Measurement of kaon and pion contributions to the T2K neutrino beam

Rebekah J Smith
Lincoln College, Oxford

Thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford

Michaelmas Term 2013

Abstract

T2K is a long-baseline neutrino oscillation experiment located in Japan, designed to observe ν_μ disappearance and look for ν_e appearance in a ν_μ beam. A near detector (ND280) and far detector (Super-K), both positioned 2.5 degrees off-axis with respect to the beam, obtain measurements of the neutrino beam flavour at baselines of 280 m and 295 km, and values of Δm^2_{23}, θ_{23}, and θ_{13} are calculated. Collisions of protons on a graphite target produce a variety of hadrons, which decay to produce neutrinos. The two backgrounds to the ν_e appearance signal, intrinsic ν_e contamination of the beam and mis-identified interactions of π^0, are dependent on the quantities of kaons and pions entering the decay pipe. The predicted backgrounds are calculated using hadron production measurements conducted by the NA61/SHINE collaboration and theoretical models, however associated uncertainties are present. This thesis tests the suitability of an intrinsic ν_e measurement using the far detector, and concludes that with an expected event rate of 10.8$\pm$3.3 intrinsic ν_e in a nominal year, and a background of 11.3$\pm$3.4 events, several additional years of data taking are required to avoid domination by statistical uncertainty. The potential use of high energy ν_μ measurements at Super-K is also investigated. A method of measuring the contributions to the T2K ν_μ beam from different parents using ND280 is developed. Features of the decay kinematics of pions and kaons are used and a 17 parameter maximum-likelihood fit is applied, with the neutrino interaction cross-section uncertainties incorporated. Using RunII and RunIIIc ND280 data, corresponding to 2.1259×10^{20} POT, we find that the measured kaon and pion contributions agree with the Monte Carlo to within 1 σ. Sources of systematic uncertainty due to energy reconstruction and magnetic field uncertainty are calculated, and the dominant source of uncertainty is found to be due to current neutrino interaction cross-section uncertainties.

To everyone who shared a dance floor with me during the DPhil years

Acknowledgements

I'd like to begin by thanking the entire T2K collaboration, whose collective hard work makes it possible to study the finest details of our world. Thank you also to everyone in Japan whose organization and hospitality made visits not only easy but a joy. I'd like to thank my supervisor, Giles Barr, for his advice and attention to detail when patiently proofreading the excessive number of pages I ended up writing. I must also thank Matthew Malek and Ryan Terri for taking the time to visit Oxford to share their knowledge of Super-K when I first joined the experiment, and the wider T2KUK community for sharing their expertise online.

I'd like to thank the many inhabitants of 664a and the Double CHOOZ office, for the daily coffee break chat ranging from exciting new results to the latest xkcd. Thanks to Phil and Chris, for being generous with their talents for making ROOT behave; and Abbey, for decorating the office with beautiful tango shoes. Special thanks to Helen O'Keeffe, postdoc extraordinaire, for every discussion we had about ND280, generally being an excellent example of a physicist, and introducing group meetings with cake!

I'd like to express my gratitude to Lincoln College for their crucial support that enabled me to finish my analysis. I am also enormously grateful to my family, who supported me throughout this endeavour, but especially towards the end, when they provided a home for a semi-nocturnal creature who did nothing but type and nap for several months straight.

Last but not least, I want to thank my friends. Thanks to my fellow T2K students who made Tokai so memorable. Thank you to the wonderful Oxford salsa community, whose company was the treat at the end of every tough day. Thanks to Jacintha, for hot chocolate and sympathy; Olive, for providing incentive to finish by demonstrating the fun I could be having; Polly, for salad dates and ludicrously positive pep talks; and Statto, for always being a green circle in a crisis, be it LaTeX or existential.

Contents

List of Figures

List of Tables

Glossary

WLS Wavelength shifting fibre

PPO Organic scintillator $C_{15}H_{11}NO$

POPOP Organic scintillator $C_{24}H_{16}N_2O$

INGRID Interactive Neutrino GRID: On-axis near detector

ND280 Off-axis near detector

PMT Photomultiplier tube

Intrinsic ν_e A ν_e which is created in the decay pipe and not the result of $\nu_\mu \rightarrow \nu_e$ oscillation. Source of background to the ν_e appearance measurement

Signal ν_e A ν_e which appears in the ν_μ beam due to oscillation

CCQE Charged current quasi-elastic interaction

CCnonQE Charged current non-quasi-elastic interaction

NC Neutral current interaction

POLfit Pattern of Light Fit. Algorithm to look for second Cherenkov ring formed by decay of π^0 in Super-K

POT Protons On Target - quantifies the events contained in a file by the number of beam protons to hit the target

MC Monte Carlo generated events

ν_μ^{fK} Muon neutrinos resulting from the decay of beam kaons produced in the interaction of protons with the graphite target

$\nu_\mu^{f\pi}$ Muon neutrinos resulting from the decay of pions. Pions are either produced by the interaction of protons with the graphite target, or produced via kaon decays

Chapter 1

Introduction

The main aims of the T2K long-baseline neutrino oscillation experiment are to search for the appearance of ν_e in a ν_μ beam, and obtain a measurement of the mixing angle θ_{13}. There are two main backgrounds to the ν_e appearance signal. Ideally the T2K neutrino beam would initially consist of pure ν_μ, however a small amount of ν_e contamination is unavoidable. These intrinsic ν_e, created in the beamline along with ν_μ and $\bar{\nu}_\mu$, are indistinguishable from the signal ν_e that appear in the beam, and therefore must be precisely quantified. The second main background is π^0-producing neutral current interactions, as the π^0 can produce a signal in the far detector which can be mistaken for the interaction of a signal ν_e.

This thesis considers the sources of these two types of background, and how our ability to predict the rates of background events is affected by current uncertainties on the production of pions and kaons at the target. Firstly, we investigate the viability of a measurement of the intrinsic ν_e in the high energy tail of the far detector beam spectrum. A unique method of measuring the pion and kaon contributions to the ν_μ beam using near detector data is devised which exploits the variation of the spectrum with off-axis angle, and a measurement is obtained using Run II and RunIIIc data from the T2K experiment.

Chapter 2 provides an overview of neutrino theory, beginning with the early detection of neutrinos and including evidence for neutrino oscillations. The chapter concludes with a discussion of some of the remaining questions in the field of neutrino physics.

Chapter 3 lists the physics goals of the T2K experiment, and describes the features of the experimental design that make it especially well-suited to achieving those goals. The production of the neutrino beam is described, followed by the structures and functions of the on-axis and off-axis near detectors, and far detector.

In chapter 4 we focus on the intrinsic ν_e background and test the possibility of making a measurement of the intrinsic ν_e interaction rate at the far detector to check the accuracy of the background predictions. Selection criteria developed by the Super-K working group are used here, and additional cuts suitable for this analysis are then considered. Chapter 5 also uses Super-K selection cuts to select high energy beam ν_μ. The uses for high energy neutrino measurements are discussed, along with the limitations of these events.

Chapter 6 discusses the second source of background to the ν_e appearance measurement, and the relationship between both background sources and the neutrino parent types. Having precise models of the quantities of kaons and pions entering in the decay pipe is essential for accurately predicting both the intrinsic ν_e present in the beam and the π^0-producing neutral current interaction rate. The difficulty of obtaining a measurement of the neutrino flux is discussed.

Chapter 7 uses existing ν_μ selection cuts to find a sample of ν_μ interactions in ND280. The dependence of the neutrino flux shape on the off-axis angle and parent type is discussed and it is demonstrated that off-axis effects are visible within ND280. Energy reconstruction of CCQE and CCnonQE events is described and separate spectra of neutrinos produced from the decays of kaons and pions are found using Monte Carlo. Chapter 8 develops a method, using these spectra, to measure the separate neutrino fluxes produced by kaons and pions, and tests are conducted using fake data samples. This method is then applied to data in Chapter 9, and the results for Run II and Run IIIc data are presented.

Conclusions follow. A derivation of the decay kinematics of pions and kaons and the relationship between the neutrino energy and the opening angle of the decay is provided in Appendix A.

Chapter 2

Neutrino physics

A brief history of the discovery of neutrinos and their properties is given, concluding with the formulation of the Standard Model description of the neutrino. Several early experimental anomalies are described and their explanations are provided. The mechanism allowing neutrino oscillation is explained and experimental evidence confirming the occurrence of oscillations is summarised. The chapter concludes with a brief discussion of outstanding questions in neutrino physics.

2.1 Standard Model neutrinos

The existence of the neutrino was first postulated by Wolfgang Pauli in 1930 to explain an observed anomaly in β-decay experiments [1]. β-decay was assumed to proceed via ${}^{A}_{Z}\mathrm{N} \rightarrow {}^{A}_{Z+1}\mathrm{N'} + e^{-}$. Neglecting the nucleon recoil, the emitted electrons should be mono-energetic, possessing all the energy made available by the transition. Instead a continuous electron energy spectrum was observed for each type of β-decay, which suggested that a fraction of the available energy was being carried away by an additional, undetected product. This additional particle would need to be neutral and spin-1/2 to conserve charge and angular momentum, and light or massless. In 1934 Enrico Fermi included this theoretical light neutral particle, now named the neutrino, into his theory of β-decay [2].

The first direct detection of (anti)neutrinos was achieved by Reines and Cowan in 1953 [3], followed by a more sensitive measurement conducted in 1956 [4]. In both cases, antineutrinos produced by nuclear reactors were detected through the reaction $\bar{\nu}_e + p \rightarrow n + e^+$. Water containing the neutron absorber $CdCl_2$ was used, along with regions of scintillator, to detect a pattern of signals indicative of a neutrino interaction. The positron quickly annihilates with electrons in the water to produce two 511 MeV γ. Absorption of the neutron by cadmium via

$$\mathrm{n} +^{108}\mathrm{Cd} \rightarrow^{109}\mathrm{Cd}^* \rightarrow^{109}\mathrm{Cd} + \gamma$$

produces an additional γ. Detection of the two annihilation γ in coincidence followed by the neutron capture γ detected 5x10^{-6} s later indicates that a neutrino interaction has occurred.

That neutrinos occur in different flavour varieties was confirmed by the discovery of the muon neutrino by Lederman, Steinberger and Schwartz in 1962 [5]. A 15 GeV proton beam was directed onto a beryllium target to produce pions (and a smaller quantity of kaons). These would decay to produce neutrinos along with muons. The neutrinos were separated and directed to a spark chamber, where interactions resulted in the production of muons. The absence of any electron production led to the conclusion that two distinct types of neutrino exist: those that couple exclusively to muons, labelled ν_μ, and those that couple only to electrons, ν_e, the antimatter equivalent of which had been observed in reactor neutrino experiments.

The total number of neutrino flavours was inferred from measurements made by the four Large Electron Positron Collider (LEP) experiments ALEPH, DELPHI, L3 and OPAL. LEP used collisions of e^+ and e^- at a range of beam energies to study the line shape of the Z^0 resonance. The total decay width of the Z^0, Γ_Z, can be expressed in terms of partial decay widths as follows

$$\Gamma_Z = \Gamma_{l^+l^-} + \Gamma_{hadrons} + \mathrm{N}_\nu \Gamma_\nu \tag{2.1}$$

where Γ_l, $\Gamma_{hadrons}$ and Γ_ν are the partial decay widths to charged lepton pairs, quark pairs and neutrinos respectively. N_ν is the number of neutrino varieties which can couple to the Z^0 (and therefore not inclusive of potential "sterile" neutrinos) with mass $\mathrm{m}_\nu < \mathrm{M}_Z/2$. The shape of the resonance peak can be modelled for various values of N_ν, and the effect on the line shape is shown in Figure 2.1. Experimental data is also shown in Figure 2.1, and clearly correlates almost perfectly with the N_ν=3 case. Analysis of the combined results yields a value of N_ν of 2.9840 $\pm$ 0.0082 [6], which matches the observed number of charged lepton and quark generations. In addition to missing energy observed in decays of the τ lepton [7], this suggested that an additional neutrino flavour, associated with the tau lepton, also existed, and the ν_τ was finally detected by the DONUT collaboration in 2000 [8].

Experiments conducted by Goldhaber et al. in 1958 set out to measure the helicity of the neutrino [9]. In contrast to all massive particles, for which the helicity of a particle is dependent on the observer's frame of reference, neutrinos were found to consistently occur with helicity of -1. All neutrinos occurring in nature were found to be left-handed, and all anti-neutrinos were found to be right-handed. Since the helicity is fixed, and cannot be altered by changing the frame of reference, it was concluded that neutrinos travel at speed c, and therefore have zero mass.

When the Standard Model of particle physics was developed in the mid-1970s, three generations of neutrino were included, matching the three lepton and quark families. Evidence such as the absence of the interaction $\nu + {}^{37}\mathrm{Cl} \rightarrow \mathrm{e}^- + {}^{37}\mathrm{Ar}$

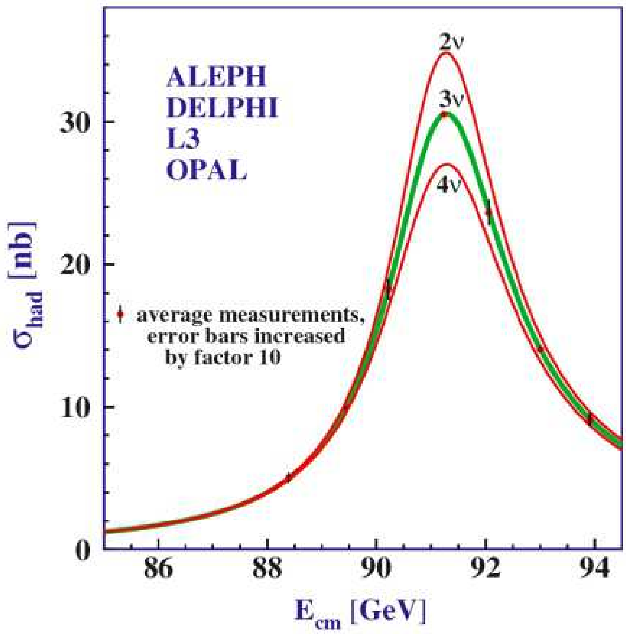

Figure 2.1: Cross-section $\sigma(e^+e^- \rightarrow Z^0 \rightarrow \text{hadrons})$ as a function of the centre of mass energy measured by The ALEPH Collaboration et al. [6] Theoretical resonance shapes for 2, 3 and 4 neutrino flavours are plotted, along with data points. Data clearly shows good agreement with N_ν=3 scenario to excellent precision.

in the presence of a reactor antineutrino flux, as observed by Ray Davis in 1955 [10], indicated that neutrinos and antineutrinos were distinct particles. Therefore the Standard Model includes six neutrinos: ν_e, ν_μ and ν_τ and their antimatter partners. All neutrinos were described as massless, with only left-handed ν and right-handed $\bar{\nu}$ occurring.

The Standard Model of particle physics has been extremely successful and provided an excellent description of the particles and interactions we observe. However, two examples of experimental measurements were gathered which could not be explained by existing physics models. In both cases, one possible explanation of the observations was that the properties of neutrinos as described by the Standard Model were in need of modification.

The first set of experimental measurements that could not be immediately explained demonstrate the Solar Neutrino Problem. The early neutrino detection experiments conducted by Ray Davis, which aimed to detect neutrinos through the inverse β decay process $\nu_e + {}^{37}\mathrm{Cl} \rightarrow \mathrm{e}^- + {}^{37}\mathrm{Ar}$, were continued at the Homestake Mine [11]. Using models of the nuclear fission chains operating in the sun and measurements of the solar energy flux, the solar neutrino flux could be predicted. The Homestake measurement released in 1968 provided an upper limit on the solar neutrino flux, and found it to be approximately 1/3 of the predicted flux. This presented three possibilities: an experimental error, an error in the Standard Solar Model, or an inaccuracy in the Standard Model description of neutrinos.

Further experiments, using different detection techniques, then confirmed the neutrino flux deficit. The Gallium Experiment (GALLEX), Gallium Neutrino Observatory (GNO) and Soviet-American Gallium Experiment (SAGE) used the interaction $\nu_e + {}^{71}\mathrm{Ga} \rightarrow \mathrm{e}^- + {}^{71}\mathrm{Ge}$ to detect solar neutrinos. This reaction was sensitive to neutrinos of lower energy than those detectable by Homestake and a flux deficit was again observed, with GALLEX/GNO [12] observing 0.58$\pm$0.05% the predicted ν_e flux, and SAGE [13] observing 0.60$\pm$0.05% [14]. A third type of measurement conducted by the Kamioka Nucleon Decay Experiment (KamiokaNDE), and later the Super KamiokaNDE (Super-K), both water Cherenkov detectors, also confirmed the existence of a deficit, with Super-K observing less than half the expected neutrino flux [15]. This collection of independent deficit observations, using different interactions and measurement techniques, made experimental errors a very unlikely explanation and indicated that a modification to either the Standard Model or Standard Solar Model was required.

The second set of anomalous observations were an unexpected feature of the atmospheric neutrino flux. Atmospheric neutrinos are produced by interactions of cosmic rays such as protons with the nuclei of molecules in the Earth's atmosphere, and result in a $\nu_\mu/\bar{\nu_\mu}$:ν_e ratio of approximately 2:1 [16]. Early measurements of atmospheric neutrinos conducted by water Cherenkov detectors found a deficit [17] of ν_μ and $\bar{\nu_\mu}$ however. Detailed results published by Super-K in 1998 revealed a dependence on the direction of travel of the neutrinos, quantified by the zenith angle θ. The flux of neutrinos detected with zenith angles $\cos\theta > 0$, corresponding to neutrinos created in the atmosphere above the detector and traveling approximately 20 km down to the Earth's surface, were as predicted. However, the flux of ν_μ and $\bar{\nu_\mu}$ detected with $\cos\theta < 0$, corresponding to neutrinos created in the atmosphere on the other side of the Earth and traveling up through the planet to the detector, over a distance of approximately 13000 km, featured a deficit. This observation suggested that neutrino flux may be dependent on the distance traveled by a neutrino between creation and subsequent detection.

The solutions to these anomalies are described in §2.3. First, a description of the necessary extension to the Standard Model neutrino is given below.

2.2 Neutrino mixing

The theoretical framework for neutrino flavour mixing, analogous to the quark mixing observed in some meson systems, was first suggested in 1957 by Bruno Pontecorvo [18]. This was developed further by Maki, Nakagawa, and Sakata [19], and later Pontecorvo, during the 1960s. The neutrinos described by the Standard Model all have zero mass. However, should the neutrino masses not be degenerate, requiring that at least one neutrino has a finite mass, it would be possible for flavour oscillations to occur.

The oscillation mechanism can be demonstrated using two neutrino flavours, which we will label ν_α and ν_β. We can also describe the two neutrinos in terms of their masses, where ν_1 has mass m_1 and ν_2 has mass m_2. Oscillation can occur if the flavour eigenstates, denoted $|\nu_\alpha\rangle$ and $|\nu_\beta\rangle$, consist of a combination of the mass eigenstates. We can describe this mixing of flavour states and mass states in terms of a rotation angle θ, where

$$\begin{pmatrix} \nu_\alpha \\ \nu_\beta \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} \tag{2.2}$$

Expanding Eq.2.2 gives each flavour eigenstate in terms of a linear combination of the mass eigenstates.

$$|\nu_\alpha\rangle = \cos\theta|\nu_1\rangle + \sin\theta|\nu_2\rangle \tag{2.3a}$$

$$|\nu_\beta\rangle = -\sin\theta|\nu_1\rangle + \cos\theta|\nu_2\rangle \tag{2.3b}$$

Neutrinos are created and detected via the weak interaction as pure flavour eigenstates. However, the propagation of a neutrino is dependent on the neutrino mass, and therefore the different mass eigenstates will propagate through space at slightly different speeds. The evolution of these mass eigenstates $|\nu_n(t)\rangle$, where n = 1 or 2, can be expressed in terms of the initial mass eigenstates $|\nu_n\rangle$ and the neutrino energies E_n using a simple plane wave solution as follows

$$|\nu_1(t)\rangle = |\nu_1\rangle e^{iE_1 t} \tag{2.4a}$$

$$|\nu_2(t)\rangle = |\nu_2\rangle e^{iE_2 t} \tag{2.4b}$$

Consider a neutrino created in a weak interaction in pure flavour state α. At time t = 0 the neutrino wavefunction is described by Eq. 2.3a. Replacing the initial mass eigenstates with the time-dependent mass eigenstates gives an expression for the composition of the neutrino wavefunction at time t, in terms of the pure mass states and neutrino energies, as follows

$$
\begin{aligned}
|\nu_\alpha(t)\rangle &= \cos\theta|\nu_1\rangle e^{iE_1t} + \sin\theta|\nu_2\rangle e^{iE_2t} \\
&= (\cos\theta|\nu_1\rangle + \sin\theta|\nu_2\rangle e^{i(E_2-E_1)t})e^{iE_1t}
\end{aligned} \tag{2.5}
$$

Expressing the energy of a particle in terms of the momentum p and mass m, we find

$$E = \left[p^2\left(1+\frac{m^2}{p^2}\right)\right]^{1/2} \approx p\left(1+\frac{m^2}{2p^2}\right) \tag{2.6}$$

$$E_2 - E_1 \approx p_2 + \frac{m_2^2}{2p_2} - p_1 - \frac{m_1^2}{2p_1} \tag{2.7}$$

Assuming the neutrinos to be ultrarelativistic[1], we state $\mathrm{p}_1 = \mathrm{p}_2$ and letting both be equal to E, find

$$E_2 - E_1 = \frac{m_2^2 - m_1^2}{2E} \tag{2.8}$$

$$= \frac{\Delta m^2}{2E} \tag{2.9}$$

where Δm^2 is the mass squared difference $m_2^2 - m_1^2$.

We now have the wavefunction of the neutrino initially created in a pure flavour state $|\nu_\alpha\rangle$, at time t after the creation of the neutrino, in terms of the pure mass eigenstates, mass difference and neutrino energy.

$$|\nu_\alpha(t)\rangle = \left(\cos\theta|\nu_1\rangle + sin\theta|\nu_2\rangle e^{\frac{i\Delta m^2 t}{2E}}\right) e^{iEt} \tag{2.10}$$

Rearranging equations 2.3a and 2.3b we find the mass eigenstates in terms of the pure flavour states.

$$|\nu_1\rangle = \cos\theta|\nu_\alpha\rangle - \sin\theta|\nu_\beta\rangle \tag{2.11a}$$

[1]The extremely small mass of the neutrino results in large Lorentz factors e.g. for a neutrino created in a nuclear reaction, with energy of order 1 MeV, γ is greater than 10^6.

$$|\nu_2\rangle = \sin\theta|\nu_\alpha\rangle + \cos\theta|\nu_\beta\rangle \tag{2.11b}$$

Substituting these into Eq. 2.10 we find the time-dependent wavefunction $|\nu_\alpha(t)\rangle$ in terms of the initial, pure flavour states:

$$|\nu_\alpha(t)\rangle = \left[|\nu_\alpha\rangle\left(\cos^2\theta + \sin^2\theta e^{i\phi}\right) + |\nu_\beta\rangle\left(\sin\theta\cos\theta\left(e^{i\phi} - 1\right)\right)\right] e^{iEt} \tag{2.12}$$

where $\phi = \frac{\Delta m^2 t}{2E}$.

The neutrino therefore propagates as a superposition of flavour states until the neutrino interacts with matter as one of the possible flavour states via a weak interaction. The probability that the neutrino, initially created in eigenstate α, will also interact in flavour eigenstate α having travelled for time t, is then given by

$$\text{Prob}(\nu_\alpha \to \nu_\alpha) = |\langle\nu_\alpha|\nu_\alpha(t)\rangle|^2 \tag{2.13}$$

Since the flavour eigenstates are orthogonal, such that $\langle\nu_\alpha|\nu_\alpha\rangle = 1$ and $\langle\nu_\beta|\nu_\alpha\rangle = 0$ we find

$$\text{Prob}(\nu_\alpha \to \nu_\alpha) = |\left(\cos^2\theta + \sin^2\theta e^{i\phi}\right) e^{iEt}|^2 \tag{2.14}$$

Substituting $|e^{iEt}|^2 = 1$ and $e^{ix} = \cos x + i\sin x$, gives

$$\begin{aligned}
\text{Prob}(\nu_\alpha \to \nu_\alpha) &= |\left(\cos^2\theta + \cos\phi\sin^2\theta + i\sin\phi\sin^2\theta\right)|^2 \\
&= \left(\cos^2\theta + \cos\phi\sin^2\theta\right)^2 + \left(\sin\phi\sin^2\theta\right)^2 \\
&= \cos^4\theta + \sin^4\theta + 2\cos\phi\cos^2\theta\sin^2\theta
\end{aligned} \tag{2.15}$$

Using the identities $\cos^2 x + \sin^2 x = 1$, $\sin(2\theta) = 2\sin\theta\cos\theta$ and

$2\sin^2 x = 1-\cos(2x)$, this reduces to

$$\text{Prob}(\nu_\alpha \to \nu_\alpha) = 1 - \sin^2(2\theta)\sin^2\left(\frac{\phi}{2}\right) \tag{2.16}$$

Since the extremely light mass of the neutrino results in a velocity of effectively the speed of light, we can relate the time elapsed between the creation of the neutrino and its subsequent detection via the weak interaction, t, with the distance travelled, L. Replacing ϕ, and adding the numerical factors necessary to convert the variables into appropriate units, we find that the survival probability of the ν_α is given by

$$\text{Prob}(\nu_\alpha \to \nu_\alpha) = 1 - \sin^2(2\theta)\sin^2\left(\frac{1.27\Delta m^2 L}{E}\right) \tag{2.17}$$

where Δm^2 is in units eV^2, the baseline L is given in km and the neutrino energy E is measured in GeV.

To find the probability that the neutrino initially created as $|\nu_\alpha\rangle$ will later be measured as $|\nu_\beta\rangle$, we instead take the amplitude formed from the pure $|\nu_\beta\rangle$ state with the $|\nu_\alpha\rangle$ wavefunction at time t. Since we are considering a scenario with only two neutrino types, the probabilities of a neutrino interacting as $|\nu_\alpha\rangle$ and $|\nu_\beta\rangle$ must add to one. Therefore the oscillation probability is given as

$$\begin{aligned}\text{Prob}(\nu_\alpha \to \nu_\beta) &= |\langle\nu_\beta|\nu_\alpha(t)\rangle|^2 &(2.18)\\ &= 1 - \text{Prob}(\nu_\alpha \to \nu_\alpha) &(2.19)\\ &= \sin^2(2\theta)\sin^2\left(\frac{1.27\Delta m^2 L}{E}\right) &(2.20)\end{aligned}$$

We know that there are in fact three neutrino flavours: ν_e, ν_μ and ν_τ. The mixing of these flavour or weak interaction eigenstates and the three mass eigenstates is described by

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = U_{PMNS} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{2.21}$$

where U_{PMNS} is the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [20]. This describes the flavour and mass mixing in terms of three rotations, parameterised as θ_{12}, θ_{23} and θ_{13}. Three complex phases are also required. If we assume that the neutrino is a Dirac particle, the two Majorana phases α_1 and α_2 may be omitted, and we instead include one complex phase δ. Using c_{ij} = cosθ_{ij} and s_{ij} = sinθ_{ij}, the PMNS matrix can be expressed as

$$U_{PMNS} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{2.22}$$

$$= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \tag{2.23}$$

Using Eq. 2.21, each neutrino flavour may be expressed in terms of the three mass eigenstates, and the method described above for the two flavour case can be used to find the survival probability for any neutrino flavour, and the oscillation probability for any pair of flavours. Eq. 2.22 is a convenient method of arranging the PMNS matrix as it emphasises the dominance of certain oscillation parameters at some specific naturally occurring baseline and energy scenarios. The formula for the survival probability of ν_μ created in our atmosphere is dominated by terms containing θ_{23}, and therefore the first matrix in Eq. 2.22 is historically referred to as the atmospheric mixing. Oscillations of solar ν_e are most sensitive to θ_{12}, therefore the last matrix represents the solar mixing. Mixing angle θ_{13} is the smallest of the mixing angles, and has only a second order effect on oscillation and survival probabilities, making it the last angle to be measured. A small dependence on θ_{13} can be seen however in the oscillations of reactor neutrinos and beam neutrinos observed over particular values of L/E, as described in more detail in §2.3.

2.3 Extensions to the Standard Model

The process of neutrino oscillation described in §2.2 was initially purely theoretical, however evidence of the atmospheric and solar neutrino anomalies continued to grow. Modelled neutrino oscillations were found to fit the experimental measurements extremely well, explaining the observed deficits of the atmospheric ν_μ flux and solar ν_e flux.

The resolution of the atmospheric neutrino anomaly was achieved in the late 1990s by the Super-K collaboration. The Super-K water Cherenkov detector was able to distinguish between interactions of $\nu_\mu/\bar{\nu_\mu}$[2] and $\nu_e/\bar{\nu_e}$. High precision measurements released in 1998 confirmed that the $\nu_\mu/\bar{\nu_\mu}$ flux was below expectations while the ν_e flux was as expected. Super-K is not sensitive to ν_τ interactions, therefore $\nu_\mu \rightarrow \nu_\tau$ oscillations would only be apparent through the disappearance of ν_μ. As shown in Eq. 2.17, the survival probability of a neutrino flavour is dependent on the distance travelled by the neutrino, L, before observation. Therefore for neutrinos of a fixed energy, the survival probabilities would differ for neutrinos created directly above the detector, and neutrinos that travel ~13,000 km through the Earth before interaction. The data were found to fit with excellent agreement to the modelled ν_μ flux assuming $\nu_\mu \rightarrow \nu_\tau$ oscillations, as demonstrated in Figure 2.2. Considering the negligible deficit or excess of ν_e detected, $\nu_\mu \rightarrow \nu_e$ oscillations were

[2] Super-K is not magnetised therefore it is not possible to determine the sign of charges on leptons produced in charged current interactions and deduce if a ν or $\bar{\nu}$ interaction occurred.

assumed to be negligible and a two flavour approximation was used to model the $\nu_\mu \rightarrow \nu_\tau$ oscillations, finding $\sin^2 2\theta > 0.82$ and $5 \times 10^{-4} < \Delta m^2 < 6 \times 10^{-3}$ eV^2. In the full three flavour system, these correspond to the atmospheric parameters θ_{23} and Δm^2_{23} [21].

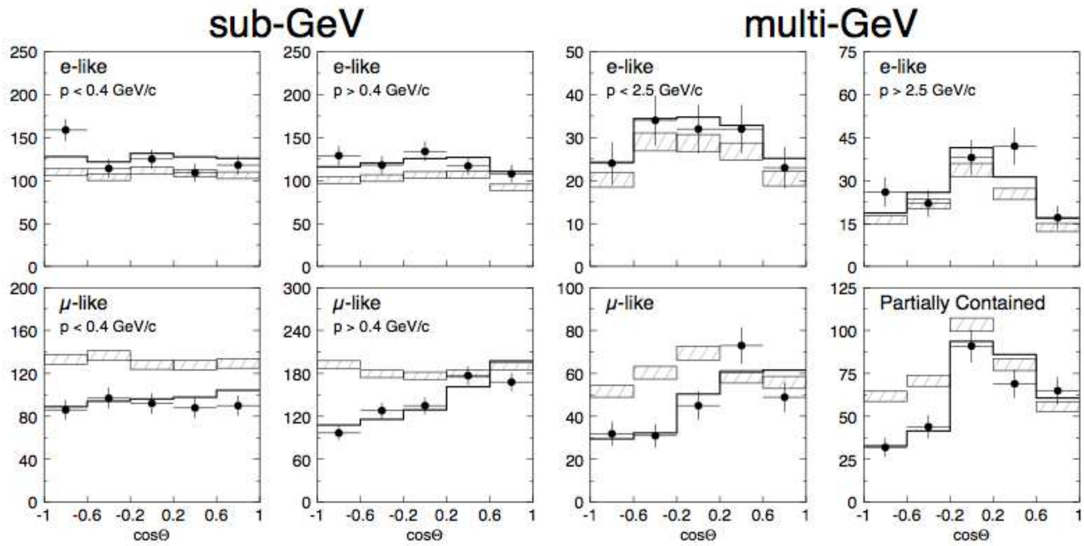

Figure 2.2: Zenith angle distributions of μ-like and e-like events detected at Super Kamiokande. Upward-going (through the Earth) particles have $\cos\Theta < 0$ and downward (created above the detector) have $\cos\Theta > 0$. The hatched region shows the Monte Carlo expectation for no oscillations. The bold line is the best-fit expectation for $\nu_\mu \rightarrow \nu_e$ oscillations. Figure from [21].

Further measurements of the atmospheric parameters have since been conducted using long baseline accelerator experiments such as the Main Injector Neutrino Oscillation Search (MINOS) [22] and K2K [23]. Such experiments provide a useful confirmation of the atmospheric neutrino observations, and can also be used to conduct more precise calculations of oscillation parameters since the neutrino energy spectrum can be measured and the baseline is constant and well known. The disappearance of ν_μ can be studied by creating a beam of pure ν_μ of a given peak energy E, and sampling the flavour of this beam once at the beam source and again at a large baseline L, such that L/E corresponds to a low survival probability. Comparison of the ν_μ interaction rates at the different distances allows θ_{23} and

Δm^2_{23} to be found. The most recent measurements of the atmospheric mixing parameters were conducted by the Super-K collaboration [24], MINOS collaboration [25], and the T2K experiment [26].

Neutrino oscillations also provided the explanation for the Solar Neutrino Problem. Results announced by the Sudbury Neutrino Observatory (SNO) in 2002 provided conclusive proof that neutrino oscillations were occurring. SNO, a heavy water Cherenkov detector, was able to measure the solar ν_e flux via the charged current (CC) interaction, which preserves the lepton flavour. It was also sensitive to neutral current (NC) interactions and elastic scattering (ES), which occur for all neutrinos without leaving an indication of the flavour, and therefore provide a measurement of the total neutrino flux inclusive of all active flavours. It was found that while a deficit of ν_e was observed, the total neutrino flux agreed with the neutrino flux prediction of the Standard Solar Model [27]. Since all neutrinos created in nuclear processes in the Sun are created as ν_e, and the total flux was as predicted, this provided proof that ν_e disappearance was occurring between the neutrinos' creation and subsequent detection on Earth. The solar mixing parameters have been refined with measurements from several experiments, including Super-K, which continued to measure the solar neutrino flux, KamLAND [28] and SNO [27].

The neutrino oscillations considered in §2.2 apply to neutrinos created in and propagating through a vacuum. Due to the high density of the environment in which the solar neutrinos are created, matter effects must also be taken into account when considering solar neutrinos. The effect of matter on oscillations was first studied by Wolfenstein [29] in 1978 and developed further by Mikheev and Smirnov [30], and is also referred to as the MSW effect. Ordinary matter contains e^- but lacks μ^- or τ^-. The result is that ν_e experience matter in a different way from ν_μ and ν_τ as they propagate through it, due to the different range of scattering interactions available to ν_e. Neutral current interactions are flavour independent, as via exchange of a Z^0, any neutrino may scatter through one of the processes

$$\nu_x + e^- \rightarrow \nu_x + e^-$$

$$\nu_x + N \rightarrow \nu_x + N$$

where N is a nucleon. For ν_e, a charged current scattering process is also possible. By exchanging a W^- with an atomic electron, coherent forward scattering can occur, proceeding as

$$\nu_e + e^- \rightarrow \nu_e + e^-$$

The result is that the mixing angles and mass-squared differences are modified in high density environments, so that the probability that a neutrino of mass ν_2 will be measured with flavour e is altered (see Figure 2.3 for vacuum mixing). At certain densities and energies, a resonant effect can produce maximal flavour mixing despite the relevant vacuum oscillation angle being small. MSW effects are not significant for neutrinos passing through a short length of relatively low density matter such as the Earth's crust (as in the case of T2K); however, the observed value of the solar ν_e flux can only be accounted for with the inclusion of matter effects. This is discussed in more detail in [31].

Measurements of the last mixing angle, θ_{13}, are ongoing. Due to this mixing angle being very small compared to θ_{12} and θ_{23}, θ_{13}-dependent oscillations were not apparent in observations of naturally occurring neutrinos. However, two experimental designs operate at the necessary L/E for θ_{13} to be measured. The reactor experiment CHOOZ, comprised of a liquid scintillator target, measured the disappearance of reactor $\bar{\nu}_e$ over a baseline of 1 km, and restricted θ_{13} to small values, placing an upper limit on the value of $\sin^2 2\theta_{13}$ of 0.10 at 90% confidence [32]. The extension to this experiment, Double CHOOZ [33], released their first results in 2011. θ_{13} can also be measured using long baseline measurements of a ν_μ beam. T2K also released their first measurements of θ_{13}, along with the first indication of ν_e appearance in a ν_μ beam, in 2011 [34], with further results of higher precision in 2013 [35]. In 2012, the Daya Bay reactor neutrino experiment released findings of a non-zero value for θ_{13} with a significance of 5.2 σ [36], with an improved result released in 2013 [37]. Also in 2012 the RENO reactor experiment observed the disappearance of reactor $\bar{\nu_e}$ to 4.9 σ [38].

Parameter	Best measurement
$\sin^2(2\theta_{12})$	0.857 ± 0.024
Δm^2_{21}	$(7.50 \pm 0.20)\ \times 10^{-5}$ eV^2
$\sin^2(2\theta_{23})$	> 0.95 (90% CL)
Δm^2_{32}	$(2.32\ ^{+0.12}_{-0.08})\ \times 10^{-3}$ eV^2
$\sin^2(2\theta_{13})$	0.095 ± 0.010

Table 2.1: Values of neutrino mixing parameters using combined experimental measurements as quoted by the 2013 Review of Particle Physics [39].

The most precise measurements of the oscillation parameters at the time of writing are listed in Table 2.1. As discussed in §2.2, neutrino oscillation is only possible if neutrinos of non-degenerate mass exist. Conclusive proof that neutrino oscillations occur, such as the comparison of the solar ν_e flux with the total solar ν flux, and the appearance of ν_e in a beam of initially pure ν_μ, requires that the Standard Model of particle physics be extended to allow for three neutrinos of different mass. Further extensions may be required. While good progress has been made measuring the mixing angles and mass-squared differences, there are still many areas of neutrino physics that require investigation. Some of these unknown features, along with their implications and plans for future study, are discussed below.

2.4 Outstanding questions

The mixing parameters can only be measured using observations of neutrino oscillations. When considering the full three flavour treatment of oscillations, we find that the oscillation and survival probabilities are often dependent on several parameters, making it impossible to isolate any one parameter for individual study. This is the case with measurements of the last mixing angle, θ_{13}. The full three flavour expression for the probability of oscillation $\nu_\mu \to \nu_e$ [40] is

$$\begin{aligned}\text{Prob}(\nu_\mu \to \nu_e) = \sin^2(\theta_{23})\sin^2(2\theta_{13})\sin^2(\frac{\Delta m^2_{23}L}{4E})+ \\ \left[\cos(\delta)\cos(\frac{\Delta m^2_{23}L}{4E}) - \sin(\delta)\sin(\frac{\Delta m^2_{23}L}{4E})\right] \times \\ \left[\cos(\theta_{13})\sin(2\theta_{12})\sin(2\theta_{13})\sin(2\theta_{23})\sin(\frac{\Delta m^2_{23}L}{4E})\sin(\frac{\Delta m^2_{12}L}{4E})\right]\end{aligned} \tag{2.25}$$

Considering the first term only, we see that $\sin^2(2\theta_{13})$ occurs multiplied by terms including θ_{23} and Δm^2_{23}. It is therefore extremely important to know these atmospheric parameters to high precision, as the precision of θ_{13} measurements will be limited by the uncertainties on the atmospheric parameters. We also note that terms including the CP violating phase δ only occur with $\sin(2\theta_{13})$, therefore if θ_{13} is zero, or extremely small, it would not be possible to measure δ. Observations of ν_e appearance and the measurements of θ_{13} made by T2K [35] and Daya Bay [37] have confirmed that θ_{13} is sufficiently large for measurements of δ to be possible.

Determining if δ is non-zero is one of the main priorities for future neutrino studies. Current cosmological models state that following the big bang, equal amounts of matter and antimatter would have been created. The excess of matter in the observable universe today conflicts with these models, unless a mechanism for creation of a matter-antimatter asymmetry can be discovered. CP violation in the lepton sector is one possible resolution of this asymmetry [41]. One method of searching for CP violation is comparing the rate of $\nu_\mu \rightarrow \nu_e$ oscillations with the rate of $\bar{\nu_\mu} \rightarrow \bar{\nu_e}$ oscillations, where any difference (after consideration of MSW effects) would imply occurrence of CP violation. A future $\bar{\nu_\mu}$ phase of T2K is planned in order to conduct such a study. The NOνA experiment [42], a ν_μ beam experiment similar to T2K but operating at higher energy over a longer baseline of 810 km, will also be sensitive to δ. Reactor $\bar{\nu_e}$ disappearance searches, such as Double CHOOZ [33] and Daya Bay [43] will also compliment these measurements.

That m_2 is greater than m_1 can be determined from measurements of solar neutrinos due to the influence of the MSW effect. The sign of Δm^2_{23} remains unknown however, since matter effects do not have a significant impact on the oscillations of atmospheric neutrinos or neutrinos which pass through the Earth's crust during low energy long baseline experiments. This results in two possible neutrino mass hierarchies, as demonstrated in Figure 2.3.

This figure also demonstrates that while the mass-squared differences can be calculated using oscillation measurements, the absolute masses of the neutrinos remain unknown. Upper limits on the neutrino masses can be calculated using cosmological data such as cosmic microwave background (CMB) observations [45] [46] and astrophysical observations such as the time distribution of neutrinos arriving from a supernova [47], and are known to be on the eV scale. The absolute mass of the ν_e could be constrained further by experiments involving beta decay. The energy distribution of electrons emitted in a given beta decay reaction would reveal the minimum energy carried away by the neutrino, corresponding to the neutrino rest mass energy. However, the small value of the rest mass makes this a very challenging measurement and no measurements of the precision required to lower the current mass limits have yet been achieved. Experiments such as the Karlsruhe Tritium Neutrino Experiment (KATRIN) aim to achieve a sensitivity

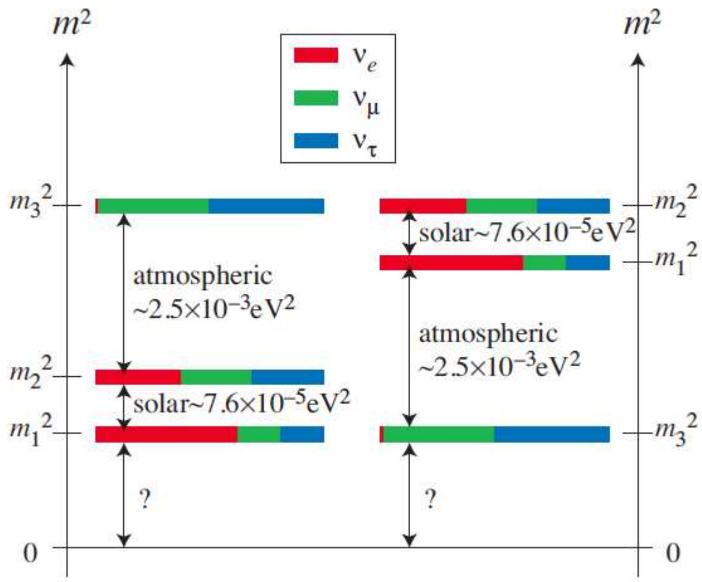

Figure 2.3: Demonstration of the two possible neutrino mass hierarchies, with normal (+ve Δm^2_{23}) on the left and inverted (-ve Δm^2_{23}) on the right. “Atmospheric” refers to Δm^2_{23} and “solar” refers to Δm^2_{12}. The flavour compositions of each mass state are shown, where the probability of measuring a ν_e is shown in red, a ν_μ in green and a ν_τ in blue. Figure from [44].

of order 0.1 eV [48].

The absolute masses could also be constrained by neutrinoless double-beta decay ($0\nu\beta\beta$-decay) searches. The Standard Model ν and $\bar{\nu}$ are distinct particles. However, that the neutrino is a neutral lepton introduces the possibility that neutrinos are Majorana fermions, as first theorised by Ettore Majorana in 1937 [49]. A massless Standard Model ν is always left-handed (LH), and a massless $\bar{\nu}$ is always right-handed (RH). However, the occurrence of oscillations and measurements of the mass-squared differences proves that for at least 2 neutrinos the neutrino mass $m_\nu \neq 0$. Therefore a small number of LH $\bar{\nu}$ and RH ν must also occur. Majorana ν and $\bar{\nu}$ would be identical in every way except for their spins. Any LH Majorana ν or $\bar{\nu}$ would participate in weak interactions as a ν, and any RH ν and $\bar{\nu}$ would behave as a $\bar{\nu}$ in weak interactions. Therefore if neutrinos are Majorana particles, should a neutrino (antineutrino) be created with RH(LH) spin, it would then be able to interact through the weak interaction as a $\bar{\nu}(\nu)$. This would not be possible with any other fermion as all other leptons and quarks also have electromagnetic charge, which must be conserved in interactions. If the neutrino is a Majorana particle, $0\nu\beta\beta$-decay would be able to proceed inside an atomic nucleus. The neutrino created in one β-decay process would be able to interact as a $\bar{\nu}$ in the second β-decay process, resulting in the emission of two electrons only. This process is shown in Figure 2.4.

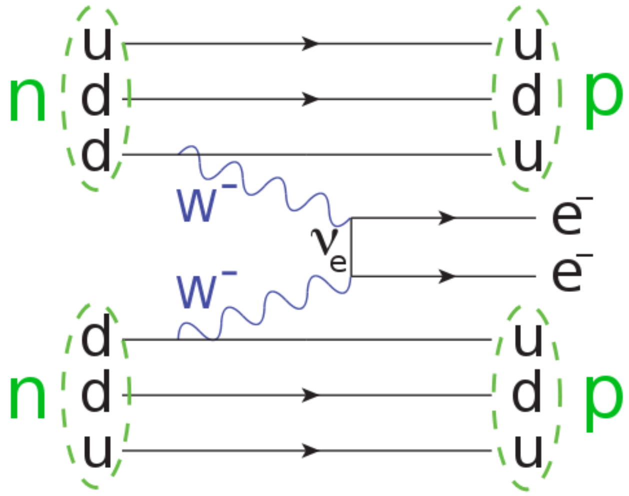

Figure 2.4: Feynman diagram demonstrating the neutrinoless double-beta decay process in the case that the neutrino is a Majorana fermion. Figure from [50].

A range of experiments are conducting $0\nu\beta\beta$-decay searches and more are under construction. Examples include MOON [51], which is currently taking data, and the planned upgrade to the Sudbury Neutrino Observatory (SNO+) [52]. While the process has an extremely low probability of occurrence, a positive result would allow the nature of the neutrino to be clarified. In addition, the $0\nu\beta\beta$-decay process is sensitive to the mass of the ν_e, therefore the detection of $0\nu\beta\beta$-decay would also allow the absolute mass of the ν_e to be constrained.

The alternative is that the neutrino is, like all quarks and other leptons, a Dirac particle, in which case the ν and $\bar{\nu}$ would be distinct particles. In this case the small number of LH $\bar{\nu}$ and RH ν would not be able to interact via the weak interaction, since only LH matter and RH antimatter can interact weakly. These neutrinos would therefore only experience gravity, and are described as 'sterile', since they cannot interact with any particles of charge, colour or flavour. It may be possible however for the three active neutrino flavours to oscillate to sterile neutrinos. Therefore it is possible to conduct sterile neutrino searches by looking for a deficit in total active neutrino flux measurements. Should such sterile neutrinos be discovered, extensions to the PMNS matrix would be required. In addition, other varieties of sterile neutrinos, including sterile neutrinos with considerable mass, also feature in many cosmological models and are potential dark matter candidates [53].

In addition to these remaining questions about neutrino properties, there are several experimental observations of neutrinos that require further investigation. One example is known as the LSND anomaly. The Liquid Scintillator Neutrino Detector (LSND) at Los Alamos National Laboratory used a tank of liquid scintillator to look for $\bar{\nu_e}$ in a beam of $\bar{\nu_\mu}$. An excess of $\bar{\nu_e}$ was observed, and while this can be explained by neutrino oscillations, the relatively large $\Delta \mathrm{m}^2_{LSND}$ required is not consistent with the 3-flavour oscillation parameters measured by other experiments [54]. Ongoing experiments such as MiniBooNE [55], ICARUS [56] and others aim to gather additional measurements which may help to resolve this.

Chapter 3

The T2K experiment

This chapter describes the specific physics goals of the T2K experiment, and how its design features are specially tailored to making the desired measurements. The full length of the baseline, from beam production to final detection is described, with the design and uses of each component given.

3.1 Aims of the experiment

The physics goals of the T2K experiment are as follows:

1. Search for the occurrence of ν_e appearance in a ν_μ beam using charged current interaction measurements. Obtain a high precision measurement of the oscillation parameter θ_{13}, with a factor of 20 improvement on the previous sensitivity [57]. Determine whether θ_{13} is non-zero and if so, sufficiently large for future measurements of the CP-violating phase δ to be possible.

2. Obtain the highest precision measurements of Δm^2_{23} and θ_{23} using observation of ν_μ disappearance. Use the measurements of charged current ν_μ interactions.

3. Measure a wide variety of interaction cross-sections using the near detector. Identification of interaction products allows specific interaction processes to

be measured individually, allowing exclusive interaction cross-section measurements.

4. Use measurements of neutral current interactions to search for oscillation to sterile neutrinos.

To study the extent of ν_μ disappearance and also search for the occurrence of ν_e appearance from a ν_μ source, an initially pure ν_μ beam is required. A high intensity neutrino beam provides extensive data for ν_μ disappearance measurements. Since θ_{13} was known to be small[1], a large flux of initial ν_μ also increases the likelihood of detecting ν_e, should appearance occur. High interaction rates resulting from a high-flux beam are also beneficial for cross-section measurements and sterile neutrino searches.

An efficient and reliable method of detecting neutrino interactions is required. It is also imperative that we can accurately distinguish between interactions of ν_μ and ν_e, so excellent particle identification capability is necessary. Super-KamiokaNDE [59] (Super-K), a neutrino observatory located under Mount Kamioka in Gifu Prefecture, Japan, is an established neutrino detector, ideal for this task. It is therefore utilised as the experiment's far detector. Super-K is described in detail in §3.4.

The Japan Proton Accelerator Research Complex (J-PARC), located in Tokai, Ibaraki Prefecture, Japan, is home to a high intensity proton accelerator, which can be used to generate a high intensity neutrino beam. The positions of Super-K and J-PARC are shown on the map of Japan presented in Figure 3.1.

J-PARC is located 295 km from Super-K, which is an appropriate baseline for the neutrino oscillations of interest. It is therefore used to generate the high intensity ν_μ beam, through a process described in more detail in §3.2. The probability

[1] The CHOOZ limit set $\sin^2 2\theta_{13} < 0.10$ [58]

Figure 3.1: Position of J-PARC facility and Super-K detector on Japanese island Honshu with baseline shown. Figure from [60].

of ν_e appearance is dependent on the energy of the neutrinos, E_ν, and the distance over which the neutrino beam has been allowed to propagate before observation, L, as well as the relevant oscillation parameters. To increase the likelihood of observing $\nu_\mu \rightarrow \nu_e$ oscillation, we must maximise the probability $\mathrm{P}(\nu_\mu \rightarrow \nu_e)$. The baseline is fixed at L=295 km, but the peak neutrino energy E is chosen such that $\mathrm{P}(\nu_\mu \rightarrow \nu_e)$ is at a maximum for L/E.

Considering 3-flavour oscillations with parameters $\sin^2(2\theta_{12}) = 0.8704$, $\sin^2(2\theta_{23}) = 1.0$, $\Delta m^2_{12} = 7.6\times10^{-5}$ eV^2, $\Delta m^2_{23} = 2.4\times10^{-3}$ eV^2, average matter density = 3.2 gcm^{-3}, and assuming $\sin^2(2\theta_{13}) = 0.1$ (the CHOOZ limit[58]), the probability of ν_μ survival at a baseline of 295 km is shown for a range of neutrino energies in Figure 3.2.

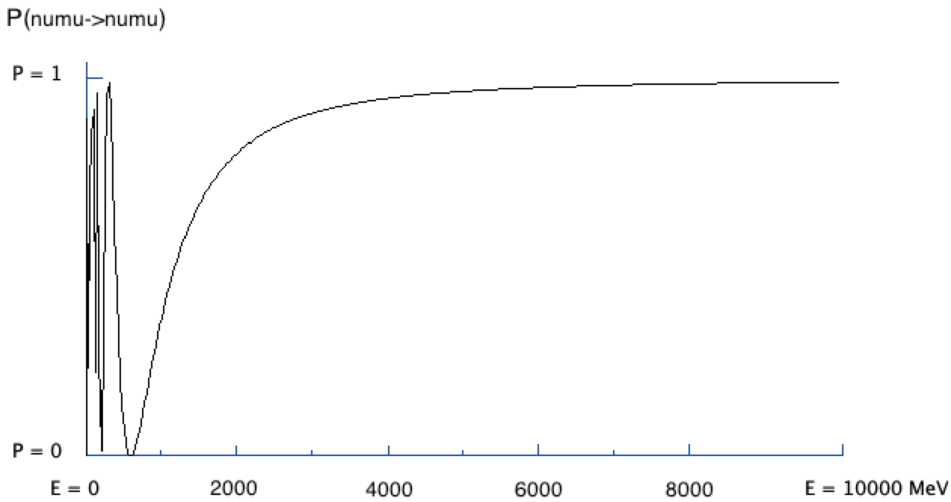

Figure 3.2: Probability of ν_μ survival at a baseline of L = 295 km for range of neutrino energies. Uses $\sin^2 2\theta_{12} = 0.8704$, $\sin^2 2\theta_{23} = 1.0$, $\Delta m^2_{12} = 7.6\times10^{-5}$ eV^2, $\Delta m^2_{23} = 2.4\times10^{-3}$ eV^2, average matter density = 3.2 gcm^{-3}, and $\sin^2 2\theta_{13} = 0.1$. Figure created using [61].

The probability drops to a minimum at numerous points, with the most pronounced dip occurring at neutrino energies of 600 MeV. ν_τ appearance is responsible for the majority of this reduction in survival probability, as demonstrated by

Figure 3.3.

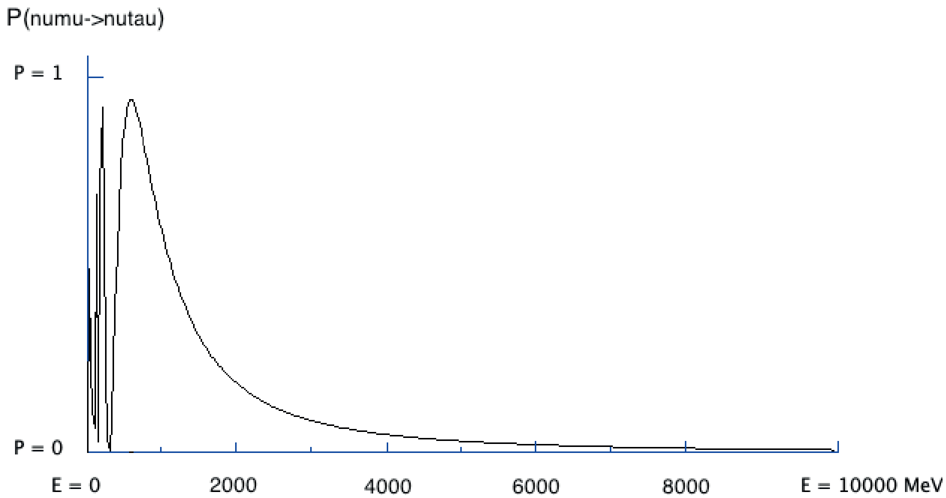

Figure 3.3: Probability of ν_μ to ν_τ oscillation at a baseline of 295 km for range of neutrino energies. Uses $\sin^2 2\theta_{12} = 0.8704$, $\sin^2 2\theta_{23} = 1.0$, $\Delta m^2_{12} = 7.6\times10^{-5}$ eV^2, $\Delta m^2_{23} = 2.4\times10^{-3}$ eV^2, average matter density = 3.2 gcm^{-3}, and $\sin^2 2\theta_{13} = 0.1$. Figure created using [61].

The appearance of ν_e also make a contribution to the low ν_μ survival probability at approximately E_ν = 600 MeV. Figure 3.4 shows that the probability of ν_μ to ν_e oscillation is known to be very low at all energies, but reaches a maximum at approximately 600 MeV, which corresponds to the ν_μ survival minimum. Using a beam of ν_μ with a peak energy of 600 MeV therefore maximises the likelihood of observing ν_e appearance at the far detector.

