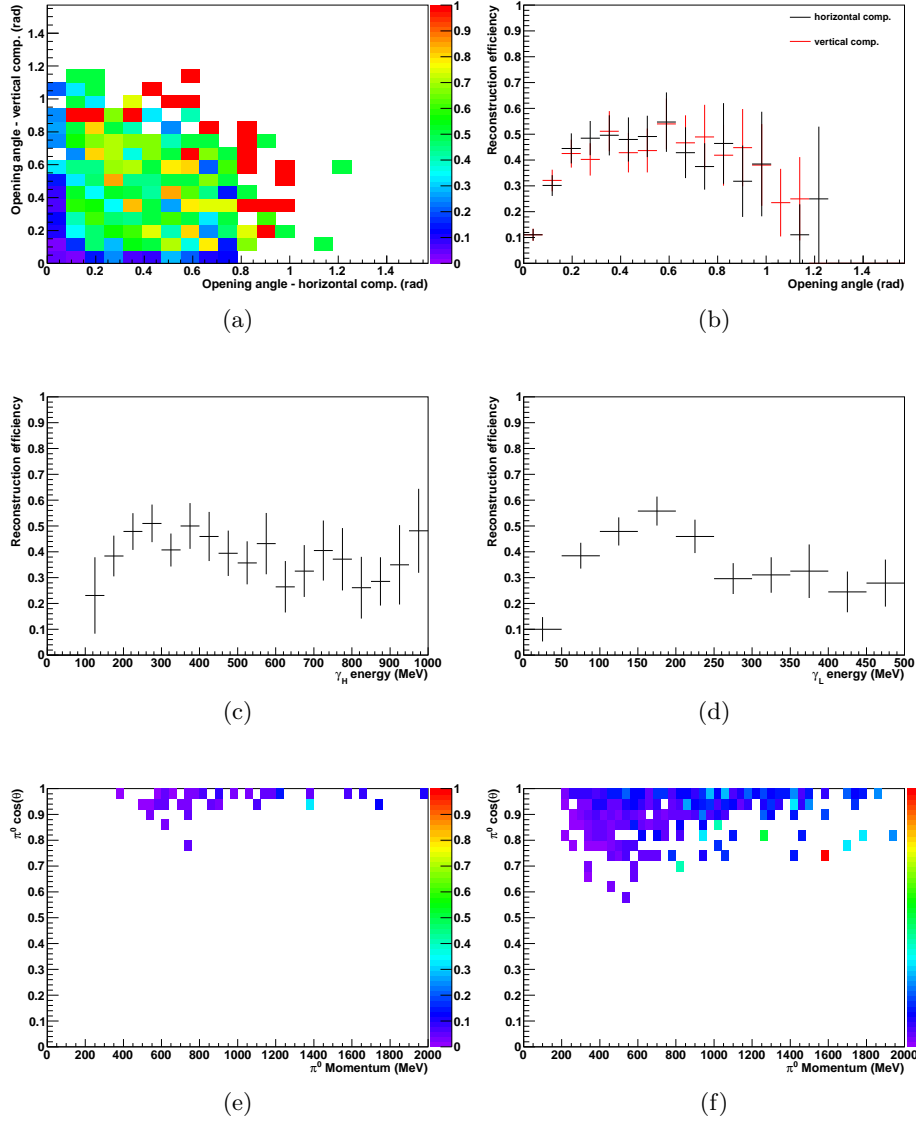


Figure 6.7.: π^0 reconstruction efficiency in the DsECal, i.e. the fraction of events where there are two reconstructed clusters, given the π^0 photons convert in the DsECal. (a) and (b) show the π^0 reconstruction efficiency as a function of the horizontal and vertical components of the opening angle. (c) and (d) show the π^0 reconstruction efficiency as a function of, respectively, the high and low photon energies. (e) shows the combined acceptance and reconstruction efficiency for π^0 s produced in FGD1, (f) in FGD2.

energy and angle of the reconstruction efficiency for single particles. For two particles, there is an additional reconstruction inefficiency if they are too close together. In this case, the hits of the clusters are likely to overlap, and reconstruct only as a single cluster. Additionally, due to the nature of the module, this only has to happen in one view, leaving the other view with a leftover 2D cluster. This can be seen in figures 6.7(a) and (b), where there is a deficiency at both low horizontal and vertical components of the opening angle. This also explains the different energy dependence of the reconstruction efficiency, shown in figures 6.7(c) and (d), compared to the single-photon reconstruction efficiency in figure 5.4(b). The higher the low-energy photon energy is, the smaller the opening angle between the two photons, leading to a lower reconstruction efficiency. The overall acceptance and reconstruction efficiency for preselected events is shown in figures 6.7(e) and (f).

When the analysis is developed further to use the whole Tracker ECal, the incidence of two photons in the same module should be relatively low, so the reconstruction inefficiencies that arise due to this should be smaller.

6.3.2. Preselection purity

After preselection, the signal purity in the Monte Carlo sample is 0.64%, with background composition as given in table 6.2. The signal is as defined in section 6.3.1. The backgrounds are categorised as follows:

NC π^0 elsewhere

Two decay photons from a NC π^0 produced in a non-FGD and non-DsECal interaction.

CC π^0

Event type	Composition (%)
NC π^0 in FGD (signal)	0.64
NC π^0 elsewhere	0.43
CC π^0	1.07
Photons from two different π^0 s	1.20
Secondary π^0 photons	0.58
Broken π^0 decay photon	1.69
At least one photon	28.46
NC DsECal interaction	9.21
CC DsECal interaction	34.74
Other	21.98

Table 6.2.: Event composition after preselection

Two decay photons from a CC π^0 produced in a non-DsECal interaction.

Photons from two different π^0 s

Photons from two separate π^0 decays (one photon from each) produced in a non-DsECal multipion interaction.

Secondary π^0 photons

Two decay photons from a secondary π^0 , e.g. one produced when a primary charged pion scatters off a nucleon in another part of the detector.

Broken π^0 decay photon

A single π^0 decay photon which, due to reconstruction inefficiencies, has been reconstructed as two clusters.

At least one photon

Any other event containing a photon from a non-DsECal neutrino interaction. These photons could be from a π^0 decay, or from other sources such as Bremsstrahlung.

NC DsECal interaction

Any NC interaction within the DsECal itself, which produces two clusters.

CC DsECal interaction

Any CC interaction within the DsECal itself, which produces two clusters.

Other


Any other event not categorised as above. Event topologies that fall into this category are: a broken neutron (forming 11% of the Other category in the preselection), a broken charged pion (10%), particles from two coincident neutrino interactions (9%), a broken muon (8%), a muon and proton (7%), a neutron and charged pion (6%), a broken proton (6%), and a muon and charged pion (5%), with the remainder of topologies under 5%.

Figure 6.8 shows the true vertex distributions for the $\text{NC}\pi^0$ signal and the background interactions.


6.3.3. Data excess

There is a data over Monte Carlo excess of 1.87 ± 0.59 in the preselected sample, when the statistics of the Monte Carlo is normalised to the POT of the data. It is believed that this excess arises from an unmodelled effect which produces extra clusters in the DsECal. There are primarily two unmodelled interactions lacking from the Monte Carlo, namely coincident cosmic rays, and neutrino interactions in the cavern walls (sand interactions).³ There

³It is planned to include sand interactions for future Monte Carlo productions.



(a) x, z distribution



(b) y, z distribution

Figure 6.8.: True vertex distributions. The $\text{NC}\pi^0$ is in green, and the background is in red. The two FGDs are the regions $-1000 \text{ mm} < x, y < 1000 \text{ mm}$, $100 \text{ mm} < z < 500 \text{ mm}$ and $1450 \text{ mm} < z < 1850 \text{ mm}$. The DsEcal is the region with $z > 2800 \text{ mm}$.

is also the issue, mentioned in section 6.2, that the noise clusters are not modelled exactly to match the data.


The probability of there being one (two) clusters in the same bunch, $P(c^{1(2)})$, is

$$P(c^1) = P(c_m^1) + P(c_u^1), \quad (6.1)$$


$$P(c^2) = P(c_m^2) + 2P(c_m^1)P(c_u^1) + P(c_u^2), \quad (6.2)$$

where $P(c_m^{1(2)})$ is the probability of one (two) modelled clusters in the same bunch, and $P(c_u^{1(2)})$ is the probability of one (two) unmodelled clusters in the bunch. Assuming that $P(c_m^1) \gg P(c_u^1)$ and $P(c_{m,u}^1) \gg P(c_{m,u}^2)$, it is possible that $P(c^1)$ (hence measurements such as in section 6.2) can be dominated by the modelled part, whereas $P(c^2)$'s contributions from modelled and unmodelled parts could be of the same order of magnitude. If the assumption is correct, then measurements that strongly depend on $P(c^1)$, for example the mean bunch times and widths of table 6.1, are well modelled, even if an excess is seen in measurements involving $P(c^2)$.


Figures 6.10(a) and (b) show the timing distributions of the data excess. The times are in units of the bunch width (σ_b), relative to the bunch mean (μ_b), which are given in table 6.1, and as mentioned above, the mean bunch times and widths should be well modelled. The plots rule out the excess being due to underestimating the amount of noise clusters, as in that case, the excess in time would look sawtoothed, peaking at values greater than $10\sigma_b$. They also rule out coincident cosmics, as in that case, the excess distributions would be flat with $t < 0$ for the earlier cluster and $t > 0$ for the later cluster. They show that the excess clusters are beam-related, systematically behind the bunch times by approximately $1\sigma_b$ (10–20 ns).