The ν_μ beam is generated using the collision of protons on carbon, which is described in detail in §3.2. A 30 GeV proton beam collides with a stationary graphite target to produce an assortment of hadrons. π^+ which will decay to produce neutrinos with energies peaking around 600 MeV at Super-K are selected.

The neutrino beam flux distribution is affected by the position of the detectors with respect to the beam axis. Commonly detectors are placed at the centre of a particle beam in order to be exposed to the highest flux of particles possible.

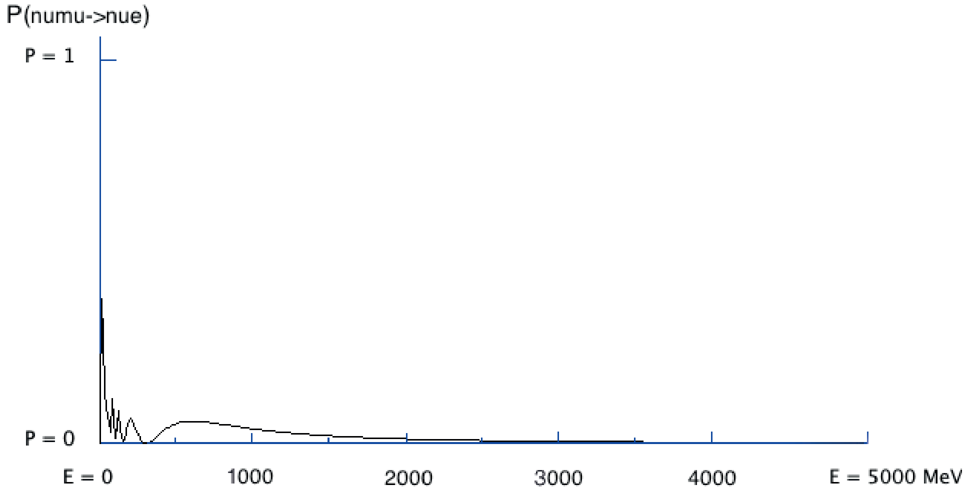

Figure 3.4: Probability of ν_μ to ν_e oscillation at a baseline of 295 km for range of neutrino energies. Uses $\sin^2 2\theta_{12} = 0.8704$, $\sin^2 2\theta_{23} = 1.0$, $\Delta m^2_{12} = 7.6\times10^{-5}$ eV^2, $\Delta m^2_{23} = 2.4\times10^{-3}$ eV^2, average matter density = 3.2 gcm^{-3}, and $\sin^2 2\theta_{13} = 0.1$. Figure created using [61].

However, the T2K beam and near detector complex have been arranged such that one of the near detectors, ND280, and the far detector, Super-K, are positioned 2.5° off-axis. This has several beneficial effects. Due to the decay kinematics of pions, the energy of the peak neutrino flux observed by a detector decreases as the detector moves further off-axis. The peak width also decreases, so that a higher proportion of the neutrinos have the peak energy. The flux for a range of off-axis angles is shown in Figure 3.5, with the peak of the flux at 2.5 degrees off-axis corresponding to the energy of the ν_μ survival minimum and the ν_e appearance maximum. This off-axis angle behaviour is discussed in detail in Chapter 7 and Appendix A.

There are two sources of background to the ν_e appearance analysis: intrinsic ν_e which contaminate the ν_μ beam, and neutral current interactions which can mimic ν_e interactions at the far detector. As discussed throughout this thesis, high energy neutrinos are responsible for the majority of the neutral current background events. Therefore reducing the flux of neutrinos in the high energy tail is desirable. Using an off-axis configuration of detectors allows us to reduce the flux of high energy neutrinos, while maintaining a high flux of neutrinos with energies around 600 MeV. This results in a high flux of neutrinos with maximised P($\nu_\mu \rightarrow \nu_e$), while also minimising the sources of background to the ν_e appearance analysis, which allows a high sensitivity to ν_e appearance.

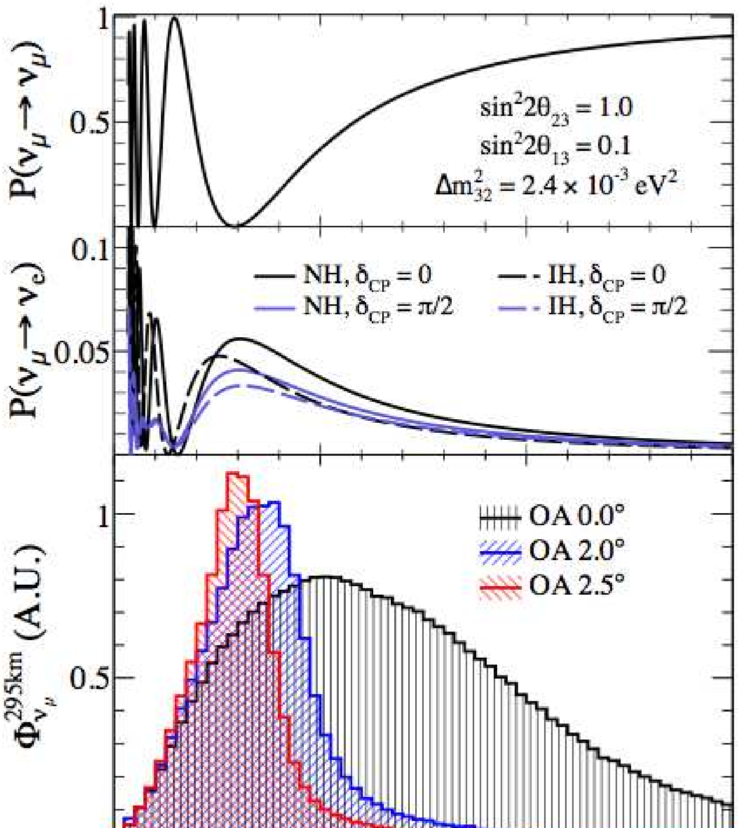

Figure 3.5: The muon neutrino survival probability (top), electron neutrino appearance probabilities (middle) at 295 km, and the unoscillated neutrino fluxes for different values of the off-axis angle (OA) (bottom). The appearance probability is shown for two values of the phase δ_{CP}, and for normal (NH) and inverted (IH) mass hierarchies. Figure from [35].

The arrangement of the main components of the T2K experiment is shown in Figure 3.6. A ν_μ beam, of peak energy 600 MeV, is generated at J-PARC.

Beam production is described in §3.2. A suite of near detectors is located at J-PARC, 280 m downstream from the beam production point. These consist of the Interactive Neutrino GRID (INGRID), which is positioned on-axis, and ND280, which is positioned 2.5° off-axis. The components of these detectors and their functions are described in §3.3. The ν_μ beam then propagates 295 km across the island of Honshu, and is observed again at Super-K, which is also located 2.5° off-axis. A description of the Super-K detector and the particle identification process used at the far detector are described in §3.4.

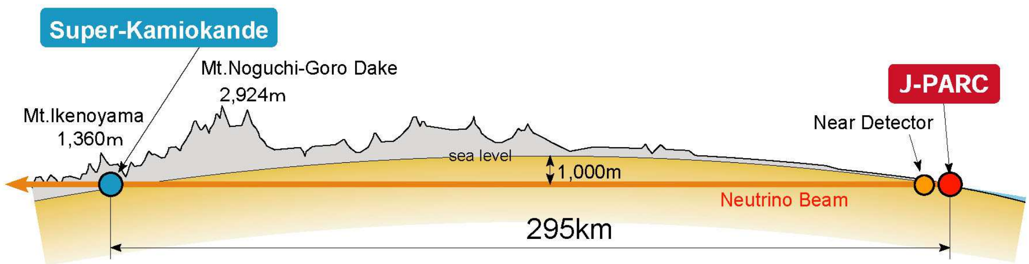

Figure 3.6: Passage of neutrino beam from point of beam production to near detector (ND280), then through the Earth's crust to the far detector (Super-K), located under Mt. Ikenoyama. Figure from [60].

3.2 Beam production

Three accelerators used in sequence generate the 30 GeV proton beam used by T2K. Their arrangement is shown in Figure 3.7. First a linear accelerator (LINAC) accelerates a beam of H^- up to 400 MeV. The H^- then pass through a charge-stripping foil, which removes the electrons to leave H^+, as they enter the Rapid Cyclic Synchrotron (RCS). The protons are accelerated up to 3 GeV by the RCS. About 5% of the bunches accelerated in the RCS are then supplied to the Main Ring (MR), where they are accelerated further to reach 30 GeV. There are 8 bunches of protons in the MR (increased from 6 before June 2010). Protons are accessed at two extraction points, with the fast extraction supplying proton bunches to the T2K neutrino beamline.

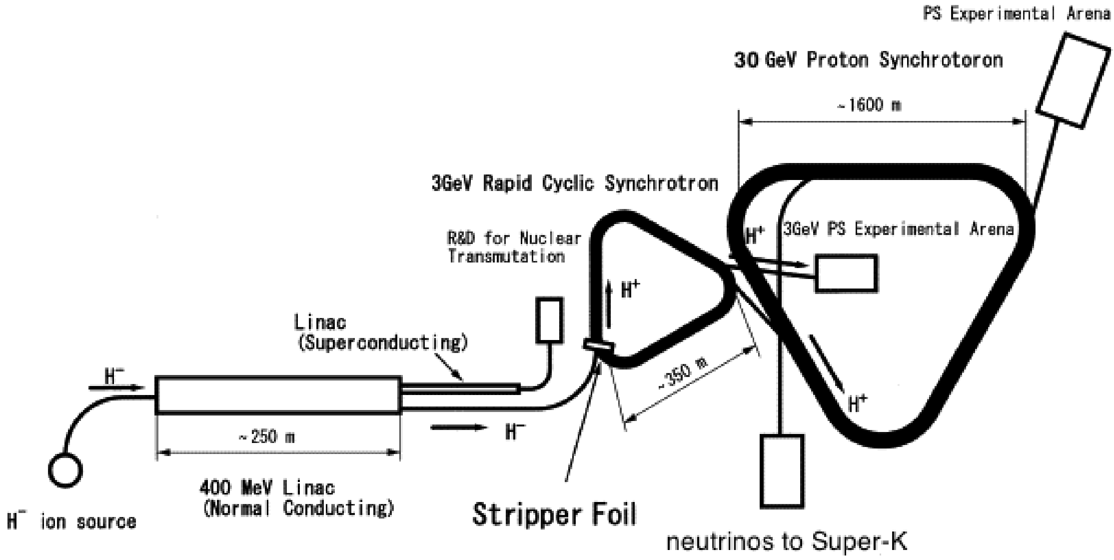

Figure 3.7: Production of protons from a H^- source using J-PARC accelerators. H^- ions pass through the LINAC, then are injected into the RCS via a stripper foil which removes the electrons to produce H^+, and finally are directed into the proton synchrotron where they are accelerated up to energies of 30 GeV. Figure from [62].

After fast extraction the proton bunches are prepared in the primary beamline for collision with the target. They pass through a 54 m preparation section, in which the beam is tuned by a series of 11 conducting magnets, followed by a 147 m arc section, in which 14 doublets of superconducting combined-function magnets alter the proton beam direction to point towards Kamioka. Finally the beam traverses the 37 m focussing section, which consists of ten normal conducting magnets which guide and focus the beam onto the target and make a final small adjustment to the beam direction so that it points down at 3.637° to the tangent to the Earth's surface. At this point they enter the secondary beam line, which is shown in Figure 3.8.

The secondary beamline consists of the target station, decay volume and beam dump. Upon entering the target station the proton beam first passes through a

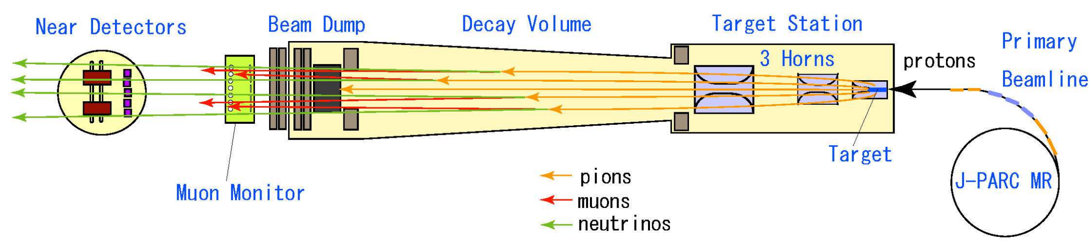

Figure 3.8: After proton bunches are extracted from the main ring (see Figure 3.7) they enter the primary neutrino beamline, in which protons are prepared and then steered to point towards the detectors. After final focussing they then enter the secondary beamline, containing the target, decay volume and beam dump. Figure from [63].

30 mm hole in a graphite baffle, which collimates the beam. After monitoring by the Optical Transition Radiation monitor (OTR) [64], the proton beam collides with the stationary target, which is in the form of a 91.4 cm long 1.8 g/cm^3 graphite rod. Helium gas flowing between the graphite core and its layers of casing cool the target.

Collision with the target produces an array of secondary hadrons, which include π^+, π^-, K^+, K^- and K^0. T2K's primary source of ν_μ is secondary pions, via the decay

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \tag{3.1}$$

Using π^+ decay to produce neutrinos results in little contamination from other decay products since the π^+ decay as in Eq. 3.1 with a branching ratio of 99.99%. Muons with momentum lower than 5 GeV/c can be removed, leaving a very pure ν_μ beam.

Immediately following production at the target, the secondary hadrons encounter the first of three magnetic horns. Each horn contains a toroidal magnetic field, where the field varies as the inverse of the distance from the horn axis. The first horn is used to collect π^+. While only π^+ are desired, relatively small numbers of π^- and kaons are also collected, and will contribute additional decay products to the beam. The majority of these decays also result in the production of ν_μ or $\overline{\nu}_\mu$; however, a small number result in the production of ν_e. As these cannot be removed, they contaminate the otherwise pure $\nu_\mu/\overline{\nu}_\mu$ beam, and act as a background to the ν_e appearance search. The decays contributing to this beam contamination are discussed in §4.1. After selection, two additional horns then focus the pions in order to increase the flux arriving at Super-K.

After focussing, the π^+ (and smaller quantities of additional hadrons) enter the decay volume. The decay volume length is carefully chosen with consideration of the pion and muon lifetimes, and set at 96 m [65]. A long decay volume would allow all pions to decay, which would produce the highest flux of ν_μ. However, the other product of π^+ decay, μ^+, will also decay with a lifetime of 2.2×10^{-6} s, via

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu_\mu} \tag{3.2}$$

Muon decay is another source of ν_e contamination in the beam, therefore minimising the occurrence of μ^+ decay is required. A short decay volume length would remove more μ^+ before they have the opportunity to decay. A length of 96 m is therefore chosen as this achieves a balance between high π^+ decay and low subsequent decay of μ^+. Located at the end of the decay volume is a beam dump, consisting of 75 tons of graphite and iron plates with a total thickness of 2.4 m. This removes any remaining hadrons and also μ^+ with momenta under approximately 5 GeV/c. This leaves a beam of neutrinos which continue on to the detectors.

μ^+ with sufficient energy to pass through the beam dump also exit the decay volume, and their distribution is profiled by the muon monitor. This is comprised of ionization chambers and silicon PIN photodiodes. Further details are provided by [65]. Measurements of these high energy μ^+ are used to monitor the neutrino beam direction and intensity.

3.3 The near detectors

The near detector complex is located 280 m downstream from the target station and contains two separate detectors: INGRID, which is centred on the neutrino beam axis, and ND280, which is positioned on the line connecting the target and Super-K, at 2.5° to the beam axis. The near detector positions with respect to the target and far detector are demonstrated by Figure 3.9.

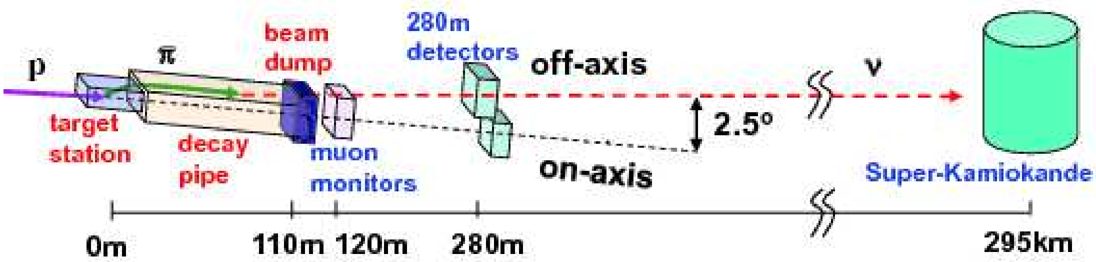

Figure 3.9: The T2K beamline, showing the positions of the near detectors and muon monitor in relation to the target and decay volume, and the far detector. The neutrino beam is directed down below the Earth's surface, with the beam centre passing below Super-K. Super-K and ND280 are located at 2.5° above the neutrino beam axis. Figure from [66].

The neutrino beam is directed below the Earth's surface, such that the far detector is positioned at 2.5° above the beam axis. The near detectors, located at 280 m from the target station, are stacked, with INGRID occupying the lower position, and ND280 above. The arrangement of the near detectors is shown in Figure 3.10.

It is the function of these detectors to measure the energy spectrum, interaction rates, and flavour content of the neutrino beam before the neutrinos can oscillate. These measurements are then used to predict the neutrino interactions that would

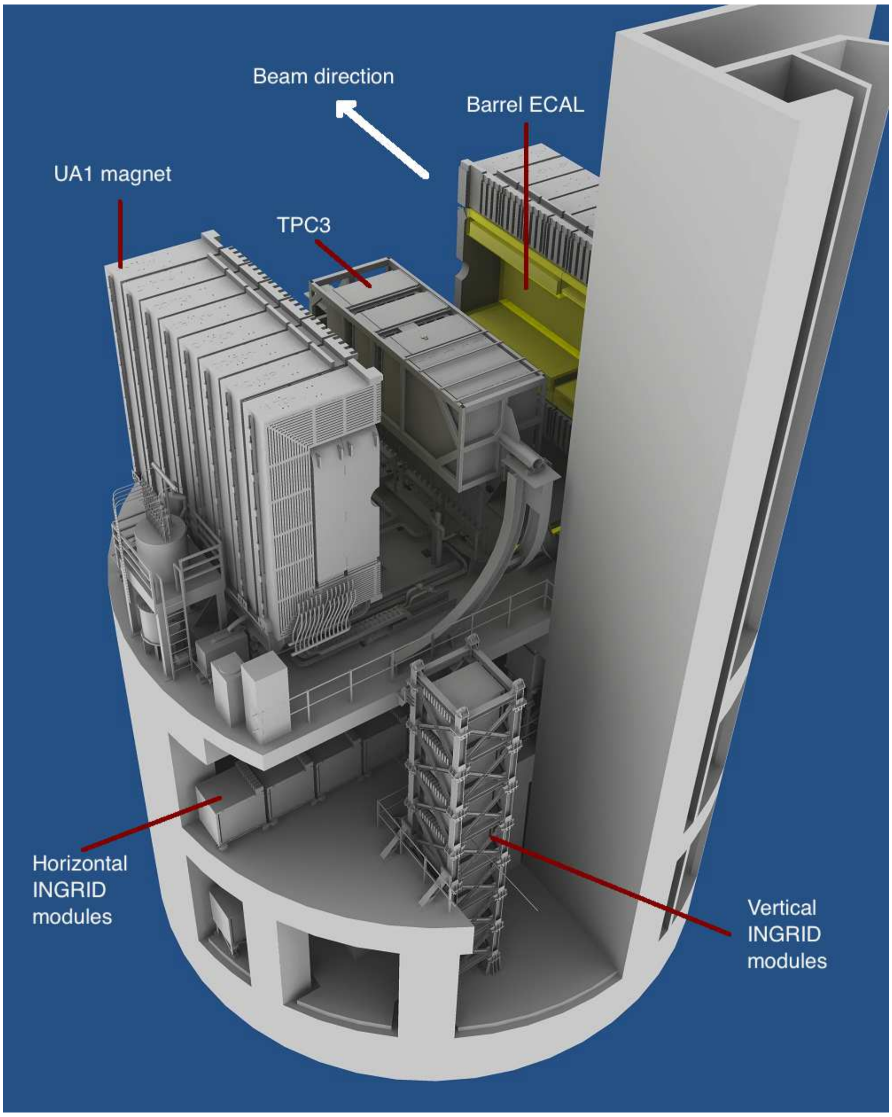

Figure 3.10: The arrangement of the near detectors, shown with the UA1 magnet open and basket and tracker components of ND280 exposed. INGRID occupies the lower levels of the pit and is centred on the beam axis. ND280 is located above INGRID, which at 280 m from the target station corresponds to 2.5° off-axis. Figure adapted from [65].

be seen at the far detector in the absence of oscillations. The far detector spectrum and rate of ν_e interactions that would be observed at Super-K when the oscillation parameters take particular values can also be modelled, and these simulations can be compared to find the combination of oscillation parameters that provide the best agreement with data.

3.3.1 INGRID

The Interactive Neutrino GRID (INGRID) is used to measure the neutrino beam direction and flux [67]. Figure 3.5 demonstrates the changes in the flux spectrum with changing position with respect to the beam axis. In order to predict the neutrino flux at the far detector accurately we must know the precise off-axis angle encountered by Super-K, and therefore must know the beam direction to a high degree of accuracy. Figure 3.11 shows the arrangement of the INGRID modules. 14 identical modules are arranged in a cross formation, with the two central modules overlapping. The beam centre, corresponding to 0° from the beam axis, passes through these two central modules. Two additional modules are located above the horizontal row of modules, at off-axis positions. Measurements of the neutrino interaction rate in each of these 16 modules allow the position of the beam centre to be calculated to within 10 cm, which is equivalent to an accuracy of 0.4 mrad in off-axis angle at 280 m.

Each module measures 124 cm × 124 cm in the xy plane, with a depth of 6.5 cm. INGRID can therefore sample the neutrino beam in a cross-section of approximately 10 m × 10 m. The modules each consist of alternating planes of iron and tracking scintillator, with 9 iron planes arranged between 11 scintillator planes. Additional veto scintillator planes cover the four faces of each module that sit parallel to the beam direction. 7.1 tons of iron per module provides the target mass for neutrino interactions. Muons produced in neutrino interactions within the iron planes are detected as they pass through subsequent scintillator planes, allowing their passage through the module to be tracked. Each tracking plane contains 24 horizontal and 24 vertical scintillator bars, which consist of polystyrene doped with 1% PPO and 0.03% POPOP by weight. Light is collected by 1 mm diameter wavelength shifting fibres (WLS), which run through the centre of each scintillator bar and are connected to Multi-Pixel Photon Counters (MPPCs) [65].

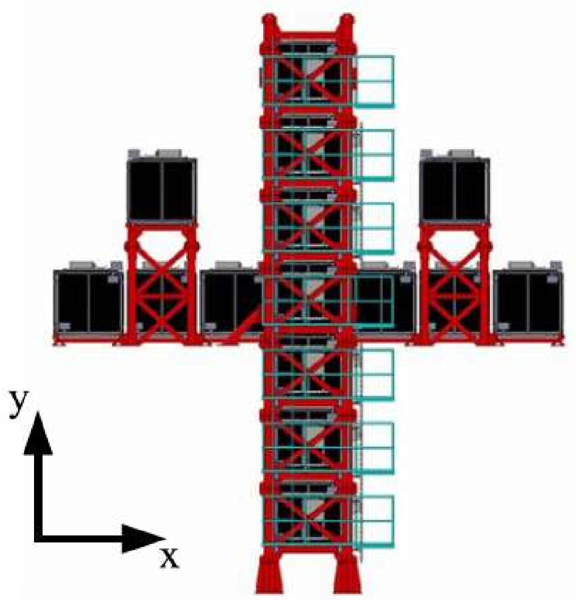

Figure 3.11: The on-axis detector, INGRID. Seven horizontal modules and seven vertical modules are arranged in a cross formation, centred on the neutrino beam. Two additional modules are placed off-axis above the horizontal row. The beam axis passes through the central modules, and the off-axis angle increases with distance from these central modules. Figure from [65].

MPPCs are used extensively by the near detectors to read out light signals generated in the scintillator regions. While use of multi-anode photomultiplier tubes (PMTs) is an effective method of light detection for experiments using scintillator and WLS fibre, and one used by many neutrino experiments, multi-anode PMTs are not suitable for use in magnetic fields. Due to the magnetic field applied across ND280, PMTs would have to be located at a distance from the detectors, which would complicate the design and calibration of ND280. MPPCs however are insensitive to magnetic fields, and are conveniently compact, making them well-suited to use throughout the near detectors for collection of signals from the various detector regions. Each MPPC holds numerous independent pixels, where each pixel functions as a Geiger micro-counter. Customized 667-pixel MPPCs were developed for T2K and manufactured by Hamamatsu. More information on the testing of these devices and their operation can be found in [68], [69] and [70].

A further "proton" module consisting only of scintillator planes is positioned at the beam centre between the central modules of the horizontal and vertical arms. The 16 standard modules are only sensitive to muons, since other products of interactions such as pions and protons will stop in the iron layer in which they were produced and not reach a scintillator bar. The Proton Module [71] is designed to detect additional interaction products and allow the identification and study of different interaction modes occurring.

3.3.2 ND280

The off-axis near detector, referred to as ND280, has several functions. Firstly, it allows the measurement of the neutrino interaction rate and energy spectrum. The neutrino flux can be calculated from this using knowledge of interaction cross-sections in ND280. Using this measurement, knowledge of the interaction cross-section in water, the drop in flux intensity over the T2K baseline and the fiducial volume of the far detector, the expected flux and energy spectrum at the far detector can be predicted.

ND280 is also used to measure the flavour composition of the neutrino beam prior to any oscillations. GPS signals sent from the beamline allow beam spills arriving at ND280 to be identified. With the development of an effective particle identification system, interactions of ν_μ and ν_e can be separated, and thus the intrinsic ν_e contamination in the beam, which forms a background to the ν_e appearance signal, can be quantified.

ND280 is designed to be able to identify a range of particles which are commonly created as neutrino interaction products. The result is that in addition to measuring inclusive interaction rates, specific interaction types can be identified and measured individually, allowing exclusive cross-section measurements to be conducted. Charged current quasi-elastic (CCQE) and charged current non-quasi-elastic (CCnonQE) interactions are of interest for the ν_e appearance and ν_μ disappearance measurements. Neutral current (NC) interaction rates must also be studied, since NC interactions are a source of background events at the far detector. Neutral current interactions resulting in the production of a single π^0 are particularly important to understand since they are the second significant background to the ν_e appearance signal at Super-K.

The off-axis detector consists of several different sections which combined allow the necessary measurements described above to be conducted. These components can be classified into 2 types: scintillator regions with MPPC light detection, and gaseous tracking volumes. Each of the sub-detectors is visible in Figure 3.12, and described in more detail below.

UA1 Magnet

The direction and degree of curvature of a charged particle's trajectory in a magnetic field is used to determine the sign and momentum of the particle respectively. In ND280 a magnetic field of 0.2 T is provided across the detectors by the recycled CERN UA1/NOMAD magnet, one half of which is shown in red in Figure 3.12.

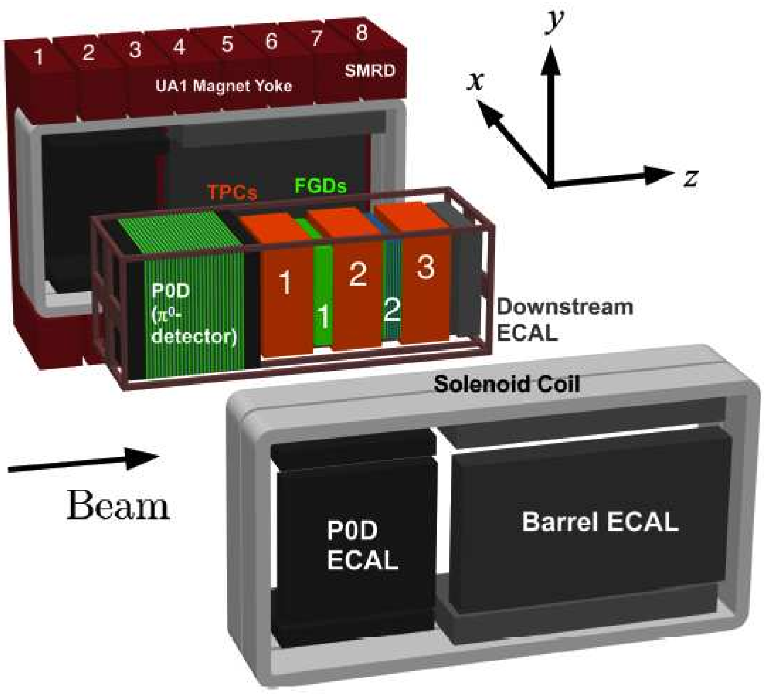

Figure 3.12: Exploded diagram of ND280 showing each component part contained within the UA1 magnet. The central basket region contains the π^0 Detector (P0D), followed by alternating Time Projection Chambers (TPCs) and Fine Grained Detectors (FGDs), which form the tracker. These components are enclosed by several separate Electromagnetic Calorimeter (ECal) units. Side Muon Range Detector (SMRD) modules are located within the magnet yoke structure (half of which is shown). Yoke elements are labelled. The z axis is defined as the beam direction, passing from the P0D to the downstream ECal. Figure adapted from [65].

The magnet consists of two halves, which are positioned together for data taking but can be separated to allow access to the various Electromagnetic Calorimeter (ECal) modules and the detectors contained within the basket. Each half comprises eight C-shaped elements constructed from steel plates, which form the magnet yoke. These 8 pairs of yoke elements are labelled from 1 at the upstream end to 8 at the downstream end, as shown. Sitting inside the steel plate structure are aluminium coils, with two elements per magnet half. The magnet is cooled by a flow of water passing through the coils.

A nominal direct current of 2900 A is directed through the magnet. The generated magnetic field was precisely measured during a B-field mapping procedure conducted in 2009. Using a system of three Hall probes, each with an intrinsic uncertainty of 0.2 G, the B-field throughout the detector region was measured, and once scaled to the nominal B-field value, is known to within 2 G (2×10^{-4} T). This high level of precision reduces the systematic uncertainty on particle momentum calculations, with a resulting uncertainty of below 2% for charged particles with momenta below 1 GeV/c [65].

π^0 Detector (P0D)

Neutral pions produced in neutral current interactions in the far detector form a significant background to the ν_e appearance measurement (see §3.4). To predict the rate of π^0-producing NC interactions at Super-K, we must know the neutral current interaction cross-sections on water. The P0D is designed to measure the interaction rate of the process

$$\nu_\mu + N \rightarrow \nu_\mu + N + \pi^0 + X \tag{3.3}$$

on water, where the shape of the neutrino spectrum measured by the P0D matches the spectrum observed by the far detector. To achieve this, alternating panels of water and scintillator are used [72].

The P0D is divided into four sections. The central sections are the upstream water target and central water target, and these consist of repeating layers of scintillator planes, brass sheets and water bags. The P0D can operate with the water target bags either filled with water, or empty. Charged particle tracks left by muons and charged pions, and electromagnetic showers generated by electrons and π^0 decay photons, can be reconstructed using signals from the scintillator planes. The upstream ECal, located at the most upstream end of the P0D, and central ECal, located at the downstream end, consist of alternating planes of scintillator and lead. These provide a veto region before and after the water target so that products of interactions occurring in other areas can be identified and rejected from water cross-section studies. The central ECal also acts to contain electromagnetic showers occurring in the P0D and prevent their detection in the adjacent TPC1.

Cross-section studies require an accurate value of the quantity of water forming the fiducial volume. This is determined by filling the water bags to fixed levels and then measuring the volume of water that is removed. The total mass of the P0D is 16.1 tons with water and 13.3 tons without, and the measured mass of the fiducial volume of the water target is 1902 ± 16 kg.

Time Projection Chambers (TPC)

The tracker section of ND280, consisting of three TPCs and two FGDs, is located downstream of the P0D, with TPC1 adjacent to the central ECal of the P0D. The near detector must be able to identify a variety of interaction types and this requires the ability to track multiple interaction products and identify them. The

TPCs provide excellent spatial resolution, which allows multiple products of an interaction to be counted, and the trajectories of charged particles to be observed as they pass through the detector and experience the applied magnetic field [73]. Information gathered about the number of particles produced in an interaction, and their types, is then used when selecting samples of specific interaction types. The momenta of charged products can also be calculated using the observed curvature in the magnetic field. This allows interaction rates to be studied as a function of the energy of the interacting neutrino.

It is necessary to be able to distinguish between different particles of the same charge in order to correctly identify all interaction products. This is especially important for measuring the intrinsic ν_e contamination in the beam. Charged current interactions of ν_μ and ν_e will result in the production of μ^- and e^- respectively, which will curve in the same manner in the applied magnetic field. The degree of ionization left by a charged particle, when combined with momentum measurements, can be used to identify the particle flavour. The TPCs are therefore designed to measure this.

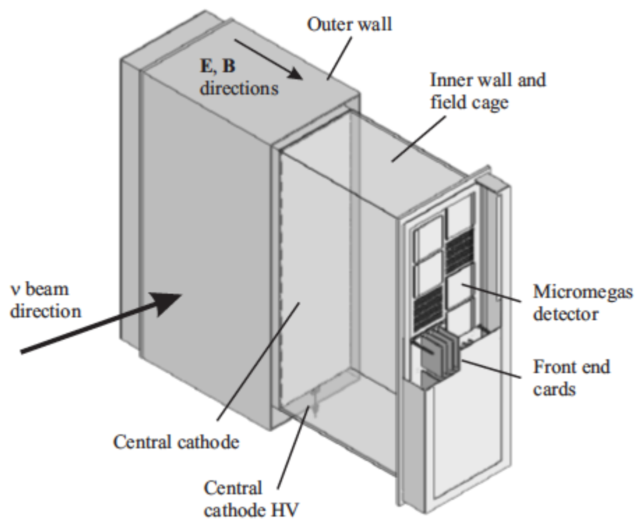

Figure 3.13: TPC structure, showing the direction of the electric drift field, the cathode position and the location of the readout planes. Figure from [73].

Each TPC consists of a chamber filled with a gas mixture (95% argon), with a central cathode panel running through the yz plane. This cathode, along with copper strips positioned on the chamber walls, produces a uniform electric drift field, aligned in the same direction as the magnetic field to run horizontally across the TPC, as shown in Figure 3.13. As a charged particle passes through a TPC, ionisation electrons are produced in the gas, which due to the electric field drift horizontally away from the central cathode towards a readout plane. On reaching one of these planes, the electrons are multiplied and generate signals. Segmentation of the readout planes allows the position of generated drift electrons to be recorded. The arrival times of drift electrons, combined with the position information, are used to calculate 3D track trajectories of any charged particles passing through the TPCs. Further TPC design and performance information can be found in [74].

Fine Grained Detectors (FGD)

Neutrino interaction cross-sections are very small, so a large target mass is required to maximise the number of interactions available for study. The TPCs provide excellent resolution for tracking of charged products, but an interaction rate of effectively zero due to the low density of their gas content. Target mass is provided by the FGDs, which combine higher density material to provide a target for interactions, with high precision tracking capability [75].

The two FGDs are placed between the three TPCs, and are labelled as shown in Figure 3.12. This positioning of the FGDs and TPCs provides information of use when forming event selections. Particles produced in interactions in either FGD can be tracked, and activity in TPC1 can be used to veto interactions not originating in an FGD, or interactions producing products that travel upstream.

Both FGDs contain layers of polystyrene scintillator bars, arranged either horizontally or vertically in the x-y plane. Each bar has a TiO_2 reflective coating and a WLS fibre running through the centre. FGD1 is constructed of scintillator only, with 30 layers of 192 bars. The orientation of the bars alternates from layer to layer, with one set of horizontal and vertical bars defined as an XY module. The scintillator both provides the target mass for interactions and tracks the interaction products. The dense arrangement of narrow scintillator bars provides detailed resolution from the point of interaction. FGD2 consists of a mixture of scintillator layers and water layers. Seven XY modules alternate with six 2.5 cm thick water modules. The presence of this second water target provides additional data for interaction cross-section studies and comparison of interaction rates in the two FGDs allows comparison of the carbon and water cross-sections.

Electromagnetic Calorimeters (P0D ECal, Barrel ECal, Ds-ECal)

The inner detectors are surrounded on five sides by the ND280 sampling electromagnetic calorimeter (ECal) [76]. There are 13 individual ECal modules, of three different types, arranged as shown in Figure 3.12. The six P0D ECal modules surround the P0D on four sides, six Barrel ECal modules surround four sides of the tracker detectors, and finally the Downstream ECal (Ds-ECal) covers the downstream face of TPC3. The majority of forward-going interaction products that are not contained by the inner detectors must therefore pass through an ECal module as they exit the detector.

The ECal modules are constructed using lead sheets between layers of plastic scintillator bars. The ECals aid the identification of charged particles created in interactions. They also measure the energy and direction of photons exiting the detector. These measurements are necessary to fully reconstruct events. In addition to detecting charged particles, the ECals are also used to detect gammas from π^0 produced in interactions in the inner detectors, and measure their energies.

Each ECal scintillator bar measures 4.0 cm $\times$ 1.0 cm, with a 2 mm diameter hole down the centre to contain a WLS fibre. The scintillator bars are polystyrene doped with 1% PPO and 0.03% POPOP. To prevent accidental signals caused by external light, each complete ECal module is sealed with thin aluminium covers for light tightness.

The Ds-ECal modules and Barrel ECal modules have the same internal structure. Bars of scintillator are arranged in layers with lead sheets of thickness 1.75 mm. The orientation of the scintillator bars in each layer is perpendicular to the orientation in the adjacent layers to allow fine spatial reconstruction of tracks. The Ds-ECal module contains 34 layers, each containing 50 scintillator bars. Each Barrel ECal module contains 31 layers.

The P0D is designed to detect and reconstruct π^0 events, and therefore it is not necessary for the P0D ECal to replicate the functionality of the ECal modules around the tracking region. Since a high degree of spatial resolution is not required, the structure of the P0D-ECal is simpler, with all scintillator bars aligned parallel to the z axis. This arrangement of scintillator is sufficient to detect photons and showers that are not contained within the volume of the P0D. The P0D ECal can also act as a veto for products of interactions occurring outside the inner detector, and confirm the presence of charged particle tracks.

SMRD - Side Muon Range Detector

The iron magnet yoke is a large, dense target mass and therefore a large number of neutrino interactions occur in its volume. Products of interactions occurring here and in the surrounding structures may travel towards the inner detectors and can be identified to prevent inclusion with interactions occurring in the ND280 target regions. Cosmic ray muons arriving at the detector can also be identified. Adding a small amount of instrumentation within the magnet yoke in order to measure particles produced in the magnet volume meets these requirements, and also allows detection of muons emitted at very large angles to the beam direction as they exit the detector region.

The steel plates that form the elements of the magnet yoke are positioned with spaces of 1.7 cm between each plate. A total of 440 scintillator modules are placed in the innermost gaps, and together form the SMRD [77]. The sides of the pairs of yoke elements numbered 1 to 5 inclusive (see Figure 3.12) contain 3 layers of modules, the 6th pair of yoke elements contain 4 layers, and pairs of yoke elements 7 and 8 contain 6 layers. All yoke pairs contain 3 modules in the top and bottom sides. The modules are constructed from scintillation counters of various dimensions which are designed to maximise the active area available between the steel plates. These scintillators consist of extruded polystyrene and dimethylacetamide with POPOP and para-terphenyl. A 1 mm diameter WLS fibre is glued into a groove in the surface of each scintillator for signal collection.

3.4 The far detector - Super Kamiokande

The role of the far detector is to measure the flavour of the neutrino beam at 295 km. By detecting interactions of ν_μ and ν_e on water, the far detector searches for the appearance of ν_e in the beam, and also measures the degree of ν_μ disappearance occurring at L = 295 km.

T2K uses an existing water Cherenkov neutrino detector located 1 km beneath the peak of Mt. Ikenoyama. Super-Kamiokande consists of a cylindrical tank of diameter 33.8 m and height 36.2 m, containing 50 kton of ultra-pure water. The inside wall of this tank is lined with 11,129 50 cm PMTs, which face into the tank and provide 40% coverage of the inner wall's surface. This volume is the inner detector (ID). Surrounding the inner detector is an additional 2 m thick layer of water, optically separated from the inner detector region. 1,885 20 cm PMTs line the inner wall of this region, facing out from the inner detector. This space forms the outer detector (OD). A 50 cm dead space separates the inner and outer detector regions. This contains plastic sheets which prevent photons from passing from one detector region to the other, and also the mounting structure and cabling for the PMTs. The inner and outer detectors and dead space combine to give a cylinder measuring 39 m in diameter and 42 m in height.

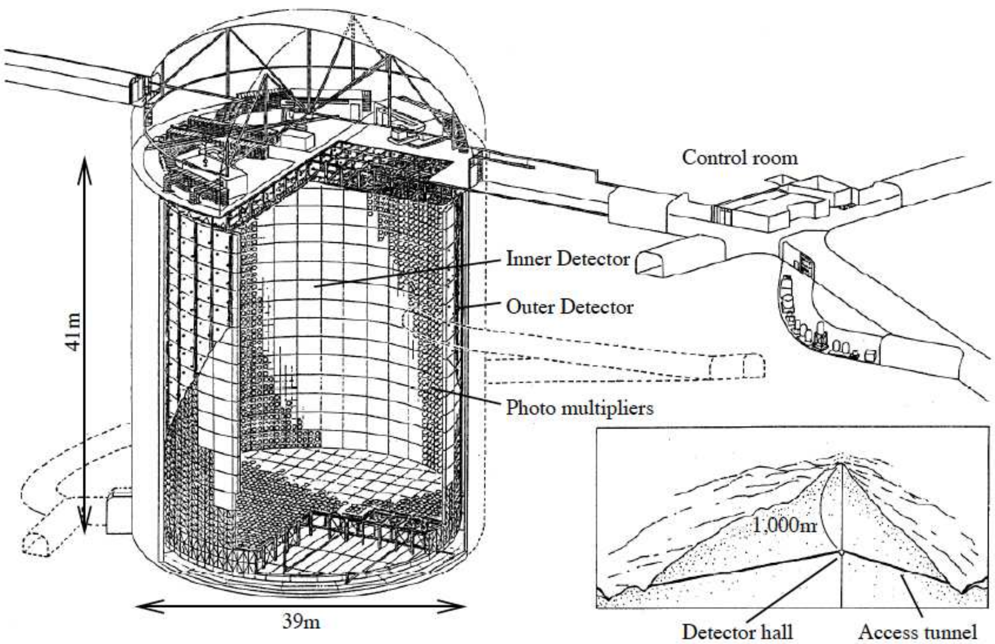

Figure 3.14: Sketch showing the structure and layout of the Super-K detector. Electronics huts are located on the lid of the water tank. The control room and dome are accessed horizontally from road level. The inner detector, inner PMTs, dead space, outer PMTs and outer detector region are demonstrated. Figure from [78].

Super-K observes neutrino interactions by detecting the Cherenkov light produced by neutrino interaction products. The primary interaction mode used for neutrino detection is charged current quasi-elastic (CCQE) interactions, which occur via

$$\nu_x + N \rightarrow N' + l_x^- \tag{3.4}$$

If the product lepton has sufficient energy that the lepton's velocity satisfies $v_l >$ c/n, where n is the refractive index of the surrounding medium, Cherenkov light will be emitted as the lepton propagates. In water, which has a refractive index of n=1.33, a charged particle will emit light if $\gamma > 1.52$. Therefore electrons with energy $E_e > 0.775$ MeV will emit Cherenkov light, as will muons of energy $E_\mu > 160.6$ MeV, and tau with $E_\tau > 2701$ MeV.

Figure 3.3 shows that the probability of ν_τ appearance in the ν_μ beam peaks at 600 MeV, and drops rapidly at higher energies. With a mass of 1776.82±0.16 MeV, the τ^- is too heavy to be produced via CC interaction of ν_τ at these energies[2], and so CC interaction of ν_τ at Super-K is very rare. In addition, the τ^- mean lifetime is very short, at $(290.6 \pm 1.0) \times 10^{-15}$ s. Should a high energy ν_τ appear in the beam and interact via the charged current, the resulting τ^- would decay, possibly to produce a muon or electron, which would emit Cherenkov light and be detected, indicating a ν_μ or ν_e interaction. This does not cause a significant background since this scenario is very rare. Therefore Super-K is not sensitive to charged current ν_τ interactions.

Muons and electrons produced by charged current interactions of ν_μ and ν_e are detected however. As a charged particle travels through the inner tank at sufficient velocity, light is emitted at an angle to the particle's path, as shown in Figure 3.15. Light emission will begin at the point of the interaction, where the lepton is created. Should the particle exit the inner tank, and be created with sufficient energy to remain above the Cherenkov threshold throughout its passage

[2] A ν_τ must have at least 3.5 GeV in order produce a τ in a CC interaction.

through the inner detector, a complete circle (or oval, depending on the particle's trajectory) of light, centred on the particle's exit point, will reach the inner wall and be detected by the PMTs.

For T2K neutrino analyses, fully contained (FC) events are selected. In this case, the energy of the lepton drops below the Cherenkov threshold before the lepton reaches the inner wall, and therefore a ring of light is detected by the PMTs, as shown on the right of Figure 3.15. Knowing the Cherenkov theshholds for various charged particles, and the rate of energy loss per unit distance traveled for different particle types, the properties of the ring pattern can be used to calculate the initial momentum of the particle. In addition, the total Cherenkov light emitted by the particle can be measured and also used to calculate the momentum of the particle.

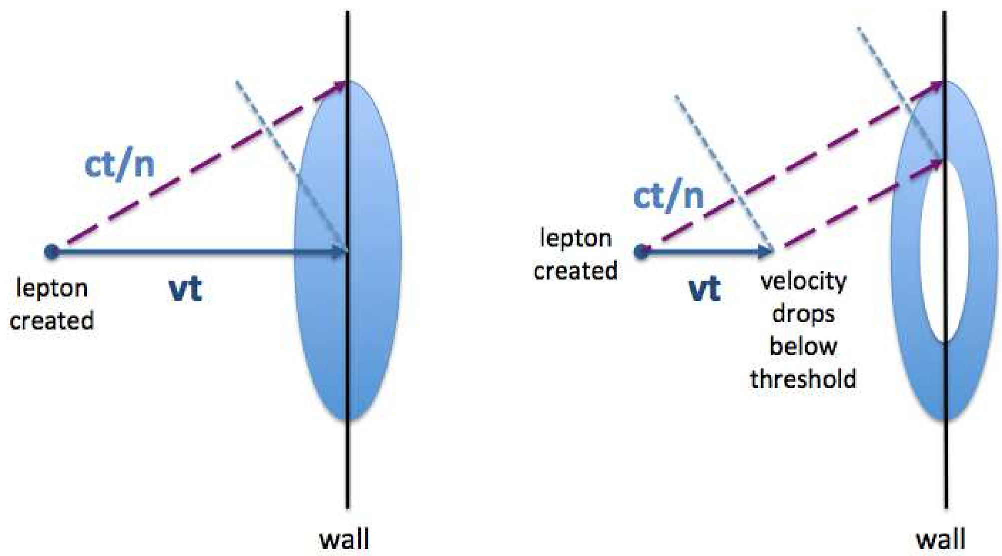

Figure 3.15: Formation of Cherenkov light patterns created by a lepton exiting the inner detector with energy above the Cherenkov threshold (left) and a lepton which ceases Cherenkov light emission within the inner detector, defined as a 'fully contained' event (right).

Properties of Cherenkov rings can also be used to identify the flavour of the lepton responsible, due to the different ways in which muons and electrons propagate through water. Due to the muon's relatively large mass, it passes through water relatively unimpeded. Any collisions with electrons in the water molecules result in little change to the muon's trajectory. Therefore the Cherenkov light cone is emitted along the straight line of the muon's path, which results in a ring pattern with a sharp, clearly defined edge. An example is given in Figure 3.16.

In contrast, the ring pattern formed by an electron is not sharp. Due to the electron's low mass, its trajectory can be altered by collisions with atomic electrons and nuclei as it propagates through water. The Cherenkov light cones produced will therefore be emitted at a range of slightly different angles as the electron travels. Changes to the direction of the electron's path also result in bremsstrahlung radiation. These photons will also be collected if they have sufficient energy to pair produce, resulting in e^- and e^+. These additional leptons can also produce Cherenkov light, or emit bremsstrahlung photons. The repetition of these processes result in a 'shower' of photons, electrons and positrons. The result of the combined photons emitted by this shower is a relatively blurred ring pattern, as shown in Figure 3.16. The different appearance of Cherenkov ring patterns produced by electrons and muons is the basis for the particle identification at Super-K, and in the case of charged current interactions, allows the flavour of the incident neutrino to be identified.

High energy photons may also pair produce, and the resulting electron and positron each result in a shower of photons, electrons, and positrons. Other than the initial particle, these showers are identical. This is the cause of a background to ν_e detection. A π^0 will decay to two photons, each of which will produce a shower. Should one of the showers not be detected, a single e-like ring will be identified. In this way neutral current interactions which produce π^0 can be mistaken for charged current interactions of ν_e. The Pattern Of Light fit algorithm (POLfit) [80], described in more detail in §4.4.2, has been developed to reduce this source

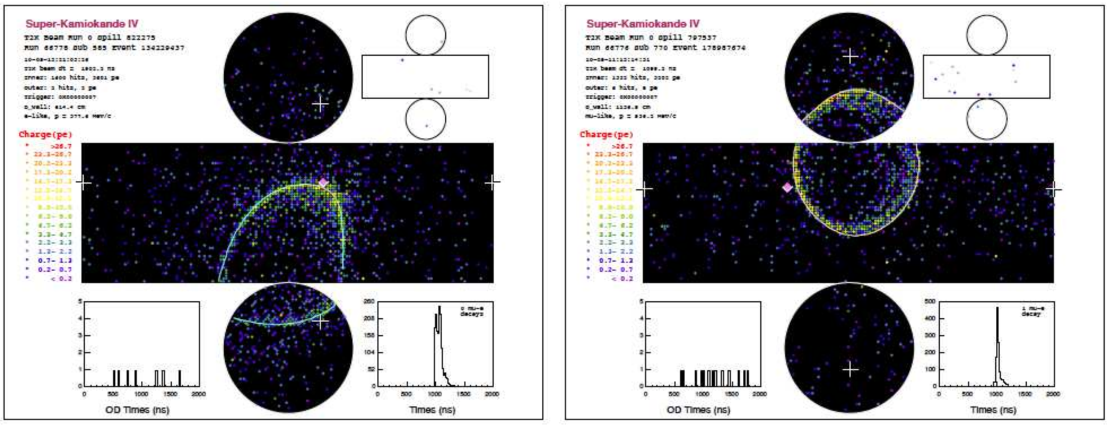

Figure 3.16: Examples of e-like (left) and μ-like (right) Cherenkov rings formed in Super-K. ID PMT map is shown by the large black cylinder projection, and OD PMT maps are shown in white. Hit PMTs are represented by coloured dots, with the colour representing the charge (timing of hit is also possible depending on event display settings). White circles show the reconstructed ring position. Figure from [79].

of background.

The beam is not the only source of neutrinos arriving at Super-K, so when counting the interactions of ν_e and ν_μ we must distinguish between interactions of beam neutrinos and non-beam neutrinos such as those originating from cosmic rays. Timing windows are used to specifically look for interactions of beam neutrinos. T2K beam spills consist of 8 neutrino bunches, distributed as shown in Figure 3.17. Each T2K beam spill is assigned a GPS timestamp which is used by the Super-K online system to define a software trigger. This records all PMT hit information recorded in a 1 ms window around the arrival time of each beam spill.

3.5 Data taking

The initial construction phase was completed in March 2009, with the first T2K beam neutrinos generated in April 2009. The first neutrino interactions were detected in INGRID in April 2009. Use of ND280 began in February 2010, with

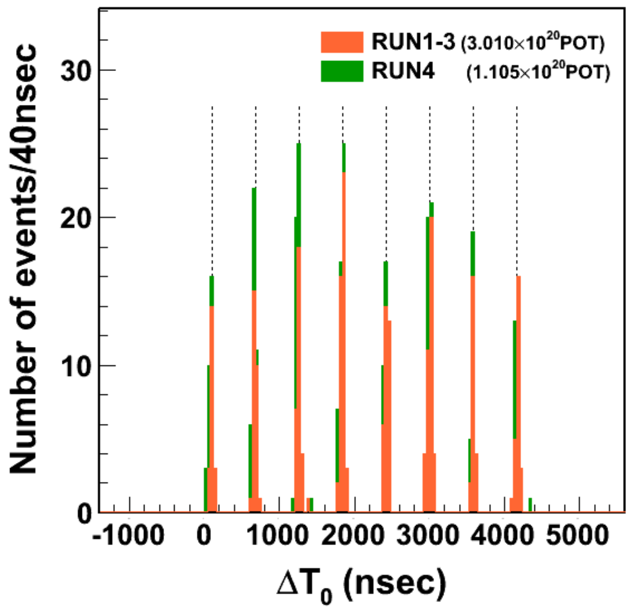

Figure 3.17: The bunch structure of a T2K beam spill. The 8 bunches (6 bunches in Run I) are distributed evenly within a window of approximately 5 μs. Figure from [81].

the collection of Run I data. The analysis presented in Ch. 9 uses Run II and part of Run III (Run details are listed in Appendix B). Data taking is ongoing.

Chapter 4

Study of intrinsic beam ν_e

The two sources of intrinsic ν_e contamination, which is one of two backgrounds to the ν_e appearance measurement, are described. The benefits of using the near and far detectors to make a direct measurement of this background are discussed, and the suitability of using the far detector is evaluated using Monte Carlo data. Expected interaction rates, along with background events, are given for one year of nominal data taking, and additional cuts designed to improve the purity of the ν_e sample are tested.

4.1 Intrinsic ν_e background

T2K is primarily a ν_e appearance experiment, where the signal is the detection of ν_e at the far detector. The number of ν_e observed for a given number of protons on target determines the $\nu_\mu \rightarrow \nu_e$ oscillation probability and the oscillation parameter θ_{13}. It is important that the sources of background to these true ν_e signal events are accounted for, for comparison with the total number of ν_e candidates seen.

One of the two main backgrounds is the intrinsic ν_e that are created in the decay pipe. Interactions of ν_e produced by the decay of hadrons in the beam pipe and ν_e resulting from ν_μ oscillations are identical and will produce exactly the same signal in our detectors. Two sources of ν_e contamination exist.

The collision of protons with the graphite target produces a variety of hadrons. T2K selects the π^+ to enter the decay pipe using magnetic horns and discards the remaining hadron types, which include π^-, K^+, K^- and K^0. However, this selection process is not perfect and some of these undesired hadrons also enter the decay pipe and contribute decay products to the beam. The predicted flux of intrinsic ν_e at Super-K is shown in Figure 4.1, and the predicted $\bar{\nu}_e$ flux is shown in Figure 4.2. In each case the total flux is shown in black and the contributions to the flux from each different parent type are shown individually.

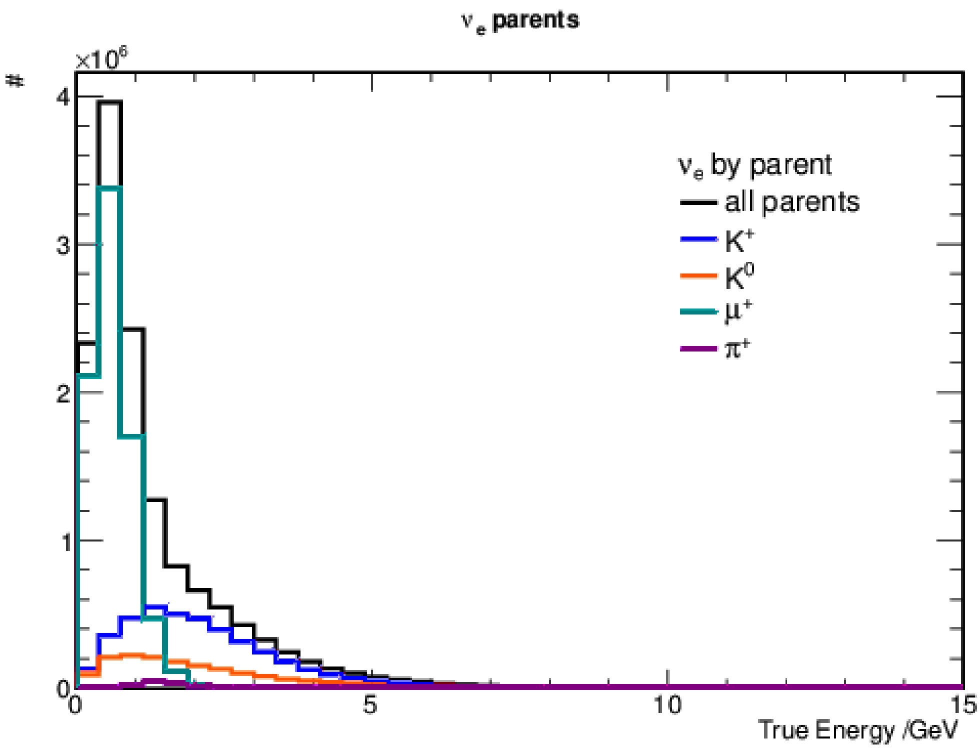

Figure 4.1: Predicted flux of intrinsic ν_e (as opposed to signal ν_e resulting from oscillations) at Super-K divided into contributions from different parents. Total flux given in black.

A π^+ decays to produce ν_μ with a branching ratio of 99.99%.

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \tag{4.1}$$

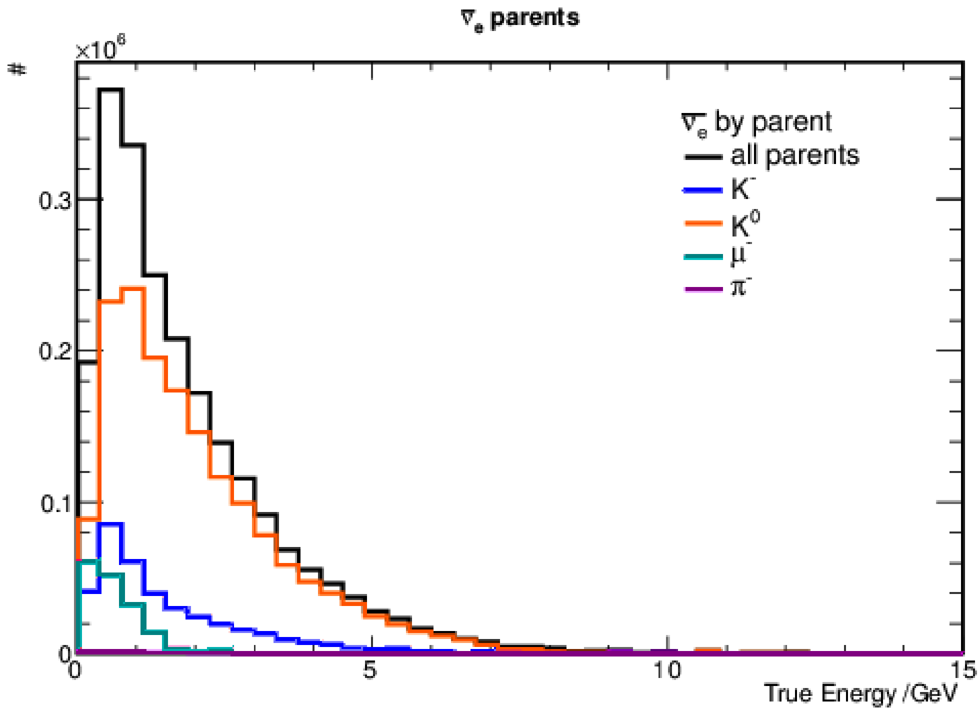

Figure 4.2: Predicted flux of $\bar{\nu}_e$ at Super-K divided into contributions from different parents. Total flux given in black.

Equivalent decays to ν_e are negligible and so do not significantly contribute to background events.

Figure 4.1 shows that μ^+ decay is the dominant source of intrinsic ν_e in the range 0–1 GeV. The decay pipe length has been optimised to allow the maximum number of pion decays to occur, in order to maximise the number of ν_μ produced, while also minimising the likelihood that a daughter μ^+ will subsequently travel far enough to also decay before it stops in the beam dump. A significant number of μ^+ will decay, as in Eq. 4.2, before being captured however. This is one source of intrinsic ν_e.

$$\mu^+ \rightarrow e^+ + \bar{\nu_\mu} + \nu_e \quad \text{(branching ratio 100\%)} \tag{4.2}$$

K$^+$ and K^0 also have several significant decay modes which produce μ^+.

$$K^+ \rightarrow \nu_\mu + \mu^+ \quad \text{(branching ratio 63.55\%)} \tag{4.3}$$

$$K^+ \rightarrow \nu_\mu + \mu^+ + \pi^0 \quad \text{(branching ratio 3.35\%)} \tag{4.4}$$

$$K^0 \rightarrow \nu_\mu + \mu^+ + \pi^- \quad \text{(branching ratio 27.04\%)} \tag{4.5}$$

The μ^+ contribution shown in Figure 4.1 includes μ^+ produced from all of these decays. Therefore, to predict the number of ν_e resulting from μ^+ decays, we must know the quantities of π^+, K$^+$, and K^0 entering the decay pipe.

Kaons also produce intrinsic ν_e directly via the following decays.

$$K^+ \rightarrow \nu_e + e^+ + \pi^0 \quad \text{(branching ratio 5.07\%)} \tag{4.6}$$

$$K^0 \rightarrow \nu_e + e^\pm + \pi^\mp \quad \text{(branching ratio 40.55\%)} \tag{4.7}$$

The intrinsic ν_e contributions from kaons peak at around 1.5 GeV, and kaons are the dominant source of ν_e above this energy.

We must be able to predict the flux of intrinsic ν_e at the far detector and hence the number of intrinsic ν_e interactions. For the appearance analysis described in [35], the uncertainty on the expected number of background ν_e events is 12%. The details of hadron production at a carbon target are not known to a high level of precision, and this introduces one of the main sources of uncertainty in the background estimations. The background estimate also requires knowledge of the decay tunnel geometry, and the mean lifetimes and branching ratios of the hadrons entering the decay pipe, which are all well-known. It is important to measure the intrinsic ν_e present in the beam to confirm the simulated ν_e contamination. Since the level of ν_e contamination in the beam is dependent on the relative quantities of pions and kaons that enter the decay pipe, a measurement of the number of

ν_e interactions could also be used to improve our knowledge of the selection of hadrons contributing to the T2K beam. One way to obtain such a measurement using a T2K detector is now discussed.

4.2 The high energy tail

A measurement of ν_e background events must be able to distinguish between the intrinsic ν_e events of interest, and signal ν_e events generated in $\nu_\mu \rightarrow \nu_e$ oscillations. Using the current measurements of the values θ_{13} and Δm^2_{23}, the probability that a neutrino created with flavour ν_μ, of energy 600 MeV, will interact as a ν_e at a baseline of 280 m is effectively zero. Therefore any ν_e detected at the near detector can only be intrinsic ν_e, produced by the decay of mesons, or in some cases subsequent leptons, in the decay pipe.

Measuring the ν_e flux at ND280 would require accurate identification of neutrino types interacting in the near detector in order to isolate the relatively low number of ν_e from the far more numerous ν_μ of the neutrino beam. The flavour of a neutrino interacting in a detector is identified by the flavour of the lepton produced in a charged current interaction. Differentiating between a μ^- and an e^- produced in a neutrino interaction in ND280 is not trivial and required the development of selection cuts to identify and separate the different interaction products. At the time of writing, such a particle identification system has been developed [82]; however, the near detector was not immediately suitable for an intrinsic ν_e measurement at the start of data taking. In contrast, as a previously established detector, Super-K has a large total run time to date and a system of particle identification was already developed before T2K data taking began. The far detector is particularly well suited to distinguishing between ν_e and ν_μ interactions. The far detector therefore met the particle identification requirements better than ND280 during the first years of the experiment.

If a ν_e background measurement using the far detector were to prove viable, it would complement the eventual background measurements made by ND280 in later years. Interaction rates are the product of a particle flux and the relevant interaction cross-sections. The neutrino interaction cross-sections in ND280 and Super-K are different, due to the different target materials. Other than easily calculable geometric factors, the intrinsic ν_e flux does not change between the near and far detectors. Measurements of the interaction rates of intrinsic ν_e at the near and far detectors could therefore provide a useful check of the cross-section values used by T2K.

T2K is designed such that $\nu_\mu \rightarrow \nu_e$ oscillations peak at 295 km, so both signal ν_e and intrinsic ν_e will be present in the beam at Super-K. We can use the relationships between energy, baseline and oscillation probability to separate the two types of ν_e.

Figure 3.4 shows the probability of ν_μ to ν_e oscillations for a baseline of 295 km and demonstrates that after a peak in probability at 600 MeV, the probability drops to zero as the neutrino energy increases. ν_μ with energies of approximately 2 GeV and above have no probability of interacting as ν_e at 295 km. Therefore any ν_e detected in the high energy tail at the far detector cannot be the result of ν_μ to ν_e appearance and so must be intrinsic ν_e.

Figure 4.3 shows the ν_e survival probability at 295 km for neutrino energies up to 10 GeV. It confirms that the probability of ν_e disappearance is small at 600 MeV and drops to zero in the high energy tail, so a negligible number of intrinsic ν_e would be lost due to oscillations.

Measuring the high energy ν_e interactions at Super-K would therefore provide a direct observation of the high energy region of the intrinsic ν_e spectrum. By only

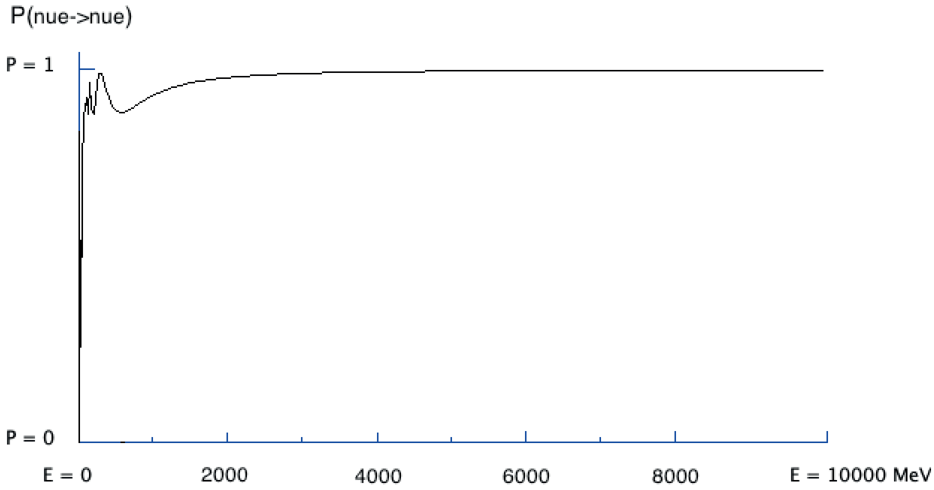

Figure 4.3: Probability of ν_e to ν_e oscillation at baseline of 295 km for range of neutrino energies. Uses $\sin^2 2\theta_{12} = 0.8704$, $\sin^2 2\theta_{23} = 1.0$, $\Delta m^2_{12} = 7.6\times10^{-5}$ eV^2, $\Delta m^2_{23} = 2.4\times10^{-3}$ eV^2, average matter density = 3.2 gcm^{-2}, and assumes $\sin^2 2\theta_{13}$ = 0.1. Figure created using [61].

selecting neutrinos above an appropriate energy, we can be confident that all the ν_e detected are intrinsic ν_e. The number of neutrinos selected can be compared to the predicted number in the same energy range to verify the predicted intrinsic ν_e flux. The high energy measurement could also be used to infer the total intrinsic ν_e present in the beam by scaling the predicted shape of the intrinsic ν_e spectrum. Setting the integral of the predicted intrinsic ν_e distribution for the high energy range equal to the number of ν_e measured in that range would scale the entire distribution.

The predicted spectrum of intrinsic ν_e is generated using current models of hadron production at the target and predictions of the relative quantities of pions and kaons that are selected and enter the decay pipe [83]. We know that the higher energy intrinsic ν_e are the products of kaons only (see Figure 4.1); therefore, comparing the observed intrinsic ν_e with the predicted ν_e in the high energy tail could identify any discrepancies present in the modelled kaon contributions.

Since the kaon and μ^+ contributions overlap at lower energies, correcting the kaon contribution would provide information about the μ^+ contribution to the ν_e spectrum also. With consideration of the known branching ratios and decay lifetimes, improving our understanding of the intrinsic ν_e generated by the different parents can also reduce the uncertainty on the total pions and kaons selected to form the neutrino beam.

4.3 Standard Super-K analysis cuts

This analysis was performed before far detector data had been collected in order to test the viability of such a measurement. The design of the T2K experiment minimises the flux of high energy neutrinos, so it is necessary to investigate whether the event rate in the high energy tail is sufficient to provide a sample that is not excessively limited by statistical uncertainty. The Super-K Monte Carlo was used to investigate the expected number of intrinsic ν_e that would be detected at high energies for a fixed period of data taking. The Monte Carlo is also used to study background events.

The flux predictions for each neutrino type, such as the distributions shown in Figures 4.1 and 4.2, are used as the basis for the Monte Carlo generation. These contain the beam simulation outputs, which for the 2010a release[1] are generated using JNUBEAM [85]. Hadron production in the target is simulated using FLUKA [86], with some additional inputs for pion production from NA61/SHINE [87] data [83]. The fluxes are then multiplied by the relevant cross-sections and files containing lists of the final state particles are produced using NEUT [88]. The interaction of the beam neutrinos in the far detector is modelled and the simulated

[1]Updated versions of Super-K Monte Carlo are generated approximately every 2 years to include the latest beam flux and cross-section information. This thesis contains results obtained using the version known as the 2010a release. Documentation is given at [84].

PMT signals are passed through the main reconstruction program. The simulated detector responses are contained in the Monte Carlo output files, along with some reconstructed quantities and truth quantities for comparison.

As described in §3.4, the properties of the cone of light produced by a moving particle are used to identify the particle type, along with the particle's trajectory through the tank and its starting position. Since this measurement of intrinsic ν_e requires neutrinos in the high energy tail only, a selection cut must be placed on the neutrino energy, and so it is important that the energy can be reconstructed accurately. Energy reconstruction is easiest in the case of charged current quasi-elastic (CCQE) interactions, such as:

$$\nu_e + n \rightarrow e^- + p^+ \tag{4.8}$$

Changes to the target nucleon are rarely observed, so in most cases the only visible product of a CCQE interaction is one charged lepton. In this case the energy of the incident neutrino can be calculated as will be described in §4.4.1. The presence of only one product particle also simplifies the particle identification process, as only one ring will be produced on the detector wall. The particle can be identified as μ-like or e-like using characteristics of the pattern of light detected, as seen in §3.4.

A set of selection cuts are defined to select single ring e-like events that are fully contained and occur in the fiducial volume of the detector. The standard cuts developed by the collaboration for the ν_e appearance analysis [79][34] are designed to select any ν_e which undergo a CCQE interaction and so are also appropriate for a selection of intrinsic ν_e. They are defined below.

1. wall>200

Fiducial volume. 'Wall' gives the shortest distance between the reconstructed vertex position and the inner wall. The reconstructed vertex position must be inside the inner detector and at least 200 cm from the inner detector wall for an event to be selected. If a neutrino interaction occurs close to the detector wall and the lepton created travels the shortest distance to the wall, reconstructing the event is more difficult as there is less PMT hit information to use. The event is also less likely to be fully contained.

Leptons created in neutrino interactions taking place outside the inner detector, and cosmic ray muons, are rejected by this cut.

'Wall' gives the shortest distance from the reconstructed vertex to the wall, and 'wallv' gives the shortest distance from the true vertex position.

2. nhitac<16

Outer hits. For an event to be classed as fully contained, there must be no hit cluster in the outer detector with more than 16 PMT hits. By limiting the allowed activity in the outer detector we reject neutrino interactions that occur outside the inner detector, interactions that produce leptons which leave the inner detector, and cosmic ray muons.

3. evis>100

Visible energy. The visible energy is the sum of energy from all rings associated with an interaction. This cut rejects low energy events, such as neutral current interactions and electrons produced via muon decay, which can be sources of background to ν_e measurements. CCQE ν_e interactions are very unlikely to occur with visible energy less than 100 MeV, so this cut increases the purity of the selected events with minimal effect on the efficiency.

4. nring=1

Number of rings. The distance traveled through the detector is used to reconstruct the energy of the lepton, so leptons which escape cannot be reconstructed as accurately. Therefore only leptons which form a ring (i.e. start and stop in the detector) are selected. Events which are not fully contained (demonstrated by Figure 3.15) and therefore do not form a ring, are rejected.

We require CCQE interactions only. The only visible product of a CCQE interaction is one charged lepton, which will produce exactly one ring.

5. ip=2

Particle identity. $\nu_\mu \to \nu_\tau$ oscillations peak at 600 MeV, so any neutrino which could interact as a ν_τ at Super-K would not have sufficient energy to produce a τ lepton. Therefore the only possible CCQE interactions at Super-K are of ν_μ or ν_e, and a single ring must correspond to a muon or electron. As described in §3.4, the distribution of light forming the ring is used to identify the lepton type (examples given in Figure 3.16), and a cut can be made on the output of the particle identification algorithm to select electron-like rings only. The *ip* variable holds the best-matching particle type, where 2 corresponds to e-like[2].

6. nummuedecay=0

No decay electrons. For every event, we search for the presence of a secondary event, occurring a short time later. If a muon decays to produce an electron with sufficient energy, this electron may produce a detectable ring. Pions may also decay via muons and generate e-like rings. To select CCQE ν_e interactions we require e-like rings which are the direct result of the neutrino interaction, and no other products. The presence of any decay electrons indicates that muons or pions were present, and so the event is not selected.

[2]The beam e-like selection efficiency is (98.9±1.1)%, from [79]

The Super-K Monte Carlo includes interactions of signal ν_e, beam ν_e, beam ν_μ and beam $\bar{\nu_\mu}$. The flux of intrinsic $\bar{\nu_e}$ are considered negligible compared to the other contributions so they are not modelled[3]. 1,000,000 beam ν_μ and 1,000,000 intrinsic ν_e simulation events are generated, and 500,000 beam $\bar{\nu_\mu}$ and 500,000 signal ν_e are generated. Analyses are performed using all available Monte Carlo and selected events are then scaled according to the expected number of interactions of each particle based on flux and cross-section values. We consider the interactions of the beam ν_μ and smaller number of $\bar{\nu_\mu}$ together, since they will both result in μ-like rings being detected. The signal ν_e events are not processed, since we will be choosing an energy range that will contain a negligible number of signal ν_e.

Table 4.1 gives the number of interactions at Super-K that pass the cuts described above and form the fully contained fiducial volume, single ring, e-like sample. Oscillation parameters $\sin^2 2\theta_{12} = 0.8704$, $\sin^2 2\theta_{23} = 1.0$, $\Delta m^2_{12} = 7.6\text{x}10^{-5}$ eV^2, $\Delta m^2_{23} = 2.4\text{x}10^{-3}$ eV^2 are applied, with the average matter density set as 3.2 gcm^{-3}, and $\sin^2 2\theta_{13}$ is set at 0.1. The true types of the events are known in the Monte Carlo, so the selected events are separated into true intrinsic ν_e that pass the cuts, and true beam ν_μ and $\bar{\nu_\mu}$ that also pass the e-like cuts and therefore form the main background to the intrinsic ν_e sample.

These numbers give the expected intrinsic ν_e in the range 0–9.95 GeV[4], and the background caused by beam ν_μ and $\bar{\nu_\mu}$ that pass the e-like selection cuts for comparison.

[3]The intrinsic beam ν_e contribute 1% of the total beam flux, and the intrinsic $\bar{\nu_e}$ flux is 10% of the beam ν_e flux, as plotted in [35]. $\bar{\nu_e}$ may therefore make a small contribution to the total e-like intrinsic beam events seen at Super-K, however these contributions are not significant enough to alter the conclusions of this chapter.

[4]This is the energy range available in some Super-K Monte Carlo data sets.

Selection cut	True ν_e	True ν_μ + $\bar{\nu_\mu}$	ϵ	π	$\epsilon.\pi$
None	72.63 $\pm$0.07	2261.81 $\pm$2.83	-	-	-
wallv>200	47.04 $\pm$0.06	1465.52 $\pm$2.28	100%	3.11%	0.0311
wall>200	45.78 $\pm$0.06	1317.00 $\pm$2.15	97.3%	3.36%	0.033
+ nhitac<16	44.41 $\pm$0.06	1158.90 $\pm$2.03	94.4%	3.69%	0.035
+ evis>100	38.51 $\pm$0.05	702.66 $\pm$1.57	81.9%	5.20%	0.043
+ nring=1	21.46 $\pm$0.04	331.56 $\pm$1.07	45.6%	6.08%	0.028
+ ip=2	20.94 $\pm$0.04	43.08 $\pm$0.39	44.5%	32.7%	0.146
+ nummuedecay=0	17.68 $\pm$0.04	33.82 $\pm$0.34	37.6%	34.3%	0.129

Table 4.1: Intrinsic ν_e expected at Super-K scaled to a nominal year of data taking (1.66×10^{21} POT, 30 GeV beam, 750 kW) which pass fully contained fiducial volume single ring e-like selection. Cuts are cumulative as of wall>200. Background due to beam ν_μ + $\bar{\nu_\mu}$ passing the selection cuts is given for comparison. ϵ gives the efficiency compared to the true fiducial volume selection (wallv>200), and π describes the purity of the selection. Statistical uncertaintiess on the selected Monte Carlo events are given.

4.4 Additional selection cuts

While the probability that a ν_μ or $\bar{\nu_\mu}$ interaction will pass the e-like selection cuts is low, the comparatively high flux of beam ν_μ and $\bar{\nu_\mu}$ results in a significant number of ν_μ and $\bar{\nu_\mu}$ passing the selection cuts and causing a background signal. There are two ways in which a ν_μ or $\bar{\nu_\mu}$ interaction can be mistaken for a ν_e interaction. Firstly, the particle identification algorithm may incorrectly identify the lepton produced in the interaction. Characteristics of the ring of Cherenkov light formed by the product lepton are used to identify the lepton flavour; however, there is a small overlap in the properties of e-like and μ-like rings which results in a small number of $\mu^\pm$ rings being mistaken for $e^\pm$ rings. The other source of background for this intrinsic ν_e measurement is the production of $\pi^0 \rightarrow \gamma\gamma$ by a ν_μ or $\bar{\nu_\mu}$ interaction. Under certain circumstances the decay products of a π^0 can create a light ring that mimics the single e-like ring indicative of a ν_e interaction.

To exclude the interactions of genuine signal ν_e which are anticipated at low energies, we first apply a cut on the reconstructed energy which only selects high energy interactions. We then consider two possible methods of reducing the background events. Since T2K is designed to minimise the number of intrinsic ν_e present in the beam in order to reduce backgrounds, the number of intrinsic ν_e present to be selected by the standard cuts is already very low. Therefore any further selection cuts implemented to improve the purity can only be considered if they result in minimal reductions of the selection efficiency.

4.4.1 High energy cut

To select only interactions of neutrinos in the high energy tail, we must first reconstruct the energy of the incident neutrinos. For a CCQE interaction of the form

$$\nu_l + n \rightarrow l^- + p^+ \tag{4.9}$$

the energy of the neutrino is given by Equation 4.10[5] [79].

$$E_\nu = \frac{2E_e(m_n - V) - m_e^2 + 2m_nV - V^2 + m_p^2 - m_n^2}{2(m_n - V - E_e + P_e\cos\theta)} \tag{4.10}$$

E_e and P_e are the reconstructed energy and momentum of the electron respectively. The neutrino beam direction is known and the lepton direction is measured using the timing information of the Cherenkov ring produced. The angle between these directions is cosθ. m_p and m_n are the masses of the proton and neutron respectively, and m_e is the mass of the electron. V represents the binding energy of a nucleon in a ^{16}O nucleus and is set equal to 27 MeV [35]. Since we will be selecting ν_e, and therefore only selecting events that pass all e-like cuts, it is appropriate to assume the observed lepton is an electron when reconstructing the

[5]The value of E$_\nu$ will not be accurate for any CCnonQE interaction which passes the CCQE selection cuts

Selection cut	True ν_e	True $\nu_\mu + \bar{\nu_\mu}$	ϵ	π	$\epsilon.\pi$
None	72.63 ±0.07	2261.81 ±2.83	-	-	
wallv>200	47.04 ±0.06	1465.52 ±2.28	100%	3.11%	0.0311
FCFV 1R e-like	17.68 ±0.04	33.82 ±0.34	37.6%	34.3%	0.129
+ $E_{recon} > 1\ GeV$	10.82 ±0.03	11.30 ±0.20	23.0%	48.9%	0.112

Table 4.2: Intrinsic ν_e expected at Super-K scaled to a nominal year of data taking (1.66×10^{21} POT, 30 GeV beam, 750 kW) which pass the fully contained fiducial volume single ring e-like selection (labelled as FCFV 1R e-like) plus the high energy selection (energy range 1.0–9.95 GeV) are given on the bottom line. Events passing the FCFV 1R e-like cuts without the added high energy cut are given for comparison above. ϵ gives the efficiency compared to the true fiducial volume selection, and π describes the purity of the selection. Statistical uncertainties on the Monte Carlo events selected are provided.

energy of both the lepton and the neutrino[6].

For each selected interaction the reconstructed energy of the neutrino is calculated using Equation 4.10. The accuracy of this estimation is limited by the accuracy of the energy and momentum reconstructions for the lepton. The reconstructed energy will not be accurate when this formula is applied to interactions other than ν_e CCQE interactions.

The number of intrinsic ν_e interactions is low, so the energy cut must be chosen such that the largest number of interactions is retained. However, the aim of this measurement is to measure intrinsic ν_e only and not include any signal ν_e in the selection, so the cut must be sufficiently above 600 MeV to ensure no signal ν_e are included in the sample. The probability of losing an intrinsic ν_e to oscillations also drops with increasing energy. Taking these effects into consideration, 1 GeV is chosen as a preliminary selection cut. The expected numbers of interactions selected for 1.66×10^{21} POT with the additional selection cut $E_{recon} > 1$ GeV are given in Table 4.2.

[6]Using these assumptions, the reconstructed energies of ν_μ will not be accurate, but interactions identified as μ-like will not pass the selection cuts and will not be included in the final sample.

NEUT mode	Description	Interaction
31	Single π from Δ resonance	ν + n $\rightarrow$ ν + n + π^0
32	Single π from Δ resonance	ν + p $\rightarrow$ ν + p + π^0
41	Multi π	ν + n/p $\rightarrow$ ν + p/n + multi π
46	Deep inelastic	ν + n/p $\rightarrow$ ν + n/p + mesons

Table 4.3: Most commonly occurring NEUT interaction modes of ν_μ and $\bar{\nu_\mu}$ which pass the intrinsic ν_e selection cuts. All are neutral current modes. (Additional modes given in Appendix C.)

The high energy tail contains proportionally fewer background events in the form of selected ν_μ and $\bar{\nu_\mu}$, and therefore this energy cut increases the purity of the sample. Since the beam is designed to peak near the oscillation maximum at approximately 600 MeV, the range 0–1 GeV contains a significant proportion of the events, so this cut almost halves the total intrinsic ν_e selected. Further reduction in the number of background events is desirable; however, any possible cuts must be tested for their effect on the selected number of intrinsic ν_e.

4.4.2 Pi0mass cut

The interaction modes of all events that pass the standard selection cuts above are shown in Figure 4.4 for a selection of Monte Carlo. The ν_μ and $\bar{\nu_\mu}$ that are mistakenly selected by the intrinsic ν_e selection cuts are shown in blue.

One of the main sources of background to the intrinsic ν_e selection is interactions of ν_μ and $\bar{\nu_\mu}$ that lead to the production of π^0 mesons. The most significant contributions to the ν_e background are listed in Table 4.3. All correspond to neutral current interactions.

The selection cuts are designed to select interactions that produce a single e-like ring. Neutral current interactions that produce π^0 mesons can have the same appearance as charged current ν_e interactions. In a neutral current interaction,

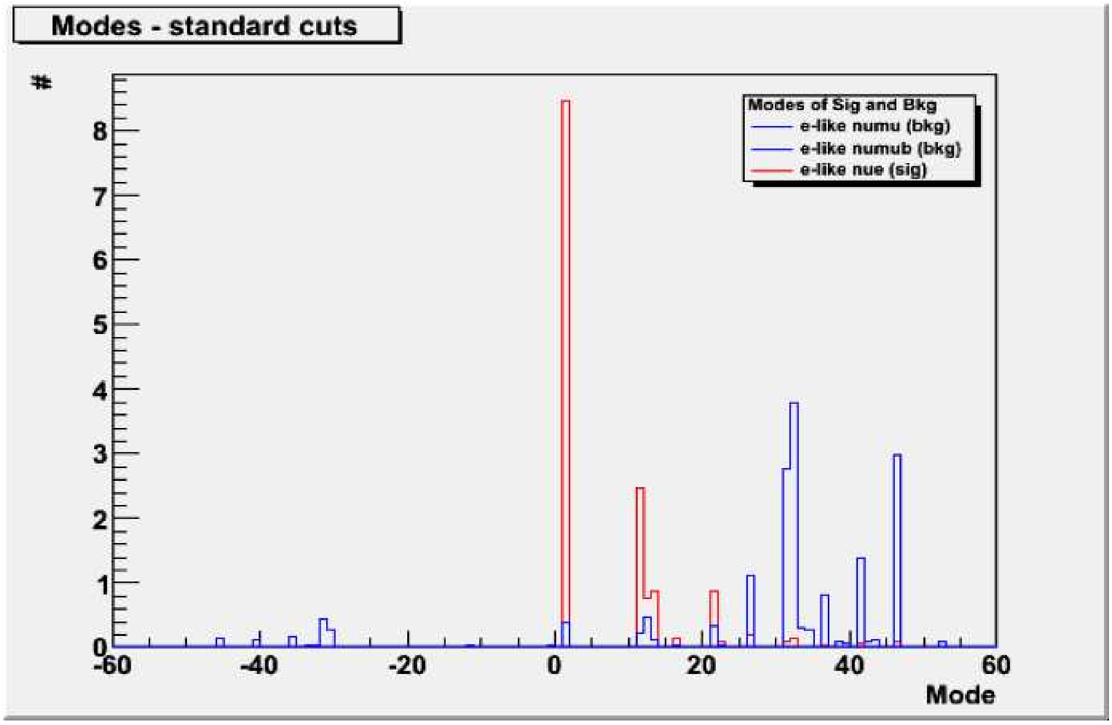

Figure 4.4: Interaction modes of interactions passing standard Super-K cuts. True intrinsic ν_e are shown in red, where the elastic mode (NEUT code of 1) accounts for approximately 59% of the true intrinsic ν_e selected. True ν_μ and $\bar{\nu_\mu}$ which pass e-like cuts and form background are plotted in blue, with the majority occurring with modes >30, corresponding to NC interactions. Full list of NEUT interaction codes is given in Appendix C.

the neutrino survives and leaves the detector without further interaction, and no leptons are produced. Changes to the nucleus are not often observed, so the only visible product of a neutral current interaction is any additional particle that may be produced, such as a pion.

A π^0 created in the interaction will decay to produce two gamma ray photons. Within a short distance, each γ may then produce leptons, for example by forming an e^+e^- pair, or scattering, resulting in the emission of an atomic electron. These e^+ or e^- will travel a short distance, emitting Cherenkov radiation, and then scatter, radiating further photons. The process continues with the result that each of the initial γ products can produce an electromagnetic shower, and the combined light emitted by each shower will form an e-like Cherenkov ring. The appearance of these rings will be identical to the rings formed by electrons produced in interactions, as the shower created is identical except for the form of the initial particle. Therefore a π^0 can be recognised by the detection of two e-like rings.

When both e-like rings are identified, the source can be correctly identified as a π^0. However, should the two gammas be emitted so that the two rings overlap, or should one of the gammas not have sufficient energy to emit enough Cherenkov radiation to produce a distinct shower, the event will appear to consist of one e-like ring. This event will pass selection cuts and be counted as a CCQE interaction of a ν_e. This is the case for the ν_μ and $\bar{\nu_\mu}$ interactions in Figure 4.4 with the NEUT codes listed in Table 4.3.

To reduce this background we can investigate all events that pass the fully contained fiducial volume, single ring, e-like selection cuts further. To do this, the probability that the single e-like ring detected is in fact one of a pair, and therefore the result of a π^0 decay, is estimated using the Pattern Of Light FIT algorithm (POLfit) [80]. POLfit assumes that a second e-like ring was produced, and using the PMT information for the event, tests all possible positions and energies for the missing ring. The optimum position and energy of the proposed second ring is found, and an invariant mass is calculated from the combination of the observed ring and the best-fitting proposed extra ring. Should the detected ring be the product of π^0 decay, and the proposed second ring be a reasonably accurate reconstruction, the invariant mass found will closely match the π^0 mass of 135 MeV. The relationship between event type and the invariant mass is demonstrated by Figure 4.5, with an example cut labelled, and shows that the invariant mass of neutral current events peaks around 135 MeV.

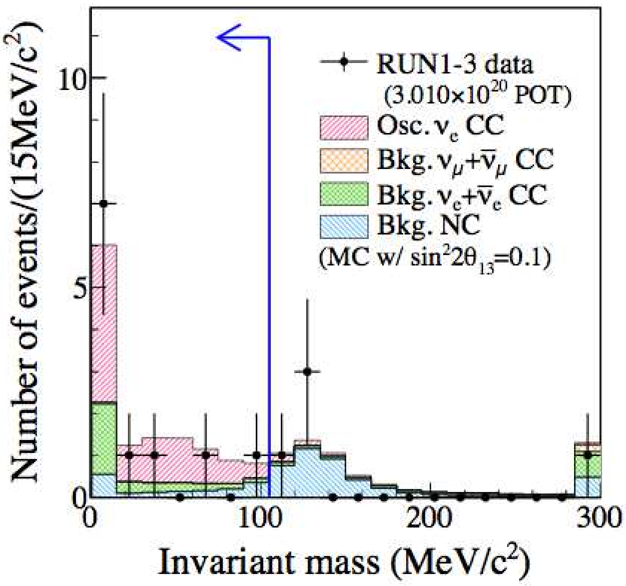

Figure 4.5: Distribution of invariant mass M_{inv} when each event is forced to be reconstructed as two photon rings. The data are shown as points with error bars (statistical only) and the MC predictions are in shaded histograms. The last bin shows overflow entries. The blue arrow shows the selection criterion $M_{inv} <$ 105 MeV/c^2. Figure from [35].

POLfit is applied to every event we select. The invariant mass is calculated and labelled in the Monte Carlo files as 'pi0mass'. For true CCQE interactions of ν_e, for which there is no true second ring, the optimum second ring suggested by

POLfit will have a very low likelihood of being a genuine ring, and the calculated value of *pi0mass* will be considerably lower than the actual π^0 mass. The *pi0mass* value can therefore be used to identify likely neutral current interactions that have been misidentified. By excluding events that have values of *pi0mass* close to 135 MeV, we can remove likely neutral current interactions that act as background events.

To demonstrate the effect of such a cut on the *pi0mass* value, Figure 4.6 shows the interaction modes of events surviving the standard selection cuts and high energy tail cut, with an additional cut of *pi0mass*<20 MeV applied. Comparison with Figure 4.4, which does not include the *pi0mass* cut, shows that the majority of the background events are removed when this cut is applied.

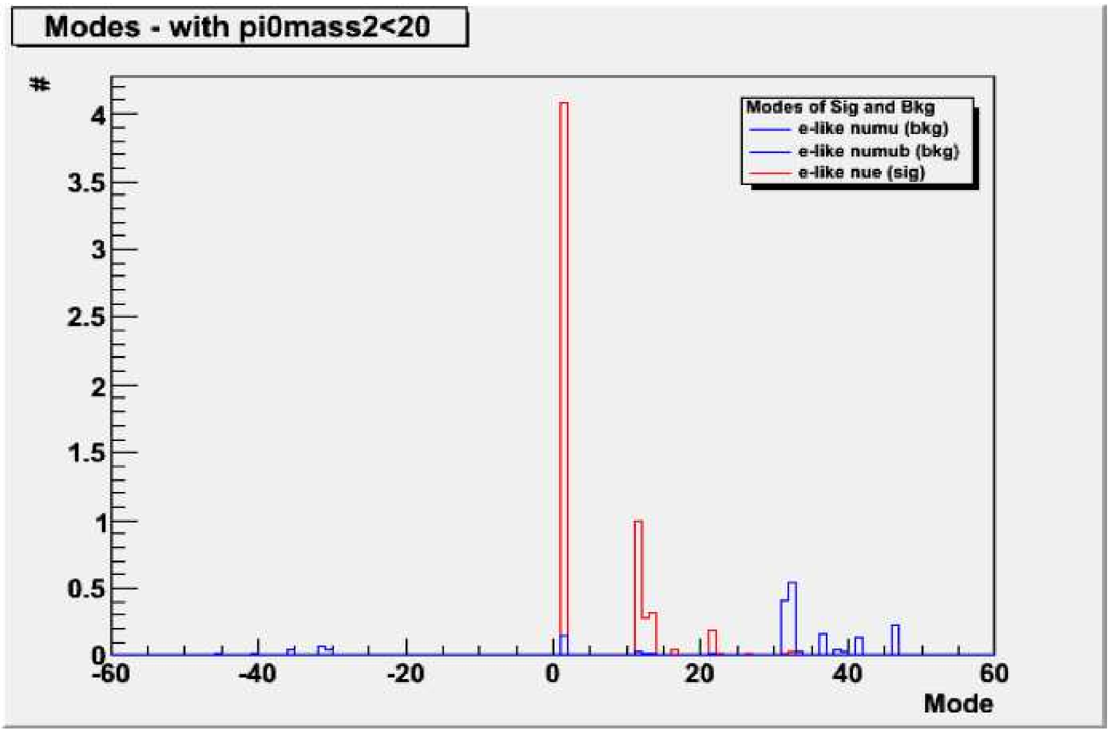

Figure 4.6: Interaction modes of interactions passing standard Super-K cuts with additional cut of pi0mass<20 MeV applied. True intrinsic ν_e are shown in red. The true ν_μ and $\bar{\nu_\mu}$ which pass e-like cuts and form background are plotted in blue. Full list of NEUT interaction codes is given in Appendix C.

Figure 4.6 shows a significant reduction in background events; however, it uses

Selection cut	**True ν_e**	**True $\nu_\mu + \bar{\nu_\mu}$**	ϵ	π	$\epsilon.\pi$
None	72.63 ±0.07	2261.81 ±2.83	-	-	-
wallv>200	47.04 ±0.06	1465.52 ±2.28	100%	3.11%	0.0311
FCFV 1R e-like	17.68 ±0.04	33.82 ±0.34	37.6%	34.3%	0.129
+ E_{recon} > 1000	10.82 ±0.03	11.30 ±0.20	23.0%	48.9%	0.112
+ pi0mass < 20	4.47 ±0.02	1.53 ±0.07	9.50%	74.5%	0.0707
or + pi0mass < 30	4.70 ±0.02	1.68 ±0.07	9.99%	73.6%	0.0736
or + pi0mass < 40	4.83 ±0.02	1.78 ±0.08	10.3%	73.0%	0.0750
or + pi0mass < 50	4.93 ±0.02	1.85 ±0.08	10.5%	72.7%	0.0763
or + pi0mass < 60	5.03 ±0.02	1.94 ±0.08	10.7%	72.2%	0.0771
or + pi0mass < 70	5.12 ±0.02	2.03 ±0.08	10.9%	71.6%	0.0780
or + pi0mass < 80	5.23 ±0.02	2.15 ±0.08	11.1%	71.0%	0.0788
or **+ pi0mass < 90**	**5.33 ±0.02**	**2.29 ±0.09**	**11.3%**	**70.0%**	**0.0794**
or + pi0mass < 100	5.45 ±0.02	2.53 ±0.09	11.6%	68.2%	0.0790
or + pi0mass < 110	5.57 ±0.02	2.82 ±0.10	11.8%	66.4%	0.0786
or + pi0mass < 120	5.69 ±0.02	3.18 ±0.10	12.1%	64.2%	0.0776
or + pi0mass < 130	5.83 ±0.02	3.66 ±0.11	12.4%	61.5%	0.0762
or + pi0mass < 140	5.97 ±0.02	4.12 ±0.12	12.7%	59.1%	0.0750
or + pi0mass < 150	6.11 ±0.02	4.52 ±0.12	13.0%	57.5%	0.0746
or + pi0mass < 160	6.25 ±0.02	4.90 ±0.13	13.3%	56.1%	0.0745

Table 4.4: Intrinsic ν_e expected at Super-K in range 1–9.95 GeV scaled to a nominal year of data taking (1.66×10^{21} POT, 30 GeV beam, 750 kW) which pass fully contained fiducial volume single ring e-like selection cuts (labelled as FCFV 1R e-like) and a pi0mass cut. Each pi0mass cut is tested separately. Background due to beam $\nu_\mu + \bar{\nu_\mu}$ passing the selection cuts is given for comparison. Statistical uncertainties on the Monte Carlo events selected are given. Efficiency ϵ and purity π are given, and the cut which results in the maximum $\epsilon.\pi$ is in bold.

a very severe cut on the *pi0mass*, which also removes approximately 50% of the intrinsic ν_e. While the fractional reduction in background events is far higher than the fractional reduction in selected ν_e, the number of intrinsic ν_e selected is already very low without any additional cuts, so any further reduction is to be avoided. To investigate the effect of various *pi0mass* cuts on the selection efficiency and the purity of the selected sample, a range of *pi0mass* cuts are tested. The results are listed in Table 4.4.

While all of these cuts result in a higher purity, the number of intrinsic ν_e selected drops significantly in each case, and the higher the purity, the lower the efficiency. The best value of the *pi0mass* to use as the boundary for selection is 90 MeV, as this cut achieves the best balance of increased purity of the sample and reasonable cut efficiency. This is demonstrated by the value of the efficiency multiplied by the purity being a maximum when a cut of *pi0mass*<90 is applied. However, the best value of $\epsilon \times \pi$ that can be achieved with a cut on the *pi0mass* is 0.0794, which is lower than the value of 0.112 which corresponds to the selection without a *pi0mass* cut. Therefore while selecting events with a *pi0mass* in a particular range does remove background events, the cut is not beneficial to the selection of the sample when the effect on the selection efficiency is taken into account.

4.4.3 Angle to beam cut

Another possible method of improving the purity of the selected sample is to select events based on the trajectory of the products. The angle θ used in Eq. 4.10 refers to the angle between the direction of the incident neutrino, and the direction of the lepton produced in the interaction. The neutrino direction is equal to the direction of the beam, which in the Super-K coordinate system is given as (0.5486540,-0.835246,0.0366437) [89].

The lepton direction is measured using the angle at which the light of the Cherenkov ring hits the detector wall, which can be determined from the order in which the PMTs are hit and the shape of the ring. In the case of a charged current ν_e interaction, the electron is most likely to continue in a direction close to the original neutrino direction. In contrast, we would expect θ to be larger in the case of a mis-identified neutral current event. A π^0 produced via a neutral current interaction will be emitted at some angle to the neutrino direction, and the two γ produced by the π^0 decay will also be at angles to the π^0 direction. If one of these γ generates a Cherenkov ring and is detected and identified as e-like, the reconstructed 'lepton' angle associated with the ring is more likely to be further from the original neutrino direction, resulting in a bigger value of θ. Selecting events with values of $cos\theta$ close to 1 would therefore select the most forward-going events and may increase the purity of the sample.

The effect of a cut on the value of $cos\theta$ is tested and the purities achieved with cuts of different severity are compared. The results are given in Table 4.5.

The results in Table 4.5 show that this cut does lead to a small improvement in the purity of the sample, with the biggest improvements achieved when the most severe cuts are applied. However, the reduction in the selection efficiency is more significant. None of the values of efficiency multiplied by purity obtained exceed the efficiency$\times$purity of 0.112 obtained without an additional cut on the $cos\theta$ value. Therefore a cut on the $cos\theta$ value does not improve the sample and will not be used.

4.5 Evaluation of intrinsic ν_e measurement

A direct measurement of the level of intrinsic ν_e contamination in the ν_μ beam would be advantageous, as it would allow us to confirm the validity of the models used to predict the backgrounds and quantify their precision. The measurement described in this chapter is in principle an appropriate way of measuring the intrinsic ν_e content present in the beam.

The interaction rate in any detector is the product of the particle flux and interaction cross-section. Since the ν_e flux at high energy is not affected by oscillations, it is consistent at the near and far detectors (other than the expected drop in intensity with distance). The relevant cross-sections at ND280 and Super-

Selection cut	True ν_e	True ν_μ + $\bar{\nu_\mu}$	ϵ	π	$\epsilon.\pi$
None	72.63 ±0.07±	2261.81 ±2.83	-	-	-
wallv>200	47.04 ±0.06	1465.52 ±2.28	100%	3.11%	
FCFV 1R e-like	17.68 ±0.04	33.82 ±0.34	37.6%	34.3%	0.129
+ $E_{recon} > 1000$	10.82 ±0.03	11.30 ±0.20	23.0%	48.9%	0.112
+ cosθ > 0.99	1.42 ±0.01	1.34 ±0.07	3.01%	51.4%	0.0155
or + cosθ > 0.97	3.54 ±0.02	3.35 ±0.11	7.53%	51.4%	0.0387
or + cosθ > 0.95	5.04 ±0.02	4.58 ±0.13	10.7%	52.4%	0.0561
or **+ cosθ > 0.93**	6.08 ±0.02	5.47 ±0.14	12.9%	**52.6%**	0.0679
or + cosθ > 0.91	6.81 ±0.02	6.16 ±0.15	14.5%	52.5%	0.0760
or + cosθ > 0.89	7.38 ±0.02	6.71 ±0.15	15.7%	52.4%	0.0822
or + cosθ > 0.87	7.82 ±0.02	7.17 ±0.16	16.6%	52.2%	0.0868
or + cosθ > 0.85	8.18 ±0.03	7.66 ±0.16	17.4%	51.7%	0.0898
or + cosθ > 0.83	8.47 ±0.03	7.99 ±0.17	18.0%	51.5%	0.0926
or + cosθ > 0.81	8.71 ±0.03	8.34 ±0.17	18.5%	51.1%	0.0946
or + cosθ > 0.79	8.92 ±0.03	8.58 ±0.17	19.0%	51.0%	0.0966
or + cosθ > 0.77	9.09 ±0.03	8.78 ±0.18	19.3%	50.9%	0.0983
or + cosθ > 0.75	9.24 ±0.03	9.00 ±0.18	20.0%	50.7%	0.0995
or + cosθ > 0.70	9.53 ±0.03	9.39 ±0.18	20.3%	50.4%	0.1020
or + cosθ > 0.60	9.91 ±0.03	9.93 ±0.19	21.1%	50.1%	0.1053
or + cosθ > 0.50	10.14 ±0.03	10.30 ±0.19	21.6%	49.6%	0.1069
or + cosθ > 0.40	10.30 ±0.03	10.55 ±0.19	21.9%	49.4%	0.1081
or + cosθ > 0.30	10.40 ±0.03	10.55 ±0.19	22.1%	49.2%	0.1088
or + cosθ > 0.20	10.48 ±0.03	10.87 ±0.19	22.3%	49.1%	0.1094
or + cosθ > 0.10	10.54 ±0.03	10.97 ±0.19	22.4%	49.0%	0.1098
or + cosθ > 0.01	10.59 ±0.03	11.02 ±0.20	22.5%	49.0%	0.1103

Table 4.5: Intrinsic ν_e expected at Super-K in range 1–10 GeV scaled to a nominal year of data taking ($1.66{\times}10^{21}$ POT, 30 GeV beam, 750 kW) which pass fully contained fiducial volume single ring e-like selection (labelled as FCFV 1R e-like) plus cosθ cuts. Each cut is applied separately. Background due to beam ν_μ + $\bar{\nu_\mu}$ passing the selection cuts is given for comparison. Statistical uncertainties on the Monte Carlo events selected are provided. Efficiency ϵ and purity π are given, and the cut which results in the maximum π is in bold.

K differ, since the target materials are scintillator and water respectively. With the development of the near detector particle identification, comparisons of the high energy intrinsic ν_e interaction rates at ND280 and Super-K could also provide indications of errors on the cross-section values used.

Using Monte Carlo data we have estimated the number of high energy intrinsic ν_e that would be selected in a nominal year of data taking, corresponding to 1.66×10^{21} POT. In this time we expect to select 10.82 ± 3.29 intrinsic ν_e. We would also expect a background of 11.30 ± 3.36 events due to misidentified interactions of ν_μ, which corresponds to a purity of 48.9%. These values are given in Table 4.2. The two additional cuts that are investigated to reduce the number of background events selected do improve the purity; however, the reductions in efficiency that they also cause are too severe. When the already low event rate is taken into consideration and the combined efficiencies and purities are compared, the additional cuts are not found to be beneficial to the selection.

The probability of $\nu_\mu \to \nu_e$ oscillation is a maximum at 600 MeV, so by measuring ν_e detected above this energy range we can avoid detecting many signal ν_e and measure only the intrinsic ν_e of interest. 1 GeV was chosen as a reasonable preliminary cut value, as this value is considerably above the energy of the oscillation maximum while also as low as possible to maximise the event rate. Since conducting this study using Monte Carlo, the first ν_e appearance measurements have been conducted and the current best measurement of θ_{13} is found to be $\sin^2(2\theta_{13}) = 0.088^{+0.049}_{-0.039}$ [35]. This is at the upper end of the estimated range for θ_{13}, and will result in numbers of signal ν_e at the higher end of the expected range at the time T2K started. This would lower the purity of the selected sample further.

Therefore while this measurement is useful in principle, in practice the effectiveness of the method is limited by the low flux of neutrinos in the high energy tail. The statistical error would lead to large uncertainties on any calculations conducted using the ν_e measurement, such as calculating the kaon contribution to the beam, and would not be an improvement on the current levels of uncertainty.

This method will be statistically limited when using data corresponding to 1 nominal year, however with additional data it would become viable. With many additional data sets combined, this method may provide a useful assessment of the models used to predict backgrounds. However, gathering sufficient data will require many additional years of running, and in that time other methods of background determination may be developed. At the time of writing, the total data collected by T2K during runs 1+2+3+4 corresponds to a beam exposure of 6.57×10^{20} POT [90] (see Appendix B), and work is ongoing at J-PARC to increase the intensity of the neutrino beam. Measuring the intrinsic ν_e at the near detector would be preferable, since the entire energy range could be included without the possibility of including signal ν_e. With the development of a reliable particle identification system, an ND280 background measurement would be superior. Therefore this method is not developed further in this thesis.

Chapter 5

Measurement of high energy ν_μ

Here we discuss the dependence of the intrinsic ν_e flux prediction on the hadron production and selection models. The parents of the ν_μ and $\bar{\nu_\mu}$ are described and we see that high energy $\nu_\mu/\bar{\nu_\mu}$ are produced exclusively by kaons. The possibility of using a measurement of high energy $\nu_\mu+\bar{\nu_\mu}$ at Super-K to test the kaon production and selection models is evaluated, and the limitations are discussed.

5.1 Uses of high energy ν_μ detected at Super-K

Accurately predicting the intrinsic ν_e background is vital when conducting a ν_e appearance experiment. This is currently performed using models and limited measurements of the fluxes of kaons and pions produced at the target and entering the decay pipe (see §6.2). The decay processes that result in the production of intrinsic ν_e of different energies are described fully in §4.1. Uncertainty on the quantities of each intrinsic ν_e parent present in the decay pipe will result in uncertainty on the intrinsic ν_e content of the beam.

A direct measurement of the intrinsic ν_e interactions could be compared to the predicted number of interactions in order to test the accuracy of the models. However, the design of the T2K experiment aims to minimise backgrounds, including the intrinsic ν_e contamination. The resulting low flux of intrinsic ν_e is

Parent	Decay	Branching Ratio
π^+	$\pi^+ \rightarrow \mu^+ + \nu_\mu$	99.99%
μ^-	$\mu^- \rightarrow e^- + \bar{\nu_e} + \nu_\mu$	100%
K^+	$K^+ \rightarrow \mu^+ + \nu_\mu$	63.55%
K^+	$K^+ \rightarrow \mu^+ + \nu_\mu + \pi^0$	3.35%
K^0_L	$K^0 \rightarrow \pi^- + \mu^+ + \nu_\mu$	27.04%

Table 5.1: Decays contributing to beam ν_μ. Only decays resulting in ν_μ with branching ratios of greater than 1% are listed. Mesons are produced in the collision of protons on the graphite target. The muon parents in this table are decay products which decay before being removed by the muon beam dump.

advantageous to the ν_e appearance measurement, but results in a very low intrinsic ν_e interaction rate in the high energy tail. Chapter 4 demonstrates that the high level of background events caused by interactions of ν_μ and the large statistical uncertainty on the small number of selected ν_e prevent a direct measurement of the intrinsic ν_e from being of use.

The intrinsic ν_e present in the beam are merely unwanted contamination, and therefore low in number. However, the beam is designed to maximise the flux of ν_μ to perform the ν_μ disappearance analysis. The rate of ν_μ, and $\bar{\nu_\mu}$, interactions at the far detector is therefore considerably higher than the intrinsic ν_e interaction rate.

T2K uses the production and subsequent decay of π^+ mesons as the primary source of its beam ν_μ, via

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \quad (99.99\%) \tag{5.1}$$

However, while this process is the source of the majority of the beam ν_μ, especially in the signal region, other decay processes also contribute to the beam. These are listed in Table 5.1.

The contributions to the total ν_μ flux spectrum from each of these processes are shown in Figure 5.1, and the $\bar{\nu_\mu}$ spectrum is also shown in Figure 5.2. These plots show the total predicted neutrino flux at the far detector, and also the predicted individual fluxes from the different parents.

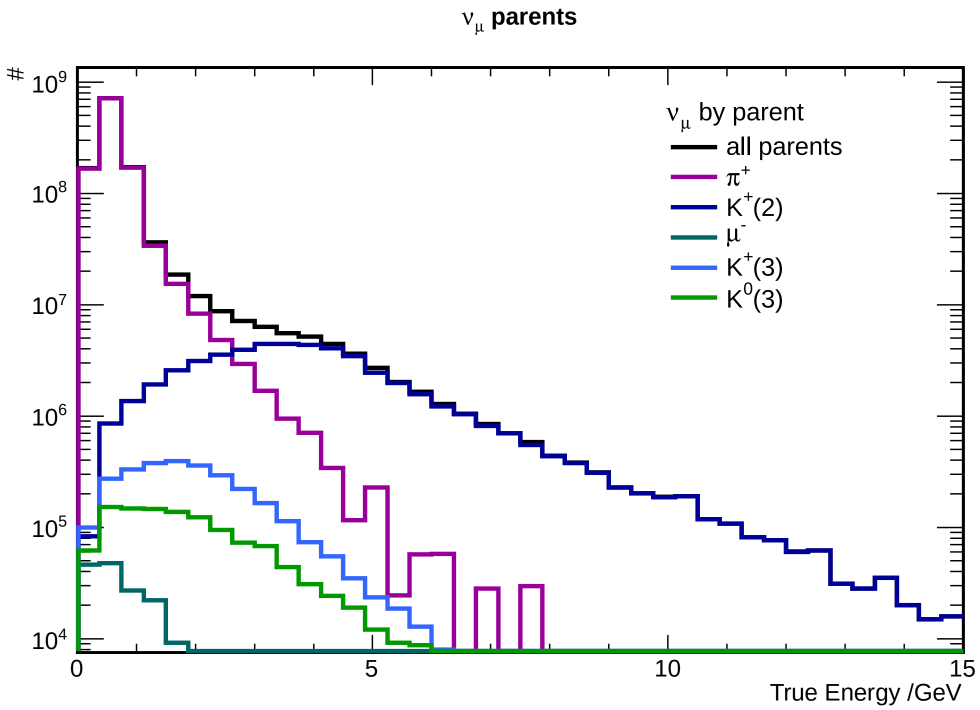

Figure 5.1: Predicted flux of ν_μ at Super-K divided into contributions from different parents. The total flux given in black. A log scale is used on y-axis to allow the smaller contributions to the flux to be seen. The numbers in brackets refer to the number of products of the decay mode, and allow identification of the corresponding decay process listed in Table 5.1.

As we can see from the plots, the dominant source of neutrinos in the signal region is the decay of pions. T2K is designed to primarily produce a ν_μ beam, and the ν_μ flux is considerably higher than the $\bar{\nu_\mu}$ flux, so we will focus on Figure 5.1. Muon decays make a small contribution to the beam at low energies, however this is several orders of magnitude smaller than the pion contribution, so we can neglect it for this analysis. The kaon decay contributions however are significant. The 2-body decay of the K^+ is the most significant, and this becomes the dominant source of ν_μ at around 2.5 GeV, and the only source above approximately 6 GeV.