(a) Early time, relative to bunch mean, in units of bunch width




(b) Late time, relative to bunch mean, in units of bunch width




(c) High fiduciality



(d) Low fiduciality




(e) High incidence angle




(f) Low incidence angle


Figure 6.9.: Distributions of the data. The green histogram is the MC normalised to the same POT as the data, and the black histogram is data (with statistical errors). The label ‘Early’ indicates the earlier of the two clusters, ‘Late’ the later of the two, ‘High’ the higher-energy cluster, and ‘Low’ the lower-energy cluster.




(a) Early time, relative to bunch mean, in units of bunch width




(b) Late time, relative to bunch mean, in units of bunch width




(c) High fiduciality



(d) Low fiduciality



(e) High incidence angle



(f) Low incidence angle

Figure 6.10.: Distributions of the excess in the data. The green histogram is the modelled distribution (N_{expected} , from the Monte Carlo, normalised to the same POT as the data), and the black histogram is the excess in the data over the modelled distribution ($N_{\text{data}} - N_{\text{expected}}$). The error bars are the sum in quadrature of the statistical errors of the data and Monte Carlo. The label ‘Early’ indicates the earlier of the two clusters, ‘Late’ the later of the two, ‘High’ the higher-energy cluster, and ‘Low’ the lower-energy cluster.

This indicates that the excess is probably formed of the sand interaction-related clusters. Figures 6.10(c) and (d) show the number of bars between the cluster and the closest edge of the DsECal module (i.e. a measure of ‘fiduciality’ – with a large value implying a cluster is in the centre of the module, and a low value implying a cluster is at the edge). These show that the data excess is greater for clusters which are closer to the edge of the module. Figures 6.10(e) and (f) show the angles of incidence of the clusters with the module, and indicate that the excess is greater for higher-angled clusters.

These distributions indicate that the excess is most likely caused by neutrons, produced in the cavern walls, and reinteracting hadronically in the magnet, producing pions, leading to particles entering the DsECal at high angles and from the sides. Assuming they are produced in neutrino interactions in the cavern upstream of the detector, they are non-relativistic because they reinteract in the detector tens of nanoseconds after the neutrino beam does. If they were relativistic, they would reinteract roughly in time with the beam, which is itself relativistic; therefore they have a mass of around 500 MeV or greater. Stability on this timescale indicates they are nucleons, rather than heavier hadrons. There is a possibility some are protons, however the range of sub-GeV protons in iron is tens of centimetres [76], so most of these would have interacted well upstream of the DsECal module.

6.4. Selection cuts

The TMVA software library [77] is used to perform a likelihood fit of the signal and background, with the aim of using the likelihood as a selection cut.

The likelihood is a product of one dimensional pdfs, and they are assumed to be uncorrelated. At this stage, only information that can be extracted from just the two clusters is used. The following variables are used as inputs to the likelihood: for both the high and low energy cluster, the cluster's incidence angle, track/shower discriminant (PID), charge skew, EM likelihood, first hit layer, number of bars to the edge of the module ('fiduciality'), number of hits, cone angle, mean position, and thrust; additionally for the high energy cluster, the cluster's 'pointing' variable, and the number of Michel tags associated with the cluster; the distance and relative angle between the two clusters; the energy asymmetry of the two clusters; and the invariant mass of the two clusters, assuming they are massless.

A description of the thrust is given in section 5.1.4, and the tangent of the incidence angle is the z -component of the thrust axis. Descriptions of the charge skew and EM likelihood are given in section 5.1.5. Descriptions of the track/shower discriminant, cone angle and mean position are given in section 5.1.6. Description of Michel tagging is given in section 5.1.7. The distance between the clusters is the magnitude of the vector $\vec{o}_h - \vec{o}_l$, where $\vec{o}_{h(l)}$ is the high (low) cluster's thrust origin. The angle between the clusters, $\theta_{hl} = \arccos(\vec{a}_h \cdot \vec{a}_l)$, where $\vec{a}_{h(l)}$ is the high (low) cluster's thrust axis. The energy asymmetry is $(E_h - E_l)/(E_h + E_l)$, where $E_{h(l)}$ is the high (low) cluster's reconstructed energy. The invariant mass is $\sqrt{2E_h E_l (1 - \cos \theta_{hl})}$.

The pointing variable (p) of a cluster is defined as the z component of the vector:

$$\vec{p} = \left(\frac{\sum_i q_i \vec{a}_{\text{PCA}} \cdot (\vec{x}_i - \vec{o}_{\text{PCA}})}{\sum_i q_i} \right) \vec{a}_{\text{PCA}},$$

where \vec{a}_{PCA} is the PCA axis, \vec{o}_{PCA} is the PCA centre, q_i the charge of the i th hit and \vec{x}_i its position, with i over all the hits of the cluster. Because a

shower has a cone-like shape, the displacement of hits upstream of the PCA centre will lie closer to the PCA axis than the hits downstream of the centre,⁴ and so will have a larger longitudinal component with respect to \vec{a}_{PCA} . The vector \vec{p} should point from the PCA centre towards the shower start, hence the name of the variable. Taking the z component of this vector is just a convenient way of reducing the pointing vector to one dimension, given that in this analysis the signal photons are entering the module from upstream. This z component should be negative for these types of showers (meaning that the cluster is pointing upstream from the PCA centre). For other types of showers this could be different; for example showers entering from the sides will have p close to 0, whereas the other components of \vec{p} should indicate the particle entrance positions. For tracks, every component of \vec{p} will be small. In the likelihood fit only the high-energy cluster's pointing is used, because as shown in section 5.2.1, the PCA axis becomes less valid for energies below around 200 MeV.

Plots of these variables, showing their signal and background separation, are in figure 6.11. The linear correlation between these variables for signal and background is shown in figure 6.12, showing that the inputs are not strongly correlated with each other, with only a few of the variables having a moderate correlation.⁵

It is difficult to verify the robustness of these distributions for application on the data, due to the excess seen and the different shapes for some of the distributions, for example as shown the incidence angle distributions in figure 6.10. All these distributions with the data overlaid are given in

⁴The terminology ‘upstream’ and ‘downstream’ in this sentence meaning, respectively, closer to and further away from the shower start, not to be confused with the same terminology used in the rest of the text.

⁵With the linear correlation coefficient ρ , strong correlation is defined as $|\rho| > 0.7$, and moderate correlation is defined as $0.4 > |\rho| \geq 0.7$.




Figure 6.11.: Input variables for the likelihood. The signal distributions are in green, the background distributions in blue. The y-axis scales are arbitrary, and the two distributions are normalised to their peaks. The label ‘High’ indicates a property of the higher-energy cluster, ‘Low’ a property of the lower-energy cluster.