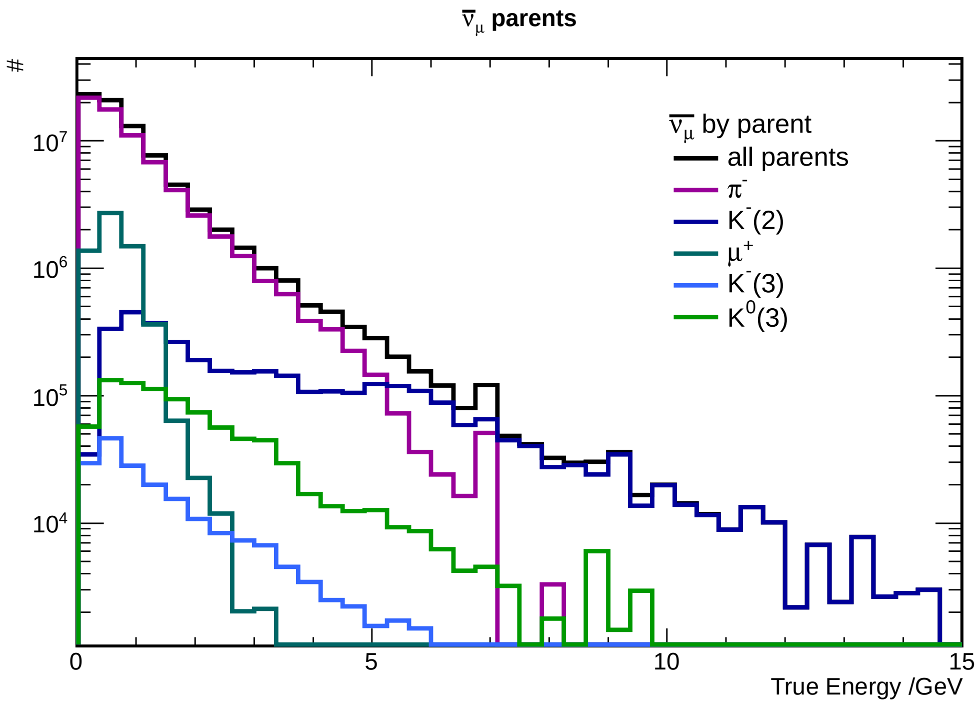

Figure 5.2: Flux of $\bar{\nu_\mu}$ at Super-K divided into contributions from different parents. The total flux given in black. A log scale is used on the y-axis to allow smaller contributions to the flux to be seen. The numbers in brackets refer to the number of products of the decay mode, and allow us to identify the (opposite charge) decay process listed in Table 5.1.

It is therefore clear, according to our models of hadron production on a carbon target and their subsequent decays, that kaons are the source of the ν_μ that form the high energy tail.

The K^+ has several hadronic decay modes which result in the production of additional π^+. Some of these π^+ may also decay to produce ν_μ before reaching the end of the decay pipe. These are labelled as π^+ contributions in Figure 5.1 since the π^+ is the direct parent. However, this demonstrates that kaons are also responsible for a fraction of the neutrinos attributed to π^+ decay, and so further emphasises their impact on the neutrino beam.

Since the high energy ν_μ are produced by various kaon decays, an accurate prediction of the high energy tail events will require correct models of the flux of each meson type present in the decay pipe. A measurement of the high energy ν_μ and $\bar{\nu_\mu}$ interactions can therefore be compared to the predicted ν_μ and $\bar{\nu_\mu}$ interactions in the same energy range in order to test the accuracy of the hadron production and selection models currently used. Should the measurement of the $\nu_\mu + \bar{\nu_\mu}$ event rate differ from the prediction, this would indicate an error in the hadron production and selection models leading to an error in the predicted neutrino flux. This information will also constrain the intrinsic ν_e prediction, since the level of intrinsic ν_e contamination is also dependent on the kaon decays occurring in the decay pipe.

5.2 ν_μ selection at Super-K

In §4.3 the standard selection cuts used to select charged current ν_e interactions at Super-K were described. A similar set of selection cuts exist to select ν_μ interactions [26]. This set of selection cuts, along with explanations of the differences required to identify μ-like events instead of e-like events, are described below.

1. wall>200

Fiducial volume. 'Wall' gives the shortest distance between the reconstructed vertex position and the inner wall. The fiducial volume definition is the same regardless of particle type, so this cut is the same for e-like and μ-like selections.

2. nhitac<16

Outer hits. For an event to be classed as fully contained, there must be no hit cluster in the outer detector with more than 16 PMT hits. This veto is the same for e-like and μ-like selections.

3. evis>30

Visible energy. The visible energy is the sum of energy from all rings associated with an interaction. A minimum of 30 MeV is required in order to reject events from radioactive decays occurring inside the detector[1].

4. nring=1

Number of rings. We require CCQE interactions of ν_μ or $\bar{\nu_\mu}$ only. The only visible product of a CCQE interaction is one charged lepton, which will produce one ring. In order to accurately reconstruct the energy of the μ^- or μ^+ produced, we require leptons that come to a stop inside the detector, therefore only rings are accepted.

5. ip=3

Particle identity. The particle identification algorithm finds the particle type that matches the properties of the Cherenkov ring most closely and returns a numerical value. 3 corresponds to μ-like[2].

6. nummuedecay≤1

[1] This cut differs from the ν_e selection cut as events including a decay electron are allowed in this case. See Cut 6.

[2] A 0.3% mis-ID rate was estimated, from [79]

0 or 1 decay electrons. For every event, we search for the presence of a secondary event occurring a short time later. A muon produced via a CCQE interaction may decay while inside the inner detector, therefore the presence of one delayed electron event is allowed. It is also possible that, after emitting Cherenkov radiation and producing a ring, muon capture occurs, removing the muon before it can decay to produce an electron. Another possibility is that having dropped below the Cherenkov threshold while in the inner detector and therefore passed the fully contained condition, the muon leaves the inner detector before decaying, producing a decay electron outside the inner detector. Zero detected decay electrons therefore also fit the CCQE $\nu_\mu/\bar{\nu_\mu}$ profile.

Using Super-K Monte Carlo data corresponding to 1.66×10^{21} POT, we find the expected number of ν_μ and $\bar{\nu_\mu}$ that would be selected, assuming that the current hadron production and selection models and interaction cross-sections are correct, and using the oscillation parameters given in §4.3. The number of true intrinsic ν_e that would also pass the μ-like selection cuts and therefore act as a background are also included. The results are listed in Table 5.2[3]

A ν_μ interaction produces a μ^-, and a $\bar{\nu_\mu}$ interaction produces a μ^+. The selection cuts described apply equally to μ^- and μ^+ Cherenkov rings and both event types are classified as μ-like. Therefore ν_μ and $\bar{\nu_\mu}$ interactions are considered together.

Unlike the intrinsic ν_e selection studied in Chapter 4, the $\nu_\mu/\bar{\nu_\mu}$ selection does not have any significant sources of background. The probability that a ν_e may be misidentified as a ν_μ is approximately equal to the probability that a ν_μ is

[3]Since we are selecting μ-like events, the neutrino energies are reconstructed assuming a muon is detected. Therefore the reconstructed muon mass, energy and momentum are used in Equation 4.10. This results in slightly different numbers of events in the $E_{recon} < 9.95$ GeV range compared to Table 4.1 prior to any cuts being made.

Selection cut	True $\nu_\mu + \bar{\nu_\mu}$	True ν_e	ϵ	π	$\epsilon.\pi$
None	2266.41 $\pm$2.83	72.30 $\pm$0.07	-	-	-
wallv>200	1468.69 $\pm$2.28	46.85 $\pm$0.06	100%	96.9%	0.969
wall>200	1320.39 $\pm$2.16	45.60 $\pm$0.06	89.9%	96.7%	0.869
+ nhitac<16	1160.63 $\pm$2.03	44.24 $\pm$0.06	79.0%	96.3%	0.761
+ evis>30	765.05 $\pm$1.64	38.97 $\pm$0.05	52.1%	95.2%	0.496
+ nring=1	387.85 $\pm$1.16	21.83 $\pm$0.04	26.4%	94.7%	0.250
+ ip=3	336.00 $\pm$1.08	0.79 $\pm$0.01	22.9%	99.8%	0.228
+ nummuedecay$\leq$1	312.26 $\pm$1.04	0.77 $\pm$0.01	21.3%	99.8%	0.212

Table 5.2: ν_μ expected at Super-K scaled to a nominal year of data taking (1.66×10^{21} POT, 30 GeV beam, 750 kW) which pass fully contained fiducial volume single ring μ-like selection. Cuts are cumulative as of wall>200. Background due to ν_e passing the selection cuts is given for comparison. ϵ gives the efficiency compared to the true fiducial volume selection, and π describes the purity of the selection. Statistical uncertainties on the Monte Carlo events selected are provided.

misidentified as a ν_e; however, the ν_e flux is far lower than the ν_μ flux, so this results in very few false ν_μ events. This lack of background and the high event rate for the ν_μ+$\bar{\nu_\mu}$ sample both indicate that a ν_μ measurement would be a more effective method of verifying the kaon content models than a direct intrinsic ν_e measurement as attempted in Chapter 4.

The numbers in Table 5.2 include selected neutrinos of energy 0–9.95 GeV. The need for a high energy selection cut and the optimum placement for this cut are discussed below.

5.3 Placement of energy cut

To remove additional sources of uncertainty, an energy range that does not contain neutrino oscillations is preferable. If ν_μ and $\bar{\nu_\mu}$ are counted in a region in which ν_μ disappearance is possible, the predicted interaction rates will have to include an estimation of the oscillations. Since the oscillation parameters are not known exactly, there will be an associated uncertainty on the predicted number of ν_μ that would disappear and therefore not be present in the ν_μ flux at Super-K. This introduces an additional source of uncertainty on the predicted event rate. Should there then be a discrepancy between the predicted event rate and observed event rate, it would not be possible to distinguish between uncertainty in the flux and cross-section models used to generate the event rate prediction, and the uncertainty on the oscillation parameters. It is also possible that such an uncertainty could mask the presence of an error in the flux or cross-section models.

The ν_μ survival probability at 295 km is shown in Figure 5.3. The black curve shows the oscillation probability using the best fit values of θ_{23} and Δm^2_{23}, which are given as $\sin^2 2\theta_{23} = 1.0$ and $\Delta m^2_{23} = 0.00243$ eV^2 in the 2011 Particle Data Group [91] release. The uncertainties on these values are then taken into consideration. The uncertainties are taken from $\sin^2 2\theta_{23} > 0.92$ and $\Delta m^2_{23} = (2.43 \pm 0.13) \mathrm{x} 10^{-3}$ eV^2. These uncertainties are combined to produce the maximum survival probabilities possible, and this curve is plotted in red. The uncertainties are then combined to produce the lowest survival probabilities, with the resulting curve plotted in blue. These curves demonstrate the range of possible survival probabilities according to the most precise measurements available in 2011.

In the ν_e-appearance signal region at approximately 600 MeV, the survival probability shows the biggest uncertainty, ranging from 0 to approximately 0.15. As the energy of the neutrino increases, the ν_μ survival probability increases and the range of possible oscillation probabilities decreases. Therefore the higher the energy cut, the smaller the uncertainty due to the oscillation parameters. However, a low minimum energy requirement would allow more interactions to be selected, and would therefore reduce the statistical error. Therefore to choose the best placement for the energy cut, we must balance the need to minimise the statistical errors and to minimise the uncertainty due to the precision of the oscillation

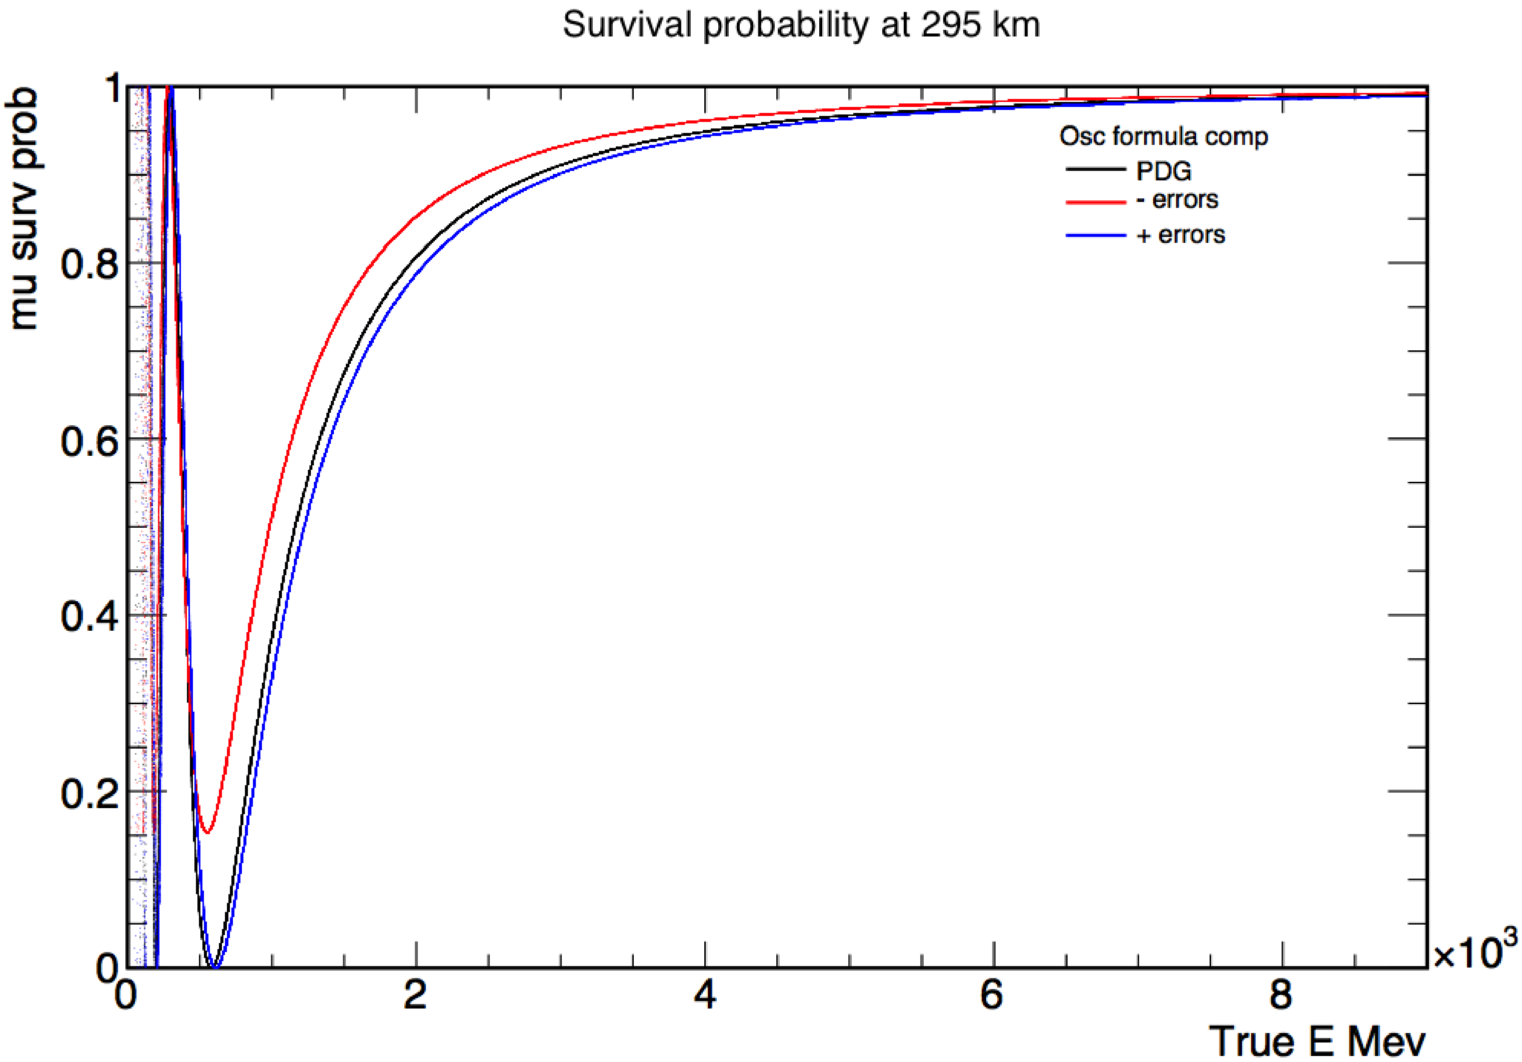

Figure 5.3: Comparison of ν_μ survival probability curves at 295 km versus energy. The probability calculated using the 2011 Particle Data Group best fit values [91] is given in black. The range of possible probabilities according to the known uncertainties on the oscillation parameters are given in blue and red.

parameters.

As in the case of the ν_e selection, the energy of each neutrino interacting in Super-K is reconstructed using Equation 4.10. However, for this measurement we ultimately select only μ-like neutrinos, and reject any e-like neutrinos. Therefore it is appropriate to use the lepton energy and momentum reconstructions calculated using the assumption that the lepton is a muon[4]. The neutrino energy is then found using [79]

$$E_\nu = \frac{2E_\mu(m_n - V) - m_\mu^2 + 2m_n V - V^2 + m_p^2 - m_n^2}{2(m_n - V - E_\mu + P_\mu \cos\theta)} \tag{5.2}$$

As before, E_μ and P_μ are the reconstructed energy and momentum of the muon respectively. The angle between the beam and lepton directions is cosθ. m_p and m_n are the masses of the proton and neutron respectively, and m_μ is the mass of the muon produced in the charged current interaction. V represents the binding energy of a nucleon in a ^{16}O nucleus and is set equal to 27 MeV.

The reconstructed energy of each interacting neutrino is calculated, and a high energy tail cut can then be applied.

7. E$_{recon}$ >$\boldsymbol{x}$ MeV

High energy tail cut. Neutrinos must have a minimum reconstructed energy in order to be selected.

The optimum energy cut will result in the lowest combined uncertainty. To find this optimum cut, we find the combined uncertainty for each possible energy cut value. A value of x is chosen and the predicted ν_μ+$\bar{\nu_\mu}$ passing all 7 cuts are

[4]The reconstructed neutrino energy will differ slightly if the lepton flavour is changed, since the lepton mass will change, along with the reconstructed energy and momentum of that lepton. This changes the number of neutrino interactions in the range 0-9.95 GeV found before e-like or μ-like selections are applied.

found, using the best fit oscillation parameters to predict the level of ν_μ+$\bar{\nu_\mu}$ disappearance. The oscillation parameters are then changed to give the maximum ν_μ and $\bar{\nu_\mu}$ survival and the predicted number is found again. The oscillation parameters are then set to correspond to the minimum ν_μ and $\bar{\nu_\mu}$ survival and the predicted ν_μ+$\bar{\nu_\mu}$ passing the cuts is found once more. These maximum and minimum predicted interaction rates are then used to find the percentage uncertainty on the ν_μ+$\bar{\nu_\mu}$ prediction for that energy cut. This uncertainty is combined with the statistical uncertainty on the ν_μ+$\bar{\nu_\mu}$ number, and the result is plotted against the value of the energy cut used. This process is repeated for all potential energy cut values, and the results are plotted in Figure 5.4.

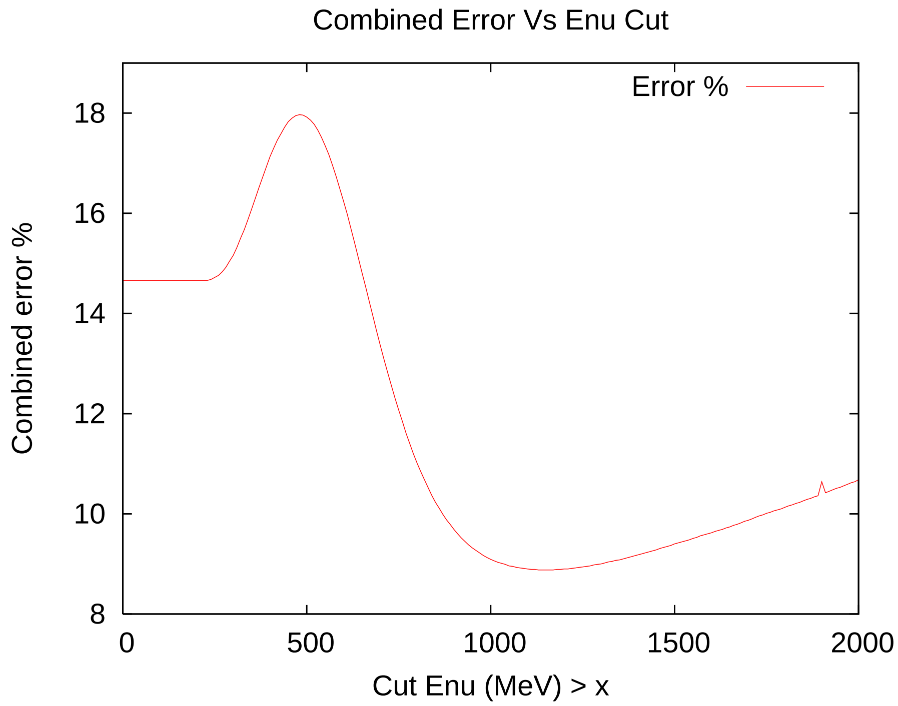

Figure 5.4: Uncertainty on number of predicted ν_μ+$\bar{\nu_\mu}$ due to current uncertainties on oscillation parameters and statistical uncertainty on sample size combined for different high energy tail cuts. Placement of cut given on x axis, where events satisfying $\mathrm{E}_{recon} > x$ are selected. Combined uncertainty is a minimum, at 8.9%, for an energy cut at 1150 MeV.

The combined uncertainty is a maximum at the energy of the oscillation maximum, where the oscillation parameter uncertainty dominates. The statistical uncertainty dominates otherwise, and becomes increasingly significant as the sample size decreases. The combined uncertainty is found to be a minimum when an energy cut of E_{recon} > 1150 MeV is used. This is a relatively low placement, as some oscillations will still occur at this energy. This shows that the sample sizes are sufficiently small that despite the oscillation parameter uncertainties, it is preferable to use a low energy cut in order to maximise the sample size.

However, oscillation parameters are not the only consideration. The aim of this analysis is to assess the accuracy of the kaon contribution to the neutrino beam, in order to ensure that the intrinsic ν_e prediction is accurate. We aim to compare a measured $\nu_\mu+\bar{\nu_\mu}$ sample with the predicted $\nu_\mu+\bar{\nu_\mu}$ sample corresponding to the same POT. If the $\nu_\mu+\bar{\nu_\mu}$ we select are produced from the decay of kaons only, then should a difference between the predicted and observed samples be found after uncertainties have been taken into consideration, we can conclude that the modelled production of kaons at the target and subsequent decay pipe selection contains errors. In contrast, if both pions and kaons contribute to the $\nu_\mu+\bar{\nu_\mu}$ sample being considered, any discrepancy must be attributed to one or both of the kaon and pion fluxes. It would not be possible to determine whether the source of the discrepancy was a result of an error in the model of the kaon flux, or the pion flux, or a combination of the two. We therefore require a sample of $\nu_\mu+\bar{\nu_\mu}$ which are the result of kaons entering the decay pipe.

Considering Figure 5.1 and Figure 5.2, we see that the high energy tail cut must be severe in order to select a $\nu_\mu+\bar{\nu_\mu}$ sample with kaon parents. A cut of E_{recon} >5000 MeV appears to be the lowest energy cut that would ensure a negligible contribution to the neutrino beam from anything other than kaon decays.

The predicted number of $\nu_\mu+\bar{\nu_\mu}$ interactions selected with different energy cuts applied are listed in Table 5.3. While the fluxes of ν_μ and $\bar{\nu_\mu}$ in the high energy tail are higher than the intrinsic ν_e flux, they are still not sufficiently high to allow a

Selection cut	True ν_μ + $\bar{\nu_\mu}$	True ν_e	ϵ	π	$\epsilon.\pi$
wallv>200	1468.69 ±2.28	46.85 ±0.06	100%	96.91%	0.969
FCFV 1R μ-like	312.26 ±1.04	0.775 ±0.008	21.3%	99.75%	0.212
+ E_{recon} >1000	171.75 ±0.75	0.116 ±0.003	11.7%	99.93%	0.117
or + E_{recon} >1500	120.76 ±0.62	0.065 ±0.002	8.22%	99.95%	0.082
or + E_{recon} >2000	90.18 ±0.54	0.042 ±0.002	6.14%	99.95%	0.061
or + E_{recon} >2500	69.09 ±0.48	0.026 ±0.001	4.70%	99.96%	0.047
or + E_{recon} >3000	53.32 ±0.42	0.019 ±0.001	3.63%	99.96%	0.036
or + E_{recon} >3500	40.08 ±0.37	0.015 ±0.001	2.73%	99.96%	0.027
or + E_{recon} >4000	28.77 ±0.31	0.012 ±0.001	1.96%	99.96%	0.020
or + E_{recon} >4500	20.00 ±0.26	0.010 ±0.001	1.36%	99.95%	0.014
or + E_{recon} >5000	13.62 ±0.21	0.008 ±0.001	0.93%	99.94%	0.009

Table 5.3: ν_μ expected at Super-K scaled to a nominal year of data taking (1.66×10^{21} POT, 30 GeV beam, 750 kW) which pass fully contained fiducial volume single ring μ-like selection, labelled as FCFV 1R μ-like. Different high energy tail cuts are tested. Background due to ν_e passing the selection cuts is given for comparison. ϵ gives the efficiency compared to the true fiducial volume selection, and π describes the purity of the selection. Statistical uncertainties on the Monte Carlo events selected are provided.

cut as severe as E_{recon} > 5000 MeV without resulting in very significant statistical uncertainties.

Therefore, while a high energy requirement provides better sensitivity to the kaon contribution to the neutrino beam, it results in a measurement which is statistically limited.

5.4 Evaluation of $\nu_\mu/\bar{\nu_\mu}$ measurement

In principle measuring the ν_μ+$\bar{\nu_\mu}$ interaction rate in the high energy part of the spectrum at the far detector and comparing this to the Monte Carlo prediction could indicate the presence of errors in the models used to generate the Super-K Monte Carlo and predict the intrinsic ν_e background. The uncertainty in the predicted interaction rate caused by the current levels of uncertainty on the relevant oscillation parameters would not be significant when compared to other sources of uncertainty, in particular, the statistical uncertainty on the sample size.

However, in order to achieve a sample of interactions of ν_μ+$\bar{\nu_\mu}$ that are exclusively products of kaon decays, a very high energy cut would be necessary. The percentage statistical uncertainty on the resulting small sample of ν_μ and $\bar{\nu_\mu}$ interactions would be comparable to any possible difference between the measured and predicted interaction rates[5], based on the current levels of uncertainty of kaon production in the models used. For example using a cut of $E_{recon} > 5$ GeV results in 13.6 selected events ± 3.7 events, corresponding to a percentage uncertainty of 27%. This would make it impossible to infer anything about the accuracy of the current kaon production models. The alternative is to use a much lower cut on the reconstructed energy, so that more events are selected. However, these events would be the products of both pion and kaon decays, and so it would not be possible to study the level of agreement between pion and kaon production models separately.

A larger sample of events would allow the high energy cut of approximately 5 GeV to be used without a large statistical uncertainty. With an energy cut of $E_{recon} > 5$ GeV we expect to detect 13.6 ν_μ and $\bar{\nu_\mu}$ interactions for every 1.66×10^{21} POT. Over an extended period of data taking this number would increase and the associated statistical uncertainties would no longer dominate.

While a large data set would reduce the dominance of the statistical uncertainty, there is one further limitation to this measurement which cannot be resolved by increasing the data taking period. To predict the interaction rate in a detector we require the flux of neutrinos arriving at the detector and also the relevant interaction cross-section values. We can observe the rate of charged current interactions of ν_μ and $\bar{\nu_\mu}$ in Super-K, and this tells us the product of the ν_μ and

[5] [92] states that the uncertainty on the number of expected ν_μ+$\bar{\nu_\mu}$ events at super-K is 8.3%.

$\bar{\nu_\mu}$ flux with the charged current cross-section on water, but we cannot separate these two factors. Therefore a comparison of the predicted and observed high energy ν_μ+$\bar{\nu_\mu}$ interaction rate at Super-K will indicate errors in the combination of the flux of neutrinos produced in kaon decays and the charged current interaction cross-sections used. It cannot, however, specifically indicate errors in the neutrino flux. For example should the observed number of events be lower than the predicted number (after consideration of known sources of uncertainty and statistical fluctuations), there are two possible causes. The flux of neutrinos produced as a result of kaon decays may be lower than the value currently assumed. It is also possible, however, that the charged current cross-section for interactions of neutrinos in the energy range of interest is lower than the value used in the Monte Carlo generation. This would also result in a lower observed event rate. Therefore the degree to which a measurement as described in this chapter can indicate the presence of errors in the kaon production models will always be limited by the current levels of uncertainty on the interaction cross-sections. To explicitly test the accuracy of the kaon production and selection models, and therefore the background estimations, a method must be devised to separate the neutrino flux and the interaction cross-section measurements. This is very challenging and is attempted in the remainder of this thesis.

Chapter 6

Neutrino parents

In this chapter we discuss how the production of π^0 via neutral current interactions of high energy ν_μ are responsible for the second source of background to the ν_e appearance measurement. We discuss the dependence of both background sources on the level of kaon production at the target. The difficulty in separating the beam flux and interaction cross-section uncertainties is described, and the possibility of using the near detector to measure the hadron production is introduced.

6.1 Importance of high energy neutrinos

As discussed in Chapter 3, T2K is designed to produce a neutrino spectrum with a narrow range of energies, peaking around the energy at which oscillations are most likely to occur. Producing a neutrino spectrum with a narrow peak by operating off-axis also minimises the number of high energy neutrinos produced. We have shown in Chapters 4 and 5, however, that a small high energy tail is still present, and these high energy neutrinos contribute to T2K's main backgrounds.

There are two main sources of background events for the T2K ν_e appearance measurement. Chapter 4 considered a direct measurement of the high energy intrinsic ν_e which contaminate the ν_μ beam. The other significant source of background is π^0 mesons, which can be mis-identified as ν_e and therefore mistaken for signal events at Super-K via the process described in §3.4. There are several methods of π^0 production possible. Figure 6.1 plots the number of true ν_μ modelled in Monte Carlo simulations according to the neutrino's type of interaction in Super-K.

The interaction types with products including a π^0 meson are listed below along with their corresponding NEUT codes. Each case involves a single π^0 produced from a delta resonance at 1232 MeV.

12 Charged current, via $\nu + n \rightarrow l + p + \pi^0$

31 Neutral current, via $\nu + n \rightarrow \nu + n + \pi^0$

32 Neutral current, via $\nu + p \rightarrow \nu + p + \pi^0$

There are two types of π^0-producing interaction. The charged current case (12) is slightly easier to identify, since a π^0 signal, seen with a charged lepton signal, is more likely to be classified correctly. Therefore π^0 produced via charged current interactions are most problematic if the lepton escapes undetected, as the π^0 could then be mistaken for the electron that would result from the charged current interaction of a signal ν_e.

π^0 produced in the neutral current interactions described above (NEUT codes 31,32) fit this scenario. The incoming neutrino scatters on a nucleon and leaves the detector again with no further indication of its presence. Changes to the target nucleon are not detected, so the only evidence that an interaction occurred is the presence of the π^0 produced. Since a π^0 may be mistaken for an electron in the far detector (see §4.4.2), this interaction can appear to have a single electron as the interaction product, which matches the outcome of a charged current quasi-elastic (CCQE) interaction of a ν_e, given by

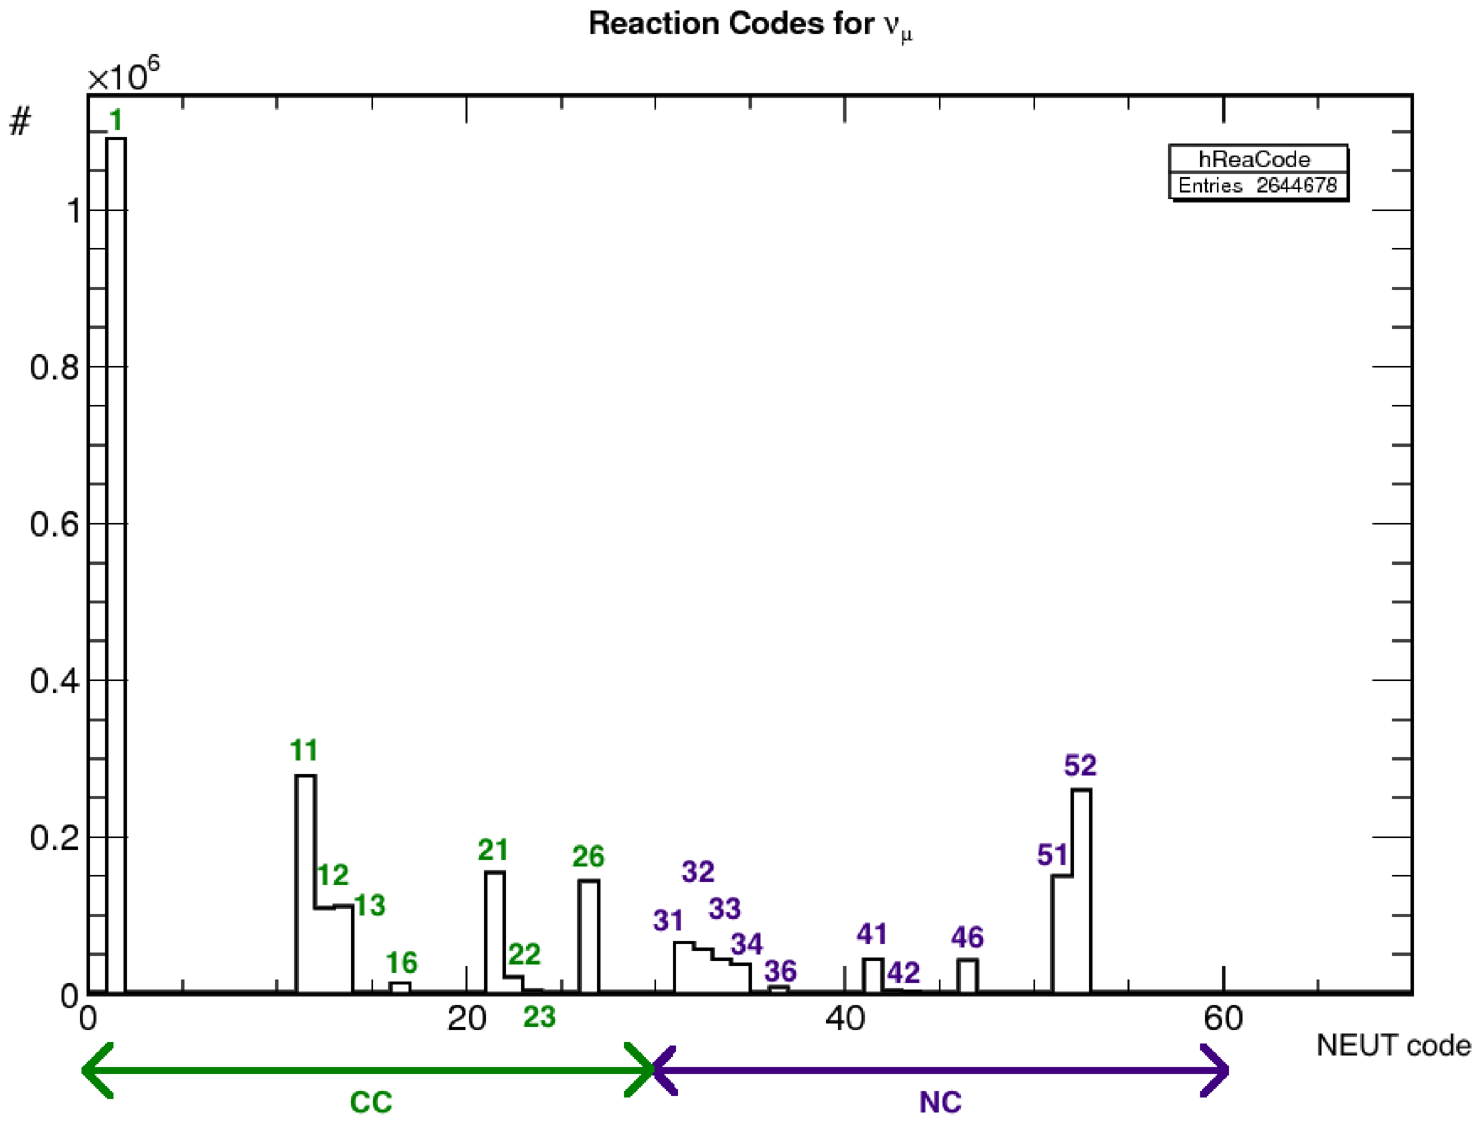

Figure 6.1: ν_μ interactions at Super-K plotted by NEUT reaction code [93] (see Appendix C for full list). Charged current reactions are (CC) given by numbers 1–30, and neutral current reactions are assigned numbers 31–60.

$$\nu_e + n \rightarrow p + e^- \tag{6.1}$$

When a CCQE interaction occurs, we can use conservation of energy, knowledge of the nucleus and measurements of the resulting lepton to obtain a reasonable estimate of the energy of the incident neutrino. In neutral current interactions however, the neutrino survives and leaves the detector with an unknown percentage of its original energy. It is possible for a high energy neutrino to interact via the neutral current and transfer only a small percentage of its energy. In this way π^0 produced in interactions of high energy ν_μ can be created with energies in the signal region, and therefore be mistaken for signal ν_e.

Since high energy neutrinos are the source of this background, to begin to predict the number of neutral current interactions of this type that will occur in a data sample, and thus the number of background events that would be produced, we must know the flux of high energy neutrinos.

6.2 Kaon production uncertainty

Predicting the rate of neutral current interactions at high energies requires two pieces of information: the neutral current cross-sections at high energies, and the flux of neutrinos. With these values, the number of expected neutral current interactions of ν_μ with sufficient energy to produce a π^0 could be calculated, and the number of π^0 that would be present in the detectors could be predicted. The neutral current cross-sections are currently not well known, but work is ongoing to improve the precision of the cross-section measurements. Once these cross-sections are well known, the neutrino flux would be required in order to calculate the backgrounds.

Quantifying the neutrino flux presents some challenges. One method is to predict the neutrino flux using the fluxes of the neutrino parents in the beam decay tunnel. The branching ratios of the decays which contribute to the neutrino beam are known, so if the spectra of kaons and pions produced at the target are well known, the spectra of neutrinos from each source can be predicted. One of the aims of the NA61/SHINE experiment [87] is to study the hadron production resulting from collisions between particles such as protons with fixed targets and data from NA61 taken using a replica of the T2K target is used by T2K to model the hadron production at the target.

The NA61 measurements are designed to cover most of the range of secondary particle energies and angles relevant to T2K. The pion production results, with uncertainties at the 5-10% level [34], have been made available to T2K. The pion production uncertainty outside the experimentally measured range is higher, at 50%. The NA61 kaon analysis is not yet complete[1] and so the current kaon production uncertainties included by T2K range from 15% to 100% across the range of kaon energies [34]. Without knowing the precise number of kaons and pions produced at each energy, we cannot apply the branching ratios and calculate how many times each decay process will occur, and therefore what decay products, including neutrinos, will be created[2]. These uncertainties on the contributions made to the neutrino beam by the different parents therefore limit our ability to predict the neutrino flux. As stated in [94], the total uncertainty for the neutrino flux prediction in the relevant energy range is evaluated to be 10-15%, where this is dominated by hadron production uncertainties.

We saw in Chapter 5 that it is possible to detect and reconstruct ν_μ interactions occurring in the high energy tail. However, while both Super-K and ND280 mea-

[1]Preliminary results were made available to T2K in September 2013.

[2]A direct measurement of the neutrino flux at T2K would therefore be a powerful cross check for the completed NA61 measurements.

sure the rate of neutrino interactions, they cannot tell us the neutrino flux directly. This is because the detectors can only detect the presence of a neutrino by observing the results of the neutrino's interaction, and therefore our measurements represent the neutrino flux multiplied by the charged current interaction cross-section. At present there are large uncertainties on the values of the interaction cross-sections[3]. Therefore while we can measure the number of neutrinos interacting in our detectors, we cannot extract the neutrino flux from these measurements without the precision being limited by the uncertainty on the cross-section. Our limited knowledge of the cross-sections mean that we cannot use observations of high energy neutrino interactions to constrain the neutrino flux in the place of precise hadron production information.

According to our beamline Monte Carlo, which contains models of hadron production and decay, we find that kaon decays are the source of the high energy tail. This is demonstrated in Figures 5.1 and 5.2, which show the ν_μ and $\bar{\nu_\mu}$ flux contributions used to generate the far detector Monte Carlo. The large uncertainty on the number of kaons produced at the target results in a large uncertainty on the flux of high energy ν_μ. Therefore, even with precise values of the neutral current cross-sections at high energy, the uncertainty on the neutrino flux limits our ability to predict the number of π^0 that may be produced in neutral current interactions.

Reducing the uncertainty on the level of kaon production at the target would help to constrain several sources of background. With a precise value for the number of kaons entering the decay pipe, the high energy neutrino flux can be more accurately predicted. Then as the precision on the neutral current cross-sections improve, the estimation of the π^0 generation will become more precise, reducing the uncertainty on that source of background.

[3][34] states that of a total uncertainty on the number of expected intrinsic ν_e at Super-K of 22.8%, cross-section uncertainties contribute 14.0%.

Improved knowledge of the relative rates of kaon and pion production would also allow better prediction of the intrinsic ν_e contamination. One source of intrinsic ν_e is the decay of kaons, so any reduction in the uncertainty on the kaon flux would also reduce the uncertainty on the number of intrinsic ν_e present in the beam. The other source of intrinsic ν_e is the decay of μ^+. The pion and muon lifetimes can be used to predict how many μ^+ will be created and subsequently decay before they can be captured at the end of the 98 m decay pipe. The main decay mode of the K^+ also results in μ^+ production. The likelihood that a μ^+ will decay before reaching the end of the decay pipe depends partly on the position of its creation, which is a function of the mean lifetime of the parent meson. π^+ and K^+ have different mean lifetimes, therefore they are not equally likely to contribute to the production of intrinsic ν_e via μ^+ decay. To predict the total number of intrinsic ν_e produced in this way, the ratio of π^+ and K^+ that enter the decay pipe must be known. Reduction of the uncertainties in the pion and kaon production rates at the target would therefore also help to predict the level of intrinsic ν_e contamination resulting from μ^+ decay.

6.3 Kaon measurement using ND280

Our knowledge of the production of pions and kaons at the target, and the contributions these hadrons make to the neutrino beam and also T2K's backgrounds, currently come from a combination of NA61 measurements and theoretical models. Using known branching ratios and decay kinematics, the composition of hadrons in the decay pipe that results in observed ν_μ fluxes from pions and kaons could be found. This hadron composition information could then be used to calculate the backgrounds. A measurement of the contributions to the ν_μ beam from pions and kaons using T2K data would therefore be beneficial to various T2K analyses.

Existing studies of the kaon contribution to the T2K beam include [95].

Chapters 4 and 5 confirm that the rate of interactions at the far detector is low, especially in the high energy tail. Since the majority of high energy neutrinos are produced via kaon decays, high energy events are very useful when making a measurement of the kaon contribution to the beam. The interaction rates at ND280 are considerably higher due to its proximity to the beam source. For this reason, near detector data is more suitable for a kaon:pion ratio measurement.

Some of the properties of the ν_μ detected in ND280 are described in Chapter 7, along with the event selection and reconstruction methods used. A method of measuring the numbers of ν_μ produced via kaon and pion decays, along with the interaction cross-sections at different energies, are then described in Chapter 8, and the results are presented in Chapter 9.

Chapter 7

Muon neutrino events in ND280

A comparison of the decay kinematics of pions and kaons is given and the differences in the neutrinos produced at T2K by different parents are discussed. An event selection for ν_μ interactions in FGD1 is described. We discuss the types of neutrino interaction that occur in FGD1 and appropriate methods of energy reconstruction for CCQE and CCnonQE events are developed. We establish that a change in the neutrino energy spectrum with changing off-axis angle is observed at ND280, and this is shown for both true and reconstructed energy.

7.1 Off-axis effects

We saw in Chapter 3 that for 2-body meson decays there is a relationship between the energy of the parent and the energy of the resulting neutrino which is a function of the opening angle of the neutrino. If we consider the example of the most common pion decay mode, $\pi^+ \rightarrow \mu^+ + \nu_\mu$, the neutrino energy is given by:

$$E_\nu = \frac{m_\pi^2 - m_\mu^2}{2(E_\pi - cos\theta\sqrt{E_\pi^2 - m_\pi^2})} \tag{7.1}$$

where θ is the angle between the neutrino and the parent pion, as shown in Figure 7.1. E_π and m_π are the energy and rest mass of the parent pion respectively, m_μ is the muon rest mass and E_ν is the energy of the produced neutrino. For a full derivation of this formula, see Appendix 1.

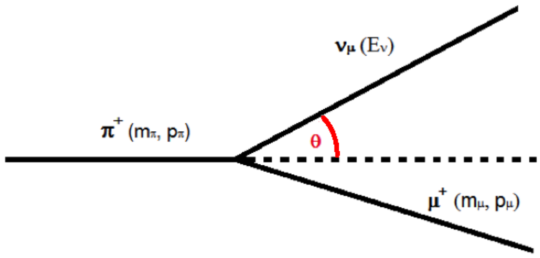

Figure 7.1: π decay showing neutrino opening angle θ

When on-axis ($\cos\theta = 1$) there is a linear relationship (see Appendix) between the pion and neutrino energies, so as the parent pion energy increases the neutrino energy increases also, with no limit. This results in a significant number of high energy neutrinos, which are less likely to oscillate and are therefore less useful to an oscillation experiment, while also contributing unwanted background events. However, as θ is increased this linear relationship disappears, and instead for a given opening angle there will be a specific maximum neutrino energy possible. This is demonstrated in Figure 7.2. For a particular θ, as the pion energy increases, the neutrino energy initially increases before peaking and then levelling out. It then stays constant or slightly decreases as the pion energy increases further. As you move to larger values of θ the value of the maximum possible neutrino energy decreases, and the pion energy at which the curve flattens out also decreases.

It is this behaviour that results in the shift of the peak neutrino energy with off-axis angle. The opening angle θ is approximately equal to the off-axis angle of the detector[1]. Figure 3.5 shows the total neutrino beam flux (from all parents combined) at various off-axis angles. The flux information from each ν_μ parent type individually, seen in Figure 5.1, shows that the largest contributor of beam

[1]Each parent meson may not be precisely on-axis when it decays, so the true off-axis angle for a ν should be a combination of the angle between the parent meson and the beam direction, ϕ, and the opening angle, θ. However the mean direction of the meson beam is on-axis, so we can approximate and say that the off-axis angle and the neutrino opening angle θ are equal.

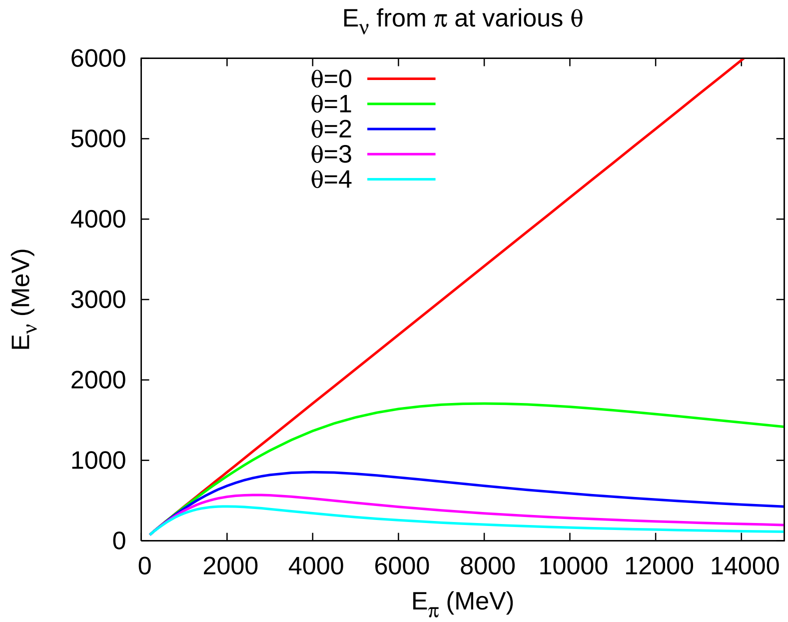

Figure 7.2: Relationship between ν_μ and parent π at different opening angles. θ measured in degrees.

neutrinos is π^+ decays, especially in the range 0-2 GeV. Therefore the movement with off-axis angle of the total beam flux peak, which sits in this range, can be attributed to the behaviour of ν from π^+ decays. Figure 3.5 shows that as you move further off-axis, the spectrum peaks at lower energies and the peak becomes narrower. This lowering of the peak position with increasing θ can be explained by Figure 7.2. The maximum neutrino energy decreases as the opening angle increases and this maximum energy occurs at lower parent pion energy. This reduces the energies of the high energy tail neutrinos and also lowers the energies of all neutrinos produced. The plateau in neutrino energies sits at a lower neutrino energy, so a higher fraction of the total neutrinos produced will have that lower plateau energy value, which causes the neutrino flux peak to shift lower in energy. The peak width is affected because at larger θ the neutrino energies become constant with respect to the parent pion energy at lower and lower pion energies, and so the range of neutrino energies produced becomes smaller, and an increased number of neutrinos are produced with the plateau neutrino energy.

Nearly all π^+ will decay via $\pi^+ \rightarrow \mu^+ + \nu_\mu$, and will obey the relationship given in Eq. 7.1. We therefore expect the spectrum of these ν_μ to be sensitive to changes in the off-axis angle.

We cannot make the same statement about the ν_μ produced via K decays. As described in Chapter 3, the collision of protons onto a carbon target produces a range of hadrons, including π^+ but also some neutral and charged K. We saw in Chapter 5 that different K$^+$ decay processes contribute ν_μ to the beam (see Table 5.1). One of those is an equivalent 2-body decay which will exhibit off-axis angle behaviour. However, there will be some ν_μ in the beam, produced from K parents through other decay modes, with energies which are not dependent on the off-axis angle in the same way. This is one way in which the spectrum of ν_μ from K is less sensitive to the off-axis angle.

For the ν_μ which are the result of the 2-body K$^+$ decay, the energies obey the same relationship as above, but with the different parent meson mass:

$$E_\nu = \frac{m_K^2 - m_\mu^2}{2(E_K - cos\theta\sqrt[2]{E_K^2 - m_K^2})} \tag{7.2}$$

The neutrino energy is plotted for a range of θ in Figure 7.3. In this case the maximum neutrino energies possible for each value of θ occur at much larger parent energies, and one must go further off-axis to see the maximum neutrino energy and eventual plateau.

To see this difference clearly, the neutrino energies with respect to parent energy are shown in Figure 7.4 for both π^+ and K$^+$ parents, at 2.5 degrees, the T2K off-axis angle.

While the maximum ν energy and plateau are clearly visible for the ν_μ from π^+, these characteristic features do not become apparent until much higher parent

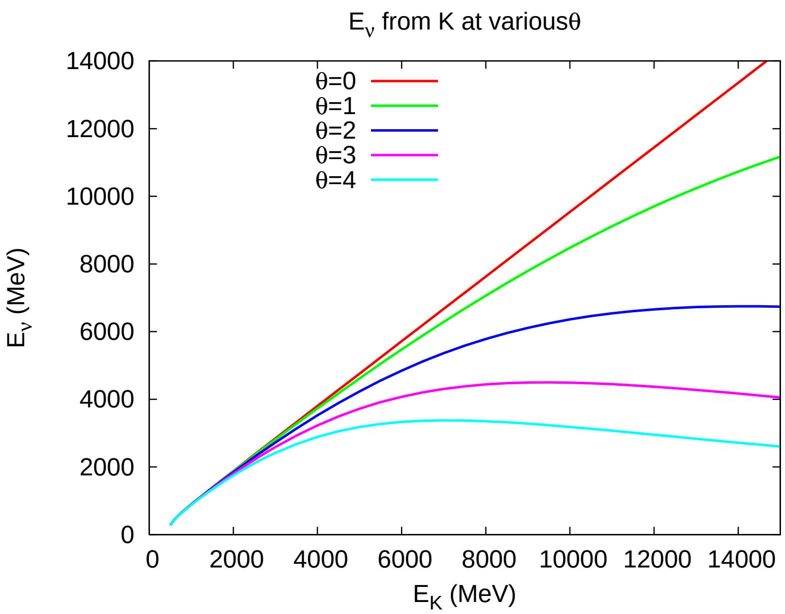

Figure 7.3: Relationship between ν_μ and parent K at different opening angles. θ measured in degrees.

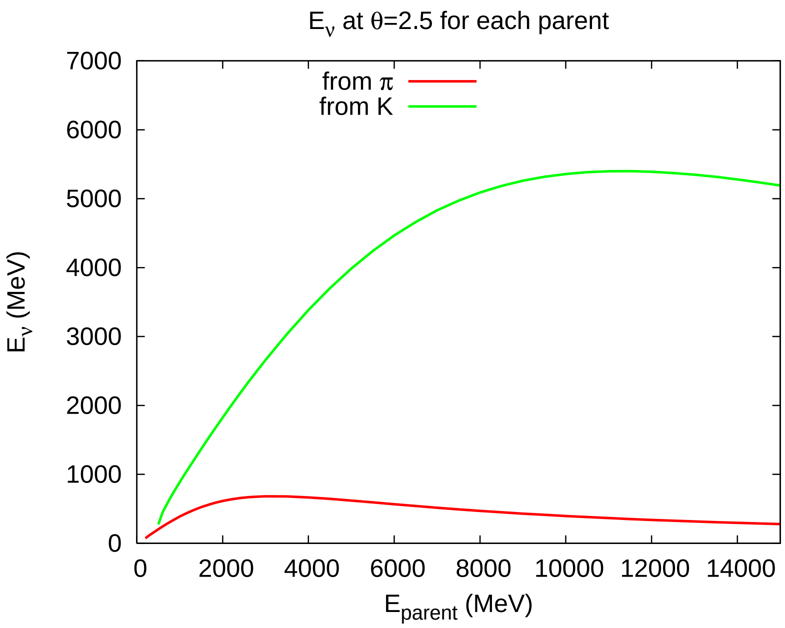

Figure 7.4: Relationship between ν_μ energy and parent energy at 2.5 degrees off-axis, where the parent can be a charged pion or kaon.

energies for the ν_μ from K^+. The ν energy is approximately linear with K energy for much of the parent energy range, before reaching a much higher maximum ν_μ energy of above 5 GeV at K^+ energy of around 10 GeV. The kaons produce a very broad ν_μ spectrum with no clear peak, so that the effect on the shape and position of the ν spectrum of a small change in off-axis angle is far smaller and more difficult to resolve.

In order to make a measurement of the relative numbers of π and K parents contributing to the neutrino beam we need a property which differs for the different types of parent. We have seen here that the $\nu_\mu^{f\pi}$ spectrum (where $\nu_\mu^{f\pi}$ is the ν_μ produced 'from pion' decays) is very sensitive to off-axis angle, and that the ν_μ^{fK} spectrum (where ν_μ^{fK} is the ν_μ produced 'from kaon' decays) is less so. The total neutrino spectrum is a combination of these, and so the degree to which the total ν_μ spectrum changes with off-axis angle should give an indication of what percentage of the total beam ν_μ are the direct decay products of π, and how many are the products of K, which in turn provides information about the relative quantities of π and K produced at the target.

We know that this off-axis behaviour exists, but due to ND280's dimensions and distance from the decay pipe, it only subtends a very small range of off-axis angles. For this off-axis behaviour to be useful for a beam parent composition measurement, the ν_μ spectra measured on different sides of the detector need to be significantly different. To investigate whether it is possible to resolve this off-axis behaviour within our detector, we will now look at ND280, and the ν_μ interactions it measures, in more detail.

7.2 Event distributions in ND280

One way to visualise the internal structure of ND280, shown in Figure 3.12, is to use Monte Carlo to plot the true vertex positions of all ν interactions in its volume. For this the 5C production version of magnet Monte Carlo is used[2]. The different detector components have different densities and therefore different interaction rates. The following plots demonstrate the internal structure of ND280, and also indicate the relative position of the beam axis, and show the variation in ν flux with position.

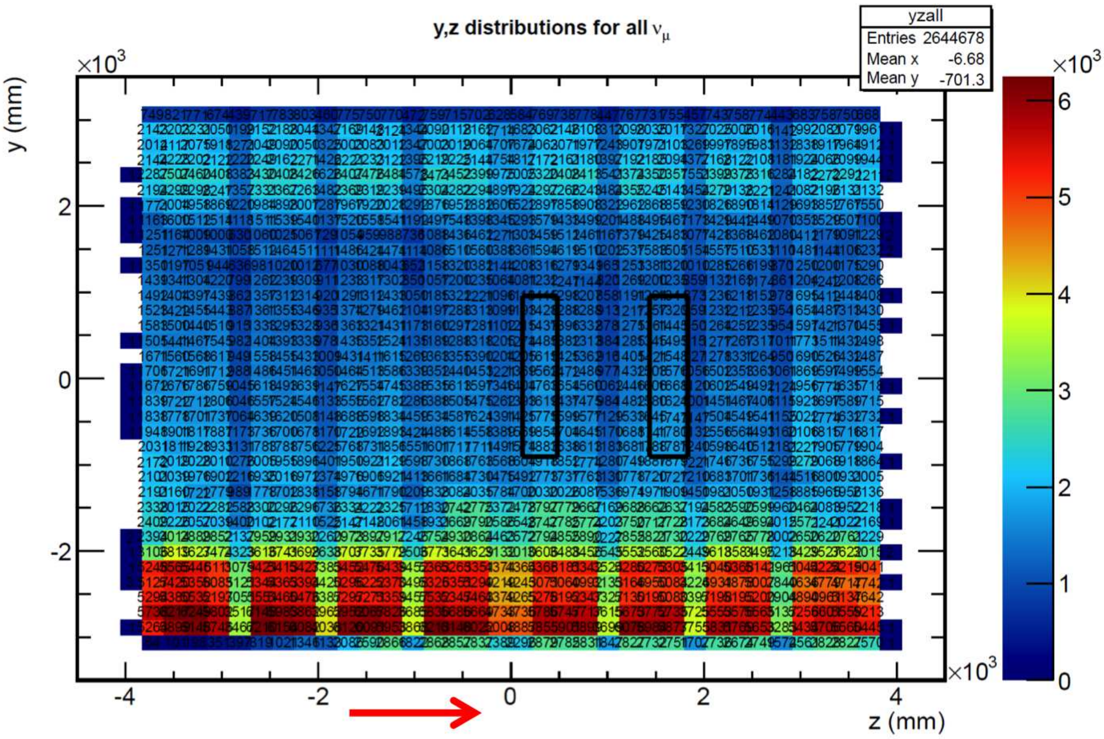

Figure 7.5: Distribution of ν interaction vertices in ND280 - side view (yz). Black boxes mark the positions of FGD1 (l) and FGD2 (r). The beam travels in the positive z direction, with the beam axis located below ND280, as shown by the red arrow (relative location not to scale). The 8 magnet yoke elements are visible.

Figures 7.5 and 7.6 show a side view and aerial view of ND280 respectively, where the beam direction is parallel to the z axis. These plots include interactions

[2]The T2K collaboration periodically releases new Monte Carlo productions, with a range of configurations and updated inputs. Monte Carlo details are provided in Appendix B.

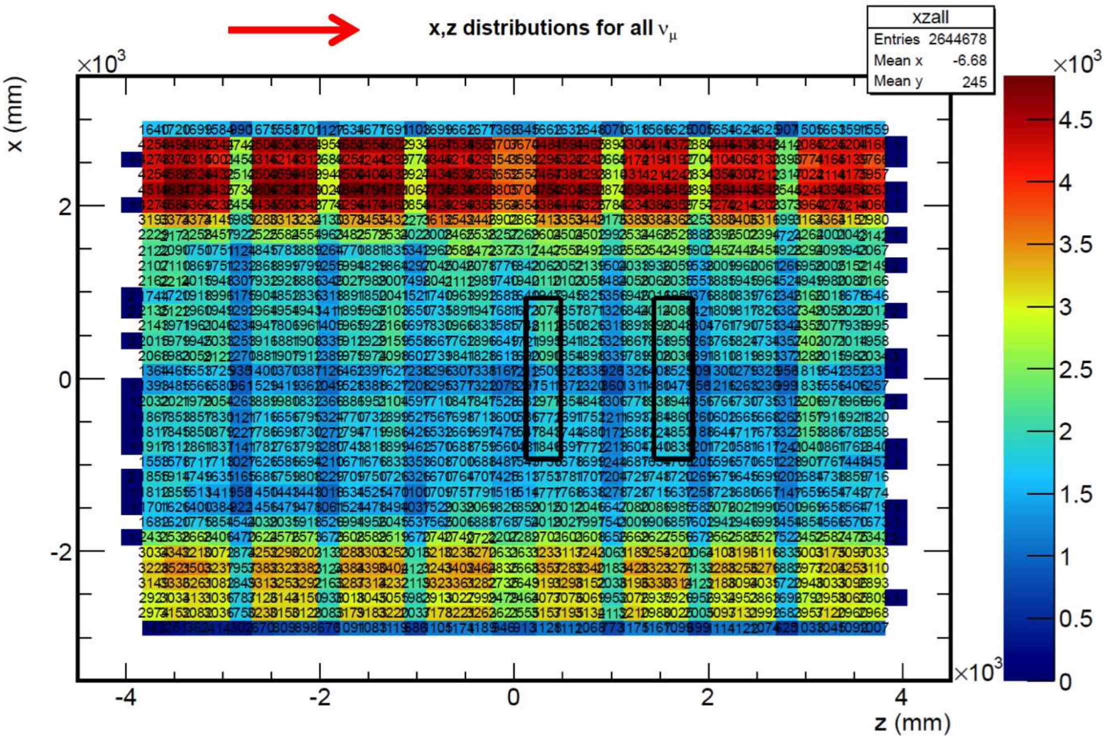

Figure 7.6: Distribution of ν interaction vertices in ND280 - aerial view (xz). Black boxes mark the positions of FGD1 (l) and FGD2 (r). The beam travels in the positive z direction, with the beam axis located on the positive x side of ND280, as shown by the red arrow (relative location not to scale). The side of the magnet yoke furthest from the beam axis (negative x) experiences approximately 72% of the interactions in the opposite side. The 8 pairs of magnet yoke elements are visible.

in the magnet yoke in addition to interactions in all detectors contained within the magnet. The structure of the 8 pairs of yoke elements is visible due to the high number of interactions in the magnet region. The highest interaction rates clearly occur in the magnet yoke, which show up as the red region at the bottom of Figure 7.5, and the red and yellow blocks at the top and bottom of Figure 7.6. The difference in the interaction rates in the magnet yoke on opposite sides of the detector, as visible in Figure 7.6, demonstrates the drop in ν flux as you move further off-axis. However, since the magnet yoke is not fully instrumented, these events cannot easily be utilised when looking at data. To test the visibility of the off-axis behaviour we need to divide the detector into regions which are at different off-axis angles and then plot separate ν_μ spectra for these regions. To do this we need a detector with good resolution of the interaction vertex position, in order to determine which region of the detector a ν_μ passes through and therefore which range of off-axis angles it belongs to. Having divided the ν_μ into off-axis angle groups, we then need to plot the energy spectra, which requires good energy reconstruction for the selected ν_μ.

Taking these requirements into consideration it is apparent that the FGDs and TPCs are best suited to detecting events of interest for this analysis. The FGDs provide significant target mass so that there is a good probability of ν interactions within their volumes, while also offering excellent vertex resolution. This allows the vertex position of a ν_μ interaction to be easily determined in a FGD, which is necessary for interactions to be grouped into bins of different off-axis angle. The products of an interaction will continue to travel downstream and most will pass through a TPC, which allows us to reconstruct any tracks and using information such as the track curvature, identify the interaction products and reconstruct their energies. While using events from both FGDs would increase the amount of available data, the separation between the two FGDs means that defining regions in FGD1 and FGD2 which would capture neutrinos with matching off-axis angles would require precise geometry information, and due to the shape and orientation of the near detector, may require omitting interactions occurring in regions of the second FGD. Therefore for this analysis we will select one FGD-TPC pair. FGD1 has twice the granularity of readout of FGD2 due to the higher number of scintillator strips present. In addition, the fiducial volume of FGD1 has been studied more extensively and is better understood, and the existing ν_μ selection has been optimised for interactions in FGD1. Therefore for this analysis, ν_μ which interact in FGD1 and produce tracks which pass through TPC2 will be selected.

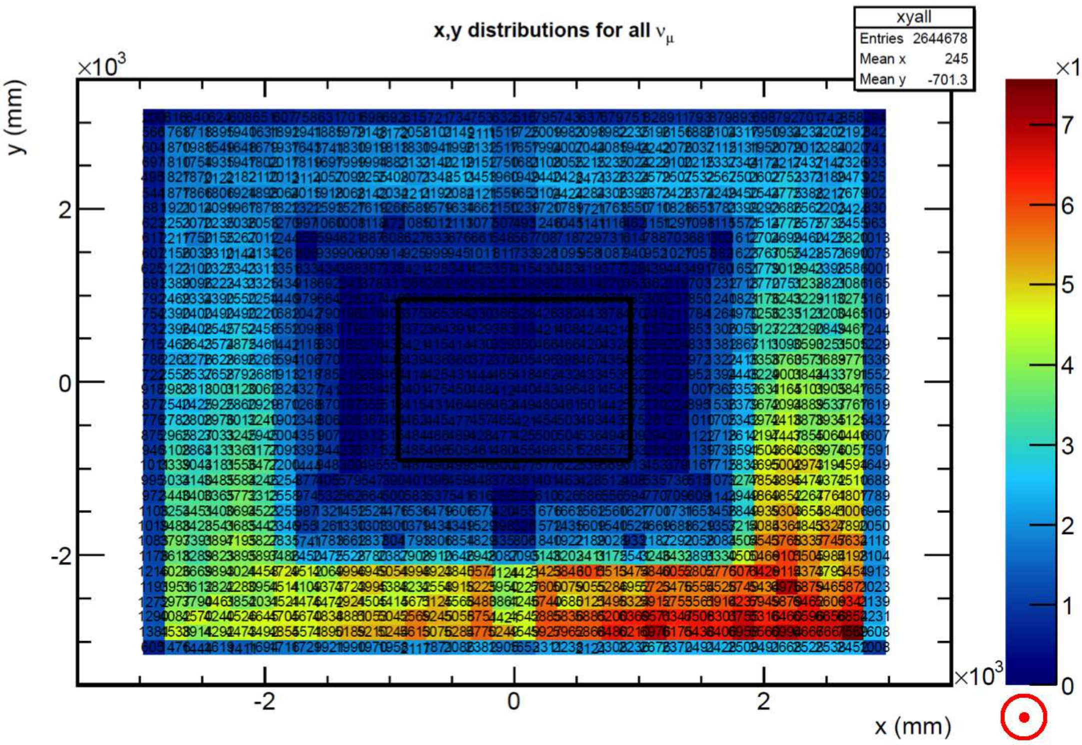

Figure 7.7: Distribution of ν interaction vertices in ND280 in the direction of the beam (xy). FGD1 boundaries are marked by the black box. Beam direction is out of the page, with the beam axis located beyond the bottom right corner of the plot, as marked in red (beam marker location not to scale).

Figure 7.7 shows an xy cross-section through ND280, looking towards the target. The boundaries of FGD1 are shown in black. The beam axis ($\theta = 0$) does not pass through ND280 but is located beyond the bottom right corner of the magnet yoke. Its position relative to ND280 can be inferred from the clear increase in interaction rate, and therefore flux, as one moves diagonally across the detector.

Using only FGD1 as the interaction volume is necessary in order to achieve optimum event selection and reconstruction; however, this reduces the width of our region of interest to 2.473 m. Figure 7.8 shows the events in FGD1 only, with the dimensions provided. The bottom right corner is closest to the beam axis, and thus ν_μ interacting here will be the products of decays with the smallest opening angles. The top left corner is the furthest from the beam axis, so these ν_μ will have larger opening angles. Since FGD1 is of limited size and subtends only a small range of off-axis angles, and considering that we must ensure a statistically significant number of interactions in each region, and a resolvable difference in the spectra from these regions, we will define only 2 regions for comparison.

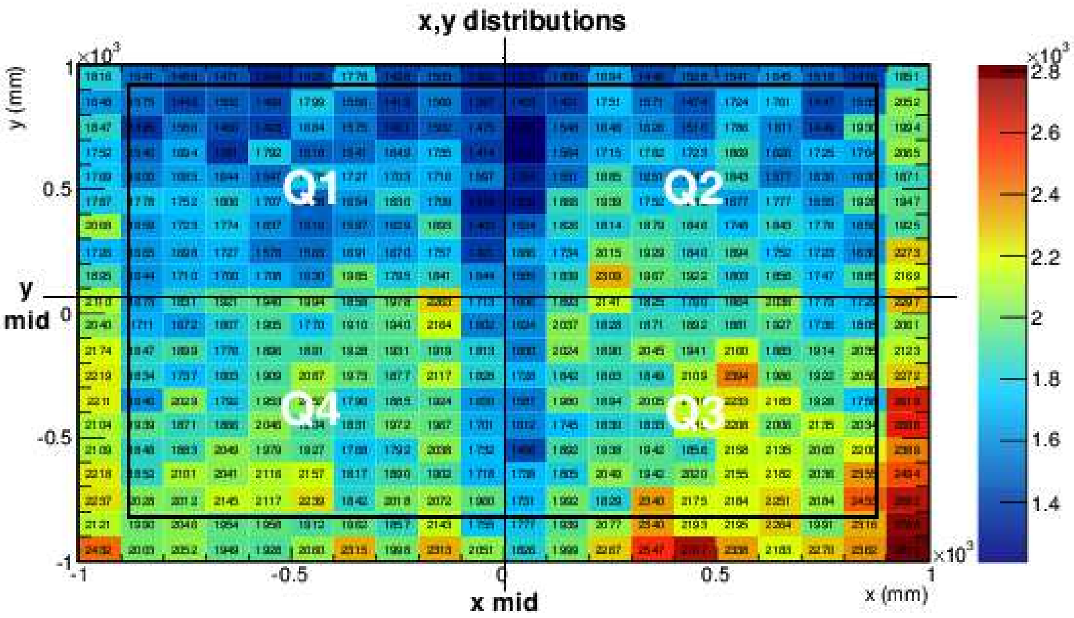

Figure 7.8: FGD1 xy cross-section with definition of quadrants. Q3 is closest to the beamline.

Figure 7.8 shows FGD1 divided into four equally-sized quadrants. These quadrants are labelled such that Q1 is furthest off-axis, and Q3 is most on-axis. The distance between the centres of these quadrants is 1.237m. Taking 232 m as the distance between ND280 and the position of neutrino production in the decay pipe, we can calculate the mean difference in off-axis angle between a ν_μ interacting in Q1 and a ν_μ interacting in Q3, and find this to be approximately 0.3 degrees, as demonstrated in Figure 7.9.

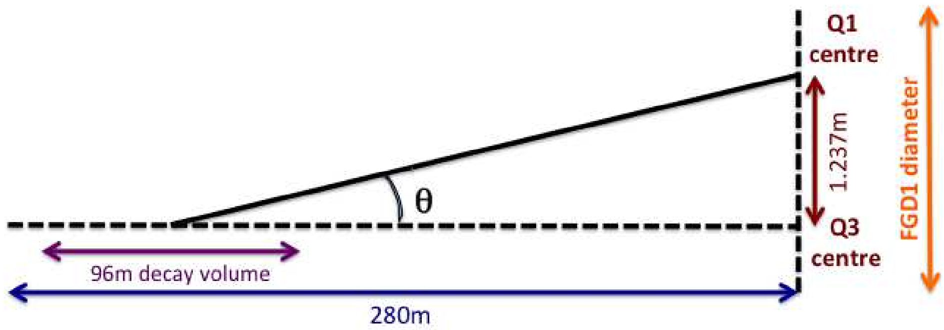

Figure 7.9: Off-axis angle difference between the mid-points of Q1 and Q3 quadrants in FGD1. Distance between mid-points of quadrants is 1.237 m. Taking the mid-point of the decay pipe as the average neutrino production position, neutrinos travel approximately 232 m to ND280. Using an approximation of the opening angle geometry, we state tan$\theta = 1.237/232$, $\theta = 0.3$ degrees.

Having defined these regions, we must now plot the energies of the ν_μ in Q1 and Q3 and determine whether a difference of approximately 0.3 degrees in θ is sufficient to resolve the changes in spectral shape.

7.3 Muon event selection

To plot the energy spectra of interest we need to identify suitable ν_μ interactions. Firstly, the many signals detected throughout the various ND280 sub-detectors must be processed and grouped into interactions. One 'event' consists of 1 beam spill, which contains 8 neutrino bunches (see Figure 3.17), and a neutrino from one or more of these bunches may interact in one of the ND280 components. An interaction will result in a number of products, which will cause tracks and showers. It is these tracks and showers which are identified, reconstructed and studied to provide information about the incident neutrino, such as its flavour and energy. It is vital then that tracks and showers within an event are assigned to the correct interaction, so that the energies are added to the appropriate totals, and the interaction types are correctly identified. Therefore some cuts are needed to ensure that we only select events in which the tracks can be grouped accurately.

The tracks themselves are subject to 'quality' cuts which guarantee a minimum level of accuracy when calculating their properties. Another factor to be considered is that the neutrino beam isn't purely ν_μ - there are intrinsic ν_e present in the beam too, as studied in Ch. 4 - so when a neutrino interaction is identified, it is necessary to select only interactions that produce μ^- tracks, rather than e^- tracks.

Many of the cuts used in this analysis are based on the selection cuts used for the main ν_μ disappearance analysis, and are described in more detail in [96]. Further cuts are then added to select ν_μ with interaction vertices at positions of interest to this analysis. We require interactions in FGD1, and forward-going tracks, so any bunches containing activity in ND280 components which sit upstream of FGD1 can be vetoed.

A full list of the selection cuts used follows. For each event we consider the bunches separately. The first step is to remove any activity which occurs outside a bunch timing window. Cuts 2 and 3 then apply on the bunch level. If a bunch passes both of these cuts, we then look at each track in that bunch and apply cuts 4 to 7 to each track.

1. Bunch timing

Ignore any track which doesn't coincide with a bunch timing window. This helps to select tracks that are products of interactions of our beam ν_μ rather than other sources of neutrinos, and removes tracks which could be the delayed products of an interaction in a previous bunch. Bunch windows of 60 ns either side of each mean bunch position are used. Timing windows for MC and data runs are given in [96].

2. **P0D, P0D ECal veto**

   Reject any bunch containing P0D or P0D ECal activity. Activity in either would mean either a backwards-going track, or an interaction that did not occur in FGD1.

3. **TPC1 veto**

   Reject any bunch that contains activity in TPC1. Activity here would be due to either a backwards-going track, or a track not originating in FGD1.

4. **TPC info**

   Ignore any track which does not contain TPC information. We require TPC information to reconstruct the track accurately.

5. **Vertex in FGD1 fiducial volume**

   We require interactions which occur in FGD1. The products of an interaction are created at the point of interaction, so an interaction in FGD1 produces tracks originating in FGD1. Therefore we only select a track if its upstream end is located in FGD1, and we define the interaction vertex location as the upstream end of the highest-momentum μ-like[3] track in the bunch.

   The cut uses the boundaries of the fiducial volume of FGD1, as follows:

   -874.51 mm < x_{vertex} < 874.51 mm

   -819.51 mm < y_{vertex} < 929.51 mm

   136.875 mm < z_{vertex} < 446.955 mm

6. **Forward going tracks**

   Only select tracks which travel downstream from their interaction point.

[3]See cut 10 for description of track identities

7. TPC2 component

Having originated in FGD1, a track must then proceed downstream and pass through TPC2, rather than escape sideways. We use the track length and curvature information gathered in the TPC to calculate properties such as charge and momentum, so TPC2 information is vital for reconstructing the type of interaction and summing the energies of the products.

Any track which passes these cuts will be used. We sort these tracks into μ-like tracks and non-μ-like tracks, where the non-μ-like tracks are the other products of a CCnonQE interaction, such as π^0. To identify the μ-like tracks we apply three further selection cuts to the existing selection of tracks. If a track passes the additional cuts 8, 9 and 10, it will be labelled as μ-like.

8. Track quality cut

Track must pass through 18 or more vertical TPC clusters. The longer the track, the more accurately we can measure its curvature and determine its momentum. A track which clips the edge of TPC2, or has low energy and doesn't pass very far into TPC2, will be harder to reconstruct and the uncertainty on the track's momentum will be larger. Also, the accuracy of the particle identification is dependent on the momentum range. Requiring that tracks travel above a minimum length removes the very low momentum tracks for which the particle identification is less reliable. Therefore we only select tracks which travel through 18 or more nodes in TPC2.

9. Track charge

Track must have negative charge, as determined by the direction of curvature in the magnetic field.

10. Track pull values

A track which passes all of the above cuts could be a muon, electron or negatively charged hadron. A ν_μ interaction will produce a μ, so we require the presence of a muon track in the bunch. For a particle of a given momentum, the energy deposited in the track, $\left(\frac{dE}{dx}\right)_{expected,i}$, is estimated for each of the possible identification hypotheses, labelled i. The energy loss for a track can also be measured using the energy deposited in the TPC, as explained in detail in [97]. We then compare the measured energy loss with the energy loss that would be expected should the particle have each of the possible identities i. The 'pull' value for a given particle hypothesis i is [96]

$$\text{Pull}_i = \frac{dE/dx_{measured} - dE/dx_{expected,i}}{\sigma_{dE/dx_{measured} - dE/dx_{expected,i}}} \tag{7.3}$$

where σ depends on the truncated mean energy deposit and its width, and the uncertainty on the momentum measurement. Figure 7.10 shows the relationships between dE/dx and the transverse momentum for various particles, and the pull variables for true electrons and muons are shown in Figure 7.11. For this selection a muon is defined as a track which passes:

-2.0 < μ pull < 2.0 and

e pull < -2.0.

Any tracks which fail to pass any of these 3 additional selection cuts are kept but labelled as non-μ-like.

Having classified all acceptable tracks as μ-like or non-μ-like, we now apply one final cut.

11. Delta Z cut

In the case that non-μ-like tracks are also present, we find the distance between the starting position of the highest momentum non-μ-like track and the starting position of the highest momentum μ-like track. If the non-μ-like track originates >150 mm upstream of the μ-like track, the event is

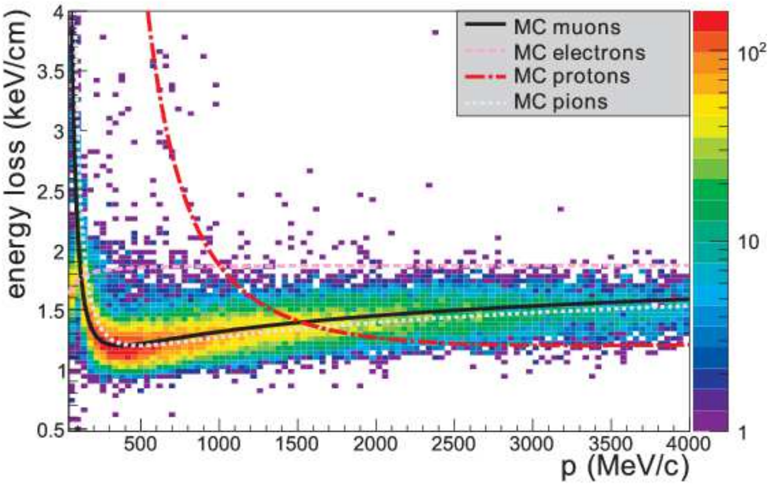

Figure 7.10: Distribution of the energy loss as a function of the momentum for negatively charged particles produced in neutrino interactions, compared to expected curves for muons, electrons, protons and pions. Figure from [73].

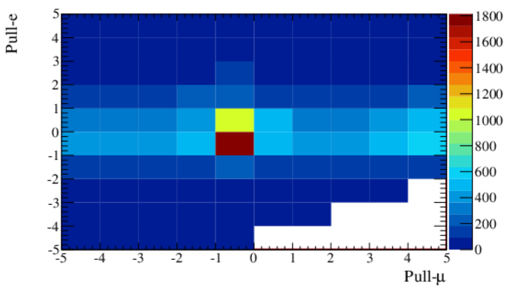

(a) True electrons

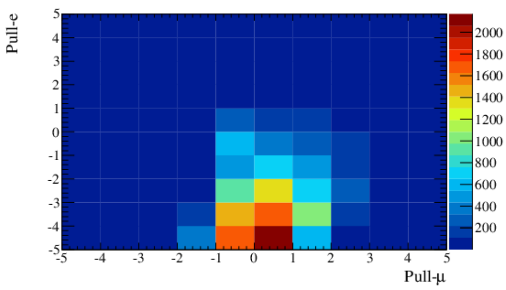

(b) True muons

Figure 7.11: Pull variables for true electrons (left) and true muons (right). The cuts on μ pull and e pull select the most true muons while minimising the number of true electrons also being selected. Plots courtesy of Dr. H O'Keeffe [98].

rejected, as there is a reasonable probability that the non-μ-like track present originated outside the FGD.

At this point we can test if the change in off-axis angle within FGD1 is large enough to be resolved. For each μ^- we have identified in the Monte Carlo, we can simply plot the true energy of its parent ν_μ. So far we have a selection of μ-like tracks which are produced in interactions in FGD1, but we can divide the sample further into Q1 and Q3 interactions. The position of the neutrino interaction vertex is taken to be the front position of the μ-like track (or if there is more than one μ-like track in a bunch[4], the highest momentum μ-like track in the bunch).

We define the centre of FGD1 in the x-direction, located at 0 mm, as x^{mid}_{FGD1} and the centre of FGD1 in the y-direction, located at 55 mm, as y^{mid}_{FGD1}, as shown on Figure 7.8. The Q1 and Q3 samples are then found using the following position cuts:

Q1 sample: $x_{vertex} < x^{mid}_{FGD1}$ and $y_{vertex} > y^{mid}_{FGD1}$

Q3 sample: $x_{vertex} > x^{mid}_{FGD1}$ and $y_{vertex} < y^{mid}_{FGD1}$

Having identified a muon track which originated in one of Q1 or Q3, we now use the Monte Carlo truth information to plot the true energy of the ν_μ which interacted. Figure 7.12 shows the spectrum of ν_μ interacting in Q1 in blue, and the spectrum of ν_μ interacting in Q3 in red.

Each individual neutrino spectrum has the shape we would expect - there is a sharp peak at around 600MeV followed by a relatively small high energy tail. Plotting the two spectra together, we see that there is indeed a small but visible difference in the neutrino spectra corresponding to opposite sides of FGD1. Q3

[4]A second mu-like track will be the result of particle misidentification, as the probability of two ν_μ interactions in one bunch is negligible. Possible additional track types are discussed in §7.4.2.

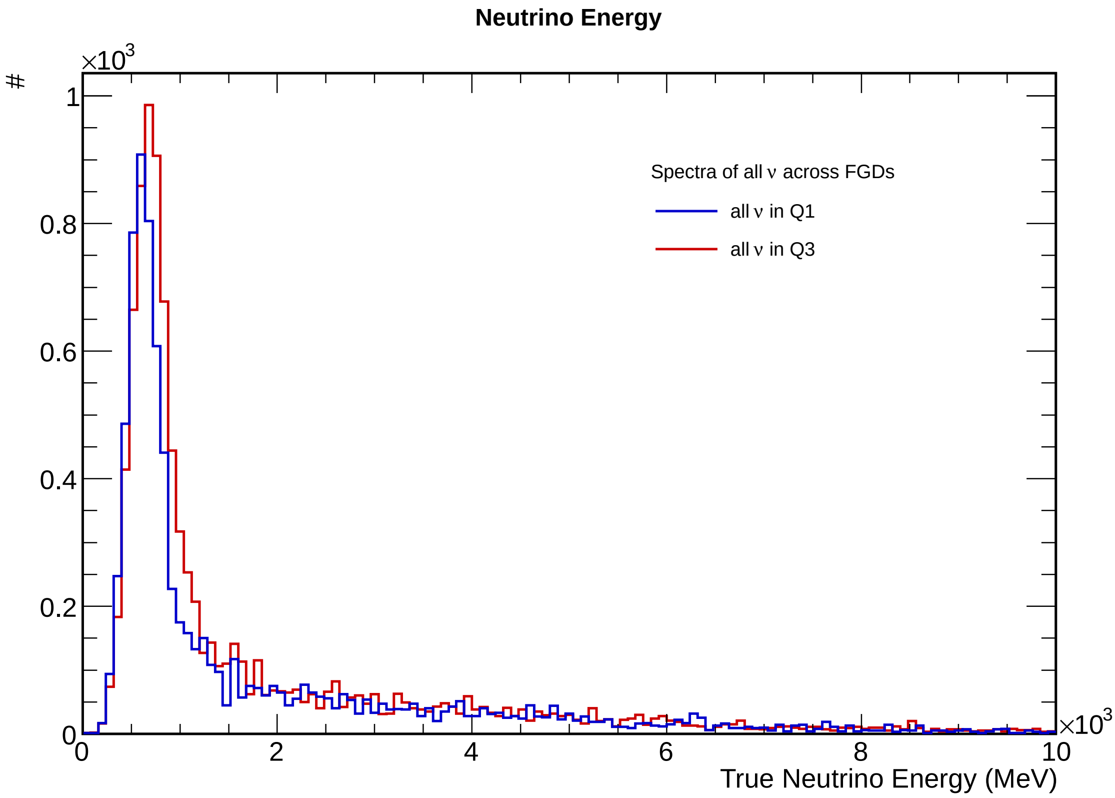

Figure 7.12: True energy spectra of ν_μ in Q1 and Q3. Shift in peak position due to off-axis angle difference is visible

is closer to the axis of the beam, so ν_μ interacting here will have smaller opening angles and as expected, the peak is at slightly higher energy than the more off-axis Q1 peak. We can also see that while Q1 and Q3 have the same fiducial volume, the Q1 spectrum contains fewer interactions due to the flux being lower at larger off-axis angles.

This confirms that, while FGD1 is only one component of ND280 and of limited size, the difference in off-axis angles between the opposite corners of FGD1 is sufficient to clearly demonstrate the effect of the off-axis angle on the neutrino spectrum shape.

7.4 Energy reconstruction

We have demonstrated that we can resolve the effects of changing off-axis angle with samples of ν_μ detected in FGD1, however these Q1 and Q3 spectra are plotted using true energy of the parent neutrinos. To use this technique with data, we must reconstruct the energy of a parent neutrino using the information we have measured for the particles produced in the interaction. We cannot expect the energy reconstruction to be perfect, since there are limitations on our ability to precisely measure the momentum of a track and identify the particle that caused it, so we must check that this off-axis angle behaviour is still visible when the spectra are plotted using reconstructed energy.

For each bunch we have processed all the tracks present and counted the number of μ-like tracks and non-μ-like tracks. We can now use these totals to infer which type of interaction occurred in the bunch. Both charged current (CC) and neutral current (NC) interactions can occur in ND280. In the case of a neutral current interaction, shown in Figure 7.13, a Z^0 boson is exchanged, and various products may be produced in the interaction and detected, but there will not be a μ^- track. If the Z^0 is exchanged with a nucleon, a number of hadronic products will be emitted (shown by 7.13a), or if the neutrino scatters off an electron (the less likely process shown by 7.13b), the electron may have enough energy to be freed and would then produce a shower.

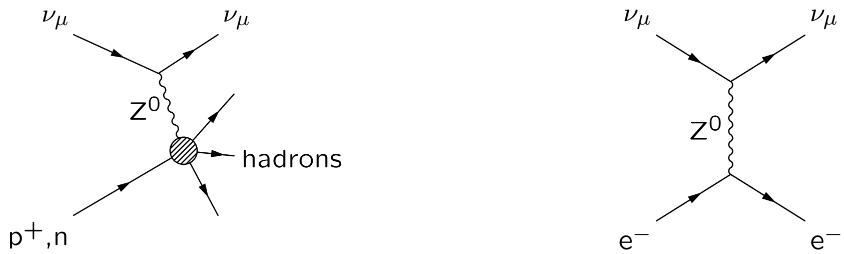

(a) Hadronic production (b) Elastic scattering

Figure 7.13: Neutral current interaction modes

We can use the detector information gathered to calculate the energies of the products of the interaction. However, in a neutral current interaction the incident neutrino is preserved and will continue on through the detector and escape without further interaction. Therefore, we cannot know how much energy the neutrino takes away, and so summing the energies of all products in order to calculate the initial energy of this neutrino is impossible. In contrast, in a CC interaction, which proceeds as in Eq.7.4, the incident neutrino is converted into a charged lepton and so all of the products are detectable.

$$\nu_x + n \rightarrow p + l_x^- \tag{7.4}$$

Therefore for this analysis we will select only CC interactions, and veto NC interactions. Since NC interactions do not produce a muon, should the number of μ-like tracks in a bunch be zero we discard the bunch. This also removes any intrinsic ν_e CC interactions, since the product in that case would be an e^- rather than a μ^-, also resulting in zero μ-like tracks.

A CC interaction proceeds by the exchange of a $W^\pm$ boson, and will always result in the production of a charged lepton of flavour matching that of the parent neutrino. Charged current interactions can be classified into 2 sub-groups; quasi-elastic (CCQE) interactions, and non-quasi-elastic (CCnonQE) interactions. Both types of interaction will always produce a μ^- from a ν_μ parent, but the presence of other additional products depends on the interaction type and the amount of energy transferred from the neutrino in the interaction. The appropriate method of parent energy reconstruction depends on the interaction type, so we will now look at each interaction type separately, and describe the energy reconstruction calculation for each.

7.4.1 CCQE

The charged current quasi-elastic interaction is the simplest ν_μ interaction we can see in ND280, as the products are limited to a single μ^- track. As shown in Figure 7.14, the neutrino interacts with the quarks in a nucleon via W$^-$ exchange, with the result that the nucleon changes type. The neutrino exchanges a W$^-$ boson with a d quark in a neutron, converting the d quark to a u quark, and thus the neutron to a proton. The negative charge is carried away by the W$^-$ boson, allowing the creation of the μ^- from the neutral ν_μ.

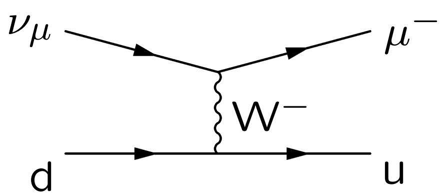

Figure 7.14: CCQE interaction: ν_μ + n $\rightarrow$ μ^- + p

In the most straightforward case, little of the neutrino's energy is transferred to the nucleon, so the proton stays in the nucleus of the atom and merely recoils. With the nuclear changes not detected, the only signal of the CCQE interaction is a single μ^- track. Therefore to identify a CCQE interaction of this type, we require that the number of tracks found in a bunch satisfy:

- Number of μ-like tracks = 1
- Number of non-μ-like tracks = 0

The momentum of the μ-like track is reconstructed using the TPC information and the energy is then calculated using the fact that the track has been identified as a particle of mass m_μ. The energy of the parent neutrino can then be calculated. Neglecting the Fermi motion, the neutrino energy can be reconstructed using equation 7.5 [26], where E_l and m_l are the measured energy and mass of the muon respectively and θ is the angle between the parent neutrino direction and the muon direction. V represents the binding energy of a nucleon in ^{16}O and is set at 27 MeV [35][5]. m_p and m_n are the masses of the proton and neutron respectively.

$$E_\nu = \frac{2E_l(m_n - V) - m_l^2 + 2m_nV - V^2 + m_p^2 - m_n^2}{2(m_n - V - E_l + P_l cos\theta)} \tag{7.5}$$

The energy spectrum, plotted using reconstructed energy, for these CCQE events only, is plotted in Figure 7.15.

We can test the accuracy of this energy reconstruction by plotting the value of E_ν calculated in this way against the true energy of the parent neutrino, found in the Monte Carlo truth information. The comparison is shown in Figure 7.16.

7.4.2 CCnonQE

There are many possible non-quasi-elastic interaction processes. The leptonic vertex is the same - a W^- is exchanged and the neutrino is converted into a μ^-.

[5] A possible improvement to this analysis may be to investigate using alternative values of V which are more appropriate for the scintillator material, such as the value for ^{12}C.

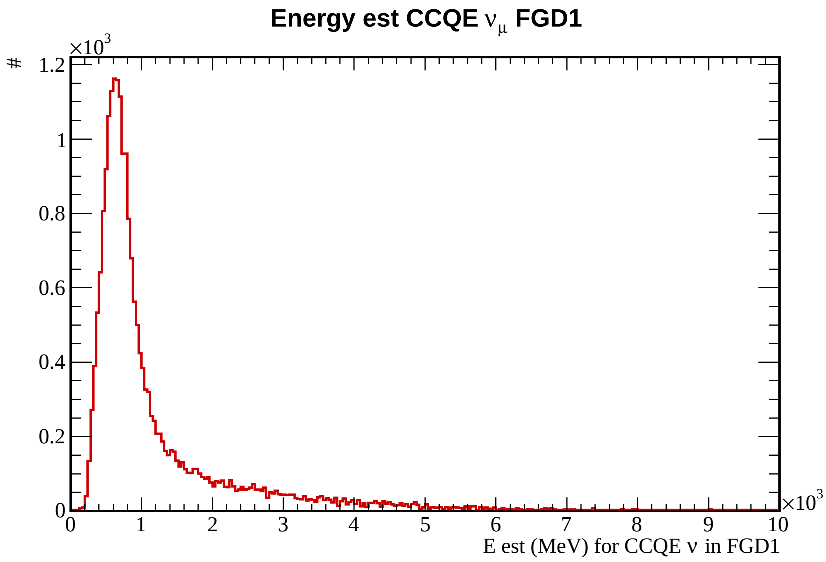

Figure 7.15: The reconstructed energies of MC ν_μ interacting via CCQE interactions in FGD1 are plotted. Energy reconstructed using Eq. 7.5.

The difference is that, in general, more energy is transferred to the nucleon, and this energy is then available for the creation of additional hadrons. While the CCQE process can be summarised as

$$\nu_\mu + N \to \mu^- + N' \tag{7.6}$$

the non-quasi-elastic processes are summarised as

$$\nu_\mu + N \to \mu^- + N' + X \tag{7.7}$$

where X are the additional products, as shown in Figure 7.17.

Events of this type are identified using the numbers of identified tracks. We select any bunch for which the following statements are true:

- Number of μ-like tracks ≥ 1 [6]

[6]The probability of an interaction in any bunch is 0.078% [96], so the probability of two interactions in the same bunch is negligible. The selection is set as ≥ 1 μ-like tracks, although

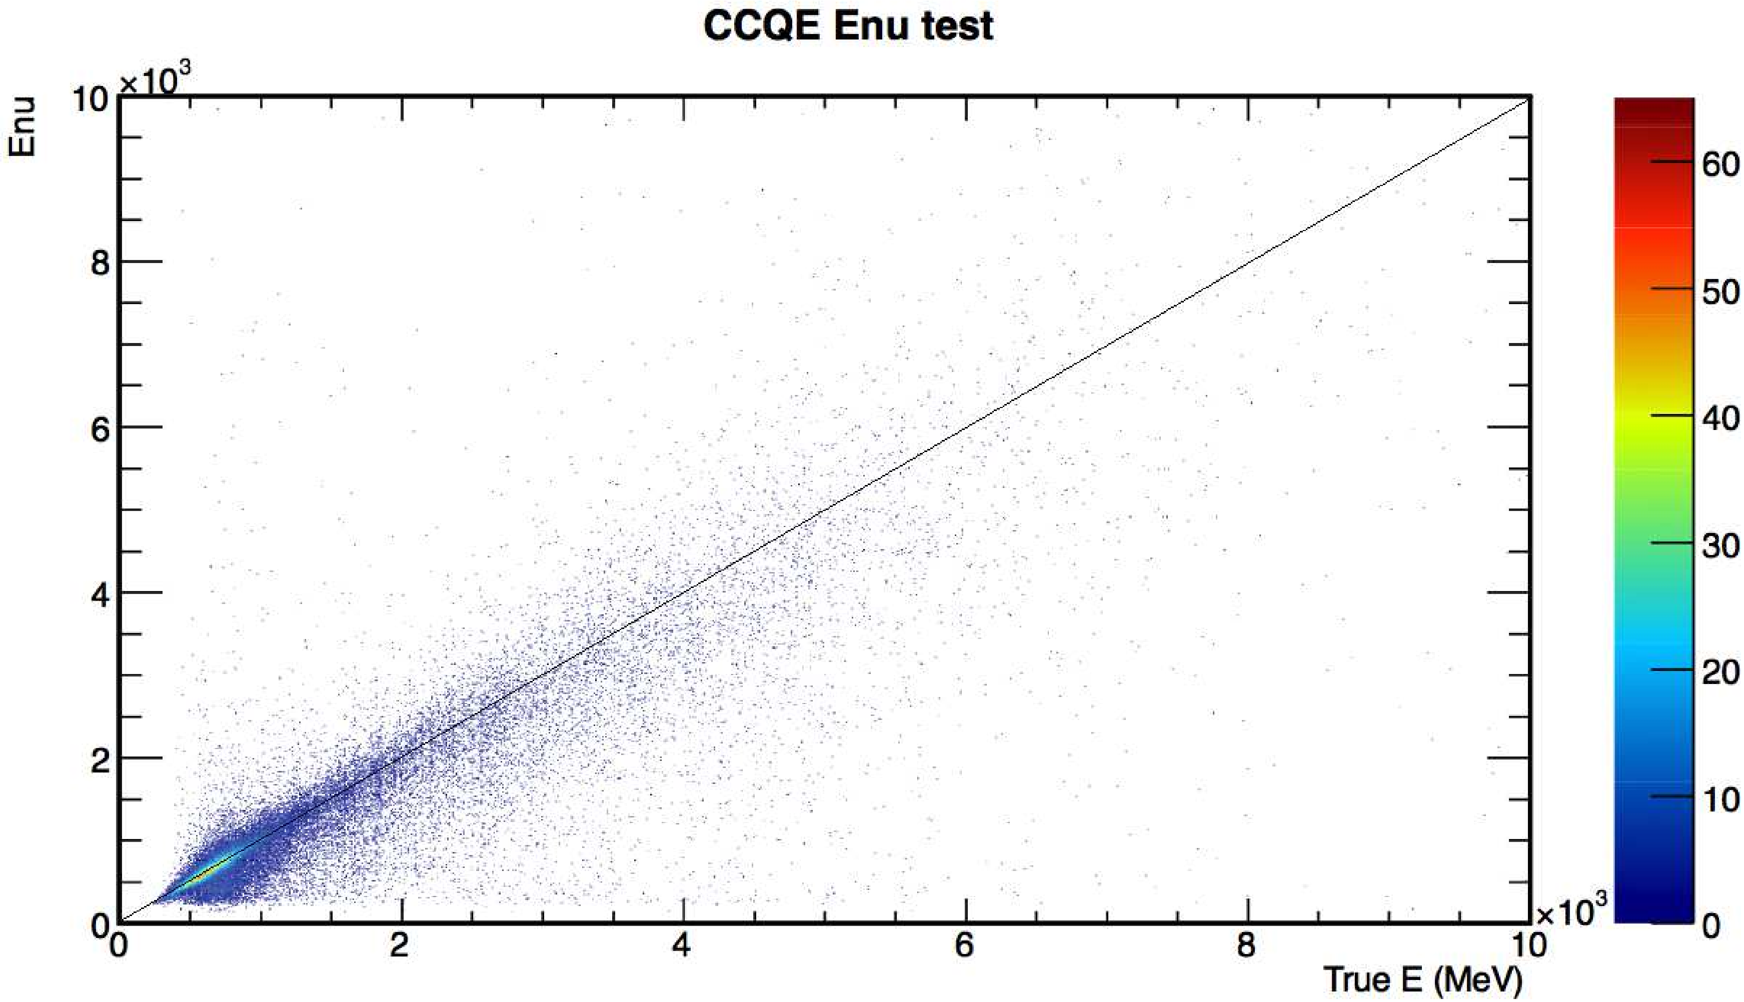

Figure 7.16: Energy reconstruction for CCQE events. For selected MC interactions the reconstructed energy (Enu) is plotted against the energy the particle was produced with (True E). A strong positive correlation is shown, indicating effective energy reconstruction. Black line showing Enu = True E plotted for reference.

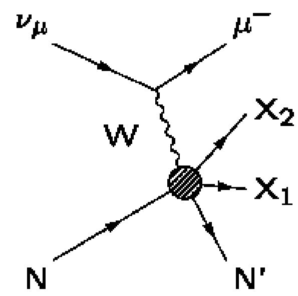

Figure 7.17: A ν_μ produces a μ^- and additional particles, labelled X_1, X_2 via a CCnonQE interaction.

- Number of non-μ-like tracks ≥ 1

To calculate the energy of the incident neutrino, we use the principle of conservation of energy. For each track in the bunch we calculate the energy of the particle using

$$E_{track} = \sqrt{p_{track}^2 + m_X^2} \tag{7.8}$$

For μ-like tracks $m_X = m_\mu$. For non-μ-like tracks we must use an appropriate mass for the track.

The energy of the neutrino is used to create the μ^-, as well as additional hadrons produced in the interaction, so the sum of the energies of the tracks detected is approximately equal to the initial energy of the neutrino. It is important that we only include the tracks emitted from the interaction directly, and not also include the later decay products, as we would then be including some energy contributions twice. By requiring that all tracks begin in FGD1, we reduce the chance of including tracks that form further downstream from the decay of hadrons already included in the sum. Should a hadron produced in the interaction decay very quickly, so that the position of decay is very close to the neutrino interaction vertex and therefore also in FGD1, the track left by the intermediate hadron would

there will be very few bunches with more than one μ-like track, and in these cases the second μ-like track is more likely to be a case of track misidentification than indicative of a second ν_μ interaction.

not have a TPC2 component, so would not be included, and the tracks caused by the decay products would be included instead, so the energy contributions would be correct. Particles of very low energy may not penetrate far enough into TPC2 to pass the track quality cut, and so their small energy contributions would be omitted. This is a potential source of inaccuracy in the reconstructed energy, along with the loss of tracks which originate near the boundaries of TPC2 and leave TPC2 after travelling too short a distance to pass the track quality cut.

To decide which particle mass to use for the non-μ-like tracks, we will now look at some examples of non-quasi-elastic interactions, and the types of particle they tend to generate.

The first case to consider is a CCQE interaction where the proton is visible. In this case the interaction proceeds as shown in Figure 7.14, but the proton receives sufficient energy from the neutrino to escape the nucleus. While the additional track is left by a proton, it would not be appropriate to use the total energy of this track in the bunch energy sum. This nucleon was merely emitted from the nucleus with additional kinetic energy gained from the neutrino, as opposed to created in the interaction using energy of the neutrino. Therefore it is only the kinetic energy that should be included in the sum of energies, and not the rest mass of the nucleon. Using $\mathrm{m}_X = \mathrm{m}_p$ to calculate E_{track} for a proton track would drastically over-estimate the energy of the incident neutrino.

Next we consider what happens when more energy is transferred to the nucleon. Consider the case where the interaction proceeds as in Figure 7.14 and a neutron is converted into a proton, but this time the proton receives enough energy that it exists in an excited state. In this case the proton will quickly drop back down to the ground state by emitting a new particle, such as a π^0, resulting in

$$\nu_\mu + n \to \mu^- + p + \pi^0 \tag{7.9}$$

We will therefore see the μ-like track, a non-μ-like track left by the proton, and another two non-μ-like tracks corresponding to the decay gammas of the π^0.

Another example is the formation of a Δ^{++} resonance. The ν_μ interacts as shown in Figure 7.14; however, instead of involving one of the d quarks in a neutron, it interacts with the one d quark in a proton. The product therefore is the Δ^{++}, consisting of (uuu), which very quickly decays to produce a proton and charged pion, as follows.

$$\nu_\mu + p \rightarrow \mu^- + \Delta^{++} \tag{7.10}$$

$$\Delta^{++} \rightarrow p + \pi^+ \tag{7.11}$$

This gives the overall result

$$\nu_\mu + N \rightarrow N' + \mu^- + p + \pi^+ \tag{7.12}$$

Therefore this example would produce three tracks in the detector. The Δ^{++} would not be accounted for directly as it would decay too quickly to leave a track, but the energies of the proton and pion would be included, so the energy of all products will still be accounted for correctly.

Rather than attempt to identify every track that we select in a bunch, which would be difficult as some particle ID selections overlap, it is reasonable to make some approximations in order to simplify the energy reconstruction in the case of extra tracks. Neutral and charged pions are commonly produced in these interactions as they are the lowest energy hadrons that can be created, so many of the non-μ-like tracks that we detect will be left by pions. While protons often escape the nucleus when there is enough energy available to them, they are more likely to leave a short track which would not pass the track quality cuts and therefore would not be counted in the bunch energy sum. Should a proton travel far enough in TPC2 that it does pass the cut, we cannot use the proton mass to calculate the energy contribution, but must estimate the energy transferred to the proton in some way. Making the assumption that every non-μ-like track is the result of a pion is a suitable approximation[7]. Also, a contribution to the neutrino energy estimation, of the correct order of magnitude, is still included for tracks left by particles which were given energy, but not created, in the interaction, such as protons.

Using this assumption, the energy of a neutrino is reconstructed as follows:

$$E_{\nu_\mu} = \sum_{\mu-tracks} \sqrt{p_{track}^2 + m_\mu^2} + \sum_{non-\mu-tracks} \sqrt{p_{track}^2 + m_{\pi^0}^2} \tag{7.13}$$

For each selected track in a bunch, the measured momentum value is used to calculate the energy of the particle, and these energies are summed. As with the CCQE reconstruction method, the accuracy will be limited by the accuracy of the momentum measurements for the tracks. There is also the possibility that the particle identification could be wrong and a track will be wrongly counted as μ-like. This method has the further limitation that contributions from short tracks and tracks that exit the TPC before traversing the required distance will be omitted. The energy carried away by any nucleons which leave the nucleus is also difficult to estimate. However, despite these challenges, the energy reconstruction accuracy achieved for CCnonQE events is good, and is shown in Figure 7.18.

The energy spectrum, plotted using reconstructed energy, for these CCnonQE events is plotted in Figure 7.19. The contrast to the spectrum of CCQE interactions in Figure 7.15 is of interest. The CCQE spectrum shows the peak at ~600 MeV that is familiar for the T2K beam, and that we know is largely formed

[7]The accuracy of this assumption depends strongly on how many protons are produced due to nuclear breakup, which is currently difficult to estimate due to the absence of measurements in neutrino interactions. Common secondary tracks are discussed in [99], which finds that protons and pions occur most frequently.

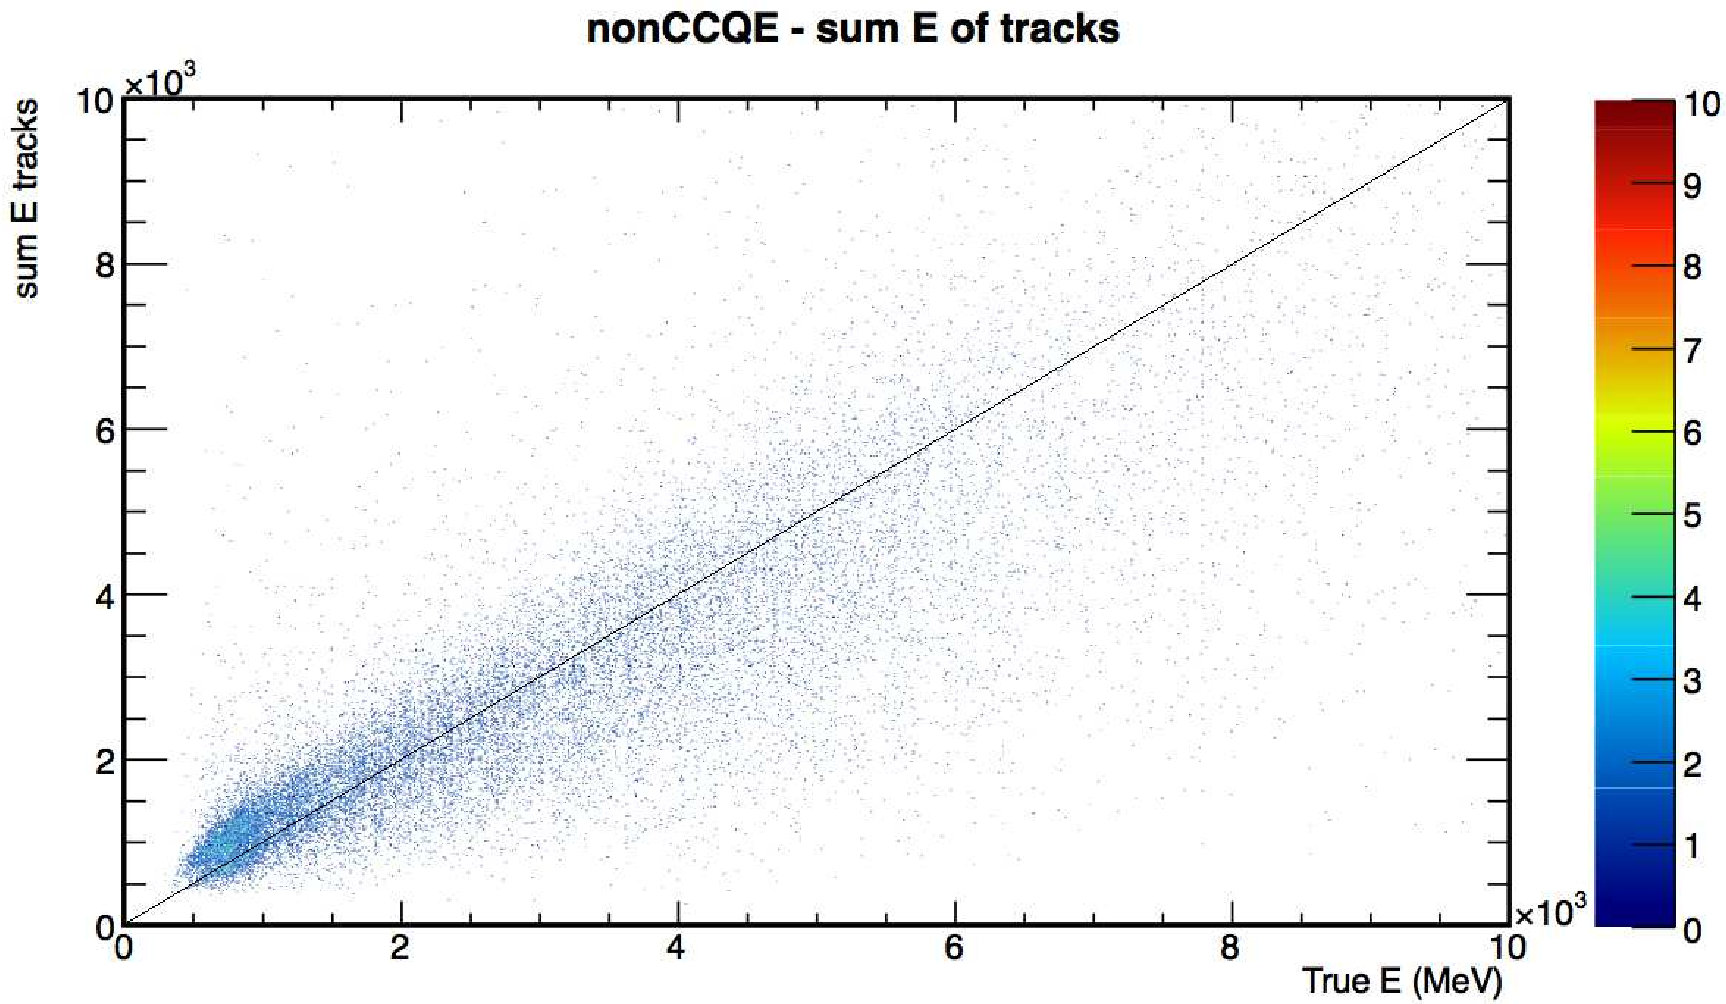

Figure 7.18: Test of the energy reconstruction for CCnonQE events. For selected MC interactions the reconstructed energy (sum E tracks) is plotted against the energy with which the particle was produced (True E). A strong positive correlation is shown, indicating that the energy reconstruction is good. Black line showing sum E tracks = True E plotted for reference.

by the neutrinos produced in pion decays. The CCnonQE spectrum does not feature this prominent low energy peak, but instead has a very broad shape and contains more high-energy neutrinos. This reflects the fact that a higher-energy neutrino is more likely to transfer the energy required to generate resonances, or lead to deep inelastic scattering, than a low energy neutrino, and so these CC-nonQE modes tend to correspond to interactions with higher energy neutrinos.