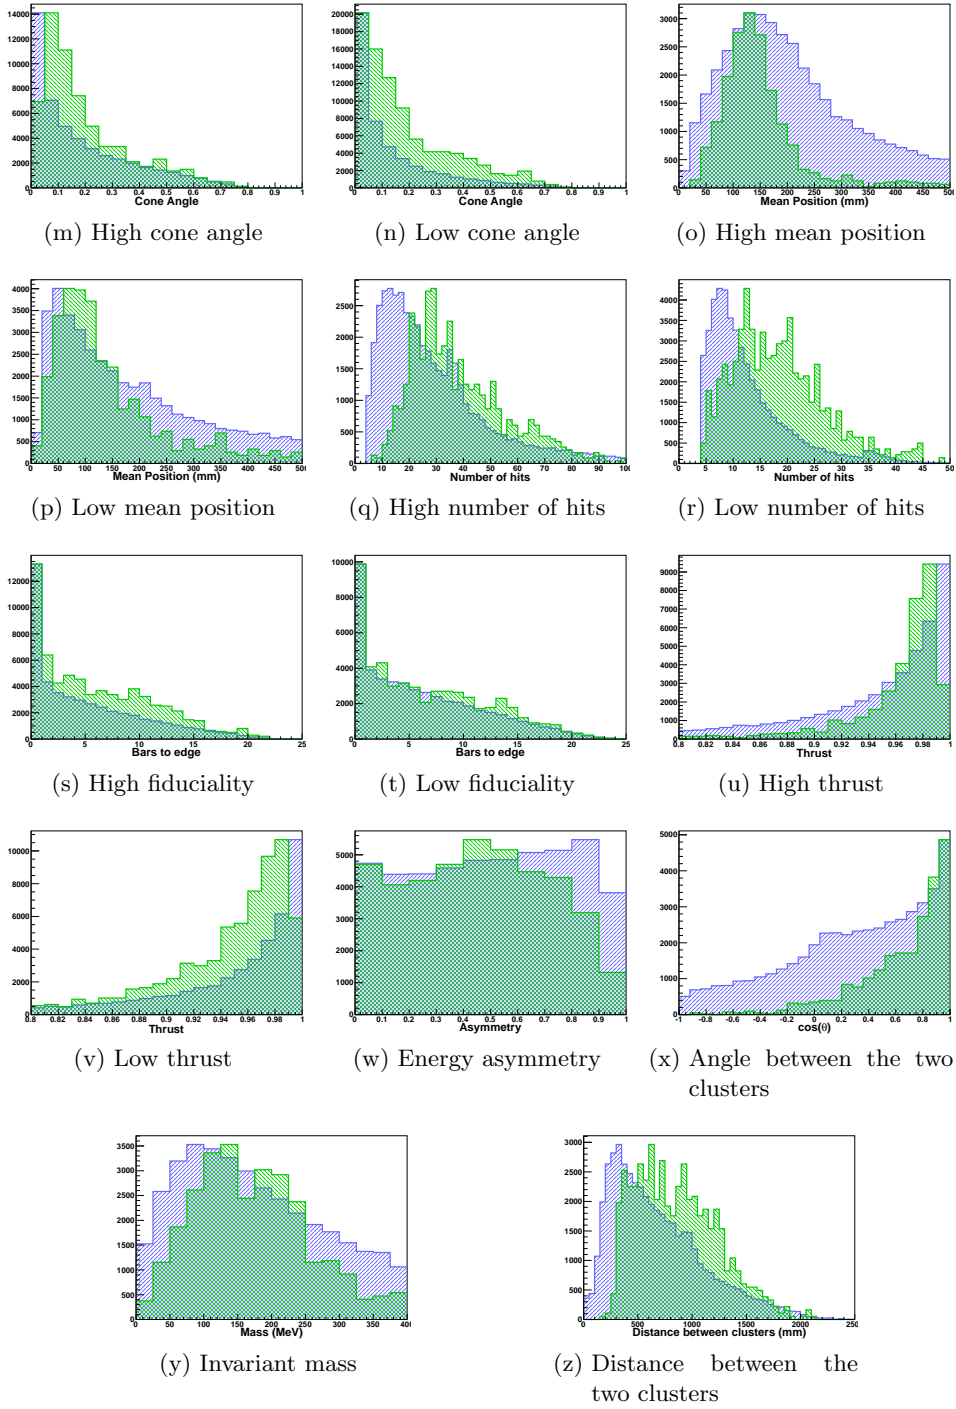


Figure 6.11.: Input variables for the likelihood, continued.

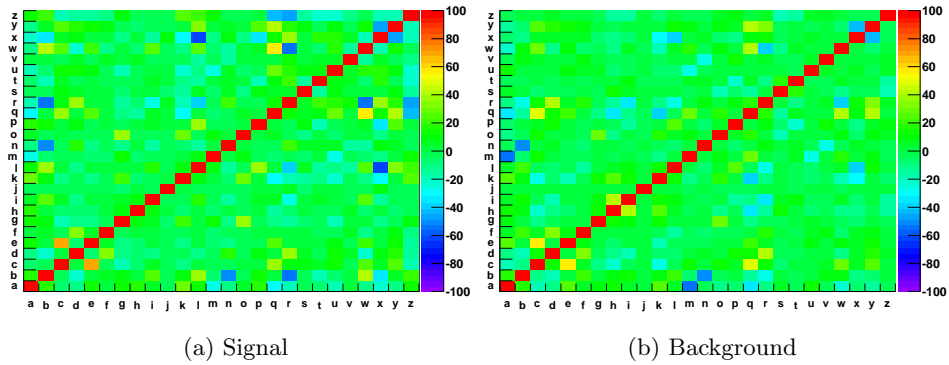


Figure 6.12.: Linear correlation coefficients (in %) between the likelihood input variables. The axis labels a–z correspond to the subfigures (a)–(z) in figure 6.11.

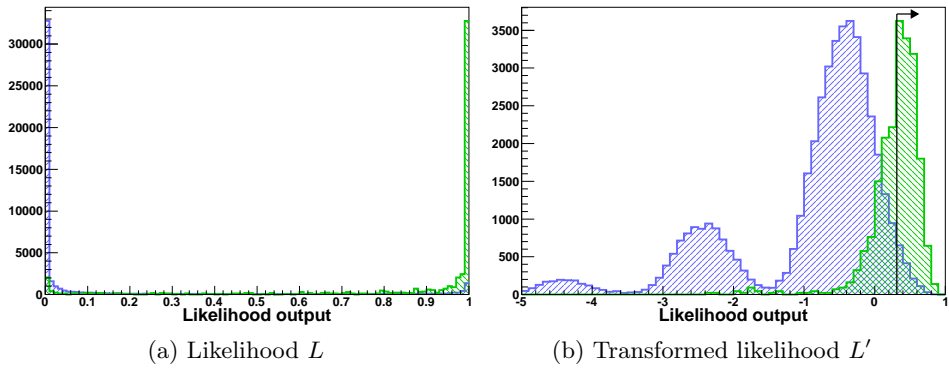


Figure 6.13.: Likelihood output, signal in green and background in blue.

appendix A.

The likelihood output produced is shown in figure 6.13(a). It is strongly peaked at 0 for background and 1 for signal, so it is transformed⁶ to $L' = -\ln(L^{-1} - 1)/15$, shown in figure 6.13(b). This transformation enhances the peak structure, and makes optimising a cut on L' easier to compute than a cut on L , because there is less variation of signal and background entries for the same variation in L' . There are strange features in the L' distribution,

⁶The transformation is implemented by the TMVA library, and is discussed in section 8.2.2 of [77].

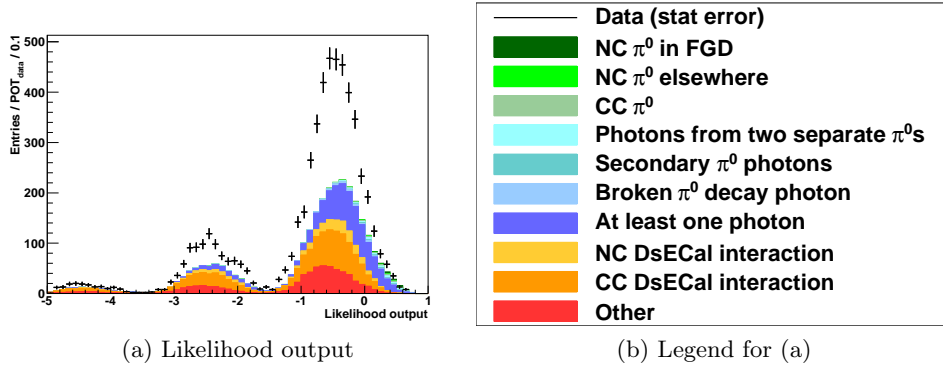



Figure 6.14.: Likelihood output for data and Monte Carlo.

Event type	Composition (%)
NC π^0 in FGD (signal)	12.0
NC π^0 elsewhere	3.2
CC π^0	12.0
Photons from two different π^0 s	9.1
Secondary π^0 photons	6.2
Broken π^0 decay photon	4.3
At least one photon	45.8
NC DsECal interaction	1.0
CC DsECal interaction	4.7
Other	1.8

Table 6.3.: Event composition after cluster-level likelihood selection

the secondary peaks at $L' \approx -2.5$ (peak B) and $L' \approx -4.5$ (peak C), which are not currently understood. There has been some investigation into these, for example finding that, if the low-energy cluster has the first hit in layers 28–34, or the high-energy cluster has its first hit in layers 20–22, 25, or 28–34, then L' will end up in peak B or C. The likelihood distribution with the data overlaid is in figure 6.14, showing that in general the shape of the data agrees with the Monte Carlo, including the secondary peak structure.

A cut is made on L' to maximise efficiency \times purity. Events are kept if




(a) Early time, relative to bunch mean, in units of bunch width (b) Late time, relative to bunch mean, in units of bunch width

Figure 6.15.: Distributions of the excess in the data, after the likelihood cut. The green histogram is the modelled distribution, and the black line is the excess in the data over the modelled distribution. The label ‘Early’ indicates the earlier of the two clusters, and ‘Late’ the later of the two.


$L' > 0.312$, with a signal selection efficiency of 58.4% and a purity of 12.0%, with the background composition given in table 6.3.

At this level of selection, the data over Monte Carlo excess is 1.18 ± 0.40 . Even though no timing information is used as an input into the likelihood, the excess of clusters systematically later than the bunch time is mostly reduced, shown in figure 6.15. If the data excess had remained high, this would have implied it was signal-like, and that the π^0 production cross sections might have been greatly underestimated in GENIE.

Figure 6.16 shows the true vertex distributions for the $\text{NC}\pi^0$ signal and the background interactions after the likelihood cut.



(a) x, z distribution



(b) y, z distribution

Figure 6.16.: True vertex distributions, after the likelihood cut. The $\text{NC}\pi^0$ is in green, and the background is in red. The two FGDs are the regions $-1000 \text{ mm} < x, y < 1000 \text{ mm}$, $100 \text{ mm} < z < 500 \text{ mm}$ and $1450 \text{ mm} < z < 1850 \text{ mm}$. The DsECal is the region with $z > 2800 \text{ mm}$.

6.5. Vertexing

Estimating the interaction vertex, hence the π^0 decay position, is an important step in the analysis. The invariant mass of a system of two massless particles, M , is given by $M = \sqrt{2E_1E_2(1 - \cos\theta)}$, where E_1 and E_2 are the two particle energies, and θ is the angle between their momenta. In a π^0 decay, this is the opening angle of the decay. This angle can be obtained using the estimated vertex position and the two cluster positions. Given that, for any single event, the two cluster positions are fixed, the opening angle will be larger (and so will the reconstructed mass) the closer the vertex position is to the DsECal module. The vertex is also important for determining which FGD the neutrino interaction occurs in, so that eventually a water-subtraction calculation can be performed to extract the $\text{NC}\pi^0$ production cross section on water.

Before trying to reconstruct a vertex, there is an activity cut. Global reconstruction is used to remove events if there is any in-time activity (within 300 ns of the π^0 candidate⁷) in the P0D or TPC1. If there are tracks in TPC1, it is far more likely that they come from interactions in the P0D, rather than high-recoil interactions in FGD1.

There are three categories of events, depending on how the vertex location is found: events with a reconstructed vertex in the tracker ('Global vertex'), events with coincident FGD activity ('FGD vertex'), and events with no other information apart from the two clusters in the DsECal ('Cluster vertex').

⁷A similar procedure as described in section 6.2 has been used to measure the bunch timings of reconstructed objects in the other subdetectors of ND280, and time offset corrections between the DsECal and the other subdetectors have been applied.




Figure 6.17.: Vertex resolution (in z) and angular deviation for the Global vertex sample. $\text{NC}\pi^0$ signal is in green, background is red, and the data is black.

6.5.1. Global vertex sample

Global reconstruction produces vertices as described in section 4.4.2. Vertices are searched for with times in a 300 ns window around the π^0 candidate time, with their z coordinate 10 cm upstream of FGD1 or greater. The spatial resolution of this type of vertex is shown in figure 6.17(a). Defining the angular deviation as $(\vec{\sigma}_{h(l)} - \vec{v}) \cdot \vec{a}_{h(l)}$, where $\vec{\sigma}_{h(l)}$ is the high (low) cluster's thrust origin, $\vec{a}_{h(l)}$ is the high (low) cluster's thrust axis, and \vec{v} is the reconstructed vertex position, this is a measure of how much the reconstructed angle of the cluster deviates from the photon hypothesis, given a photon would have travelled in a straight line from the vertex to the cluster. This is comparable to the method of working out the angular resolution in section 5.2.1, where the MC-truth photon direction was used instead of a $(\vec{\sigma}_{h(l)} - \vec{v})$ hypothesis. The angular deviation of the two clusters in the sample is shown in figure 6.17(b). This information could be used in a further selection cut, as the signal events have deviations closer to one than the background. Using knowledge of the angular resolution, this variable could

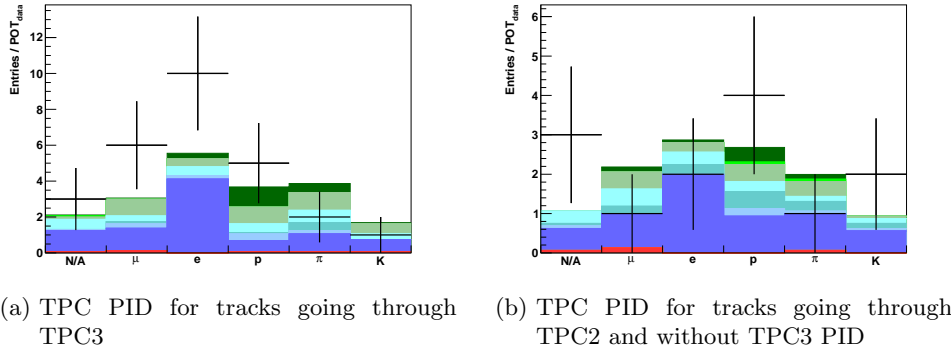


Figure 6.18.: TPC PID information for the highest momentum track in the ‘Global vertex’ sample. N/A denotes no TPC PID information available for the track. For legend see figure 6.14(b).

be turned into how many σ the angles deviate from the expected true angle.

Events with this type of vertex also have additional PID information from the TPCs. The highest momentum track originating at the vertex is examined: if it goes through TPC3, it is assigned the PID for which the pull is closest to 0 (see section 4.4.1). Sometimes the pull is not available, as for some reason the PID algorithm failed. If the track does not have TPC3 PID information (either it does not traverse that TPC or the pull is not available), and it goes through TPC2, then likewise it is assigned the PID for which the pull in this TPC is closest to 0. Figure 6.18 shows the TPC PID information if the track goes through either of the two TPCs. An event is selected if the track is identified as a proton in this manner, or it has no PID information at all. Figure 6.19 shows the reconstructed mass spectrum of events in this sample.




Figure 6.19.: Reconstructed π^0 mass in the ‘Global vertex’ sample. For legend see figure 6.14(b).

6.5.2. FGD vertex sample

For those events lacking a global vertex, a search is made for unused FGD hits, again in a timing window of 300 ns around the bunch time. Within the FGD, the position of the most upstream hit is selected as the vertex. If there is FGD activity in both FGDs, then the FGD2 vertex is chosen, given that the signal acceptance is more than ten times greater for that FGD. The spatial resolution of this type of vertex is shown in figure 6.20(a), and the angular deviations of the two clusters is shown in figure 6.20(b). Figure 6.21 shows the reconstructed mass spectrum of events in this sample.

6.5.3. Cluster vertex sample

For the remaining events, only the reconstructed cluster positions and directions can be used to estimate a vertex position. One such method is to look for the intersection of the cluster directions. In general, they will not intersect in three dimensions, so the problem is broken down into the two




Figure 6.20.: Vertex resolution (in z) and angular deviation for the FGD vertex sample. $\text{NC}\pi^0$ signal is in green, background is red, and the data is black.




Figure 6.21.: Reconstructed π^0 mass in the 'FGD vertex' sample. For legend see figure 6.14(b).




Figure 6.22.: Unconstrained vertex resolution (in z) for the Cluster vertex sample.

views of the DsECal module. In each view, the two-dimensional vertex position is worked out from the intersection of the 2D cluster directions, and the third, unknown coordinate is calculated using just the direction of the high-energy cluster at that 2D vertex position. The reconstructed vertex that is chosen, from the two options, is the one closest to an FGD module. The spatial resolution, in figure 6.22, shows that this vertexing method performs worse than the other two, due to the poor angular resolution of the cluster directions, and the large ‘lever arm’ to the vertex position.

An alternative method is to artificially constrain the vertex to be in FGD2. The z coordinate is chosen to be in the centre of the FGD, and the x, y coordinates are the mean of the two clusters’ x, y positions. The spatial resolution in figure 6.23(a) shows that this performs better, however there is a bias in the z coordinate when using this vertex position. The bias arises due to the fact that any interaction activity, for example recoil protons, cannot produce hits in the FGD or other tracks (as these events would be




Figure 6.23.: Constrained vertex resolution (in z) and angular deviation for the Cluster vertex sample. $\text{NC}\pi^0$ signal is in green, background is red, and the data is black.

placed in the first two categories). Hence this is biasing the interaction to be as close to the downstream end of the FGD, where there is less material for the activity to produce hits. There is also less material for one of the π^0 photons to convert within the module itself. Placing the vertex at the centre of the FGD is biasing it upstream of where the majority of the true vertices actually are, for those events which are accepted. However, this bias in the reconstruction is preferable to placing the constrained vertex at the mean of the true interaction distribution, because that option could be model-dependent (true interaction positions for the GENIE and NEUT Monte Carlos are shown in figure 6.24). This bias in the position leads to a bias in the reconstructed mass, because of the opening angle depending on this position. The angular deviation using this constrained vertex is shown in figure 6.23(b). Figure 6.25 shows the mass spectrum of events in this sample.

The mass resolution of the Cluster vertex sample can be checked using the other two samples. Their cluster-based mass (M_{reco}) can be calculated in the same way as for the Cluster vertex sample (by ignoring the Global or FGD




Figure 6.24.: True interaction positions (z -coordinates) for two different Monte Carlo generators, for $\text{NC}\pi^0$ interactions in the Cluster vertex sample.




Figure 6.25.: Reconstructed π^0 mass in the 'Cluster vertex' sample. For legend see figure 6.14(b).





Figure 6.26.: Mass resolution cross-check, calculating the Cluster-vertex based mass on the Global- and FGD-vertex based samples. For legend see figure 6.14(b).


reconstruction information), and then this mass can be compared to their reconstructed mass ($M_{\text{glob,FGD}}$) using the Global or FGD reconstruction information. Figure 6.26 shows the mass resolution ($M_{\text{reco}} - M_{\text{glob,fgd}}$) of this cross-check sample. The secondary peak in the Global-vertex sample is due to the vertex being constrained to FGD2 whereas the Global vertex is in FGD1. The resolution on the Global-vertex sample has a bias of $21 \pm 3 \text{ MeV}/c^2$ with a resolution of $16 \pm 2 \text{ MeV}/c^2$. The resolution on the FGD-vertex sample has a bias of $14 \pm 2 \text{ MeV}/c^2$ with a resolution of $11 \pm 2 \text{ MeV}/c^2$.

Category	Mass (MeV/c^2)		Momentum (MeV/c)	
	Resolution	Bias	Resolution	Bias
Global vertex	12.3 ± 1.9	5.5 ± 2.8	98.9 ± 21.6	51.4 ± 21.1
FGD vertex	22.9 ± 10.9	-1.9 ± 5.0	59.5 ± 14.9	12.4 ± 12.6
Cluster vertex	32.4 ± 4.2	16.8 ± 4.7	62.8 ± 10.9	24.0 ± 10.1


Table 6.4.: Reconstructed π^0 mass and momentum resolutions and biases.




(a) Global mass




(b) Global momentum




(c) FGD mass



(d) FGD momentum



(e) Cluster mass



(f) Cluster momentum

Figure 6.27.: Reconstructed π^0 mass and momentum biases (reconstructed values – true values) for the three event categories in green, with Gaussian fits in black.




Figure 6.28.: Mass spectrum with perfect vertex resolution

6.5.4. Resolution

Figure 6.27 shows the reconstructed π^0 mass and momentum resolutions for each category of events, and table 6.4 lists their resolutions.

The mass spectra for the three categories (figures 6.19, 6.21, and 6.25) show a peak for the signal around $135 \text{ MeV}/c^2$, the mass of the π^0 [32]. The width of the mass spectrum is best for the global vertex category, and worst for the cluster vertex category.

As a final selection cut, only those events where the mass is more than $100 \text{ MeV}/c^2$ and less than $170 \text{ MeV}/c^2$ are selected. In general, each category can have its own mass cut values, as they have different mass resolutions and different signal purities; however for this analysis a single cut is used, which has not been optimised.