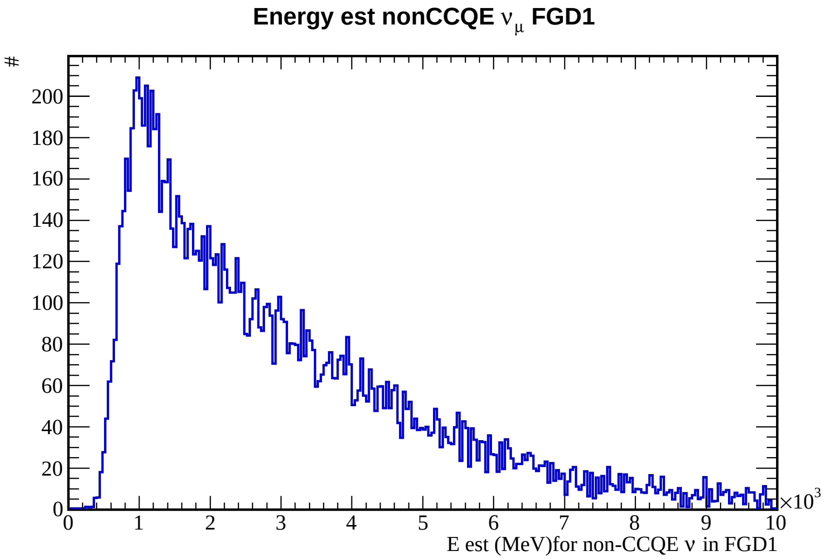

Figure 7.19: The energy spectrum of selected ν_μ MC events plotted using the reconstructed energy as calculated using Eq. 7.13. CCnonQE interactions only. The peak corresponding to the Δ^{++} resonance is clearly visible at 1232 MeV.

7.4.3 Reconstructed energy spectra

Using these methods of energy reconstruction described above, we can now plot the energies of the ν_μ interacting in the regions of interest in FGD1. As in §7.3 we can use the vertex position measurements to select ν_μ interacting in the more off-axis region, Q1, and ν_μ interacting in the more on-axis region, Q3, and plot these spectra separately, this time using the reconstructed neutrino energies. These spectra are shown in Figure 7.20, and show that the difference in the two spectra is still clearly resolvable when using reconstructed energy. A difference of approximately 80 MeV is observed between the peak positions, which is in agreement with the expected difference in peak energy due to the average off-axis angle difference between Q1 and Q3 of ∼0.3 degrees[8].

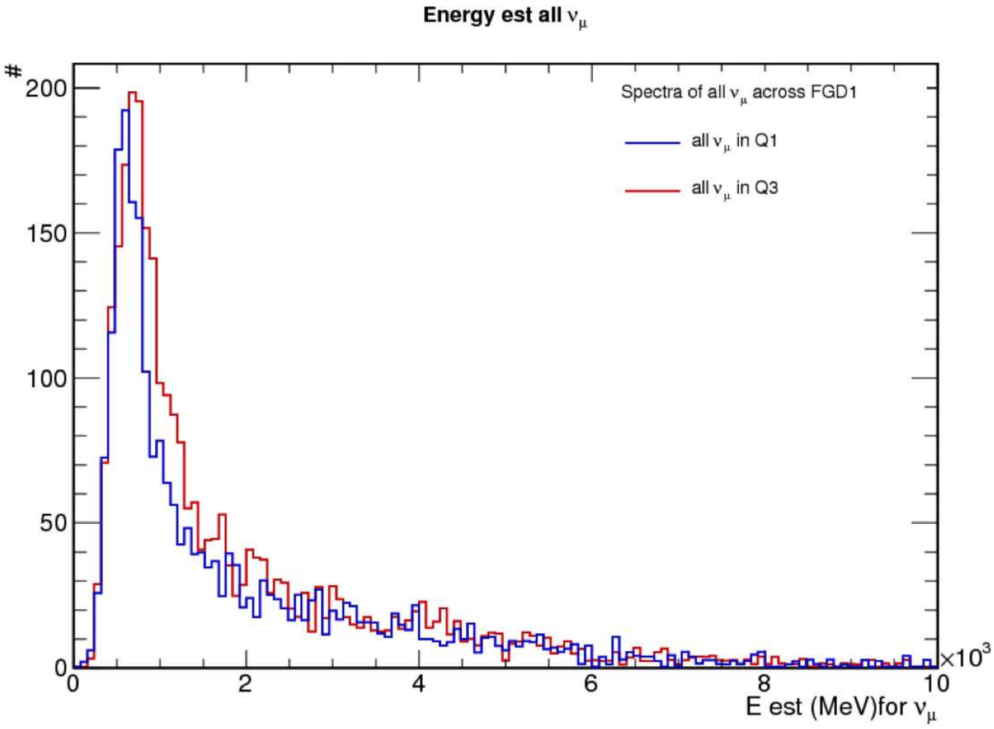

Figure 7.20: Comparison of the reconstructed energy spectra of selected ν_μ (CCQE and CCnonQE combined) interacting in Q1 and Q3. The shift in peak position due to the off-axis angle difference between Q1 and Q3 is visible, with the Q1 peak at ∼570±30 MeV and the Q3 peak at ∼650±30 MeV.

7.5 Summary

This chapter has discussed the differing behaviour of ν_μ produced from K decays, ν_μ^{fK}, and ν_μ produced from π decays, $\nu_\mu^{f\pi}$, and the reasons for the differences. A ν_μ selection has been defined, and methods for reconstructing the energy of ν_μ interacting in FGD1 have been developed. We have examined the changes to the neutrino spectrum with off-axis angle, and confirmed that while ND280 subtends

[8] 0.3 degrees is equivalent to 5.2 mrad. An off-axis angle shift of 1 mrad results in an energy shift of ∼16 MeV [100].

a relatively small range of off-axis angles, the sensitivity of the neutrino spectrum to off-axis angle changes is strong enough that the shift in position of the energy spectrum peak is visible within the detector.

The mean difference in off-axis angle between a ν_μ interacting in Q1 and a ν_μ interacting in Q3 is $\sim$0.3 degrees. By plotting the spectra of the ν_μ interacting in these regions separately we can utilise this off-axis angle behaviour to tell us more about the neutrino beam. Several methods of using this characteristic of the beam neutrinos to extract information about the neutrino parents have been considered and tested, and the method that is most effective for this level of spectrum shift and with the available statistics is described in the following chapter.

Chapter 8

Measuring K and π beam content using the ND280

Using the event selections and energy reconstruction described in Chapter 7, along with Monte Carlo truth information, we create probability density functions for ν_μ^{fK} and $\nu_\mu^{f\pi}$ interacting in different regions of FGD1. A fitting technique is developed to find the combinations of these distributions which provide the best fit to data, where both detector regions are taken into account. A method of finding corrections to the interaction cross-sections at different energies is included. Several tests are conducted using toy Monte Carlo, including statistical bias tests and investigating the effects of potential systematic uncertainties.

8.1 Introduction

8.1.1 Changes in energy spectrum with off-axis angle and neutrino parent

Using the event selection criteria described in Chapter 7, along with the energy reconstruction developed for suitable ν_μ events, we now have the tools to plot the reconstructed energy spectra for ν_μ interacting in regions of FGD1. Using the quadrant boundary definitions given in §7.3, the Q1 and Q3 spectra can be plotted separately.

In Chapter 5 the beam fluxes according to parent type were plotted and this demonstrated the different parents contributing to the ν_μ beam. These decays are listed in Table 5.1. Counting the 3 types of kaon decay together, there are contributions from kaons, pions, and muons. We see from Figure 5.1 that the contribution from muon decays is very small compared to the kaon and pion contributions. In approximately 40,000 selected Monte Carlo ν_μ interactions, only 39 involve a ν_μ which is the product of a muon decay. Therefore we will consider this to be a negligible contribution to the beam, and consider the total ν_μ spectrum to be merely a combination of the ν_μ from all K parents (ν_μ^{fK}) and the ν_μ from π^+ parents ($\nu_\pi^{f\pi}$).

When processing Monte Carlo interactions, it is possible to use the truth information with which the MC event was generated to divide these spectra according to parent type. ν_μ interacting in Q1 and Q3 are selected as before, but now before plotting their reconstructed energies, we access the parent type information. For each quadrant, the ν_μ^{fK} and $\nu_\mu^{f\pi}$ are plotted separately, resulting in four energy spectra. These are plotted together in Figure 8.1 for comparison.

As expected, the familiar peak at 600 MeV is formed largely from $\nu_\mu^{f\pi}$, and the slight difference in peak position between the Q1 and Q3 $\nu_\mu^{f\pi}$, caused by the change in off-axis angle as discussed in Chapter 7, is visible here. These $\nu_\mu^{f\pi}$ spectra tail off quickly and at approximately 2.5 GeV the two $\nu_\mu^{f\pi}$ spectra and the two ν_μ^{fK} spectra cross, and the ν_μ^{fK} spectra dominate thereafter. The ν_μ^{fK} spectra have the very broad peak that we would expect based on our knowledge of the kaon off-axis decay behaviour. The Q1 and Q3 ν_μ^{fK} spectra do not show a noticeable difference in shape, as we expected from our discussion of the off-axis behaviour of kaon decays in Chapter 7. We see that the ν_μ^{fK} spectra peak at approximately 4 GeV

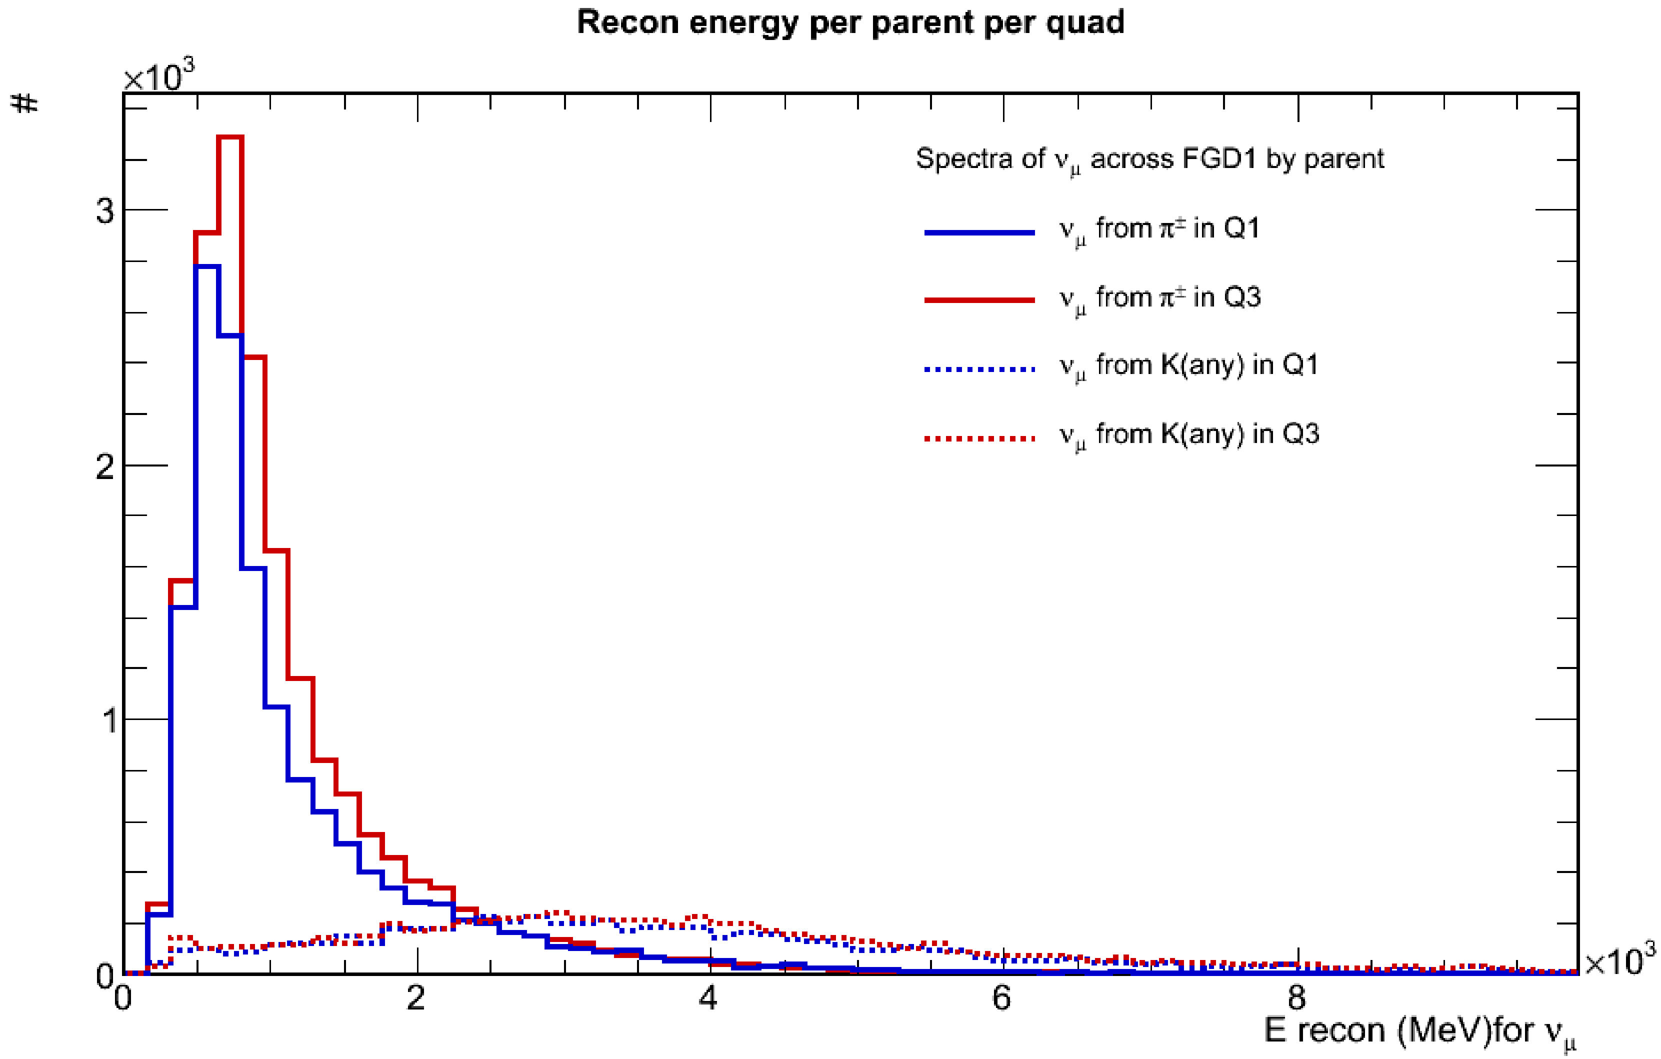

Figure 8.1: Reconstructed energy spectra for ν_μ, divided according to parent type, with pion contributions given by solid lines and kaon contributions represented by dashed lines. Interactions in Q1 are shown in blue, and interactions in Q3 are shown in red.

and extend all the way up to 10 GeV, so this confirms that the vast majority of the high energy neutrinos, which are responsible for many of T2K's background events, come from kaon parents. The more accurate our knowledge of these ν_μ^{fK}, the better we can understand our sources of background.

8.1.2 Analysis aim

The aim of this analysis is to measure the fraction of beam ν_μ which are produced from the decay of kaons, as opposed to pions. We have established that, omitting the negligible contribution from muon decays, the total supply of beam ν_μ is a combination of the ν_μ^{fK} and $\nu_\mu^{f\pi}$. Therefore the total beam spectrum observed must be the addition of some amount of the ν_μ^{fK} spectrum and a complementary amount of the $\nu_\mu^{f\pi}$ spectrum. This is true for the whole of FGD1, and it is also true for each of the quadrants Q1 and Q3 individually, as both the total beam spectrum and the parent-specific spectra will change slightly as we move between detector regions. If we know the shape of the separate ν_μ^{fK} and $\nu_\mu^{f\pi}$ spectra, we can generate different combinations of these spectra and check their compatibility with the observed total beam spectrum. By fitting test combinations to the observed data spectrum in this way, we can find the combination that gives the best match, and the relative scaling of each of the ν_μ^{fK} and $\nu_\mu^{f\pi}$ spectra that gave this optimum match will therefore tell us the contribution from each parent type.

In order to make a measurement of the kaon contribution to the beam, it is important to understand all of the sources of uncertainty in our current kaon and pion ratios. As discussed in Chapter 6, NA61/SHINE provides information about the fluxes of π and K produced at a target; however, the uncertainties on these are larger for kaons, and especially at high energies, where they are as high as 100%. The spectrum shapes are constructed using the Monte Carlo simulation.

Extracting the separate ν_μ^{fK} and $\nu_\mu^{f\pi}$ spectra from our Monte Carlo as done here provides spectra which incorporate the best current theoretical and experimental information available to T2K. However, the uncertainty in the relative contributions from pions and kaons, and uncertainty in the exact shapes of these parent specific spectra, are still present in this Monte Carlo information.

When fitting spectra in this way we are sensitive to both of these sources of uncertainty. Testing all combinations of potential parent-specific spectrum shapes as well as the relative contributions of these spectra is possible using a fitting method, but beyond the scope of this thesis. Since we think that errors in the $\nu_\mu^{f\pi}$ and ν_μ^{fK} spectral shapes could be reasonably small, and that the main source of uncertainty is the number of ν_μ^{fK} present in our beam, we will use the shapes provided by the Monte Carlo truth information, as shown in Figure 8.1, and only vary the contributions of each of these spectra. Testing the effect of making modifications to the shapes of these spectra is a possible extension to this work, and is discussed in Chapter 9.

8.1.3 Combining Monte Carlo samples

The Monte Carlo spectra should be as precise in shape as possible, so the higher the number of interactions used to form the spectra, the smoother the spectral shapes will be. There are several different productions of Monte Carlo available to use, and by combining three of the appropriate sets of Monte Carlo available, we can achieve a total of 39946 interactions in Q1 and Q3 combined, where 17976 of these interactions fall in Q1, and 21979 fall in Q3.

The three Monte Carlo sets used have some slight differences in configuration. The two largest sets are labelled ‘beam b’, and the smallest is ‘beam c’, where the beam labels refer to the beam specifications (beam configurations and MC production details are provided in Appendix B). The number of bunches, bunch separation and bunch duration are all the same for beam b and beam c; however, the beam power, beam repetition and protons on target per spill change. The 'b' configuration includes a beam of 120 kW, with repetition time 3.2 s, and contains 7.9891 $\times 10^{13}$ POT per beam spill. The 'c' configuration includes a higher intensity of 178 kW, the lower repetition time of 2.56 s, and 9.463 $\times 10^{13}$ POT per beam spill. By increasing the beam intensity and POT it is possible that a higher number of interactions could occur in each beam spill, and the possibility of multiple interactions in a bunch increases, which may affect the event selection. The other difference between the productions is the contents of the P0D, which can be either air or water. The larger of the two 'beam b' Monte Carlo sets is configured to have a water-filled P0D, while the smaller 'beam b' set includes an air-filled P0D. The one 'beam c' Monte Carlo set is configured with air. This difference will change the cross-section for interactions in the P0D, but since we are selecting events which interact in FGD1, and discarding any events with P0D activity, this difference should not affect the spectrum shapes used by this analysis. However, it is imperative that the parent and quadrant-specific spectral shapes found by these three Monte Carlo sets are identical if they are to be combined, therefore to check, each set is compared to each other set before combining.

For each Monte Carlo production being used, we have 4 spectra; $\nu_\mu^{f\pi}$ in Q1, $\nu_\mu^{f\pi}$ in Q3, ν_μ^{fK} in Q1 and ν_μ^{fK} in Q3. The three different productions correspond to different numbers of protons on target, so in order to perform comparisons between the different productions all spectra are normalised, enabling comparison of shape without the need to consider scale. First a visual comparison is performed. As examples, the comparisons for each of the 4 spectra between 'beam b, air' and 'beam c, air' are plotted in Figures 8.2, 8.3, 8.4 and 8.5.

In each plot the bin contents from the two different Monte Carlo sets are

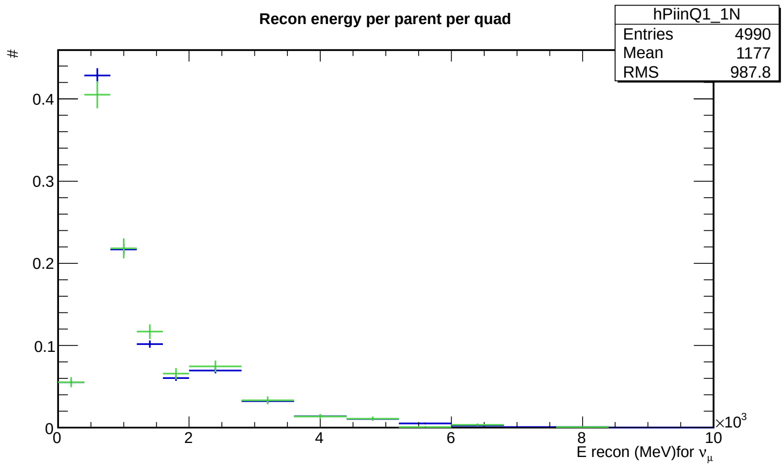

Figure 8.2: Monte Carlo comparisons of $\nu_\mu^{f\pi}$ interacting in Q1, where 'beam b, air' MC is plotted in blue and 'beam c, air' MC is plotted in green. Error bars show statistical uncertainties.

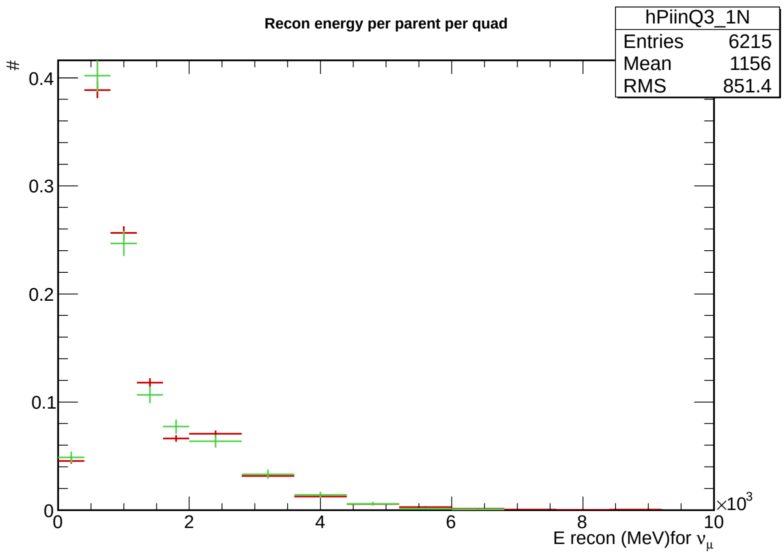

Figure 8.3: Monte Carlo comparisons of $\nu_\mu^{f\pi}$ interacting in Q3, where 'beam b, air' MC is plotted in red and 'beam c, air' MC is plotted in green. Error bars show statistical uncertainties.

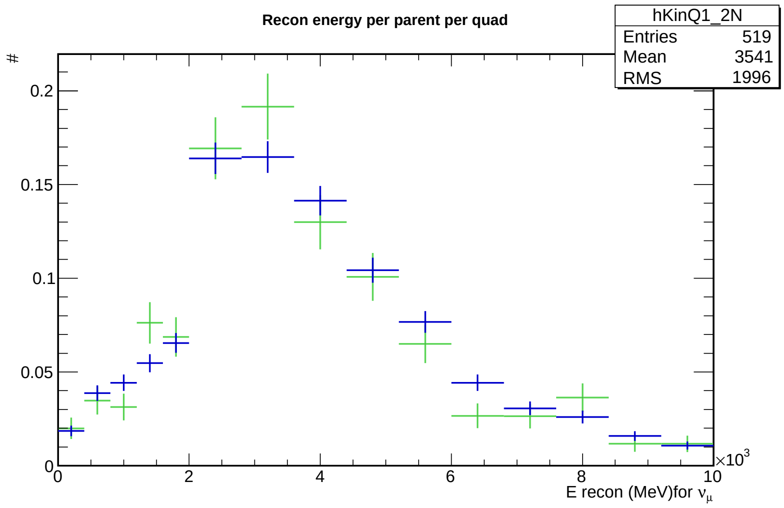

Figure 8.4: Monte Carlo comparisons of ν_μ^{fK} interacting in Q1, where 'beam b, air' MC is plotted in blue and 'beam c, air' MC is plotted in green. Error bars show statistical uncertainties.

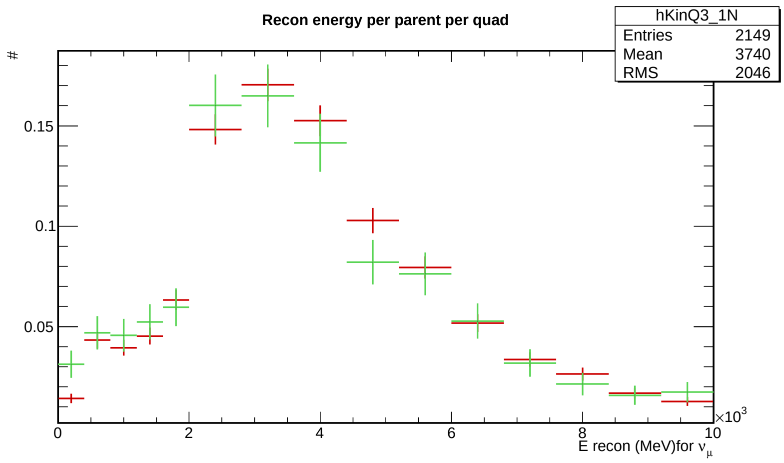

Figure 8.5: Monte Carlo comparisons of ν_μ^{fK} interacting in Q3, where 'beam b, air' MC is plotted in red and 'beam c, air' MC is plotted in green. Error bars show statistical uncertainties.

displayed in different colours, with statistical uncertainties included. Allowing for statistical fluctuations, it is clear that the bin contents are in agreement in all plots. The error bars overlap in the majority of bins, and an examination of the colour coding shows that the distribution with the greater number of events switches from bin to bin as you move across each distribution range, supporting the conclusion that any small differences in bin contents are due to statistical fluctuations. It is possible that some small variation in shape could be hidden by the statistical fluctuations however, so in addition to a visual confirmation we also perform a more quantitative test. A Kolmogorov-Smirnov test, or KS test, is particularly well suited to testing the compatibility of two distributions as rather than simply comparing individual bin contents, it also compares the cumulative bin contents across the range of the distributions, and in this way is sensitive to differences occurring in adjacent bins. This makes the test particularly sensitive to changes in shape between the distributions being compared. Each set of matching spectra was compared using the ROOT TMath function KolmogorovTest [101], and the probability outputs are given in Table 8.1. For compatible histograms, the outputs should be uniformly distributed between 0 and 1, while for incompatible histograms, the distribution is peaked close to 0. Only histograms with a test value of greater than 0.05 should be considered compatible, which is true for all of the spectra tested here.

Using the described event selection and energy reconstruction, the ν_μ spectrum can be plotted using T2K data. When ND280 data is processed, the position of the interaction vertex is reconstructed. Using this information and the quadrant definitions defined above, it is possible to plot the ν_μ data spectra separately for Q1 and Q3. We can therefore add extra sensitivity to a fitting programme by considering Q1 and Q3 separately. The fact that the Q1 and Q3 spectra differ allows us to test the success of each combination of ν_μ^{fK} and $\nu_\mu^{f\pi}$ spectra twice -

MC set 1	MC set 2	Spectrum	Kolmogorov Output
beam b, air	beam b, h2o	ν_μ^{fK} in Q1	0.999953
		ν_μ^{fK} in Q3	0.983415
		$\nu_\mu^{f\pi}$ in Q1	0.454657
		$\nu_\mu^{f\pi}$ in Q3	0.442377
beam c, air	beam b, h2o	ν_μ^{fK} in Q1	0.518468
		ν_μ^{fK} in Q3	0.613705
		$\nu_\mu^{f\pi}$ in Q1	0.64724
		$\nu_\mu^{f\pi}$ in Q3	0.862395
beam c, air	beam b, air	ν_μ^{fK} in Q1	0.370129
		ν_μ^{fK} in Q3	0.291873
		$\nu_\mu^{f\pi}$ in Q1	0.544137
		$\nu_\mu^{f\pi}$ in Q3	0.825266

Table 8.1: Kolmogorov test outputs for each combination of Monte Carlo spectra.

once in Q1 and once in Q3. Test quantities of ν_μ^{fK} in Q1 and the $\nu_\mu^{f\pi}$ in Q1 can be added together and compared to the ν_μ data in Q1, and a parameter quantifying the agreement can be recorded. This can be repeated for Q3, using the same test ratio of ν_μ^{fK} and $\nu_\mu^{f\pi}$, and a second agreement parameter recorded. These two tests of the compatibility can then be combined for each ratio being tested, thus making the fit more powerful.

8.1.4 Choice of binning for energy spectra

For both the data and Monte Carlo we plot the reconstructed energies of interactions in the range 0-10 GeV. These spectra need to be binned appropriately. The lower end of the spectrum contains the highest number of entries since the $\nu_\mu^{f\pi}$ peak is located here, and it is this peak which shows the clearest change between off-axis quadrants, which is a detail we wish to preserve. The event rate drops as you move higher in energy. Bins with no entries should be avoided during the fit, however the fine detail of the $\nu_\mu^{f\pi}$ peak must be retained, therefore bins of equal width are not suitable. Using the POT corresponding to the available data sets, we can scale the number of interactions selected in the Monte Carlo to estimate the number of interactions that would be present when looking at the data. Considering these numbers, several binning regimes were tested and a set of 15 bins was chosen, with 5 bins of width 400 MeV in the range 0 - 2 GeV, and 10 bins of width 800 MeV in the range 2 - 10 GeV. Using this binning system the shift in peak position is still visible, and there is a suitable number of entries in the highest energy bins.

8.2 The likelihood function

Consider an energy bin i of one of our off-axis quadrant spectra. The number of observed entries in this bin, n_i, is fixed, as it is the number of events observed in our data set. We wish to construct a test spectrum using a combination of our parent-specific spectra to compare to this data spectrum. A ratio of the number of ν_μ^{fK} to the number of $\nu_\mu^{f\pi}$ to be tested is picked and the ν_μ^{fK} and $\nu_\mu^{f\pi}$ spectra are combined according to this ratio, to create a test total ν_μ spectrum. The expected events in our selected bin, denoted by μ_i, is then the number of entries in energy bin i, of this test spectrum. Assuming that this test spectrum is an accurate description of the true spectrum, the probability that an observation run would find n_i entries in our data can be found by the Poisson probability P, where

$$P_i(n_i;\mu_i) = \frac{\mu_i^{n_i} e^{-\mu_i}}{n_i!} \tag{8.1}$$

If the test spectrum we create is very unrealistic, the bin contents of the test spectrum will not be a close match to the corresponding data bin contents, and so in each bin the probability that the test spectrum represents the distribution observed by our data measurements will be very small. The more accurate our test spectrum becomes, the more closely the test bin contents and data bin contents will match, and the larger the Poisson probabilities P_i will be. We wish to measure the compatibility of the whole spectrum, so to do this we combine the values of P calculated for each bin by multiplying them. The likelihood, L, is defined as the product of the Poisson probabilities for each of the 15 bins in our spectra, such that

$$L = \prod_i P_i = \prod_i \frac{\mu_i^{n_i} e^{-\mu_i}}{n_i!} \tag{8.2}$$

In the case of a good match, each probability included in the product will have a relatively large value, resulting in a larger value of L. The larger the value of L, the closer our test spectrum is to the observed spectrum. Therefore to find the most accurate quantities of ν_μ^{fK} and $\nu_\mu^{f\pi}$ in the data, we look for the combination of the ν_μ^{fK} and $\nu_\mu^{f\pi}$ spectra which give the largest value of L. This is called the maximum likelihood technique.

For each bin the value of P will be less than 1 and often rather small, so that the product of the values of P for the 15 bins will be a very small number. It is convenient to take the natural logarithm of the likelihood, as this aids mathematical manipulation and avoids small numbers when using a computer. Using the relation

$$ln(ab) = ln(a) + ln(b) \tag{8.3}$$

it follows that the logarithm of the product becomes the sum over the energy bins, as follows.

$$lnL = \sum_i lnP_i = \sum_i ln\frac{\mu_i^{n_i} e^{-\mu_i}}{n_i!} \tag{8.4}$$

Using Eq. 8.3, along with the following relations

$$ln(\frac{a}{b}) = ln(a) - ln(b) \tag{8.5}$$

$$ln(a^b) = b.ln(a) \tag{8.6}$$

$$ln(a!) = a.ln(a) - a \tag{8.7}$$

we can expand Eq. 8.4 as follows[1]:

$$\begin{aligned} lnL &= \sum_i ln\frac{\mu_i^{n_i}e^{-\mu_i}}{n_i!} && (8.8)\\ &= \sum_i [ln(\mu_i^{n_i}) + ln(e^{-\mu_i}) - ln(n_i!)] && (8.9)\\ &= \sum_i [n_i ln(\mu_i) - \mu_i - (n_i ln(n_i) - n_i)] && (8.10)\\ &= \sum_i [n_i - \mu_i - n_i ln(\frac{n_i}{\mu_i})] && (8.11) \end{aligned}$$

We now have an expression for the log-likelihood calculated over all 15 energy bins in our 10 GeV range. However, so far this only takes into account one set of corresponding test and data spectra. We have the observed ν_μ spectra, and the parent-specific Monte Carlo spectra, separately for each of Q1 and Q3. Therefore, we can construct test distributions and compare these to data twice - once in each off-axis quadrant. There are two sets of 15 energy bins to calculate the Poisson probabilities for, and these 30 values should all be included in the log-likelihood sum. We therefore extend Eq. 8.11 to sum over the quadrants.

[1]Stirling's Approximation is used for Equation 8.7. Higher order terms can be neglected since the n_i (where this approximation is applied in Eq. 8.10) are fixed, therefore the same constants are added for each test, which does not affect our result.

$$lnL = \left(\sum_i [n_i - \mu_i - n_i ln(\frac{n_i}{\mu_i})]\right)_{Q1} + \left(\sum_i [n_i - \mu_i - n_i ln(\frac{n_i}{\mu_i})]\right)_{Q3} \tag{8.12}$$

$$= \sum_{i=1}^{15} \sum_{Q1,Q3} [n_{iQ} - \mu_{iQ} - n_{iQ} ln \left(\frac{n_{iQ}}{\mu_{iQ}}\right)] \tag{8.13}$$

where we define n_{iQ} and μ_{iQ} to be the observed and expected numbers of events respectively in energy bin i ($i = 1...15$) of the spectrum corresponding to Q = Q1 or Q3. The values of n_{iQ} can be read directly from the binned data histograms. The values of μ_{iQ} are dependent on the combination of ν_μ^{fK} and $\nu_\mu^{f\pi}$ we are testing.

8.3 Finding the expected events

The contributions to the neutrino beam from each parent type can be described in terms of the total numbers of ν_μ^{fK} and $\nu_\mu^{f\pi}$ in our data sample. Using the number of interactions found in Q1 and Q3 in our Monte Carlo selection and comparing the POT corresponding to the Monte Carlo used and the POT of our available data, we find that we would expect to select approximately 3500 ν_μ interactions in Q1 and Q3 of our data sample. We will define the total ν_μ selected in Q1 and Q3 in our data sample as $N_{\nu_\mu}^{data}$. $N_{\nu_\mu}^{data}$ must be the sum of the quantities of ν_μ^{fK} and $\nu_\mu^{f\pi}$. We define the total number of ν_μ^{fK} in Q1 and Q3 combined as T_K and the total number of $\nu_\mu^{f\pi}$ in Q1 and Q3 combined as T_π. Nominally

$$N_{\nu_\mu}^{data} = T_K + T_\pi \tag{8.14}$$

but in this fit each of T_π and T_K are allowed to vary independently.

Constructing the different test spectra for comparison with data can be done in terms of these T values. By allowing each of T_K and T_π to vary between 0 and

$N_{\nu_\mu}^{data}$, we can calculate the log-likelihood of all possible beam compositions that could occur.

8.3.1 Probability density functions

To find the number of entries in a particular bin of a spectrum, we multiply the total number of entries in that spectrum by the probability of an entry being in that bin. We can use this principle to find the expected number of entries in a bin in the Q1 and Q3 test spectra. To do this we need to generate probability distribution functions for our parent-specific distributions.

The four theoretical spectra for ν_μ^{fK} and $\nu_\mu^{f\pi}$ in Q1 and Q3 are shown in Figure 8.1. If we normalise each of these spectra we will obtain the probability distribution functions for each of the spectra individually. However, we would then need separate totals for the different quadrants, since the drop in flux with increasing off-axis angle means that Q1 always contains fewer interactions than Q3.

Rather than divide the T values into quadrant totals, we keep one total for the ν_μ^{fK} in both Q1 and Q3 (T_K), and one total for the $\nu_\mu^{f\pi}$ in both Q1 and Q3 (T_π), and instead include the changing flux information in the probability distribution functions. This can be achieved by normalising the Monte Carlo spectra in pairs according to parent type. This is demonstrated in Figure 8.6.

After scaling, the integrals of the Q1 and Q3 spectra for each parent type add to 1. The new Q1 integrals will always be slightly under 0.5 due to the relatively lower flux, and the new Q3 integrals will be slightly higher than 0.5, reflecting the fact that an interaction is more likely in Q3 than in Q1 due to the higher flux.

The probabilities contained in these pdfs are labelled as P_{iQ}^{p}, where p indicates

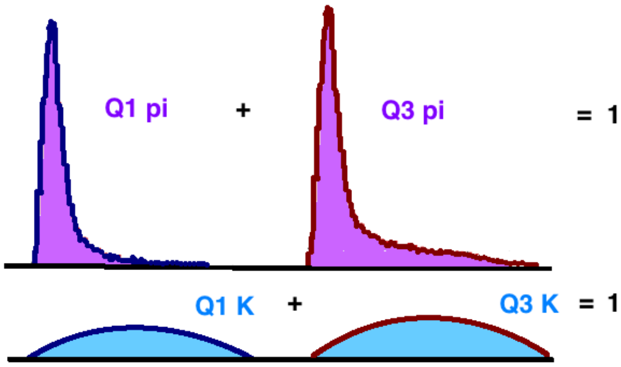

Figure 8.6: Linked probability density functions. Distributions of events in Q1 and Q3 are shown in blue and red respectively. Notice the integrals of the Q3 spectra are slightly larger than the Q1 integrals due to the difference in neutrino flux experienced by these quadrants. The upper distributions, featuring the prominent narrow peak, represent the $\nu_{\mu}^{f\pi}$ spectra and are scaled to give a combined integral of 1. The lower distributions, consisting of very broad peaks, represent the ν_{μ}^{fK} spectra, and are also normalised as a pair.

the parent type, which will be one of K or π. As an example, the probability of a neutrino produced by the decay of a pion, with reconstructed energy in the range 2.0 - 2.8 GeV, which corresponds to energy bin 6, interacting in Q1, would be described as $\mathrm{P}^{\pi}_{6,Q1}$. The total number of $\nu_{\mu}^{f\pi}$ interactions in Q1 in that energy range is then given by $\mathrm{P}^{\pi}_{6,Q1}.\mathrm{T}_{f\pi}$.

The total number of ν_{μ} interactions in an energy range and in a specific quadrant is the combination of neutrinos from pion decays and neutrinos from kaon decays. Therefore the expected number of events, per energy bin i and off-axis quadrant Q, is given by the following expression.

$$\mu_{iQ} = [P^{\pi}_{iQ}T_{\pi} + P^{K}_{iQ}T_{K}] \tag{8.15}$$

The probabilities will be read from the 4 probability distribution functions generated from the Monte Carlo and will not change. The values of T_{π} and T_{K}

are fixed appropriately for each test being conducted, and it is in this way that the expected events will represent the test distribution being studied, so that it can be compared to data and a value of the log-likelihood can be calculated.

8.3.2 Cross-section uncertainties

Equation 8.15 is correct assuming that the interaction cross-sections used in the Monte Carlo generation are exactly correct. Should the cross-section have a different value in a certain energy range, the number of interactions we predict in that energy range will not be accurate, and this would affect the shape of the neutrino spectrum. For example, consider the high energy tail, where we know that the spectrum is dominated by neutrinos produced via kaon parent decays. If the interaction cross-section has been underestimated in this region, we will anticipate too few interactions, and our probability distribution functions for kaons at these energies will also be too low. When comparing test spectra to data in that region, which will have the higher, true number of interactions, tests which include erroneously high values of T_K will appear to be a better match since they will have a higher number of entries in the higher energy bins. Therefore uncertainties in the cross-section values used in the Monte Carlo generation would be carried through to the pdf shapes, and that would cause the fitting programme to favour incorrect values of T_K and T_π.

The expression for μ_{iQ} can be modified to take account of this potential source of error in the Monte Carlo. The expected events are constructed using the Monte Carlo probability density functions, however we introduce an additional factor, c, which scales the entries in a given energy bin up or down as necessary. The interaction cross-section at a given energy will not change between Q1 and Q3, as these quadrants are located in the same sub-detector and consist of identical materials.

The neutrino interaction cross-sections are also not sensitive to the parent of the neutrino interacting, so we do not need separate factors for different parent types. Since interaction cross-sections are energy dependent, and the uncertainties are also energy dependent, we could potentially require different correction factors in different regions of the energy spectrum. Therefore we must have a different value of c for each energy bin i. The modified version of Eq. 8.15 is therefore as follows:

$$\mu_{iQ} = c_i[P^{\pi}_{iQ}T_{\pi} + P^{K}_{iQ}T_{K}] \tag{8.16}$$

where $c_i = 1$ when the cross-sections used in that bin i are correct, and is allowed to vary from 1 if the entries in the bin need to be scaled up or down to compensate for an incorrect cross-section used in the Monte Carlo. This changes the fit from having 2 parameters and 30 bins to 17 parameters and 30 bins.

Everything in this expression is either known or set by us, except the values of c_i. In order to calculate lnL for each test setting of T_π and T_K we must first find the correction factors for each bin. Since the aim is to maximise lnL, in each bin we desire the value of c_i that would give the largest (most positive) value of lnL. We can find this value analytically by finding the partial derivative of lnL with respect to c_i and setting the first derivative equal to zero and solving for c_i.

However, before differentiating, we can add an extra term. While some freedom to alter the cross-sections must be allowed, this should be secondary to finding the most accurate values of T_π and T_K. We know that the cross-sections may include uncertainties, but the current uncertainties in the numbers of ν_μ^{fK} in the beam are larger. Therefore we would preferentially let T_π and T_K vary from the Monte Carlo ratio rather than include large corrections to the cross-section values. If we differentiate the expression for lnL as given by Eq. 8.13 (with μ_{iQ} defined as in Eq. 8.16), excessively large values of c_i may be returned, as any discrepancies between the observed events n_{iQ} and the expected events μ_{iQ} will be accounted for by assuming an incorrect cross-section. This would also result in inaccurate values of T_π and T_K being presented as the most favourable matches, when in fact the T_π and T_K are not good matches at all, but the lnL result is artificially large due to large correction factors altering the bin contents.

Therefore we include an extra term which influences the log-likelihood to reflect the extent of the suggested corrections to the cross-sections. We do this by adding a penalty term to the expression, such that the full expression for the log-likelihood is

$$lnL = \sum_{i=1}^{15} \sum_{Q1,Q3} [n_{iQ} - \mu_{iQ} - n_{iQ} ln(\frac{n_{iQ}}{\mu_{iQ}})] - \sum_{i=1}^{15} \frac{(c_i - 1)^2}{2\sigma_i^2} \tag{8.17}$$

where σ_i is the known uncertainty on the cross-section in energy bin i. The form of the penalty term is influenced by the form of the pdf of a Gaussian distribution, given by

$$\mathrm{P}_{gaus} = \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) \tag{8.18}$$

Therefore

$$ln\mathrm{P}_{gaus} = \left(-\frac{(x-\mu)^2}{2\sigma^2}\right) \tag{8.19}$$

The penalty term will be zero if the Monte Carlo cross-sections are correct, and therefore $c_i = 1$, and gets larger as a value of c strays further from 1. This has the effect of reducing the overall log-likelihood of any scenario in which the cross-sections are found to need large corrections. Consider an example where a very extreme and unlikely combination of T_π and T_K result in a high log-likelihood, but only due to unrealistically large cross-section corrections altering the expected bin contents. This is not a desirable outcome as the likelihood value is only so high due to the cross-sections being altered considerably, and we know that the cross-section uncertainties are reasonably small, so this is not a probable solution. In this case, the large corrections would cause the penalty term to be large, and when this is subtracted from the sum of probabilities, it reduces the final value of the log-likelihood. This test would therefore no longer appear to be a good match, reflecting the fact that the bin contents only match because the cross-sections have been altered beyond allowed levels.

The scale of the penalty values is set using our knowledge of the accuracy of the cross-sections. The GENIE group have made a compilation of measurements of the uncertainties on the neutrino interaction cross-sections and these are shown in Figure 8.7.

An average value for the fractional uncertainty in each of our defined energy bins can be read from this plot, and the values, labelled σ_i, are given in Table 8.2.

Energy bin / GeV	σ_i	Energy bin / GeV	σ_i
0 - 0.4	0.22	4.4 - 5.2	0.32
0.4 - 0.8	0.34	5.2 - 6.0	0.30
0.8 - 1.2	0.37	6.0 - 6.8	0.28
1.2 - 1.6	0.37	6.8 - 7.6	0.26
1.6 - 2.0	0.38	7.6 - 8.4	0.25
2.0 - 2.8	0.39	8.4 - 9.2	0.24
2.8 - 3.6	0.38	9.2 - 10.0	0.21
3.6 - 4.4	0.35		

Table 8.2: Fractional cross-section uncertainties by energy bin

These uncertainties are included in the denominator of the penalty term. If the uncertainties are very large, the penalty is small, since it would be reasonable for the cross-sections to need large corrections, and so the log-likelihood should not be affected severely by a large correction being needed. If, conversely, we knew the cross-sections to a extremely high degree of accuracy, we would not expect the cross-sections to change at all, and so any deviation from 1 for any value of c should cause the log-likelihood to drop significantly.

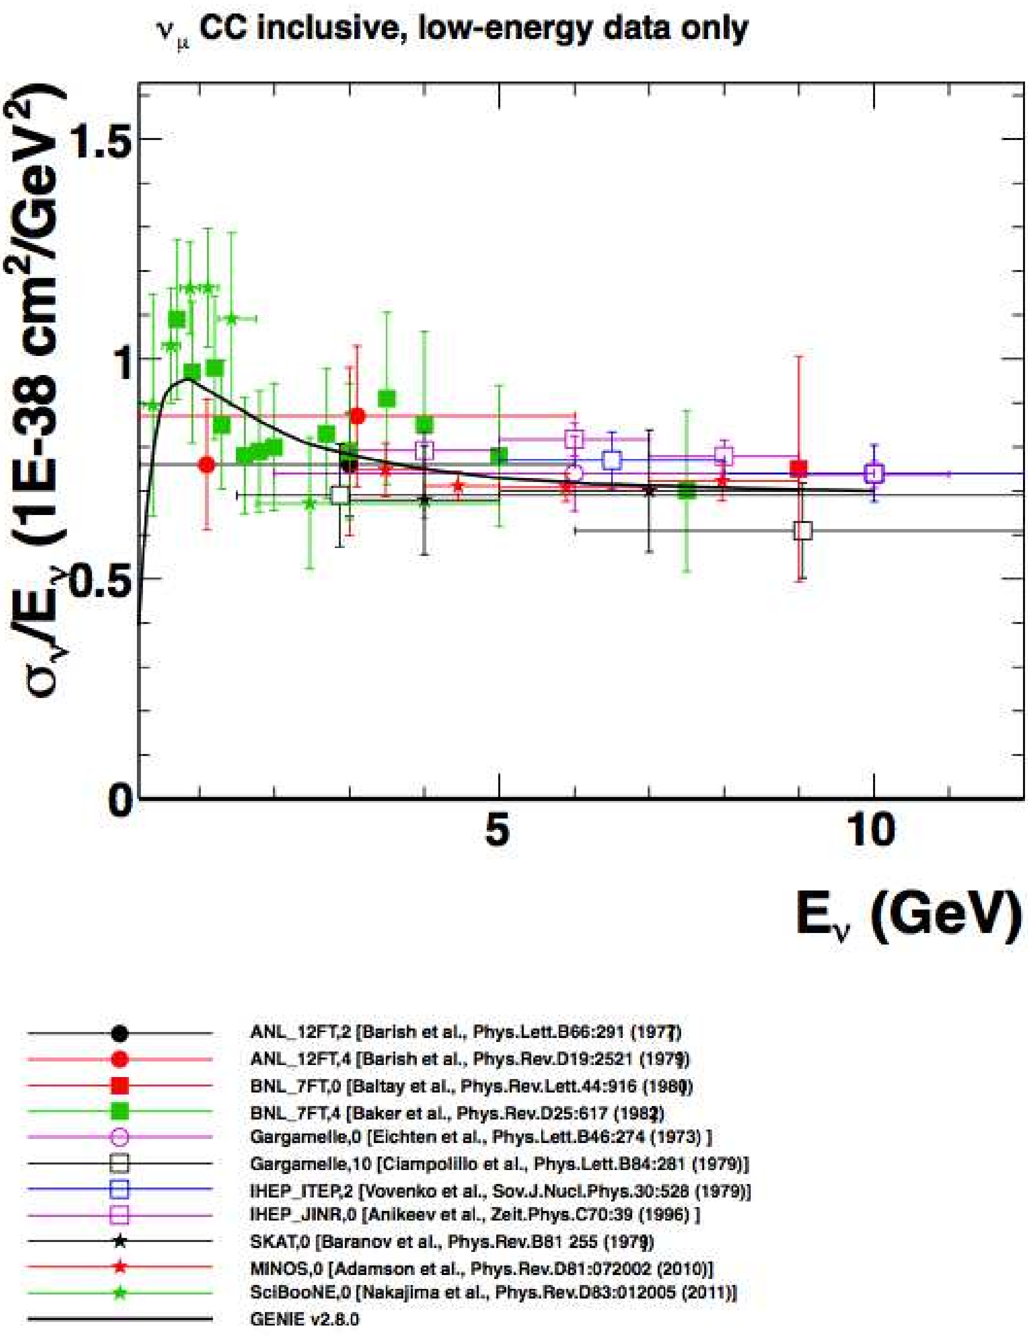

Figure 8.7: Cross-section uncertainties divided by neutrino energy for ν_μ interactions. Experimental measurements are shown, along with the GENIE values. Figure from [102].

In this way we favour combinations of T_π and T_K that best match the data spectra, with the added feature that while small alterations to the cross-sections are allowed, scenarios that only work if the cross-sections are very wrong will be penalised, and the extent to which the current cross-sections would need to be incorrect is another factor in the final likelihood values calculated.

Using the full expression given in Eq. 8.17 we can now maximise lnL with respect to each c_i. This can be done analytically by partial differentiation with respect to one of the c_i values, for example c_j, while holding all the other c_i ($i \neq j$) and T_π and T_K constant. The only terms in Eq. 8.17 which contribute are from the energy bin j.

$$lnL_j = n_{jQ1} - \mu_{jQ1} - n_{jQ1} ln\left(\frac{n_{jQ1}}{\mu_{jQ1}}\right) + n_{jQ3} - \mu_{jQ3} - n_{jQ3} ln\left(\frac{n_{jQ3}}{\mu_{jQ3}}\right) - \frac{(c_j - 1)^2}{2\sigma_j^2} \tag{8.20}$$

The n_{jQ} terms are constants read from the data spectra. The μ_{jQ} terms are given by Eq. 8.16, where the P terms are constants read from the probability distribution functions and the T values correspond to the test being conducted. Therefore

$$\frac{\partial lnL}{\partial c_j} = -[P^{\pi}_{jQ1} T_\pi + P^{K}_{jQ1} T_K] + \frac{n_{jQ1}}{c_j} - [P^{\pi}_{jQ3} T_\pi + P^{K}_{jQ3} T_K] + \frac{n_{jQ3}}{c_j} - \frac{(c_j - 1)}{\sigma_j^2} \tag{8.21}$$

To maximise lnL, we set Eq. 8.21 equal to zero and solve for c_j. Doing this gives an expression which is quadratic in c_j, as follows.

$$[P^{\pi}_{jQ1} T_\pi + P^{K}_{jQ1} T_K] + [P^{\pi}_{jQ3} T_\pi + P^{K}_{jQ3} T_K] - \frac{n_{jQ1}}{c_j} - \frac{n_{jQ3}}{c_j} + \frac{(c_j - 1)}{\sigma_j^2} = 0 \tag{8.22}$$

Using the analytical solution for quadratic equations, given by

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a} \tag{8.23}$$

where

$$ax^2 + bx + c = 0 \tag{8.24}$$

we can rearrange Eq 8.22 to find

$$a = \frac{1}{\sigma_j^2} \tag{8.25}$$

$$b = [P^{\pi}_{jQ1}T_{\pi} + P^{K}_{jQ1}T_{K}] + [P^{\pi}_{jQ3}T_{\pi} + P^{K}_{jQ3}T_{K}] - \frac{1}{\sigma_j^2} \tag{8.26}$$

$$c = -n_{jQ1} - n_{jQ3} \tag{8.27}$$

The quadratic solution for the first derivative implies that there are two values of c_j that correspond to turning points. Since we require values of c_j that result in maximum values of lnL, we require the local maxima on plots of lnL against c_j, and these can be found by investigating the second derivative of lnL with respect to c_j.

$$\frac{\partial^2 lnL}{\partial c_j^2} = \frac{-1}{c_j^2}(n_{jQ1} + n_{jQ3}) - \frac{1}{\sigma_j^2} \tag{8.28}$$

The different c_j solutions were investigated and it was found that the c_j values calculated using the positive square root in Eq. 8.23 always satisfied the inequality

$$\frac{\partial^2 lnL}{\partial c_j^2} < 0 \tag{8.29}$$

which is satisfied in the case of a local maximum. Therefore to find each value of c_i we take the solution to Eq. 8.23 (along with with Eq. 8.25, Eq. 8.26 and Eq. 8.27) with the positive root. By marginalising the c_i values in this way we find the values of the cross-section corrections which will result in the largest $\ln L_i$ values for each energy bin, whilst restricting the extent of the corrections by a degree consistent with the known uncertainties.

8.4 The fitting program

The 17 parameter fit to find T_π, T_K and $c_1 \dots c_{15}$ proceeds by maximising the likelihood using the analytical procedure from §8.3 to find $c_1 \dots c_{15}$ and scanning over all possible values for the remaining parameters T_π, T_K. The fitting procedure is implemented in a computer program. A brief outline of the program's mechanism is described here.

- The program requires six input spectra. The four parent-specific Monte Carlo spectra divided into Q1 and Q3 selections are read in first. Using integrals these four spectra are normalised in pairs, as described in section 8.3.1, to form the four required probability distribution functions. These p.d.f.s are the source of the P_{iQ} values used in the μ_{iQ} terms from Eq. 8.16.

- The data selections for Q1 and Q3, plotted using the predetermined binning, are then accessed. The values of n_{iQ} are found from these spectra.

- The integrals of the Q1 and Q3 data spectra are also taken and added together to find $N_{\nu_\mu}^{data}$. Two loops are then set up to test all the possible combinations of T_π and T_K, where each of T_π and T_K are allowed to be any integer between zero and $N_{\nu_\mu}^{data}$.

- A pair of T values is selected to be tested. First the 15 c_i values are found using the marginalisation process outlined above. These c_i values are the allowed corrections that would result in the highest possible value of $\ln L$, and therefore the best match between data and test spectra, for the T values being tested. The c_i values are found using Eq. 8.23, where a, b, and c are defined in equations 8.25 - 8.27, and the positive square root is used. We double check that each solution of c_i found corresponds to the maximum lnL using Eq. 8.28 and Eq. 8.29.

- These 15 values of c_i are then used to calculate the μ_i terms for Q1 and Q3, according to Eq. 8.16. The c_i values are also used to calculate the ith elements of the penalty term.

- The Q1 expected and observed components, Q3 expected and observed components and penalty term for each energy bin are then available to calculate $\ln L_i$ according to Eq 8.20. This process is repeated for each of the 15 energy bins and the $\ln L_i$ components are summed to provide one final value of lnL. This is the likelihood that the current pair of T values describe the data.

- This process is repeated for all possible combinations of T_π and T_K. For each combination of T_π and T_K the value of lnL is plotted. The largest value of lnL, defined as $\ln L_{max}$ is identified and the value of T_π and T_K which correspond to this maximum log-likelihood are then taken as the beam parent composition which provide the closest fit to the data.

- To find the 1 sigma uncertainty on the optimum values of T_π and T_K, we find the values of T_π and T_K which resulted in a log-likelihood of value ($\ln L_{max}$ - 0.5) [2]. A contour linking all lnL results which have this value is drawn and the highest and lowest values of T_π and T_K that sit on this contour are the 1 sigma uncertainties on the best fit T values. This is demonstrated by Figure 37.5 in [103].

[2]Using Eq. 8.19, the change in lnL caused by a 1 σ shift can be found by setting x = $(\mu + \sigma)$. This results in lnP = -0.5 therefore a 1 σ shift reduces lnP by 0.5.