For the π^0 mass resolution, the vertex position resolution dominates, as the angular resolution of the photon clusters have no input to the final reconstructed mass calculation. Figure 6.28 shows the mass spectrum achievable with perfect vertex resolution, with the spread being just due to the energy

resolution of the clusters.

6.6. Systematic errors

Only one source of systematic errors has been taken into account in this analysis. This is the uncertainty in the flux prediction, as shown in figure 3.5, and a simplified model of that uncertainty was used to calculate the systematic uncertainties in the analysis. For events with true neutrino energy (E_ν) below 2 GeV, the flux uncertainty was estimated as 15%; for events with $2 \text{ GeV} < E_\nu < 5 \text{ GeV}$, the uncertainty was estimated as 30%; and for events with $E_\nu > 5 \text{ GeV}$, the uncertainty was estimated as 45%.

Other sources of systematic error, which have not been taken into account, include uncertainties of the cross section models of the interaction generators, reconstruction inefficiencies (such as the connection of TPC3 tracks to DsECal clusters), and variations of cut values.

6.7. Selection summary

Table 6.5 gives the $\text{NC}\pi^0$ selection efficiency ($N_{\text{final}}/N_{\text{presel}}$) and purity in the final selection, along with the expected and observed numbers in the data. The background composition is listed in table 6.6.

The overall mass spectrum, before the final mass cut, is shown in figure 6.29. It shows good agreement between data and Monte Carlo. The reconstructed momentum spectrum is shown in figure 6.30, and the π^0 angular distribution is in figure 6.31. As expected, the momentum is relatively high, and the angle is very forward-going.

Category	Selection		Expected MC $\pm \sigma_{\text{stat}} \pm \sigma_{\text{syst}}$		
	Efficiency	Purity	All	NC π^0	Data
Global	$33 \pm 7\%$	$30 \pm 6\%$	$7.2 \pm 0.7 \pm 2.2$	$2.2 \pm 0.4 \pm 0.7$	9
FGD	$40 \pm 9\%$	$43 \pm 10\%$	$3.9 \pm 0.5 \pm 1.0$	$1.7 \pm 0.3 \pm 0.4$	7
Cluster	$34 \pm 6\%$	$39 \pm 7\%$	$6.6 \pm 0.6 \pm 1.7$	$2.6 \pm 0.4 \pm 0.6$	5
Overall	$34 \pm 4\%$	$37 \pm 4\%$	$17.7 \pm 1.1 \pm 4.9$	$6.5 \pm 0.6 \pm 1.7$	21

Table 6.5.: Final selection efficiencies and purities

Event type	Composition (%)
NC π^0 in FGD (signal)	36.62
NC π^0 elsewhere	4.93
CC π^0	18.66
Photons from two different π^0 s	6.69
Secondary π^0 photons	4.58
Broken π^0 decay photon	4.58
At least one photon	19.37
NC DsECal interaction	0.70
CC DsECal interaction	3.87
Other	0.00

Table 6.6.: Event composition in the final selection

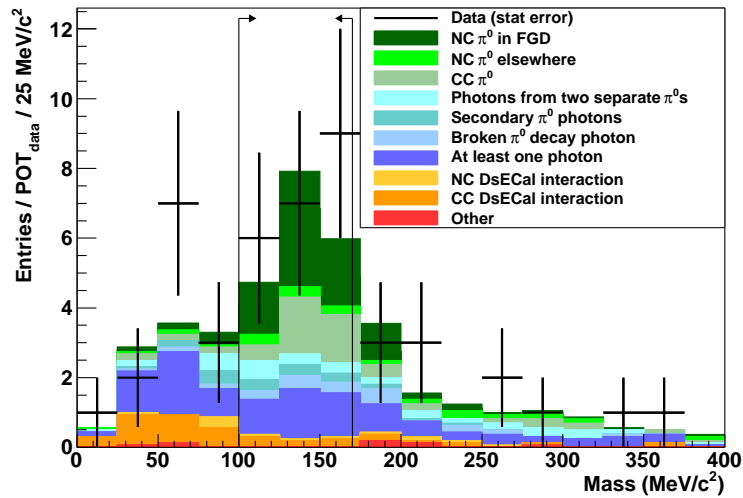


Figure 6.29.: Reconstructed π^0 mass in the final selection, before applying mass cut




Figure 6.30.: Reconstructed π^0 momentum in the final selection

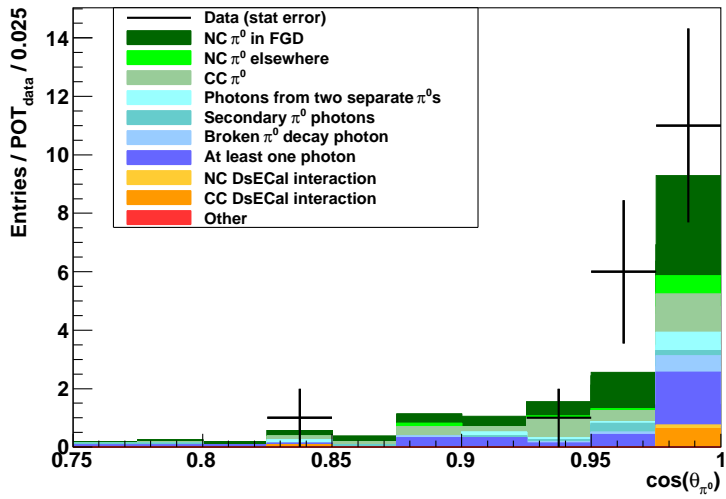


Figure 6.31.: Reconstructed π^0 angle in the final selection

6.8. Crosscheck with NEUT

As a crosscheck, the same analysis was performed using Monte Carlo produced with the NEUT neutrino interaction generator [58]. Statistics of 1.68×10^{21} POT was produced, again with 33% in the Run I configuration. The analysis proceeds in an identical manner as with the GENIE production, using the exact same cuts.

In NEUT, the rate of NC π^0 production in the FGDs is 87% of that in GENIE, while the totally inclusive rate over the whole detector is 111% that of GENIE.

After the preselection, the data/MC ratio is 1.76 ± 0.54 , which is slightly smaller than the excess seen in the GENIE production. This difference can possibly be attributed to the increased overall cross section in NEUT. The signal purity is 0.51%.

The cluster-level likelihood output is shown in figure 6.32, and the same cut as with GENIE is applied. This cut for the NEUT production performs slightly worse than in the GENIE production, with a selection efficiency of 53% and purity of 9.9%.

After applying all the selection cuts, the efficiencies and purities for the various event categories is shown in table 6.7. These show that is not significantly different to the results with GENIE, with the most significant difference being a 2σ difference between the purities in the Global vertex sample. The background composition of the final selection is given in table 6.8, and shows no major differences to that of the GENIE production.

The overall reconstructed π^0 mass in the final selection is shown in figure 6.33, and the reconstructed π^0 momentum and angle are given in figures 6.34 and 6.35.

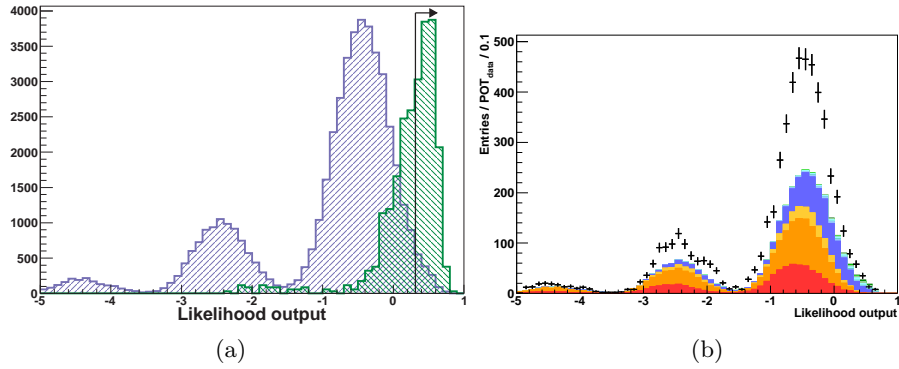


Figure 6.32.: Likelihood output for the NEUT production. (a) shows the shape of the distributions, with signal in green and background in blue. (b) shows the data overlaid on the distribution; for legend see figure 6.14(b).




Figure 6.33.: Reconstructed π^0 mass, before mass cut, for the final selection in the NEUT production

Category	Selection		Expected MC $\pm \sigma_{\text{stat}} \pm \sigma_{\text{sys}}$			Data	Data/MC
	Efficiency	Purity	All	NC π^0	Data		
Global	$23 \pm 6\%$	$16 \pm 4\%$	$7.2 \pm 0.7 \pm 2.3$	$1.1 \pm 0.3 \pm 0.4$	9	1.3 ± 0.5	
FGD	$39 \pm 10\%$	$37 \pm 9\%$	$3.7 \pm 0.5 \pm 1.0$	$1.4 \pm 0.3 \pm 0.4$	7	1.9 ± 0.7	
Cluster	$37 \pm 6\%$	$43 \pm 8\%$	$5.9 \pm 0.6 \pm 1.5$	$2.5 \pm 0.4 \pm 0.5$	5	0.9 ± 0.3	
Overall	$32 \pm 4\%$	$30 \pm 4\%$	$16.8 \pm 1.0 \pm 4.7$	$5.0 \pm 0.6 \pm 1.2$	21	1.3 ± 0.4	

Table 6.7.: Selection efficiencies and purities in the NEUT production.

Event type	Composition (%)		
	Preselected	Likelihood	Final
NC π^0 in FGD (signal)	0.51	9.93	30.19
NC π^0 elsewhere	0.28	3.13	5.28
CC π^0	0.90	12.54	22.64
Photons from two different π^0 s	1.33	14.18	7.55
Secondary π^0 photons	0.56	6.64	5.28
Broken π^0 decay photon	1.40	3.58	3.40
At least one photon	26.11	43.06	21.13
NC DsECal interaction	9.41	1.19	0.00
CC DsECal interaction	37.14	3.58	2.64
Other	22.36	2.16	1.89

Table 6.8.: Event composition after preselection, after the likelihood cut, and in the final selection, for the NEUT production.




Figure 6.34.: Reconstructed π^0 momentum for the final selection in the NEUT production

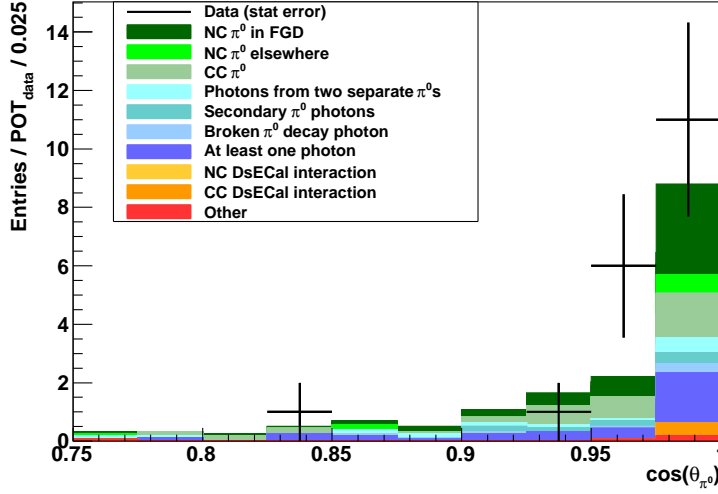


Figure 6.35.: Reconstructed π^0 angle for the final selection in the NEUT production

6.9. Summary and remarks

Both the GENIE and the NEUT sample are statistically consistent with the data, at the level of 0.5σ and 0.7σ respectively. The data/MC ratio for GENIE is 1.2 ± 0.4 , and for NEUT is 1.3 ± 0.4 .

The reconstructed NC π^0 s are very likely to have been produced in coherent scattering, as they are highly forward-going, although this has not been checked with the Monte Carlo truth.

A main background to the NC π^0 signal in the final selection is CC π^0 s, most of which are produced in the FGDs. It should be possible to greatly reduce this component by further development of the TPC PID cuts in the ‘Global vertex’ sample, for example by looking at all the tracks in the event, not just the highest momentum one. There have also been improvements made to Tracker reconstruction, since the version used for this analysis, which should remove the cases where the TPC PID is unavailable. There is

also FGD PID available for tracks that don't go through any TPCs (they leave the FGDs at high angles).

The other major background is the 'At least one photon' category of events. Of these types of events that end up in the final selection, almost all (87%) have a π^0 decay photon producing one of the two clusters, with 70% of these having an electron- or positron-related deposit as the other cluster. Whether these are Bremsstrahlung photons, or the result of a showering process beginning upstream of the DsEcal, and whether they are related to the other decay photon, has not been investigated. However, with refinement of the reconstruction, some of these events might be recoverable into the signal or CC π^0 event categories.

7. Conclusions

This thesis has shown that reconstruction of neutral current π^0 s with the Tracker ECal, one of its design purposes, is possible. A π^0 mass peak is well reconstructed, albeit with low statistics and only around 33% purity.

The analysis is applied to two separate Monte Carlo simulations, based on GENIE and NEUT. The expected number of events with GENIE is $17.7 \pm 1.1(\text{stat}) \pm 4.9(\text{sys})$, which is consistent with the number observed in the data by 0.5σ , and the scaling factor is 1.2 ± 0.4 . The expected number of events with NEUT is $16.8 \pm 1.0(\text{stat}) \pm 4.7(\text{sys})$, consistent with the data at the 0.7σ level, with a scaling factor of 1.3 ± 0.4 .

The analysis presented in the thesis will be expanded upon in the near future, to include the whole Tracker ECal. The gain in statistics by doing this is around an order of magnitude.

This work will hopefully lead to new cross section measurements for $\text{NC}\pi^0$ production, of which there are very few as of writing.

A. Distributions of likelihood input variables in data

This appendix contains plots of distributions of the likelihood input variables used in the analysis, described in section 6.4. Each figure is a likelihood input, shown before the likelihood cut and after the cut.