In order to perform a blind analysis the data set should not be analysed using this fitting method until the code has been fully tested and all systematic uncertainties have been considered and calculated. To provide a check that the code is functioning correctly, the Monte Carlo can be compared to itself. The Monte Carlo provides the four probability distribution functions and the data provides the two reconstructed energy spectra for all ν_μ in Q1 and Q3. However, we can also plot the Q1 and Q3 reconstructed energy spectra from the same Monte Carlo sample to provide a pair of spectra to take the place of the real data spectra. Since these 'data' spectra will be an exact addition of the p.d.f.s, the code should be able to find the exact values of T_π and T_K present in the Monte Carlo spectra.

Figure 8.8 shows the format of the fit output. The range of possible T_π are plotted on the x-axis and the possible T_K values fill the y-axis. The log-likelihood value of each combination is given by the colour at that point. The extreme low values of T_π and T_K result in values of lnL that approach -∞, so the axes are restricted to omit the very low T values. The white cross shows the position of the largest value of lnL and therefore the best-fitting values of T_π and T_K.

The Monte Carlo spectra used for this test contained 1999.21 ν_μ^{fK} and 4116.07 $\nu_\mu^{f\pi}$ events distributed between Q1 and Q3. Therefore T_π and T_K were each allowed to vary from 0 to 6115. The optimum values of T_π and T_K returned by the fit were $T_\pi = 4116$ and $T_K = 1999$, which are the closest integer values of T_π^{test} and T_K^{test} to the true values, so this confirms that the code functions properly.

8.5 Statistical tests with the fitting procedure

The behaviour of the fitting program must be studied extensively before it can be used to analyse our one data set. Sources of systematic uncertainty must also be considered and their effect on the fit result must be calculated. Some of these

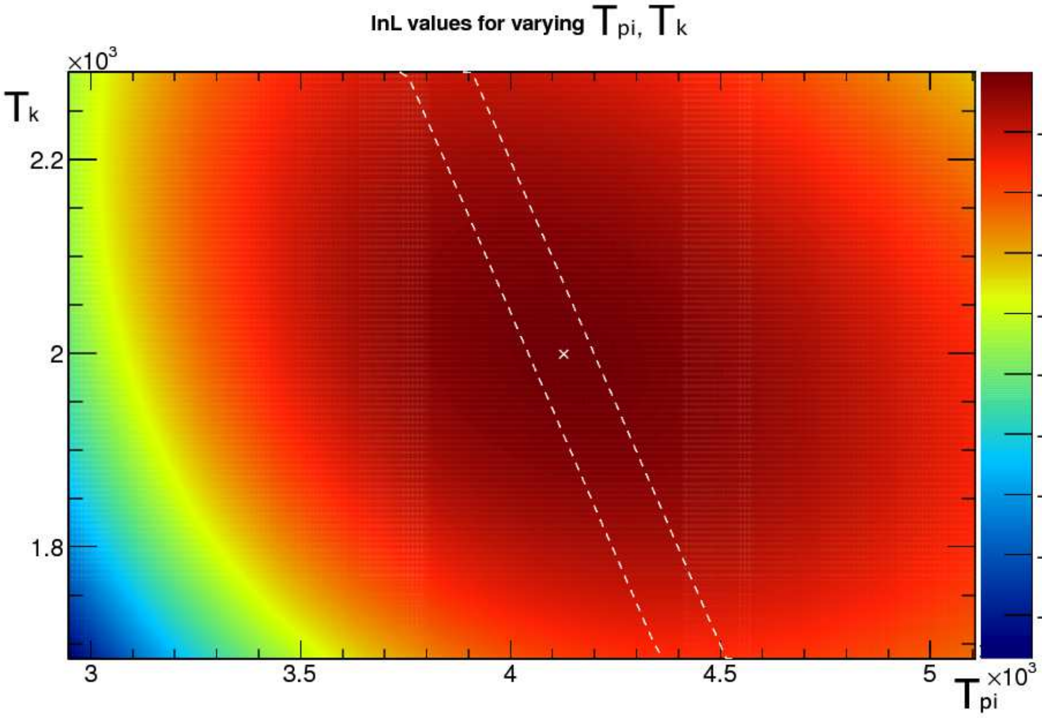

Figure 8.8: Typical output showing log-likelihood values for all combinations of T_π and T_K, with the best fit point marked. Dashed lines mark the region of T value sums that equal the true total number of events present ± the 1 σ statistical uncertainty on that number.

require running the fitting code in different scenarios. In order to do this, multiple sets of fake data are generated using a toy Monte Carlo. The various tests with and without this fake data are described in the following sections.

The fake data sets are generated to contain required numbers of ν_μ^{fK} and $\nu_\mu^{f\pi}$. The reconstructed energy spectra plotted from the Monte Carlo selections for ν_μ^{fK} and $\nu_\mu^{f\pi}$ in Q1 and Q3 are used as the basis for the fake data, since these distributions are the form we expect the data will take. A combination of ν_μ^{fK} and $\nu_\mu^{f\pi}$ to study is selected, and to represent this combination, test values of T_π and T_K, labelled T_π^{test} and T_K^{test}, are set. First we consider the $\nu_\mu^{f\pi}$ distributions in Q1 and Q3. An energy in the range 0-10 GeV is randomly selected, from either the Q1 or Q3 distribution, where the probability of selecting a certain energy value in either of the quadrants is greater for bins with higher bin contents. Therefore we are more likely to select an energy value from within the range of one of the $\nu_\mu^{f\pi}$ energy peaks, and we are slightly more likely to select an energy value from the Q3 distribution, since there are more entries in the Q3 quadrant. One event, at this selected energy, is then added to either the Q1 or Q3 fake data spectrum, to match the quadrant from which the energy value was chosen. This process occurs T_π^{test} times, so that the fake data spectra gain T_π^{test} entries, split between the fake Q1 spectrum and the fake Q3 spectrum. The same process is then applied for the ν_μ^{fK}, so that T_K^{test} entries, of energies randomly selected from the ν_μ^{fK} Monte Carlo distributions, are also added to the fake Q1 and Q3 data spectra. The result is two fake ν_μ spectra - one for Q1 and one for Q3 - containing a combined total of $\mathrm{T}_\pi^{test} + \mathrm{T}_K^{test}$ entries.

There are two different ways of using the fake data sets, as outlined in sections 8.5.1 and 8.5.2.

8.5.1 Fit convergence with different input values

It is important that the fitting code functions correctly and to the same degree of accuracy regardless of the actual ratio of ν_μ^{fK} to $\nu_\mu^{f\pi}$ present in the data. Therefore, the first use for samples of fake data is to test the stability of the fitting code when analysing spectra with extreme ratios of ν_μ^{fK} and $\nu_\mu^{f\pi}$ to identify if there are any combinations that the fitting program cannot process. We can generate fake data spectra containing different combinations of ν_μ^{fK} and $\nu_\mu^{f\pi}$ and process these using the fitting code.

In order to do this we create several different sets of 100 fake data samples, with vastly different ratios of ν_μ^{fK} and $\nu_\mu^{f\pi}$ in each set. These are used in the place of the real Q1 and Q3 data spectra. The standard probability distribution functions obtained from the Monte Carlo selections are used as the distributions to be combined and compared to the fake data. Of the 39946 ν_μ selected in the Monte Carlo we find that 10,583 are ν_μ^{fK} and 29363 are $\nu_\mu^{f\pi}$, giving a ν_μ^{fK} to $\nu_\mu^{f\pi}$ ratio of approximately 1:2.8. Proton on target comparisons tell us that we expect to select approximately 3500 events in the data samples, therefore a range of fake data spectra containing around 3500 events, split between ν_μ^{fK} and $\nu_\mu^{f\pi}$, are generated. Each of these fake spectra are processed using the fitting programme. Tested combinations of ν_μ^{fK} and $\nu_\mu^{f\pi}$ are listed in Table 8.3.

Since a ν_μ^{fK} to $\nu_\mu^{f\pi}$ ratio of approximately 1:2.8 is predicted by current models, and the uncertainty on the parent contributions is limited, the test ratios listed in Table 8.3 are sufficient to confirm that the fitting code will function properly in and beyond the range of ν_μ^{fK} to $\nu_\mu^{f\pi}$ ratio results we could expect to measure. There are no combinations that result in the fit failing to calculate a log-likelihood value. The tests involving the lowest numbers of events also confirm that the code functions appropriately in the case that some bins contain no entries.

T_K^{test}	T_π^{test}	ν_μ^{fK}:$\nu_\mu^{f\pi}$ ratio	Code functions?
1000	2000	1:2	Yes
900	2000	1:2.22	Yes
800	2000	1:2.5	Yes
700	2000	1:2.86	Yes
600	2000	1:3.33	Yes
500	2000	1:4	Yes
50	2000	1:40	Yes
2000	2000	1:1	Yes
2000	1000	2:1	Yes
2000	1	2000:1	Yes
500	1	500:1	Yes

Table 8.3: Ratios of ν_μ^{fK}:$\nu_\mu^{f\pi}$ tested. Fitting code functions normally for all reasonable ratios of parent contributions.

8.5.2 Statistical spread and fit bias

The one data set available corresponds to a particular number of protons on target, and contains $N_{\nu_\mu}^{data}$. If a second set of data was collected, using exactly the same number of protons on target, we would not expect to observe exactly the same $N_{\nu_\mu}^{data}$, but instead would expect to select a number in the range $N_{\nu_\mu}^{data} \pm \sqrt{N_{\nu_\mu}^{data}}$. Therefore to model this behaviour, for each T_π^{test} and T_K^{test} combination of interest, we create 100 fake data sets, where T_π^{test} and T_K^{test} are each allowed to vary by their square roots from set to set. The result is 100 sets of fake data where the ratio of ν_μ^{fK} and $\nu_\mu^{f\pi}$ is approximately the same, but the total number of events fluctuates from set to set, in the same way that real data sets would behave.

For a given combination of base values of T_π^{test} and T_K^{test} being tested, the fitting programme can be applied to each of the 100 fake data sets in turn, and the best-fitting values of T_π and T_K can be calculated and plotted for each set. The means of the values of T_π and T_K that produce the highest log-likelihood result for each fake data set can be found. Comparing the mean best fitting T_π and T_K values found for a set of 100 fake data samples with the base T_π^{test} and

T_K^{test} values used in the generation of that 100 samples will reveal how accurate the fitting programme is in that test scenario. The variation in optimum T_π and T_K values returned for each of these individual fake data sets reflects the natural variation we can expect with data sets of limited size, and the way in which the optimum T_π and T_K values move also illustrates the uncertainty on the fit results.

The three Monte Carlo samples used to create the probability density distributions correspond to 2.50×10^{21} protons on target and approximately 40,000 ν_μ interacting in Q1 and Q3 pass the selection cuts. The total data used for this analysis contains 2.13×10^{20} POT, which we would therefore expect to produce approximately 3400 events in quadrants Q1 and Q3, using the same selections. If 100 data samples corresponding to 2.13×10^{20} POT were collected and processed, we would expect the individual ν_μ^{fK} and $\nu_\mu^{f\pi}$ totals to fluctuate in each sample, and the total number of events to vary accordingly.

To observe this natural fluctuation, and the variation in best fit values of T_π and T_K that would be calculated by the fitting programme for each different sample, we can generate fake data which mimics the real data. The base ratio of ν_μ^{fK} to $\nu_\mu^{f\pi}$ is set to be the same for all 100 fake data samples and to be approximately the same as the ratio predicted by our models and used in the Monte Carlo. This is achieved by setting $T_\pi^{test} = 2570$ and $T_K^{test} = 930$ as the base values for the number of ν_μ^{fK} and $\nu_\mu^{f\pi}$ to be included in the samples, resulting in a total of 3500 events. These test values are then allowed to fluctuate from sample to sample however, and it is very unlikely that two fake data samples will contain exactly the same numbers of ν_μ^{fK} and $\nu_\mu^{f\pi}$. Instead, each sample will contain numbers of T_π and T_K which sit in the expected range of T_π and T_K, centred on 930 ν_μ^{fK} and 2570 $\nu_\mu^{f\pi}$, and the total events will fluctuate around 3500 accordingly. These 100 samples of simulated Q1 and Q3 data are then processed by the fitting programme, and the best fitting values of T_π and T_K are calculated for each set. The results are shown in Figure 8.9.

We can see from this figure that the mean T_{π}^{opt} is 2597, which is 27 events higher than the base T_{π}^{test} of 2570, or accurate to 1%. The mean T_{K}^{opt} is 949.2, where the fake data was generated with a base T_{K}^{test} of 930. This difference of +19.2 means T_{K}^{opt} is accurate to 2%.

While the level of accuracy on both the T_{π} and T_{K} found by the fitting program is good, it must be noted that both the mean T_{π}^{opt} and T_{K}^{opt} values returned are slight over-estimates, which causes the mean total events identified to be 3546.2, where the pre-fluctuation total events was set to be 3500. The uncertainties are listed in Table 8.4. The source of this over-estimation is of interest. There are two possible causes for this. One is that the random fluctuations introduced to the T_{π}^{test} and T_{K}^{test} values as part of the fake data generation happen to have been positive for the majority of this particular 100 samples, so that the mean fluctuated T_{π}^{test} and T_{K}^{test} values, and therefore also the mean total events, are slightly larger than the base values for this 100 data sets. The other possibility is that the fitting code features a small inherent bias that favours results with slightly higher T values.

Parent type	Base T^{test}	mean T^{opt} (σ)	$\sigma/\sqrt{n}$	Accuracy	Significance
π	2570	2597 ±123.4	12.34	+27 (1.05%)	2.19 $\sigma/\sqrt{n}$
K	930	949.2 ±71.86	7.19	+19.2 (2.06%)	2.67 $\sigma/\sqrt{n}$
All	3500	3546.2 ±142.80	14.28	+46.2 (1.32%)	3.24 $\sigma/\sqrt{n}$

Table 8.4: Fit outputs for 100 fake data sets, with uncertainties. Base T values are the event totals before fluctuation. The accuracy of the fit gives the difference between the mean optimum T value returned and the base value, as well as the % difference. The significance of this difference in terms of $\sigma/\sqrt{n}$, where n=100 (the number of fake data sets considered), is also provided.

To test this, the sums of the fluctuated T_{π}^{test} and T_{K}^{test} values used for each of the 100 fake data samples are found. The 100 fluctuated values of T_{π}^{test} + T_{K}^{test}, referred to as T_{tot}^{fluc}, are plotted in Figure 8.10, and demonstrate the Gaussian distribution of event totals used. While the 100 sets used 3500 events as the base

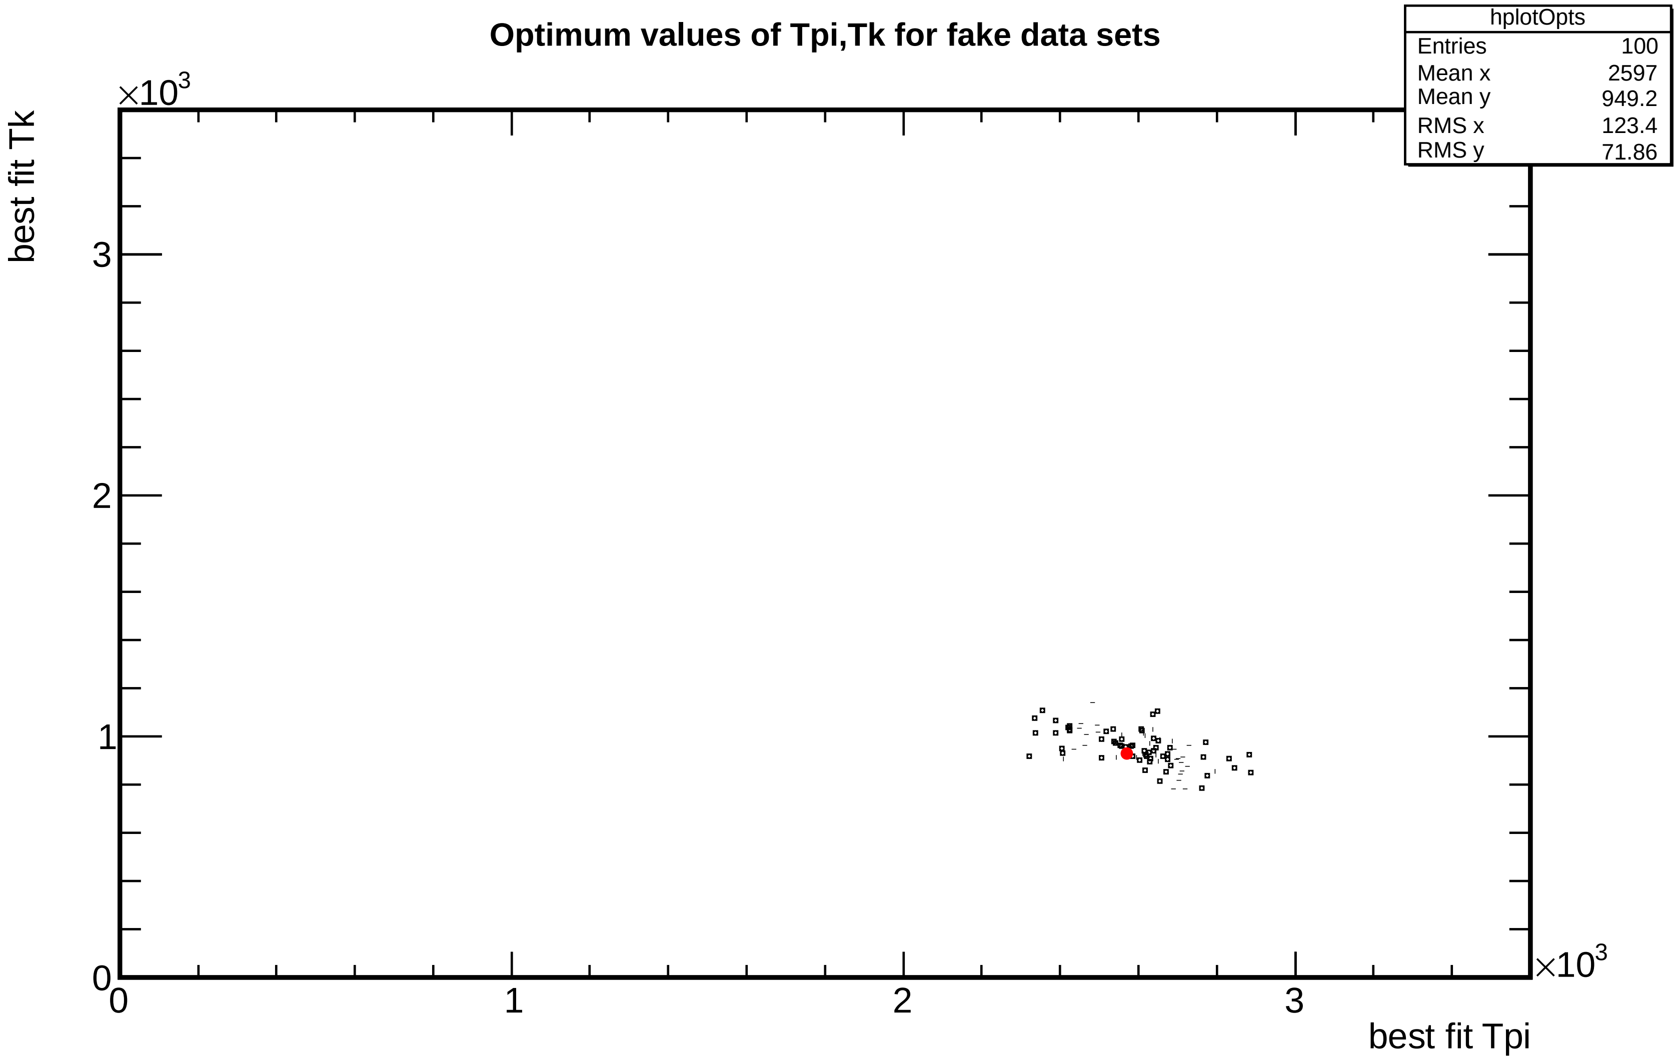

Figure 8.9: Fit output for 100 fake data samples. The best fitting combinations of T_π and T_K are marked in black for each of the 100 sets. The red dot shows the base configuration of the fake data before fluctuations, corresponding to $T_\pi = 2570$ and $T_K = 930$.

value for the combined Q1 and Q3 spectra contents, the mean total events in the 100 fake data sets is 3506.7 events. The standard deviation of 58.22 results in a $\sigma/\sqrt{n}$ of 5.822, so the slight increase of 6.7 events is equivalent to 1.15 $\sigma/\sqrt{n}$. This is a reasonable shift and therefore not indicative of a bias in the Poisson fluctuation mechanism used to generate the T_{tot}^{fluc} values. However it does confirm that with a limited number of fake data experiments, the effective T values and total events can be slightly higher than the base settings on average, and this could be a small contribution to the fit overestimation seen for these 100 samples.

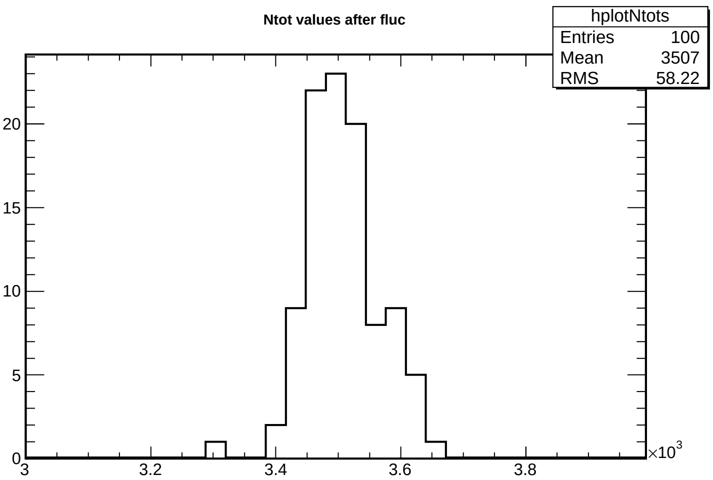

Figure 8.10: Plot of the total events in Q1 and Q3 combined used in each of the 100 fake data samples generated. Totals are 3500 ± appropriate statistical fluctuation, and therefore the distribution should be centred around 3500.

The other consideration is a built-in bias present in the fitting programme. One way to investigate this further is to create another 100 fake data sets, but this time without allowing the statistical fluctuation on the individual T_{π}^{test} and T_{K}^{test} values. The 100 pairs of fake Q1 and Q3 spectra will still differ between samples as the events randomly selected from the Monte Carlo distributions to form the fake spectra will still be different in each fake experiment, resulting in different samples. The total number of events contained in each pair of fake Q1 and Q3 spectra will always be constant however, as will the T_{π}^{test} and T_{K}^{test} values.

For this test a smaller set of fake data was generated: the 100 fake data samples each contain a total of 1520 events, split as $T_{\pi}^{test} = 1120$ and $T_{K}^{test} = 400$. These 100 fake data samples are processed and the optimum T_{π} and T_{K} values for each fake experiment are plotted in Figure 8.11. The spread of values returned by the fitting programme is still present as the small differences in the Q1 and Q3 spectral shapes between each fake experiment, magnified by the lower number of entries in each spectrum, will cause the programme to find that slightly different combinations of ν_{μ}^{fK} and $\nu_{\mu}^{f\pi}$ provide the best fit to each different pair of fake data spectra.

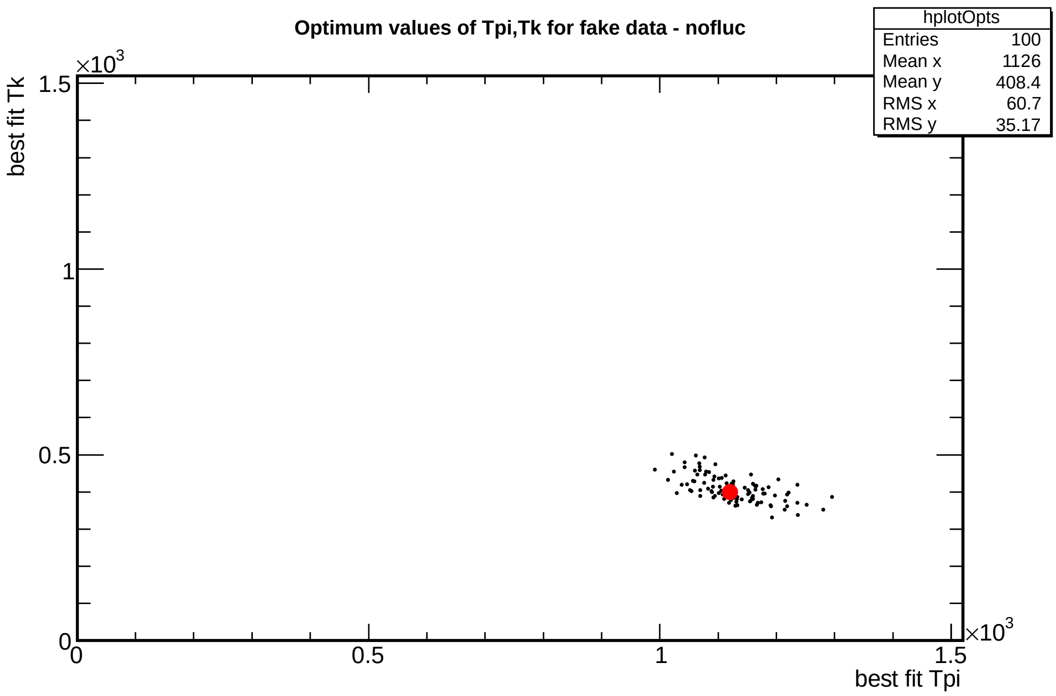

Figure 8.11: Fit output for 100 fake data samples generated without statistical fluctuation of event totals. The best fitting combinations of T_{π} and T_{K} are marked in black for each of the 100 sets. The red dot shows the configuration of the fake data, with $T_{\pi}^{test} = 1120$ and $T_{K}^{test} = 400$

The results for this test are summarised in Table 8.5.

Once again the optimum T values found by the fitting programme are both slight overestimates. The mean total events found by the programme is 2.05

Parent type	Base T^{test}	mean T^{opt} (σ)	$\sigma/\sqrt{n}$	Accuracy	Significance
π	1120	1126 $\pm$60.7	6.07	+6 (0.54%)	0.99 $\sigma/\sqrt{n}$
K	400	408.4 $\pm$35.17	3.517	+8.4 (2.1%)	2.39 $\sigma/\sqrt{n}$
All	1520	1534.4 $\pm$70.15	7.015	+14.4 (0.95%)	2.05 $\sigma/\sqrt{n}$

Table 8.5: Fit outputs for 100 fake data sets generated without natural fluctuation of event numbers, with uncertainties. Base T values are the event totals. The accuracy of the fit gives the difference between the mean optimum T value returned and the base value, as well as the % difference. The significance of this difference in terms of the quantity $\sigma/\sqrt{n}$, where n=100 (the number of fake data sets provided), is also provided.

$\sigma/\sqrt{n}$ from the input mean events. We would expect a fitting code with no bias to produce a mean N_{tot} of this value or higher approximately 5% of the time, so this is not sufficient evidence to conclude that the fit definitely contains a bias. The presence of a slight bias is a possibility which must be taken into consideration when fitting to real data; however, we have shown that the overestimation is very small, so having observed and recorded this effect it is not a cause for concern.

8.6 Systematic effects

The standard 100 fake data experiments processed and plotted in Figure 8.9 demonstrate the expected behaviour of the fitting programme in the case that the Monte Carlo distributions are correct and there are no sources of systematic uncertainty present. Several potential effects need to be considered, and the results shown in Figure 8.9 provide the basis for comparison for the investigations described below.

8.6.1 Magnetic field uncertainty

The momenta of particles detected in ND280 are reconstructed using information about the tracks produced. The momentum of a particle passing through ND280 is a function of the magnetic field the particle travels through, therefore an accurate evaluation of the momentum of a particle requires precise knowledge of the magnetic field in the detector.

When a particle of charge e and mass m travels with a velocity v in a direction perpendicular to a magnetic field B, it will follow a circular path of radius r, such that

$$F = evB = \frac{mv^2}{r} \tag{8.30}$$

This can be rearranged to show that

$$mv = erB \tag{8.31}$$

$$p \propto B \tag{8.32}$$

From Eq. 8.32 it follows that the uncertainty in the momentum of any particle is proportional to the uncertainty in the magnetic field the particle is travelling through. Therefore the uncertainty on the momentum values calculated for each of the particles selected for this analysis is proportional to the uncertainty on the magnetic field applied across ND280. The magnetic field is measured accurately with the aim of ensuring that the uncertainty related to the magnetic field is below 2% [65].

As described in detail in Chapter 7, the energies of the ν_μ of interest are reconstructed using the momenta of the particles produced by their interactions. An error in the calculated momenta would result in an error in the reconstructed energies of the ν_μ. If the magnetic field experienced by particles passing through ND280 differs from the value used in the Monte Carlo simulations, there will be a difference in the energy reconstruction between the Monte Carlo and the data. This means that the probability distribution functions for ν_μ^{fK} and $\nu_\mu^{f\pi}$ in Q1 and Q3 generated from the Monte Carlo will be a less precise representation of the real contributions forming the data spectra, and the fitting program may struggle more to find a combination of ν_μ^{fK} and $\nu_\mu^{f\pi}$ that match the data well. The combination of these incorrect probability distribution functions which provides the best match to the data will not necessarily be the true ν_μ^{fK} to $\nu_\mu^{f\pi}$ ratio.

Fake data experiments were used to investigate the effect of an uncertainty in the magnetic field on the fit results. If the magnetic field is modelled incorrectly in the Monte Carlo, it would result in the probability distribution functions used by the fitting program being slightly scaled up or down with respect to the true distributions. We therefore create new probability distribution functions with the effect of a magnetic field uncertainty incorporated. Once a Monte Carlo interaction is selected, the energy of the incident neutrino is reconstructed. To create new, altered probability distribution functions, the momenta of all tracks included in the energy reconstruction calculations in the Monte Carlo are scaled up or down, to reflect possible errors in the magnetic field value. Fitting these altered probability distribution functions to the standard set of 100 fake data samples, generated by the original Monte Carlo, will then demonstrate the effect of our Monte Carlo actually containing an incorrect magnetic field value.

Three altered sets of probability distribution functions are created to investigate this effect. Since the magnetic field uncertainty in ND280 is expected to be less than 2%, one set of Monte Carlo is created with all track momenta multiplied by a factor of 0.98, and another set is created with track momenta scaled by 1.02. A third set featuring a scaling of 1.10 is created, despite this being an unrealistic error, to clearly demonstrate the effect of an incorrectly modelled magnetic field. Each of these altered sets of probability distribution functions are then used as inputs to the fitting programme when applied to the standard 100 fake data samples. The results are shown in Figures 8.12, 8.13, and 8.14 and summarised in Table 8.6.

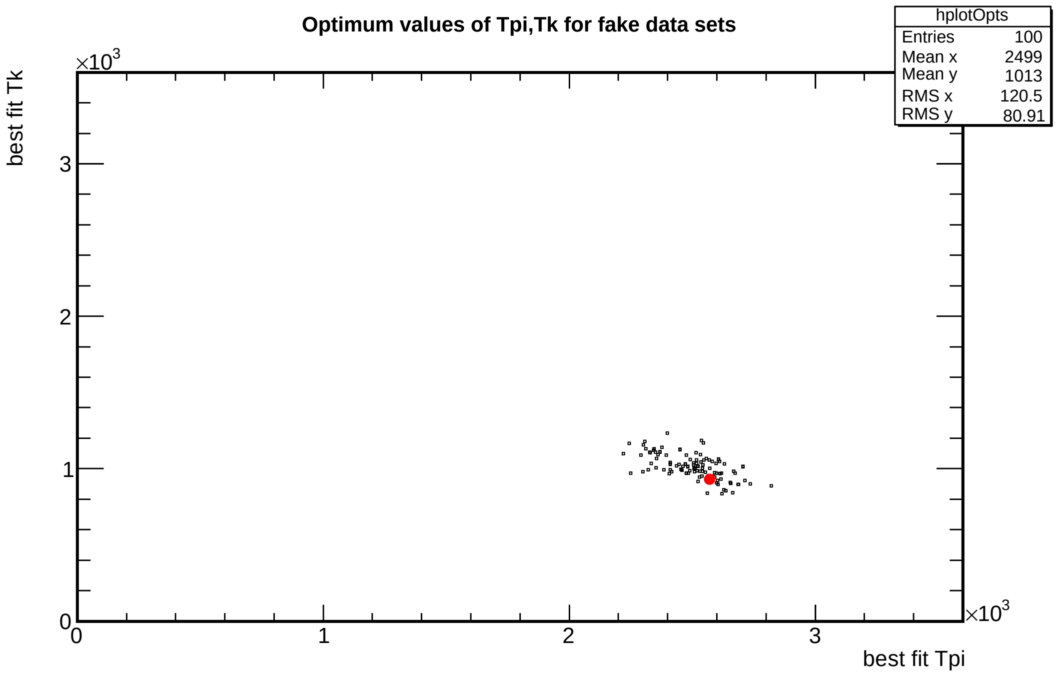

Figure 8.12: Fit output for 100 standard fake data samples using probability distribution functions with scale factor of 0.98 applied to track momenta. The red dot shows the fake data base configuration of $T_{\pi}^{test} = 2570$ and $T_{K}^{test} = 930$ and the black marks show the optimum T_{π} and T_{K} values found for each set.

First we consider the case that the magnetic field experienced by particles in the detector is higher than modelled in the Monte Carlo. This is investigated using the probability distribution functions formed with the scale factor of 0.98 applied to all track momenta. Reducing the momenta has the effect of reducing all of the reconstructed energy values, and therefore shifts the reconstructed energy spectra, and the probability distribution functions, lower in energy. This will cause a slight offset between the probability distribution functions and the data distributions. The probability distribution functions will be shifted lower in energy than the data

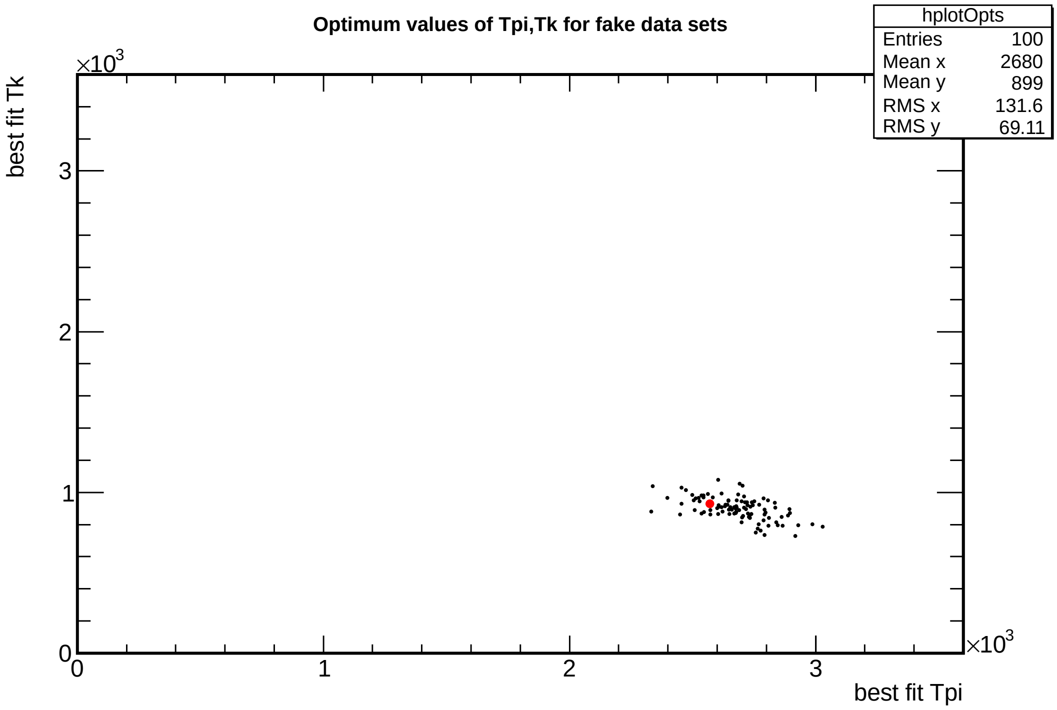

Figure 8.13: Fit output for 100 standard fake data samples using probability distribution functions with scale factor of 1.02 applied to track momenta. The red dot shows the fake data base configuration of T_{π}^{test} = 2570 and T_{K}^{test} = 930 and the black marks show the optimum T_{π} and T_{K} values found for each set.

p scaling	lnL	Mean T_{π}^{opt} $\pm\sigma$	Shift in mean T_{π}^{opt}	Mean T_{K}^{opt} $\pm\sigma$	Shift in mean T_{K}^{opt}
None	-9.008	2597 $\pm$123.4	-	949.2 (71.86)	-
110%	-12.51	3213 $\pm$169.2	+616 (23.7%)	734.3 $\pm$61.04	-214.9 (22.6%)
102%	-10.23	2680 $\pm$131.6	+83 (3.2%)	899 $\pm$69.11	-50.2 (5.3%)
98%	-10.09	2499 $\pm$120.5	-98 (3.8%)	1013 $\pm$80.91	+63.8 (6.7%)

Table 8.6: Results of magnetic field uncertainty tests. Standard, unscaled results are provided for comparison. The shifts in mean T values caused by the scale factors, compared to the standard, unscaled case, are given, along with the % changes seen.

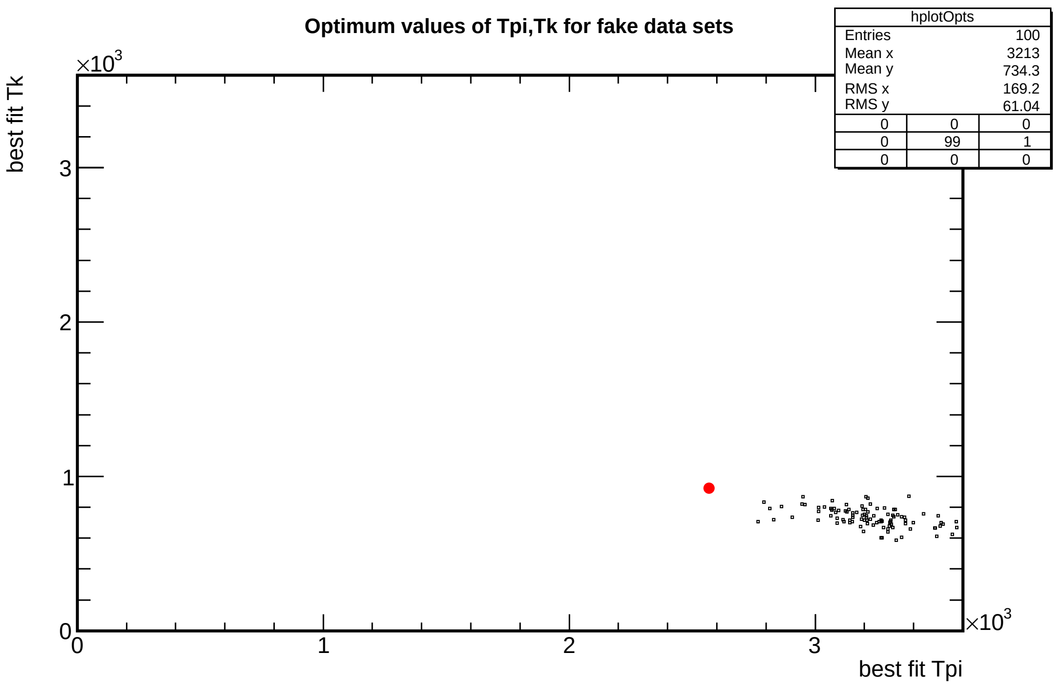

Figure 8.14: Fit output for 100 standard fake data samples using probability distribution functions with scale factor of 1.10 applied to track momenta. The red dot shows the fake data base configuration of $T_\pi^{test} = 2570$ and $T_K^{test} = 930$ and the black marks show the optimum T_π and T_K values found for each set.

distributions. At low energy the $\nu_\mu^{f\pi}$ is the main contributor to the data spectrum, and so this contribution is underestimated in order to fit the rising edge of the peak to the data spectrum. At high energy the ν_μ^{fK} spectrum dominates, so to achieve sufficient bin contents in the high energy tail the ν_μ^{fK} contribution is overestimated. Therefore when the probability distribution functions sit at lower energies than the data spectra, the fitting program compensates, resulting in low best-fitting T_π (T_π^{opt}) values and high best-fitting T_K (T_K^{opt}) results. The movement of the 100 results with respect to the T_π^{test} and T_K^{test} values is shown in Figure 8.12.

The increased difficulty in finding a combination of ν_μ^{fK} and $\nu_\mu^{f\pi}$ distributions that fit the data is demonstrated by the highest lnL value, corresponding to the T_π^{opt} and T_K^{opt} values, being lower than the standard case. This indicates that the best combination of the probability distribution functions that could be found is not as close a fit as the best fitting combination in the standard case.

To consider the case of the true magnetic field being lower than modelled in the Monte Carlo, we apply a scale factor of greater than 1. In this case the probability distribution functions will be scaled up in energy, so that the probability distribution functions will sit at higher energies than the data spectra. The fitting code will now find that an overestimation of the $\nu_\mu^{f\pi}$ will provide the required event numbers in the low energy edge of the peak, so T_π^{opt} is found to be high. In contrast, the ν_μ^{fK} contribution must be reduced in order to give the appropriate number of events in the high energy bins, therefore T_K^{opt} is found to be low. This effect is demonstrated most clearly in Figure 8.14, where the track momenta are scaled up by 10%, and the 100 T_π^{opt} and T_K^{opt} values show a very clear shift away from the T_π^{test} and T_K^{test} values used in the fake data generation.

The uncertainty on the magnetic field is less than 2%, therefore the shifts summarised in Table 8.6 are the uncertainties on the fit results due to an incorrectly modelled magnetic field. These systematic uncertainties will be compared to the

1 sigma statistical uncertainty obtained for the fit to the real data set.

8.6.2 Energy reconstruction uncertainty

An uncertainty in the magnetic field strength is just one possible cause of a difference in the energy reconstruction applied to Monte Carlo events and data events. Rather than investigate each potential cause of an energy systematic separately, we can study the effect in terms of a resultant change in the energy reconstruction. To do this we use the same Monte Carlo selections but this time apply a scale factor to the reconstructed energy calculated for each ν_μ before plotting the ν_μ^{fK} and $\nu_\mu^{f\pi}$ distributions. The fitting program is then applied to the same 100 unaltered fake data samples, using these new scaled distributions as the probability distribution functions. To make the effect clear, we use scalings of 90% and 110%. The results are plotted in Figures 8.15 and 8.16.

The results are summarised in Table 8.7. They show that a 10% uncertainty in the energy reconstruction would cause a significant inaccuracy in the fit results. Therefore when fitting to any real data set, an important check is to compare the data spectra and the Monte Carlo spectra. Should any offset be present, it would indicate that the reconstructed energies are not identical for the data and Monte Carlo, and the level of offset must be measured in order to calculate the effect on the fit values.

8.6.3 Uncertainties in the neutrino cross-sections

This analysis makes the assumption that the distribution shapes for ν_μ^{fK} and $\nu_\mu^{f\pi}$ given in the Monte Carlo are accurate and focusses on finding the contributions to the beam from these distributions.

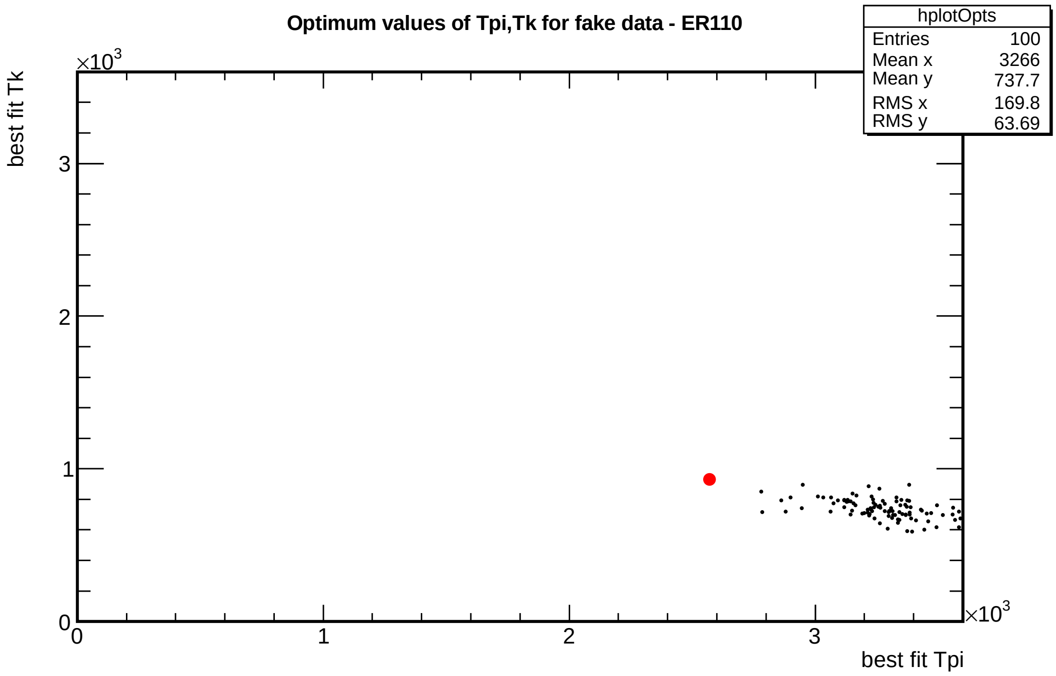

Figure 8.15: Fit output for 100 standard fake data samples using probability distribution functions with scale factor of 1.1 applied to reconstructed energy values. The red dot shows the fake data base configuration of $T_{\pi}^{test} = 2570$ and $T_{K}^{test} = 930$ and the black marks show the optimum T_{π} and T_{K} values found for each set.

E_{recon} scaling	lnL	Mean $T_{\pi}^{opt} \pm\sigma$	Shift in mean T_{π}^{opt}	Mean $T_{K}^{opt} \pm\sigma$	Shift in mean T_{K}^{opt}
None	-9.008	2597 ±123.4	-	949.2 (71.86)	-
110%	-12.61	3266 ±169.8	+669 (25.8%)	737.7 ±63.69	-211.5 (22.3%)
90%	-11.32	2249 ±112.4	-348 (13.4%)	1232 ±90.46	+282.8 (29.8%)

Table 8.7: Results of energy reconstruction uncertainty tests. Standard, unscaled results are provided for comparison. The shifts in mean T values caused by the scale factors, compared to the standard, unscaled case, are given, along with the % changes seen.

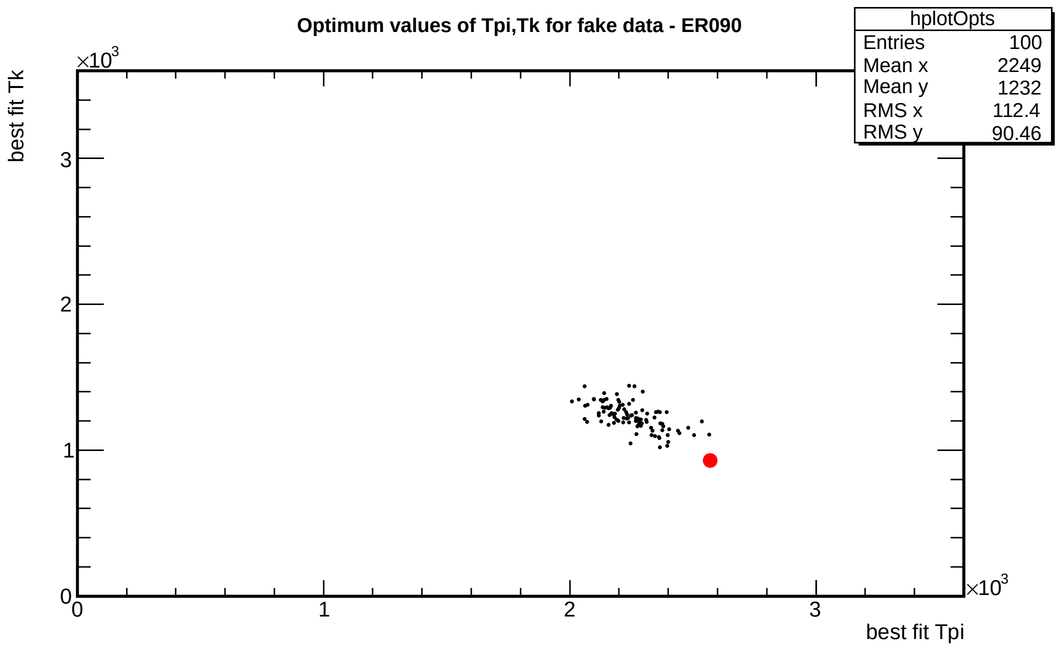

Figure 8.16: Fit output for 100 standard fake data samples using probability distribution functions with scale factor of 0.9 applied to reconstructed energy values. The red dot shows the fake data base configuration of $T_\pi^{test} = 2570$ and $T_K^{test} = 930$ and the black marks show the optimum T_π and T_K values found for each set.

Errors on some of the cross-section values used in the Monte Carlo would result in errors on the bin contents for certain energies in the neutrino spectra. This would have the effect of distorting the shapes of the probability distribution functions used. The cross-section correction factors, c_i, introduced in §8.3.2, are designed to allow small alterations to the expected bin contents, within limits allowed by the known cross-section uncertainties corresponding to an energy bin i, in order to provide a better fit to the data. It is important to test if these correction factors are functioning as intended.

Applying scale factors of different magnitudes to the contents of the different energy bins models the case that the cross-sections at some energies are inaccurate. We investigate the effect of fitting altered probability distribution functions to the standard, unaltered fake data sets. Several different alterations to the probability distribution functions are tested. For each test, new ν_μ^{fK} and $\nu_\mu^{f\pi}$ spectra in Q1 and Q3 are created by applying the stated scaling to the standard Monte Carlo spectra. Corresponding probability distribution functions are created, and these altered functions are fitted to the standard 100 fake data samples which are generated with the original, unscaled Monte Carlo spectra. The scale factors applied to the bin contents at different energies are given below.

- Test 1 = 0% at 0 Gev, increasing uniformly to 10% at 10 GeV
- Test 2 = 10% at 0 GeV, decreasing uniformly to 0% at 10 GeV
- Test 3 = 0% at 0 GeV, increasing uniformly to 50% at 10 GeV
- Test 4 = 50% at 0 GeV, decreasing uniformly to 0% at 10 GeV

Tests 1 and 3 apply no change to the lowest energy events and scale up the bin contents in the high energy tail. The high end of the spectrum contains relatively few events, but is dominated by ν_μ^{fK}, so these tests are designed to investigate the effect on the T_K^{opt} results. The majority of neutrinos interacting have energies around 600 MeV, so we also apply a scale factor which is largest at 0 GeV and reduces to zero at high energies. Altering bin contents in this way does cause the distributions to change; however, it does not move the position of the peaks, so it does not create offsets between the probability distribution functions and the fake data.

Test 1, which increases the bin contents by up to 10% at high energies, creates the spectra that we would expect if the cross-sections are in fact 10% higher at high energies than currently modelled in the Monte Carlo. When compared to the standard fake data, this tests the effect of the Monte Carlo cross-sections being higher than the true cross-sections at high energies. Small changes to the cross-sections of this type can be compensated for to some extent by the c_i values.

The results are given in Table 8.8. The results of fitting to the unaltered probability distribution functions are provided for comparison.

Scaling	Mean $T_\pi^{opt} \pm\sigma$	Shift in mean T_π^{opt}	Mean T_K^{opt} $\pm\sigma$)	Shift in mean T_K^{opt}
None	2597 ±123.4	-	949.2 (71.86)	-
Test 1	2631 ±124.5	+34 (1.31%)	916.2 ±69.48	-33 (3.48%)
Test 2	2563 ±122.4	-34 (1.31%)	985.2 ±74.52	+36 (3.79%)
Test 3	2746 ±129.5	+149 (5.74%)	815.8 ±63.07	-133.4 (14.05%)
Test 4	2460 ±120.5	-137 (5.28%)	1101 ±84.94	+151.8 (15.99%)

Table 8.8: Results of altered probability distribution function tests. Scaling values used for each test are listed separately. Shifts in mean T values are given, along with the % changes.

The results for tests 1 and 2 show that even a 10% alteration to the bin contents will cause a shift in the best fit values of T_K and T_π to occur. The value of $\sigma/\sqrt{n}$ for T_π^{opt} is 12.45 in Test 1, therefore the shift of +34 observed is 2.73 $\sigma/\sqrt{n}$. However, this corresponds to a 1.31% shift in the value of T_π^{opt}, which is small compared to the average shift applied to the bin contents in this case. The shift observed in the mean T_K^{opt} value, of 3.48%, equivalent to 4.75 $\sigma/\sqrt{n}$, is more significant. However, it is not unexpected that these changes would have a larger effect on the T_K values since there are regions of the spectrum that consist purely of ν_μ^{fK}, so that changes to those bin contents would have to be accounted for purely by changes to T_K. Similar results are found for test 2.

Tests 3 and 4, which feature bin content scale factors of up to 50%, show much larger shifts in the mean T_K^{opt} and T_π^{opt} values. Due to the current levels of uncertainty on the cross-sections and the influence of the penalty term, there is a limit to the extent to which the c_i parameters can stray from 1 and adjust the expected bin contents to provide closer matches to the data spectra. While the c_i values can be set to compensate for small differences in the bin contents, their effect on the results of tests 3 and 4, which applied scale factors larger than the cross-section uncertainties (given in Table 8.2), will be minimal. We therefore see larger shifts, of approximately 5% on the pion content and 15% on the kaon content. However, these are still considerably smaller than the alterations of 50% that were made to the bin contents.

This source of uncertainty must be taken into consideration when evaluating the final fit results using data.

8.6.4 Cathode plane interference

Each TPC contains a cathode panel which produces an electric drift field. These cathodes sit on the x=0 plane, as shown in Figure 8.17.

Due to the direction of the B field across the TPCs, μ^- produced in neutrino interactions in FGD1 will curve upwards as they travel through TPC2. The presence of the cathode panel causes a narrow plane of dead space in the TPC. It also divides any tracks which traverse this plane in two, since the ionisation electrons produced by a track will always travel away from the plane and be detected at the

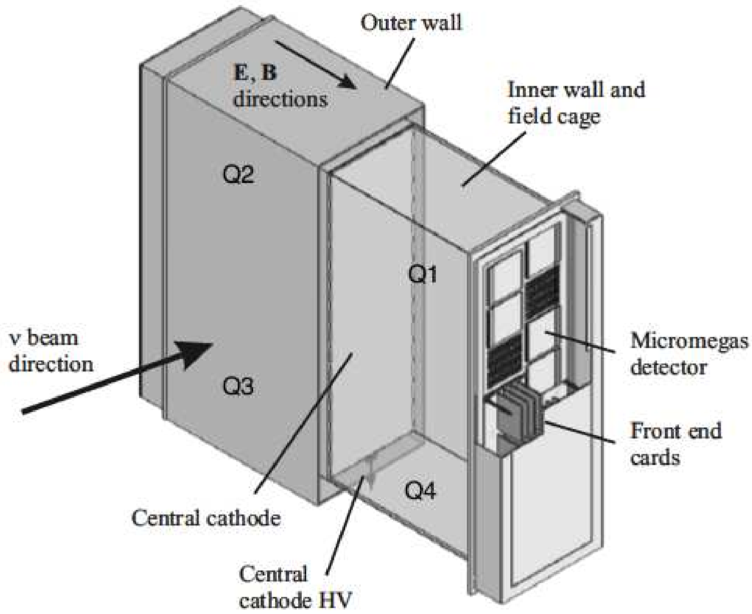

Figure 8.17: TPC structure demonstrating position of cathode panel in relation to neutrino beam axis. Ionisation electrons move away from the cathode panel and towards the readout planes. Q labels describe the FGD quadrant to which each corner of the TPC is adjacent. Figure adapted from [65].

edges of the TPC. The track quality cut requires that tracks travel a minimum distance through TPC2 in order to guarantee a minimum level of accuracy when reconstructing the trajectory. A track which crosses the cathode panel may be less likely to pass this cut, since the components of the track detected separately on each side of the cathode must each also be of a minimum length to be counted, and discarding one section of a track may reduce the overall length to lower than the quality cut threshold. Interactions occurring at x=0 may produce charged muons which curve up entirely in the plane of the cathode and produce very few ionisation electrons, and so these events could be lost. This has the potential to be a source of systematic error, since the track reconstruction for Monte Carlo and data may differ.

To investigate whether the presence of a cathode panel in TPC2 could affect the event selection, we plot the x coordinate of the interaction vertices in FGD1.