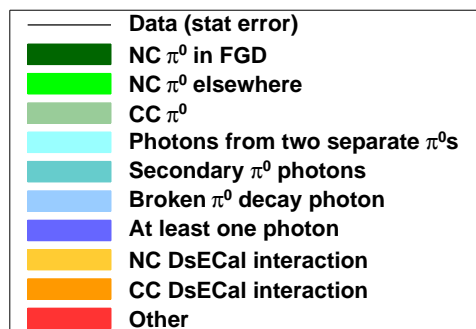


Figure A.1.: Legend for the plots in the appendix




Figure A.2.: High track/shower discriminant




Figure A.3.: Low track/shower discriminant




Figure A.4.: High charge skew

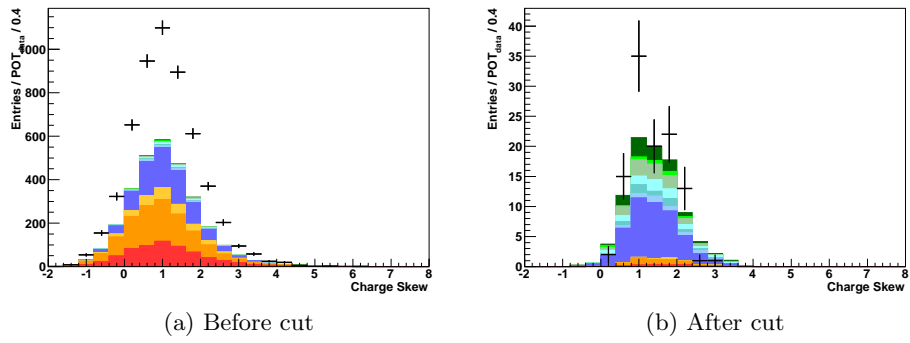


Figure A.5.: Low charge skew

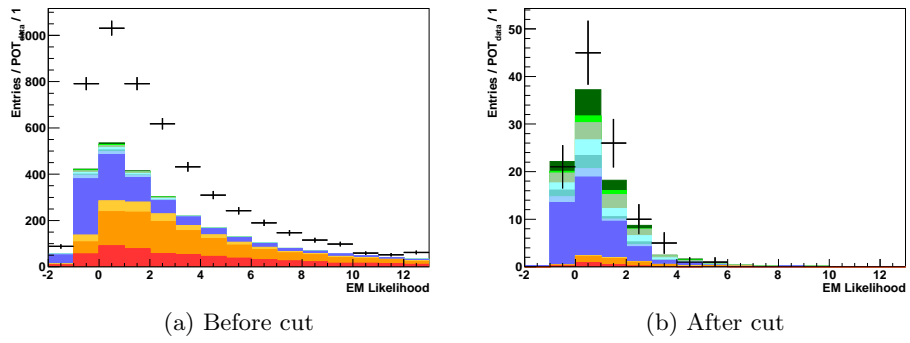


Figure A.6.: High EM likelihood

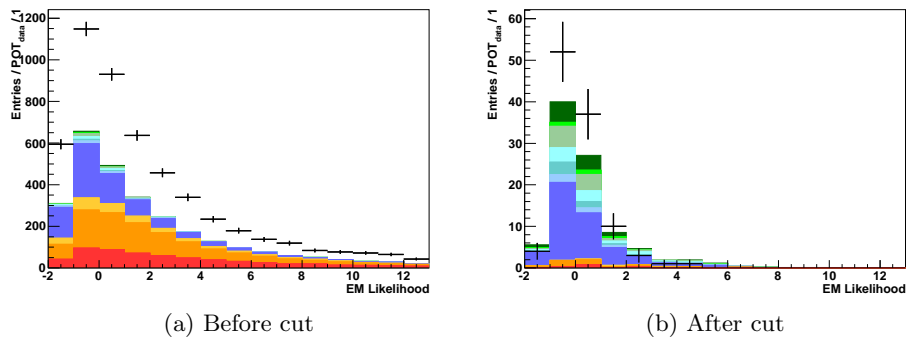


Figure A.7.: Low EM likelihood

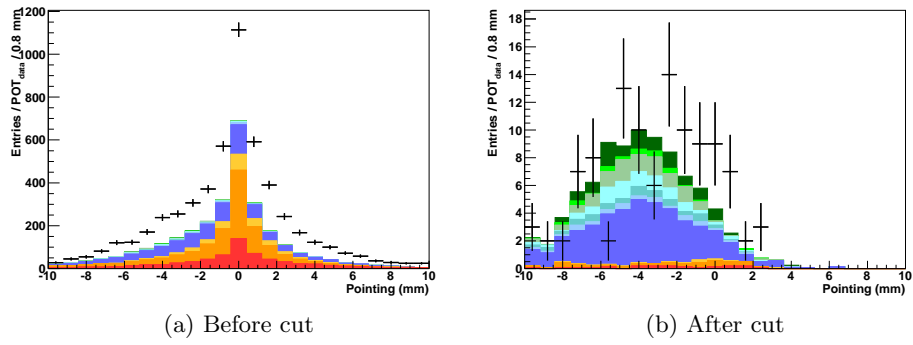


Figure A.8.: High pointing

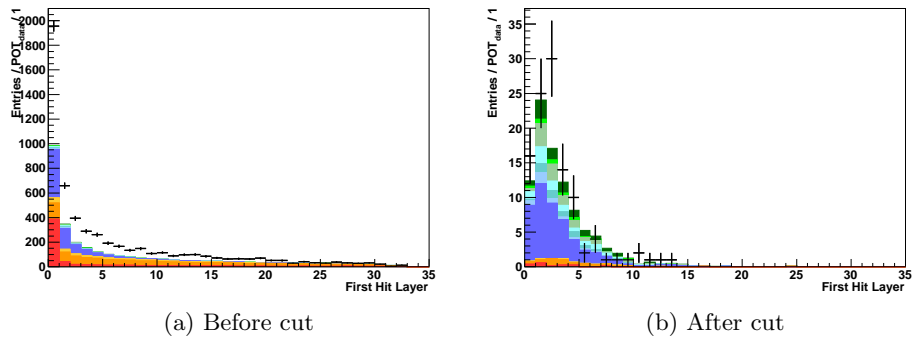


Figure A.9.: High first layer

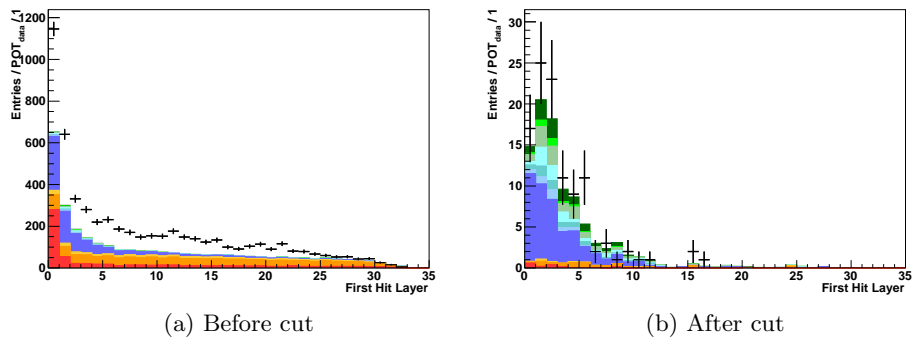




Figure A.10.: Low first layer




(a) Before cut




(b) After cut

Figure A.11.: High number of Michel tags




(a) Before cut




(b) After cut

Figure A.12.: High incidence angle



(a) Before cut



(b) After cut

Figure A.13.: Low incidence angle

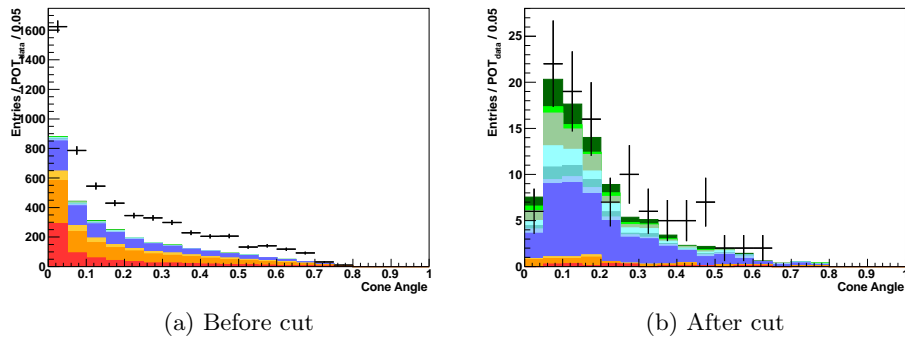


Figure A.14.: High cone angle




Figure A.15.: Low cone angle

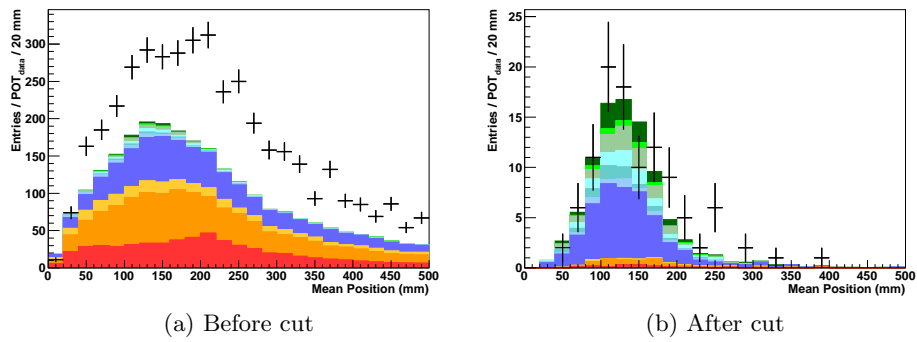




Figure A.16.: High mean position




(a) Before cut




(b) After cut

Figure A.17.: Low mean position




(a) Before cut




(b) After cut

Figure A.18.: High number of hits



(a) Before cut



(b) After cut

Figure A.19.: Low number of hits




Figure A.20.: High fiducialty

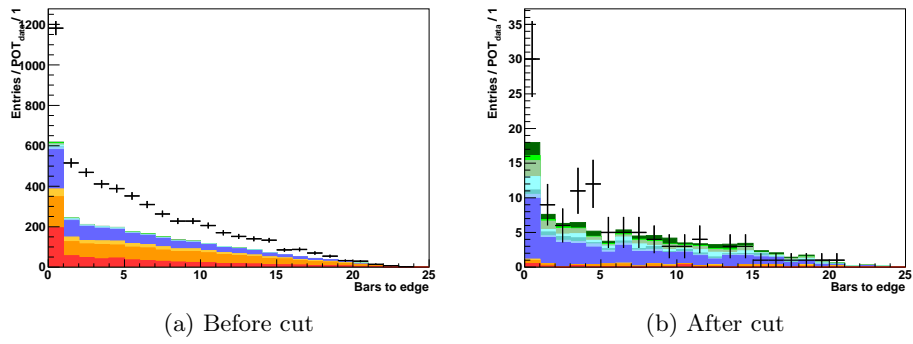


Figure A.21.: Low fiducialty




Figure A.22.: High thrust

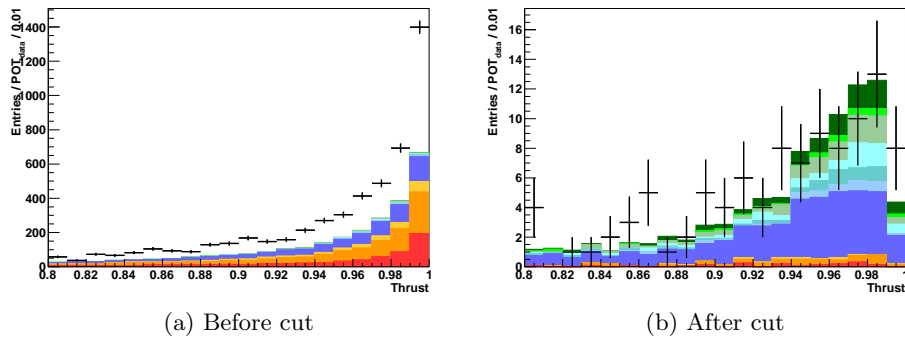


Figure A.23.: Low thrust






Figure A.24.: Energy asymmetry




(a) Before cut




(b) After cut

Figure A.25.: Angle between the two clusters




(a) Before cut




(b) After cut

Figure A.26.: Invariant mass



(a) Before cut



(b) After cut

Figure A.27.: Distance between the two clusters

Bibliography

- [1] C. D. Ellis and W. A. Wooster, Proceedings of the Royal Society of London. Series A **117**, 109 (1927), doi:10.1098/rspa.1927.0168.
- [2] L. Meitner and W. Orthmann, Zeitschrift f. Physik **60**, 143 (1930), doi:10.1007/BF01339819.
- [3] W. Pauli, Pauli letter collection: letter to Lise Meitner, CERN-ARCH-PLC:meitner_0393, 1930.
- [4] E. Fermi, Zeitschrift f. Physik **88**, 161 (1934), doi:10.1007/BF01351864.
- [5] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, Science **124**, 103 (1956), doi:10.1126/science.124.3212.103.
- [6] F. Reines and C. L. Cowan, Nature **178**, 446 (1956), doi:10.1038/178446a0.
- [7] R. Davis, Phys. Rev. **97**, 766 (1955), doi:10.1103/PhysRev.97.766.
- [8] G. Danby *et al.*, Phys. Rev. Lett. **9**, 36 (1962), doi:10.1103/PhysRevLett.9.36.
- [9] G. S. Abrams *et al.*, Phys. Rev. Lett. **63**, 2173 (1989), doi:10.1103/PhysRevLett.63.2173.

- [10] B. Adeva *et al.*, Physics Letters B **231**, 509 (1989), doi:10.1016/0370-2693(89)90703-X.
- [11] D. DeCamp *et al.*, Physics Letters B **231**, 519 (1989), doi:10.1016/0370-2693(89)90704-1.
- [12] M. Akrawy *et al.*, Physics Letters B **231**, 530 (1989), doi:10.1016/0370-2693(89)90705-3.
- [13] P. Aarnio *et al.*, Physics Letters B **231**, 539 (1989), doi:10.1016/0370-2693(89)90706-5.
- [14] K. Kodama *et al.*, Physics Letters B **504**, 218 (2001), doi:10.1016/S0370-2693(01)00307-0.
- [15] R. Davis, D. S. Harmer, and K. C. Hoffman, Phys. Rev. Lett. **20**, 1205 (1968), doi:10.1103/PhysRevLett.20.1205.
- [16] B. T. Cleveland *et al.*, The Astrophysical Journal **496**, 505 (1998), doi:10.1086/305343.
- [17] P. Anselmann *et al.*, Physics Letters B **314**, 445 (1993), doi:10.1016/0370-2693(93)91264-N.
- [18] J. Abdurashitov *et al.*, Physics Letters B **328**, 234 (1994), doi:10.1016/0370-2693(94)90454-5.
- [19] T. J. Haines *et al.*, Phys. Rev. Lett. **57**, 1986 (1986), doi:10.1103/PhysRevLett.57.1986.
- [20] K. Hirata *et al.*, Physics Letters B **205**, 416 (1988), doi:10.1016/0370-2693(88)91690-5.

- [21] Y. Fukuda *et al.*, Physics Letters B **335**, 237 (1994), doi:10.1016/0370-2693(94)91420-6.
- [22] K. S. Hirata *et al.*, Phys. Rev. Lett. **65**, 1297 (1990), doi:10.1103/PhysRevLett.65.1297.
- [23] Super-Kamiokande Collaboration, Y. Fukuda *et al.*, Phys. Rev. Lett. **81**, 1562 (1998), doi:10.1103/PhysRevLett.81.1562.
- [24] Q. R. Ahmad *et al.*, Phys. Rev. Lett. **89**, 011301 (2002), doi:10.1103/PhysRevLett.89.011301.
- [25] M. Apollonio *et al.*, The European Physical Journal C - Particles and Fields **27**, 331 (2003), 10.1140/epjc/s2002-01127-9.
- [26] KamLAND Collaboration, K. Eguchi *et al.*, Phys. Rev. Lett. **90**, 021802 (2003), doi:10.1103/PhysRevLett.90.021802.
- [27] The KamLAND Collaboration, S. Abe *et al.*, Phys. Rev. Lett. **100**, 221803 (2008), doi:10.1103/PhysRevLett.100.221803.
- [28] K2K Collaboration, M. H. Ahn *et al.*, Phys. Rev. D **74**, 072003 (2006), doi:10.1103/PhysRevD.74.072003.
- [29] MINOS Collaboration, P. Adamson *et al.*, Phys. Rev. Lett. **101**, 131802 (2008), doi:10.1103/PhysRevLett.101.131802.
- [30] Z. Maki, M. Nakagawa, and S. Sakata, Progress of Theoretical Physics **28**, 870 (1962), doi:10.1143/PTP.28.870.
- [31] B. Pontecorvo, Sov. Phys. JETP **26**, 984 (1968).
- [32] Particle Data Group, K. Nakamura *et al.*, J. Phys. G **37**, 075021 (2010), doi:10.1088/0954-3899/37/7A/075021.

- [33] L. Wolfenstein, Phys. Rev. D **17**, 2369 (1978), doi:10.1103/PhysRevD.17.2369.
- [34] S. Mikheyev and A. Smirnov, Sov. J. Nucl. Phys. **42**, 913 (1985).
- [35] P. Lipari, Nuclear Physics B - Proceedings Supplements **112**, 274 (2002), doi:10.1016/S0920-5632(02)01783-8.
- [36] D. Rein and L. M. Sehgal, Annals of Physics **133**, 79 (1981), doi:10.1016/0003-4916(81)90242-6.
- [37] D. Rein and L. M. Sehgal, Nuclear Physics B **223**, 29 (1983), doi:10.1016/0550-3213(83)90090-1.
- [38] G. Zeller, (2003), arXiv:hep-ex/0312061 [hep-ex].
- [39] W. Krenz *et al.*, Nuclear Physics B **135**, 45 (1978), doi:10.1016/0550-3213(78)90213-4.
- [40] The MiniBooNE Collaboration, A. A. Aguilar-Arevalo *et al.*, Phys. Rev. D **81**, 013005 (2010), doi:10.1103/PhysRevD.81.013005.
- [41] Provided by the T2K collaboration.
- [42] Adapted from <http://j-parc.jp/Acc/en/layoutE.html>.
- [43] S. van der Meer, (1961), CERN-61-07.
- [44] V. Galymov, Predicting neutrino flux for T2K, NuFact'11 XIIIth International Workshop on Neutrino Factories, Super beams and Beta beams, 2011.
- [45] The NA61/SHINE Collaboration, N. Abgrall *et al.*, Phys. Rev. C **84**, 034604 (2011), doi:10.1103/PhysRevC.84.034604.

- [46] G. Arnison *et al.*, Physics Letters B **126**, 398 (1983), doi:10.1016/0370-2693(83)90188-0.
- [47] J. Altegoer *et al.*, Nucl. Instr. Meth. A **404**, 96 (1998), doi:10.1016/S0168-9002(97)01079-6.
- [48] Y. Giomataris, P. Rebourgeard, J. Robert, and G. Charpak, Nucl. Instr. Meth. A **376**, 29 (1996), doi:10.1016/0168-9002(96)00175-1.
- [49] D. Renker and E. Lorenz, Journal of Instrumentation **4**, P04004 (2009), doi:10.1088/1748-0221/4/04/P04004.
- [50] A. Vacheret *et al.*, Nucl. Instr. Meth. A **656**, 69 (2011), doi:10.1016/j.nima.2011.07.022.
- [51] J. Estrada, C. Garcia, B. Hoeneisen, and P. Rubinov, FERMILAB-TM-2226 (2003).
- [52] The T2K Collaboration, Y. Itow *et al.*, KEK-REPORT-2001-4, ICRR-REPORT-477-2001-7, TRI-PP-01-05 (2001), arXiv:hep-ex/0106019 [hep-ex].
- [53] S. Nakayama *et al.*, Phys. Lett. B **619**, 255 (2005), doi:10.1016/j.physletb.2005.05.044.
- [54] Official plot of the T2K collaboration, unknown author.
- [55] R. Brun and F. Rademakers, Nucl. Instr. Meth. A **389**, 81 (1997), doi:10.1016/S0168-9002(97)00048-X.
- [56] A. Ferrari, P. Sala, A. Fasso, and J. Ranft, (2005), CERN-2005-10, INFN/TC_05/11, SLAC-R-773.

- [57] R. Brun, F. Bruyant, M. Maire, A. C. McPherson, and P. Zancarini, *GEANT 3: user's guide Geant 3.10, Geant 3.11; rev. version* (CERN, Geneva, 1987), CERN DD/EE/84-1.
- [58] Y. Hayato, *Acta Physica Polonica B* **40**, 2477 (2009).
- [59] C. Andreopoulos *et al.*, *Nucl. Instr. Meth. A* **614**, 87 (2010), arXiv:0905.2517 [hep-ph], doi:10.1016/j.nima.2009.12.009.
- [60] S. Agostinelli *et al.*, *Nucl. Instr. Meth. A* **506**, 250 (2003), doi:10.1016/S0168-9002(03)01368-8.
- [61] R. Frühwirth, P. Kubinec, W. Mitaroff, and M. Regler, *Computer Physics Communications* **96**, 189 (1996), doi:10.1016/0010-4655(96)00040-9.
- [62] H. Maesaka, *Evidence for Muon Neutrino Oscillation in an Accelerator-based Experiment*, PhD thesis, Kyoto University, 2005.
- [63] F. Di Lodovico, T2K results and future plans, NuFact'11 XIIIth International Workshop on Neutrino Factories, Super beams and Beta beams, 2011.
- [64] M. Day, Measuring the beam electron neutrino component in the T2K experiment, Lake Louise Winter Institute, 2011.
- [65] C. Giganti, *Particle Identification in the T2K TPCs and study of the electron neutrino component in the T2K neutrino beam*, PhD thesis, l'Université Paris-Sud 11, 2010.
- [66] A. Cervera-Villanueva, J. Gomez-Cadenas, and J. Hernando, *Nucl. Instr. Meth. A* **534**, 180 (2004), Proceedings of the IXth International

Workshop on Advanced Computing and Analysis Techniques in Physics Research, doi:10.1016/j.nima.2004.07.074.

- [67] V. Aleksakhin, Y. Bedfer, S. Gerasimov, and A. Korzenev, *Physics of Particles and Nuclei Letters* **4**, 350 (2007), doi:10.1134/S1547477107040103.
- [68] I. Taylor, *Development of T2K 280m Near Detector Software for Muon and Photon Reconstruction*, PhD thesis, Imperial College, University of London, 2009.
- [69] S. Brandt, C. Peyrou, R. Sosnowski, and A. Wroblewski, *Physics Letters* **12**, 57 (1964), doi:10.1016/0031-9163(64)91176-X.
- [70] E. Farhi, *Phys. Rev. Lett.* **39**, 1587 (1977), doi:10.1103/PhysRevLett.39.1587.
- [71] A. Carver, *Electron identification in and performance of the ND280 Electromagnetic Calorimeter*, PhD thesis, University of Warwick, 2010.
- [72] L. Michel, *Proc. Phys. Soc. A* **63**, 514 (1950), doi:10.1088/0370-1298/63/5/311.
- [73] C. Bouchiat and L. Michel, *Phys. Rev.* **106**, 170 (1957), doi:10.1103/PhysRev.106.170.
- [74] T. Kinoshita and A. Sirlin, *Phys. Rev.* **108**, 844 (1957), doi:10.1103/PhysRev.108.844.
- [75] Official plot of the T2K collaboration, produced by G. Christodoulou.
- [76] M. Berger, J. Coursey, M. Zucker, and J. Chang, *Stopping-Power and Range Tables for Electrons, Protons, and Helium Ions: Physical Reference Data; electronic version* (NIST, Gaithersburg, MD, 1999).

[77] A. Hoecker *et al.*, PoS **ACAT**, 040 (2007), CERN-OPEN-2007-007.