Particles produced at these interaction points will travel downstream, with a vertical curvature due to the B field. A small additional horizontal displacement may also occur in TPC2 due to the electric field; however, the effect on the neutrino interaction products would be very small compared to the effect on the electrons produced via ionisation. If the presence of the cathode significantly lowers the probability of a track being selected in the TPC, we would expect to see a lower event rate in the region surrounding the x=0 plane. The x coordinates of the interaction vertices for selected muon-like tracks and non-muon-like tracks are plotted in Figures 8.18 and 8.19.

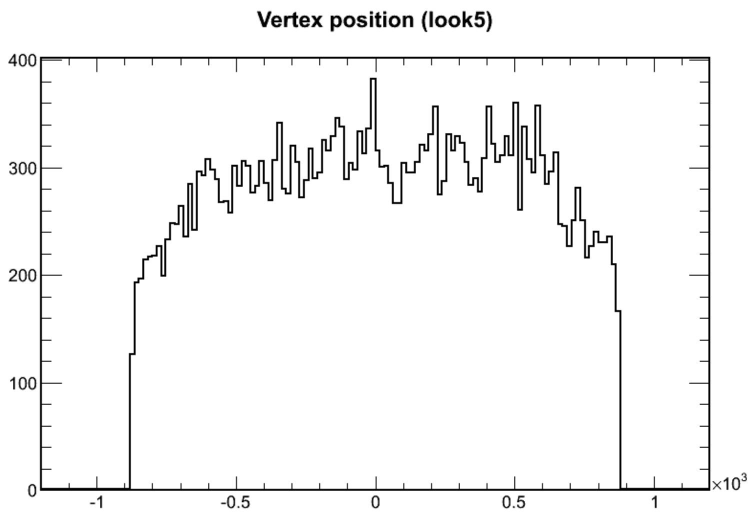

Figure 8.18: Plot shows the x positions of the upstream end of Monte Carlo tracks which pass the muon selection cuts. Lower numbers of selected tracks at largest $|x|$ values demonstrate that tracks originating near the edges of the FGD are less likely to travel a sufficient distance within the TPC to pass the track quality cut.

These plots demonstrate that any potential reduction in the selection of events near the x=0 plane is less significant than the statistical fluctuation in the selection of events at different x positions. With a considerably higher number of Monte Carlo events, the event selection rate would become more smooth across the FGD and any effect on the likelihood that a track passes the quality cuts would become more obvious. With this level of statistics however, which is a factor of 10 more

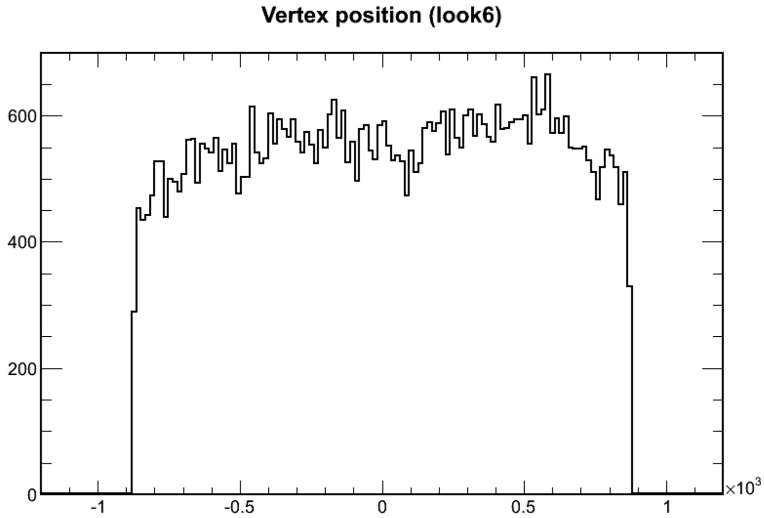

Figure 8.19: Plot shows the x positions of the upstream end of Monte Carlo tracks which pass the track quality cuts but are identified as non-μ-like.

than is expected in the data, there does not appear to be a significant effect.

Chapter 9

Analysis of data

The method of measuring the hadron contributions to the ν_μ beam developed in Chapter 8 using the event reconstruction described in Chapter 7 is applied to ND280 data here. Run II and Run IIIc are used. Results are given using a standard fitting procedure, and some additional fits are performed to demonstrate the effects of cross-section uncertainties. An evaluation of the fitting method and possible future extensions are discussed.

9.1 Data

Chapter 8 described the development of a technique which can in principle measure the contribution of kaon decay to the neutrino beam, while also allowing for uncertainties in the cross-section values used by T2K. The effects of systematic errors have been investigated (§8.6) and the fitting procedure has been checked for the presence of any bias (§8.5). We now use this fitting technique to examine T2K data.

For this analysis, Run II and Run IIIc ND280 data sets are used, which correspond to 7.8378×10^{19} POT and 1.3421×10^{20} POT[1] respectively (see Appendix B for data information). The data is processed using ND280 software version 5c,

[1]These numbers give the POT detected specifically by ND280, in contrast to the total POT provided in Appendix B.

using the event selection described in Chapter 7, and the energy reconstruction methods discussed there are applied to the selected events. We select 1602 events in Q1, and 1922 events in Q3, giving 3524 ν_μ in total.

The resulting ν_μ spectra observed in quadrants Q1 and Q3 are plotted. Before fitting to our data, we must check that the energy reconstruction has not malfunctioned in any way when applied to the data files instead of the Monte Carlo files. To check the reconstruction, the Monte Carlo spectra and data spectra are plotted together for comparison. Since the fitting program will be run on coarsely binned spectra, we will compare the spectra with the final binning applied. The Monte Carlo spectra are scaled to correspond to the same protons on target as the data, in order to give the spectra approximately equal integrals. The Q1 and Q3 comparisons are plotted in Figure 9.1 and Figure 9.2 respectively.

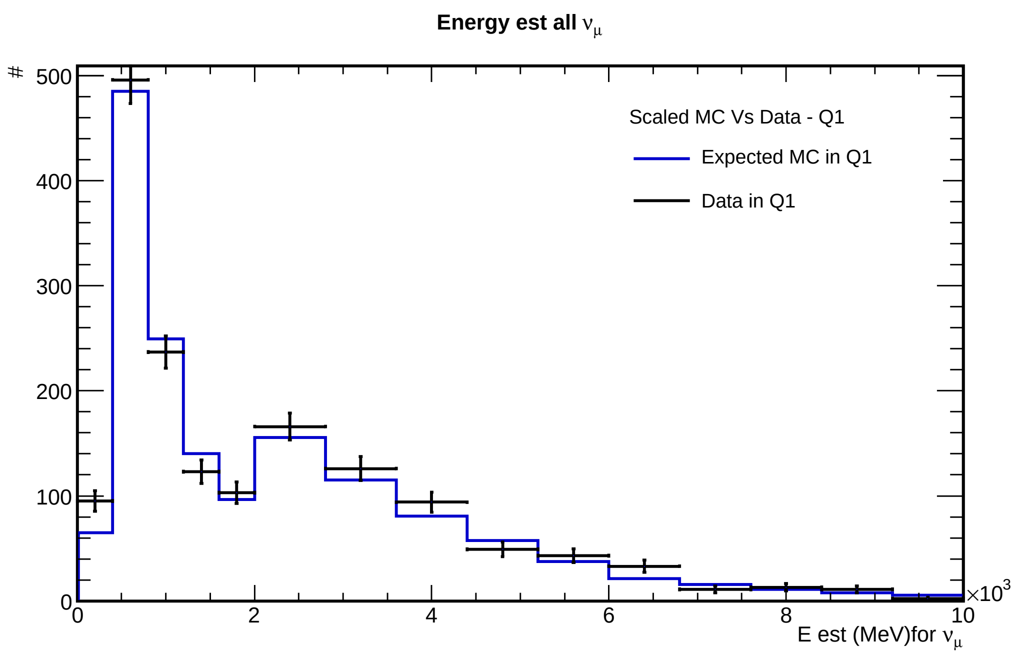

Figure 9.1: Histogram of ν_μ energies plotted by reconstructed energy for data and MC in Q1. Data plotted in black with statistical error bars given. Monte Carlo selection, scaled to match the POT of the data, plotted in blue. Note the change in binning at 2 GeV.

The data and Monte Carlo agree well. Some statistical fluctuation of bin

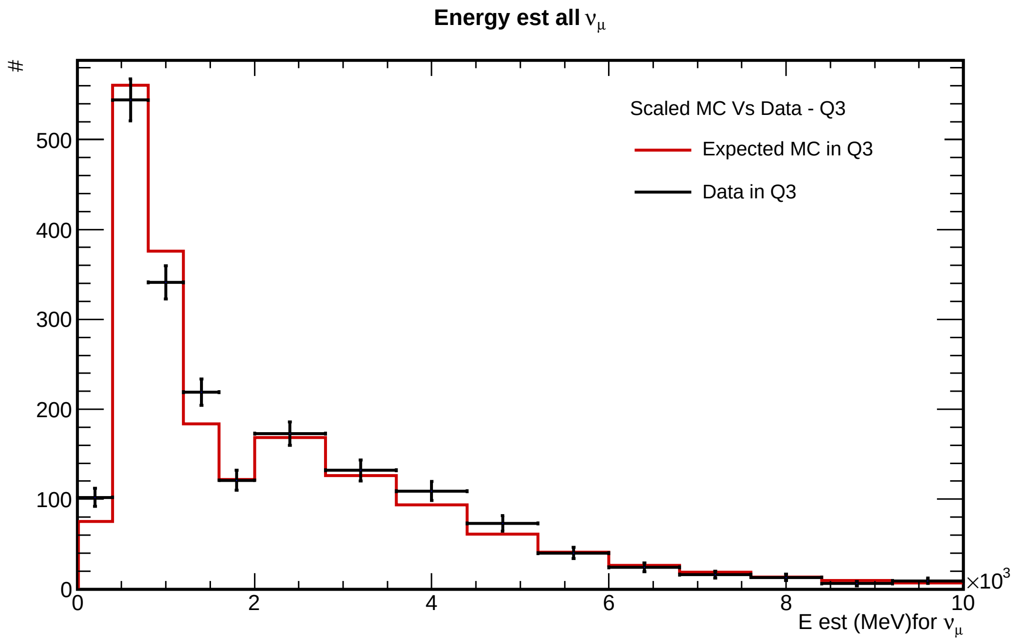

Figure 9.2: Histogram of ν_μ energies plotted by reconstructed energy for data and MC in Q3. Data plotted in black with statistical eror bars given. Monte Carlo selection, scaled to match the POT of the data, plotted in blue. Note the change in binning at 2 GeV.

contents is expected, so when the error bars are considered, and we notice that the spectrum with the higher bin content varies from bin to bin without an obvious pattern, the agreement between the spectra is satisfactory and of the level we would expect.

We therefore proceed and apply the fitting program. Since the data contains 3524 events, we allow each of T_π and T_K to vary from zero to 3524 + 59, and find the lnL value for each combination. The results are shown in Figure 9.3.

The combination of ν_μ^{fK} and $\nu_\mu^{f\pi}$ which generate the largest lnL value and are therefore the best fit to the observed data spectra is given by:

- $T_\pi = 3014^{+18.5\%}_{-14.5\%}$
- $T_K = 955^{+13.9\%}_{-11.9\%}$

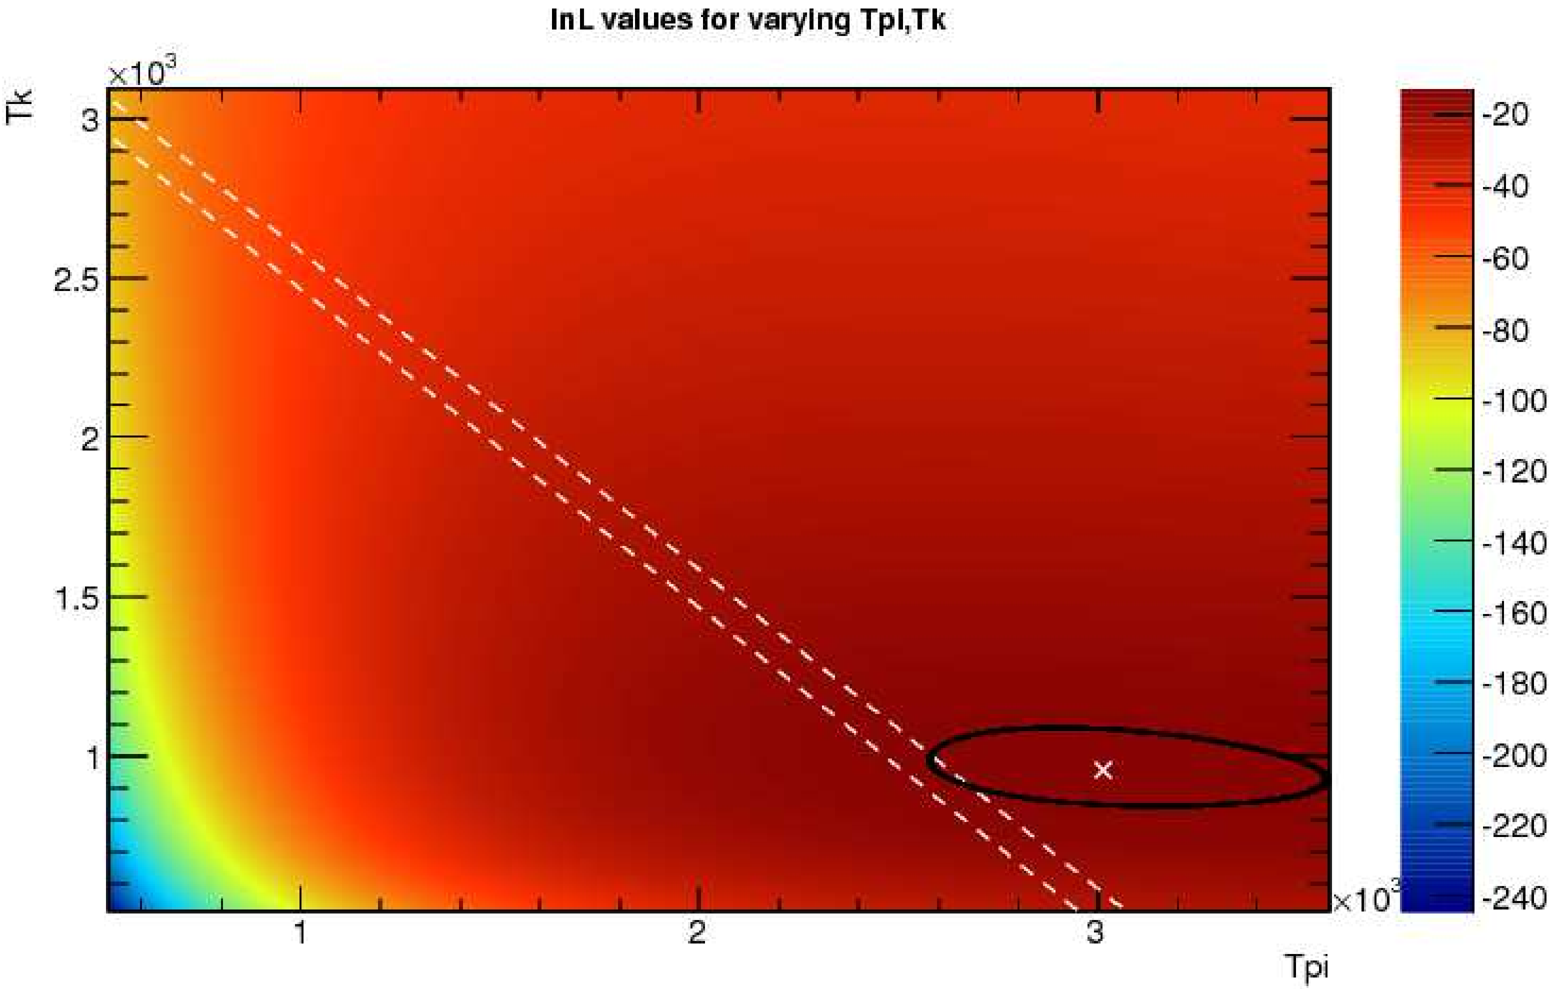

Figure 9.3: lnL values for all T_{π} and T_K combinations tested on Run II + Run IIIc data sets. lnL values for very small T values approach -∞, therefore plot is zoomed in on the contour region to remove the most unlikely T value combinations near T=0 and provide a more detailed colour scheme for the region of interest. The white dashed lines mark the region in which T_{π} and T_K sum to 3524$\pm$59. If additional data samples corresponding to the same number of POT were processed, the best fit points would fall within this contour 39.4% [103] of the time.

The uncertainties are provided by the four edges of the contour shown in Figure 9.3. To find the 1 σ uncertainty we use the lnL values, as described in §8.4. The best-fitting combination of T values are the values that correspond to the highest lnL value, which we label lnL_{opt}. All T value combinations which result in a lnL value of (lnL_{opt}-0.5) form the contour. If additional data samples corresponding to the same number of protons on target were collected and processed in the same way, the best fit points would fall within a contour this size 39.4% [103] of the time[2]. This contour includes several sources of uncertainty. The fluctuation in the number of events observed in different data samples corresponding to the same number of protons on target is included, and the statistical uncertainty on the sample size effects the size of the contour. The size of the uncertainties on the neutrino interaction cross-sections, given in Table 8.2, also affect the size of the contour through the penalty term. Reducing the magnitude of the cross-section uncertainties also reduces the size of the uncertainty on the fit. The precision of the fit can therefore be increased by increasing the size of the data sample used, and by improving the accuracy of the cross-sections measurements.

The dashed lines mark the range of T_π and T_K values whose sums are equal to 3524 $\pm$ 59. The fitting programme includes sources of uncertainty, so if the 1 σ contour and the line of $\mathrm{T}_\pi + \mathrm{T}_K = \mathrm{N}_{data}$ overlap this is considered an acceptable result. In this case however the contour touches the edge of the N_{data} range, and the sum of the T_π^{opt} and T_K^{opt} values is 3969, which is 445 above the true number of data events. This could indicate the presence of a source of systematic error. Investigations of possible systematics showed that high N_{tot} values can be a symptom of the probability distribution functions and the data spectra not matching. Figures 9.1 and 9.2 show that any discrepancy between the Monte

[2]The number of parameters of interest in a fit must be taken into account when plotting an error contour. To find the contour containing 68.3% of results, when fitting for 2 parameters using maximum likelihoods, a contour joining points with lnL_{opt} - 1.15 must be plotted. See Table 37.2 in [103] for factors.

Carlo spectra and data spectra can not be very large as it would have been visible in the comparisons. However a difference small enough to be hidden by the coarse binning and the effect of statistical fluctuations could be present. To look for this, Figures 9.1 and 9.2 are replotted using very fine binning, and are presented as Figures 9.4 and 9.5.

These more detailed plots show that the agreement between the reconstructed data events and Monte Carlo events is very good. However, focussing on the pion peaks between 0 and 1 GeV reveals a pattern in the bin content differences. Along the lower-energy edges of the pion peaks, the bin contents of the data spectra are always larger than Monte Carlo bin contents. Conversely, on the higher-energy edges of the peaks, each bin contains higher Monte Carlo contents than data contents. This pattern implies that the differences in bin content values are not purely due to statistical fluctuations and that the peaks of the data spectra occur at slightly lower energies than the peaks of the Monte Carlo spectra. An offset of this type in the peak positions would cause the fit to overestimate the number of pions required, as higher T_π values would provide the necessary events to match the data bin contents on the low energy edges of the peaks, and the resulting excess events on the high-energy peak edge could be compensated for by low T_K values. This explanation is consistent with the results of the energy reconstruction systematic error studies conducted in §8.6.2.

9.2 Energy correction

A more reliable fit result can be attained by compensating for this offset between the reconstructed energies of the data and Monte Carlo events before fitting. We cannot use the Q1 and Q3 spectra for this, as they are used to provide the fit result. However, our treatment of FGD1 provides two additional quadrants, Q2

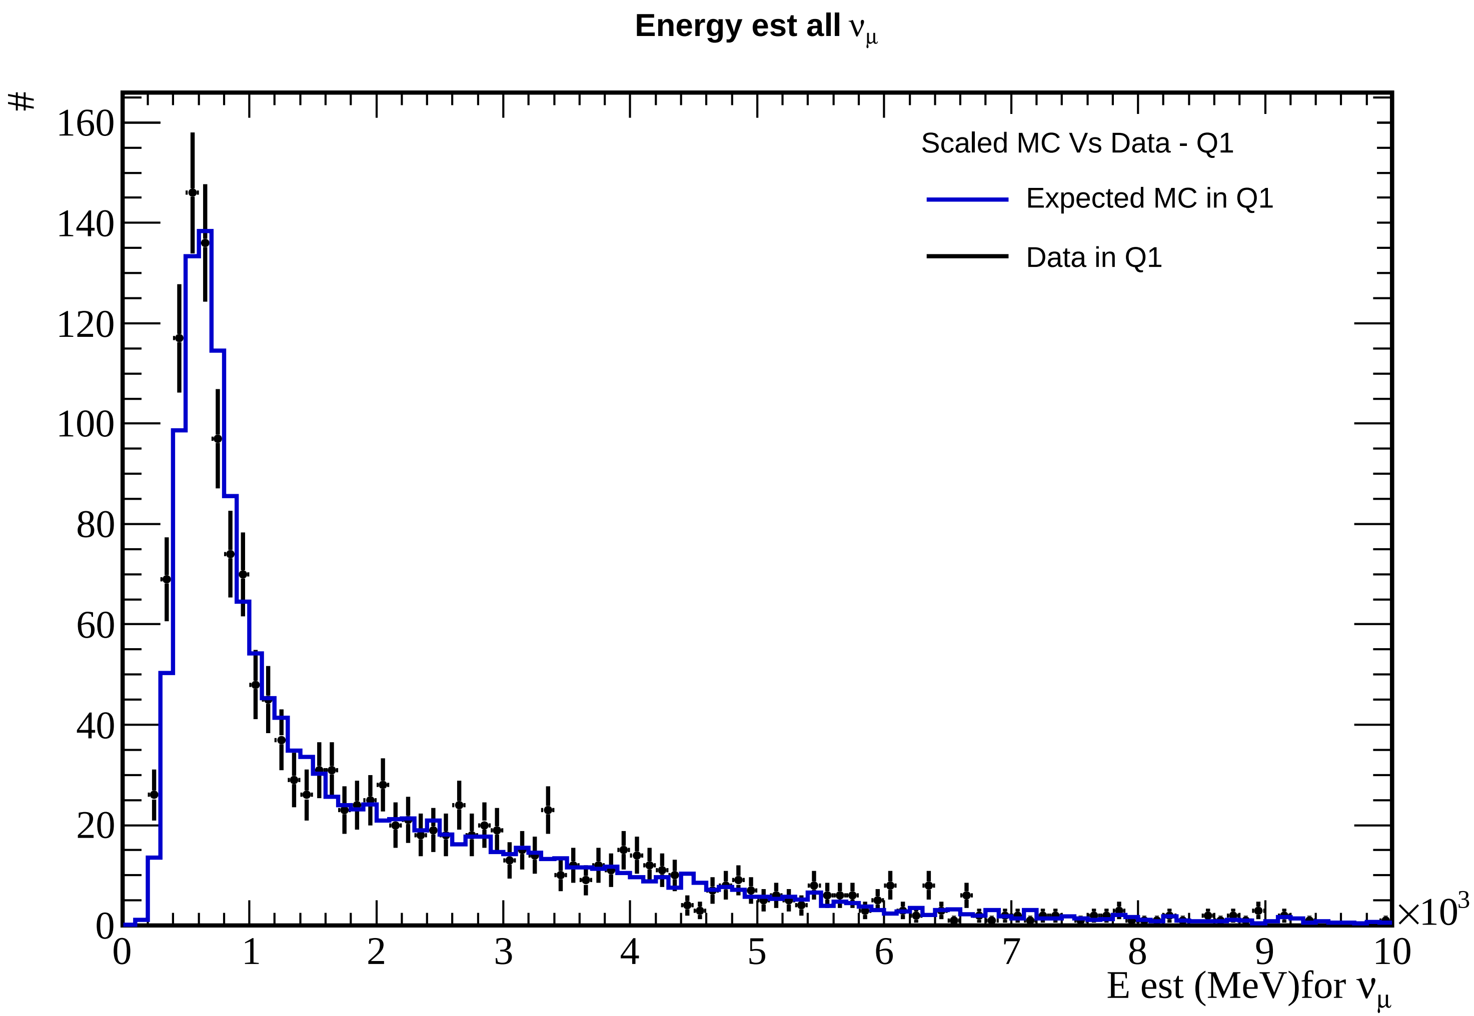

Figure 9.4: Reconstructed energy spectrum of ν_μ interactions in Q1. Monte Carlo, scaled to match the protons on target of the collected data, is plotted in blue and data is plotted, with error bars, in black. The plot shows a small offset between the data and Monte Carlo peaks, with the data spectrum leading the Monte Carlo spectrum.

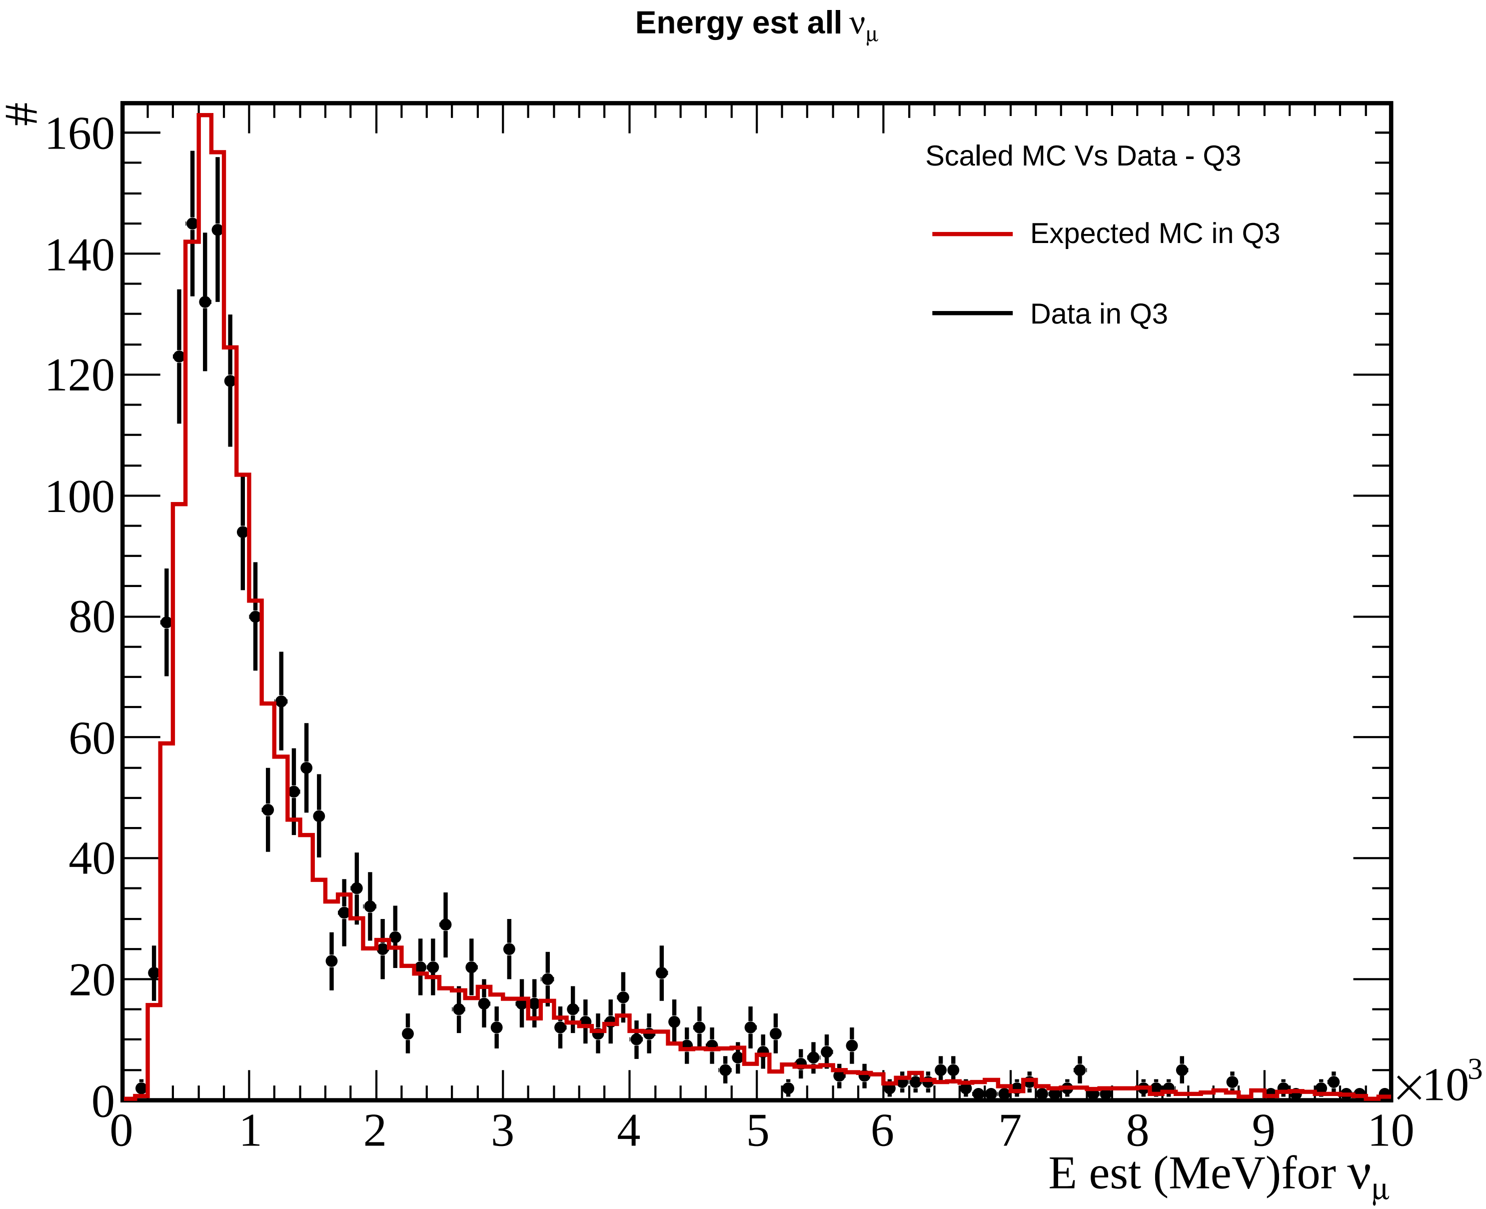

Figure 9.5: Reconstructed energy spectrum of ν_μ interactions in Q3. Monte Carlo, scaled to match the protons on target of the collected data, is plotted in red and data is plotted, with error bars, in black. The plot shows a small offset between the data and Monte Carlo peaks, with the data spectrum leading the Monte Carlo spectrum.

and Q4 (defined in Figure 7.8), which are not used by this analysis. The shift in reconstructed spectrum position between Monte Carlo events and data events is present in both the Q1 and Q3 comparisons, so the offset does not appear to be dependent on off-axis position. Neutrino interactions in Q2 and Q4 are identical to interactions in Q1 and Q3 in every way, since the entire FGD consists of the same material, with the same granularity. The only difference between the FGD quadrants is off-axis position. The event selections and energy reconstruction applied to interactions in Q2 and Q4 will function in the same way as when applied to Q1 and Q3 interactions, so the cause of the offset between Monte Carlo and data spectra should effect the Q2 and Q4 reconstructed spectra in the same way. We can therefore use the Q2 and Q4 spectra to find the size of the reconstructed energy offset, and then use this knowledge to apply a compensating correction to the probability distribution functions obtained for Q1 and Q3.

The same event selection and reconstruction schemes as described in Chapter 7 are applied, however the vertex position selections are changed to select neutrinos interacting in the Q2 and Q4 quadrants. The Monte Carlo spectra obtained are normalised to match the protons on target appropriate for the data sample, and in this way versions of Figures 9.4 and 9.5 are created for quadrants Q2 and Q4. χ^2 tests can now be conducted, in both Q2 and Q4 separately, to quantify the compatibility of the data and Monte Carlo spectra.

The Q2 and Q4 reconstructed spectra each contain just under 2000 events, so binning of 100 MeV is chosen to ensure that the bin contents are sufficient for a valid χ^2 calculation. The energy offset is most apparent in the region of the ν_μ beam peak, so bins in the range 300 MeV to 1500 MeV are chosen for inclusion, as this range contains the peak and these bins contain sufficient entries. For each of the Q2 and Q4 comparison plots, a value of χ^2 is calculated, where

$$\chi^2 = \left(\sum_i \frac{(N_i^{data} - N_i^{MC})^2}{N_i^{data}} \right) /k \tag{9.1}$$

where k is the number of degrees of freedom, which in this case is equal to one fewer than the number of bins included in the sum[3]. N_i^{MC} is the number of events in bin i of a POT normalised Monte Carlo spectrum, and N_i^{data} is the number of events in matching bin i of the data spectrum for the same quadrant. These values - one for each of Q2 and Q4 - are used as reference points. We now take the POT normalised Monte Carlo spectra and shift them down in energy by 10 MeV. χ^2 values are found again, this time comparing the data spectra and the shifted Monte Carlo spectra. This process is repeated until χ^2 values are obtained for Monte Carlo energy shifts up to 100 MeV. The values of χ^2 obtained for each of Q2 and Q4 are shown in Figure 9.6.

The best match of Monte Carlo spectra and data spectra occurs when the χ^2 value is a minimum. It is generally seen that the same shift is required for both Q2 and Q4, which have been analysed independently. When the Monte Carlo is initially processed the reconstructed energies are plotted into 10 MeV bins. Therefore when forming the shifted spectra, 10 MeV shifts are the finest shifts possible, and so this limits the precision with which we can find the peak energy offset. Performing the same test with smaller shifts would be possible with additional Monte Carlo processing. The minimum χ^2 value is obtained when a shift of 50 MeV is applied to the Q4 spectrum. In the case of the Q2 spectrum, two minima appear to be present, at 40 MeV and 60 MeV, although more detail is needed to resolve this clearly. Taking these results into consideration, a value of 50 MeV is chosen as the best measurement of the offset between the data and Monte Carlo reconstructions, with an uncertainty of 10 MeV due to the size of the shift increments tested.

[3]12 bins are considered, therefore k=11

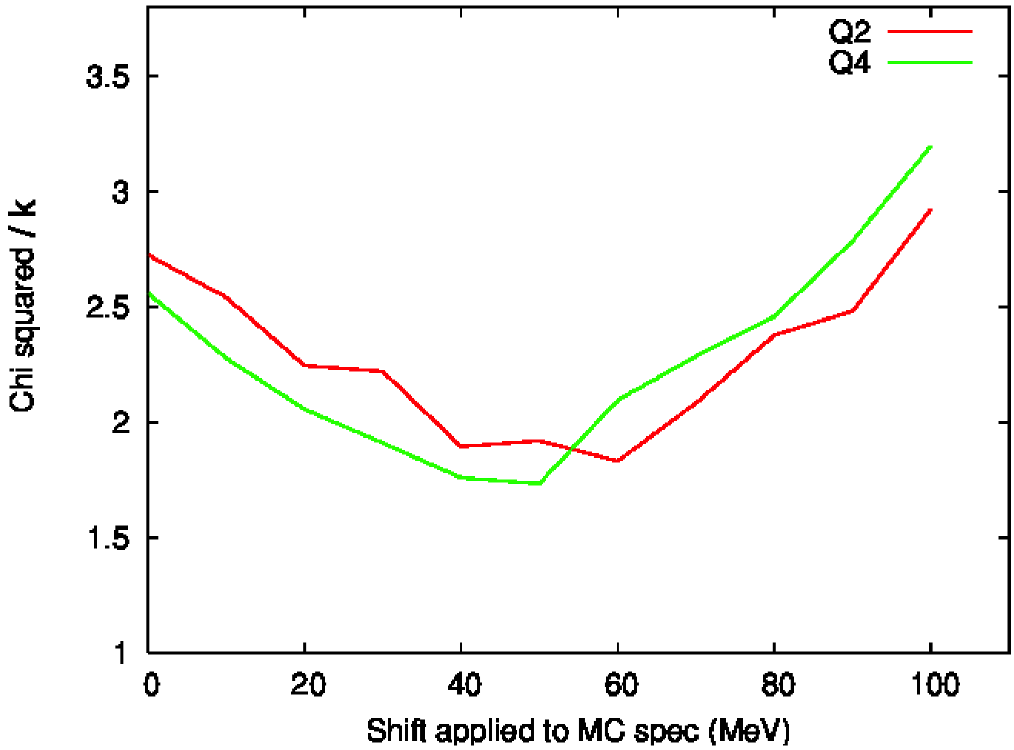

Figure 9.6: Results of comparisons between data spectra and Monte Carlo spectra. χ^2/k values plotted for comparisons with Monte Carlo spectra shifted by various energies. Magnitude of shift applied is given on the x axis.

An alternative measure of the uncertainty on the size of the energy offset can be found using the values of χ^2. As described in [103], the 1 σ uncertainty corresponds to the change in energy offset that results in a change of +1 in χ^2, and a 2 σ uncertainty corresponds to a change of +2 in χ^2. Figure 9.6 plots χ^2 per degree of freedom, therefore an increase of 2/11 from the best fit values will reveal the 2 σ uncertainty windows for each of Q2 and Q4. We find the 2 σ uncertainties for Q2 are 50^{+21}_{-16} MeV and for Q4 the 2 σ offset window is $50^{+5.1}_{-30}$ MeV, which is consistent with a precision of 10 MeV.

Using this value of 50 MeV, we can now compensate for the offset between the data and Monte Carlo peaks[4]. The original finely binned (10 MeV) ν_μ^{fK} and $\nu_\mu^{f\pi}$ Monte Carlo event distributions found in Q1 and Q3 are used to create new shifted distributions. For each bin i, the bin contents are replaced with the contents of bin i+5. This has the effect of shifting the distributions lower in energy by 50 MeV. The distributions are then rebinned according to the coarse binning scheme used for this analysis, as described in §8.1.4. These new distributions are then normalised in pairs, as described in §8.3.1, to form the shifted probability distribution functions.

9.2.1 Final result with energy correction

The new, corrected probability distribution functions described above are used as the inputs for the fitting program, and the data sample is processed once more. The results are shown in Figure 9.7.

The combination of ν_μ^{fK} and $\nu_\mu^{f\pi}$ which generate the largest lnL value and are therefore the best fit to the observed data spectra is given by:

[4]This makes the assumption that the offset is a constant. Other possibilities include an energy-dependent offset, and can be investigated along with the cause of the offset.

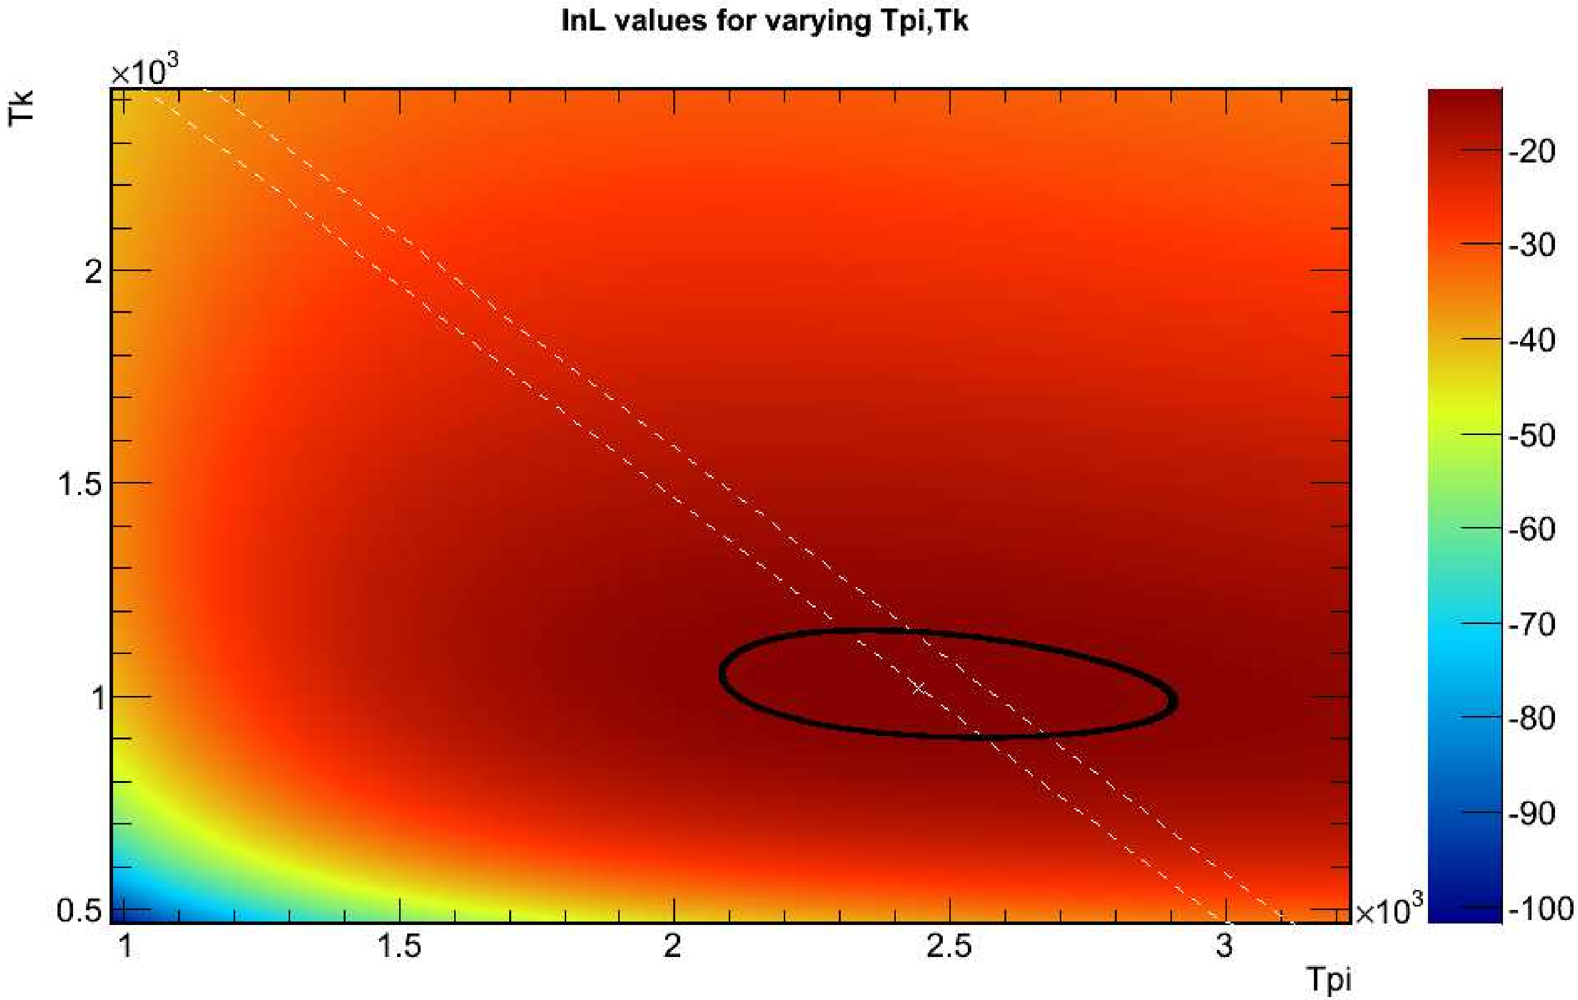

Figure 9.7: lnL values for all T_π and T_K combinations tested on Run II + Run IIIc data sets. Four probability distribution functions obtained from Monte Carlo are shifted down by 50 MeV to compensate for the difference between the data and MC energy reconstruction. lnL values for very small T values are extremely negative, therefore the display axes are restricted to the contour region to remove the most unlikely T value combinations and provide a clearer colour scheme for the region of interest. The white dashed lines mark the region in which T_π and T_K sum to 3524 ± 59. If additional data samples corresponding to the same number of POT were processed, the best fit points would fall within this contour 39.4% [103] of the time.

- $T_\pi = 2442^{+19.0\%}_{-14.6\%}$
- $T_K = 1018^{+13.6\%}_{-11.3\%}$

Energy bin / GeV	i	c_i	σ_i	Energy bin / GeV	i	c_i	σ_i
0 - 0.4	1	0.888	0.22	4.4 - 5.2	9	1.001	0.32
0.4 - 0.8	2	0.981	0.34	5.2 - 6.0	10	1.036	0.30
0.8 - 1.2	3	1.030	0.37	6.0 - 6.8	11	1.103	0.28
1.2 - 1.6	4	1.114	0.37	6.8 - 7.6	12	0.828	0.26
1.6 - 2.0	5	1.041	0.38	7.6 - 8.4	13	1.033	0.25
2.0 - 2.8	6	1.033	0.39	8.4 - 9.2	14	0.954	0.24
2.8 - 3.6	7	1.061	0.38	9.2 - 10.0	15	0.969	0.21
3.6 - 4.4	8	1.120	0.35	-	-	-	-

Table 9.1: Optimised cross-section correction factors found by fitting program per energy bin. Uncertainties on the cross-sections, obtained from Figure 8.7, labelled as σ_i, given for comparison.

The position of the best fit point and contour are now in agreement with the number of events in the data sample, which indicates that the effect of a systematic error in the energy reconstruction has been reduced.

The fit also considers corrections to the cross-sections and these corrections are given in Table 9.1. We see that the corrections are always within the levels of known uncertainty, as listed in Table 8.2. For example the largest correction is found in the 6.8 GeV - 7.6 GeV range, where a correction factor of 0.828 is found, which is a change of 17.2% on the cross-section. However the uncertainty on the cross-section on this range is 26%, therefore this correction is consistent.

Using the values of the cross-section uncertainties used by the fit as listed in Table 8.2 and the marginalised c_i values returned by the fitting program, we can calculate the value of the penalty term added to the lnL calculation to allow for the favoured changes to the cross-sections. The penalty term subtracted from the lnL value is calculated by

$$P = \sum_{i=1}^{15} \frac{(c_i - 1)^2}{2\sigma_i^2} \tag{9.2}$$

as included in Eq. 8.17. Calculating the contribution to the penalty term for each energy bin and summing these values, we find that the total penalty applied to this optimum value of lnL due to correction factors is -0.5962.

These correction factors are influenced by two effects. Their function is to allow the expected events in an energy bin, denoted μ_{iQ} in Eq. 8.17 and calculated from the probability distribution functions and T values being tested, to be scaled in order to match the observed events, n_{iQ}, more closely. This scaling could compensate for an erroneous cross-section value being used in that bin in the Monte Carlo generation. However it could also improve the agreement of the expected events with a statistical fluctuation in the observed events in that bin. Should the number of observed events in a bin happen to be particularly low or high in the data sample we have, and therefore not match the smoother distributions obtained from the more abundant Monte Carlo events, the c_i value for that bin will attempt to correct for that. As the amount of data available increases, the effect of statistical fluctuations in the bin contents will become less significant, and the c_i values returned will become increasingly sensitive to necessary corrections to the cross-sections.

Since the 1 σ statistical uncertainties are found by reducing lnL_{opt} by 0.5, and the penalty term reduces lnL by an amount larger than 0.5, this indicates that in this case the c_i terms are correcting for more than just the statistical fluctuations expected for a data sample of this size, and that they are therefore also making some small changes to some bin contents to alter the cross-sections.

Systematic effects on the fit results must also be considered.

Magnetic field uncertainty

The uncertainty on the magnetic field strength has an upper limit of 2%, therefore the effect on the fit results observed when a 2% change is applied to the track momenta, as found in §8.6.1, is the maximum change we could expect to see. These uncertainties are quoted in Table 9.2 for comparison with other error sources.

Energy reconstruction discrepancy

The same event selection and energy reconstruction scheme results in slightly different reconstructed energy spectra when applied to Monte Carlo events and data. The extent of this effect has been measured and is found to be a shift of 50 MeV, to an accuracy of 10 MeV, and this has been compensated for before applying the final fit. However, due to the limited degree of accuracy available, an offset of up to 10 MeV may still be present (this is consistent with the 2 σ uncertainties). The ability of the χ^2 test to find the offset is also limited by the number of entries in the spectra used, as what appears to be a difference in bin contents could be statistical fluctuations. The effect of an energy reconstruction systematic was studied in §8.6.2 by scaling the reconstructed energies up or down by 10%. We can estimate the remaining systematic error due to energy reconstruction using these results, with the approximation that there is a linear relationship between the energy offset and the resulting shift in T values. At the peak energy of 600 MeV, which is the region in which the offset is measured, a 10 MeV offset corresponds to a scale factor of 1.67% being applied to the reconstructed energies. Using this value and the results for the effect of 10% scale factors (given in Table 8.7), we can find the potential shifts in the fit results that would be caused by the presence of a remaining 10 MeV offset between the probability distribution functions and the data[5]. The results are given in Table 9.2.

We find that with the current cross-section uncertainties and using the events

[5]This uncertainty could be reduced in future analyses if the Monte Carlo and data are more finely binned (requiring higher event numbers), allowing the magnitude of the reconstructed energy offset to be found to a higher degree of precision.

-	T_{π}^{opt}	T_{K}^{opt}
Result	2442	1018
Uncertainty from fit	+19.0%, -14.6%	+13.6%, -11.3%
Magnetic field uncertainty	+3.2%, - 3.8%	+5.3%, -6.7%
Energy reconstruction	+4.31%, -2.24%	+3.72%, -4.98%

Table 9.2: Summary of uncertainties on the final fit results obtained using probability density functions corrected for the energy offset. Uncertainty from the fit is obtained from the one σ contour and includes statistical uncertainty and cross-section uncertainties combined. Fit uncertaintiess dominate.

available in data samples RunII and RunIIIc, the combined cross-section and statistical errors dominate.

The fit finds a total of 3460 events and tells us the percentage of these events that are ν_{μ}^{fK} and $\nu_{\mu}^{f\pi}$. We can compare these percentages to the relative contributions we expect based on the models used in the Monte Carlo. Of the 45366 neutrinos contained in the Monte Carlo spectra, 12671 are ν_{μ}^{fK} and 32695 are $\nu_{\mu}^{f\pi}$. This produces percentages of 72.1% $\nu_{\mu}^{f\pi}$ and 27.9% ν_{μ}^{fK}. These percentages can be used to find the numbers of ν_{μ}^{fK} and $\nu_{\mu}^{f\pi}$ the Monte Carlo models predict to find in a sample of 3460 events, and these numbers can be compared to the fit results. The results are summarised in Table 9.3.

Parent type	Data events	Contour limits	Predicted MC events
π	2442	min 2086, max 2906	2494.66
K	1018	min 903, max 1156	965.34

Table 9.3: Comparison of final fit results with Monte Carlo expectations

Therefore the fitting program observes 98% of the predicted $\nu_{\mu}^{f\pi}$ and 106% of the expected ν_{μ}^{fK}. The predicted number of $\nu_{\mu}^{f\pi}$ according to the Monte Carlo ratios is well within the 1 σ statistical uncertainty on the observed number of $\nu_{\mu}^{f\pi}$. The fit finds a slightly higher number of ν_{μ}^{fK} than predicted by the Monte Carlo, however the expected number of ν_{μ}^{fK} is comfortably within the 1 σ statistical uncertainty on the fit.

9.3 Additional data studies

Some additional fits are performed using the data sample and are described below.

9.3.1 Altering the influence of the correction factors

The correction factors c_i are calculated for each combination of T_π and T_K such that the highest value of lnL can be achieved. These factors will therefore make small changes to the expected bin contents that are compared to the data if a small change can provide a better match. However, the extent to which bin contents can be altered by these factors is limited by the values of sigma given in Table 8.2, since there is a known level of uncertainty on the cross-section measurements. By altering the sigma values used we can investigate how the fitting program would function if these correction factors were allowed to behave differently.

Firstly we effectively remove the correction factors by forcing each value of c_i to be 1, thus assuming that the cross-sections used in the Monte Carlo generation are exactly correct[6]. This is achieved by setting the sigma value for each bin to be almost zero[7], which corresponds to the absence of any uncertainty. This would produce a penalty value, given in Eq.8.17, approaching infinity should c_i be anything other than 1 in any energy bin i. Performing the fit to the data sample replacing the σ_i values listed in Table 8.2 with $\sigma_i = 0.00001$ gives the result shown in Figure 9.8.

Fitting to the data using these σ_i settings gives the results

- $T_\pi = 2438^{+2.42\%}_{-2.30\%}$
- $T_K = 1086^{+4.24\%}_{-3.87\%}$

[6] This corresponds to using Eq. 8.15 instead of Eq. 8.16 to find the expected events in each energy and quadrant bin

[7] A value of precisely 0 in the denominator would cause an error so $\sigma = 0.00001$ is used.

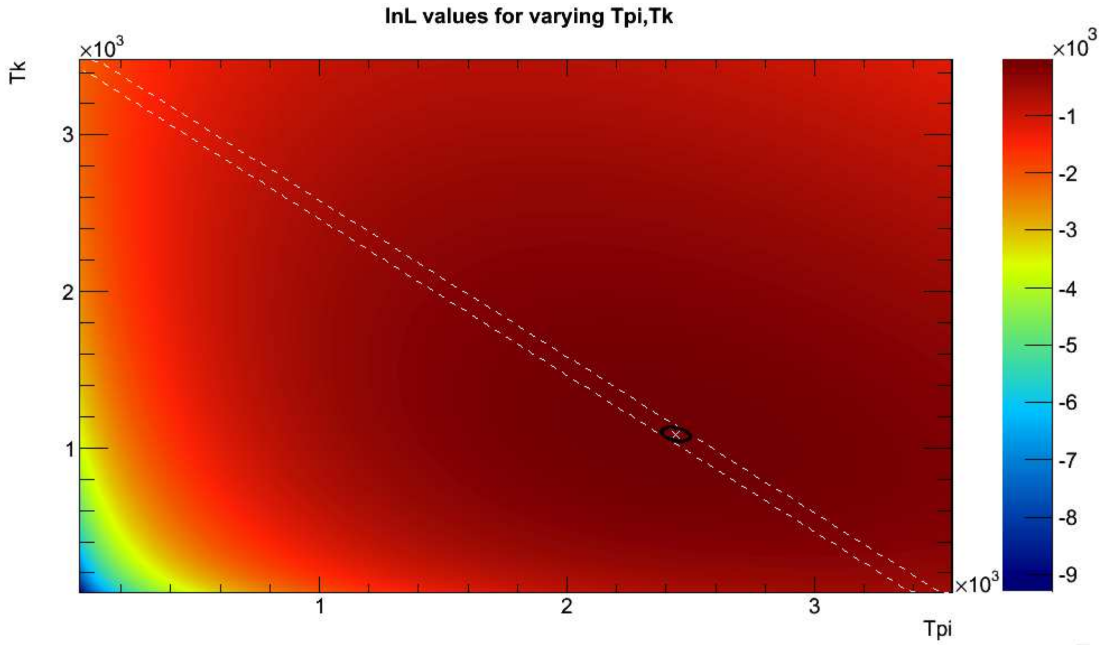

Figure 9.8: Results of fit to the data sample using $\sigma_i = 0.00001$. This result assumes that the cross-sections used in the Monte Carlo generation are correct to 0.001% and is the result that would be obtained with no compensation for potential cross-section uncertainties. The white dashed lines mark the region in which T_π and T_K sum to 3524±59. If additional data samples corresponding to the same number of POT were processed, the best fit points would fall within this contour 39.4% [103] of the time.

The 1 sigma uncertainties given on the final fit results in §9.2.1 combine the statistical uncertainty due to the size of the data sample, and the uncertainties on the cross-sections. By fitting with the assumption that the cross-sections are correct and not including their associated uncertainties, we can separate the contributions to the uncertainties on the fit results. The uncertainties given by the 1 sigma contour in this test do not include any cross-section uncertainty, and so represent the statistical uncertainty due to the limited data sample only. They are considerably smaller than the standard uncertainties, which demonstrates that the biggest challenge when measuring beam contributions is acquiring precise measurements of the cross-sections with minimal uncertainties. If the fit were performed in this way, without allowing for small changes to the cross-sections, the uncertainty due to the cross-section uncertainties would have to be evaluated separately and included as an additional source of systematic uncertainty.

To illustrate the uncertainty caused by the cross-sections, we can also fit to the data using the opposite scenario of increased cross-section uncertainties. For this test we set the σ_i values to 1000, which corresponds to the case that the cross-sections are effectively unknown. The probability distribution functions are created assuming certain cross-section values. By increasing σ_i to a very large value, the penalty term will always be very small. This effectively removes the penalty associated with altering the Monte Carlo cross-sections. Therefore the c_i values will be allowed to change by any amount that is necessary to provide the best match between the expected events and the data, without the lnL value being effected. Fitting to the data sample with $\sigma_i = 1000$ produces the result shown in Figure 9.9.

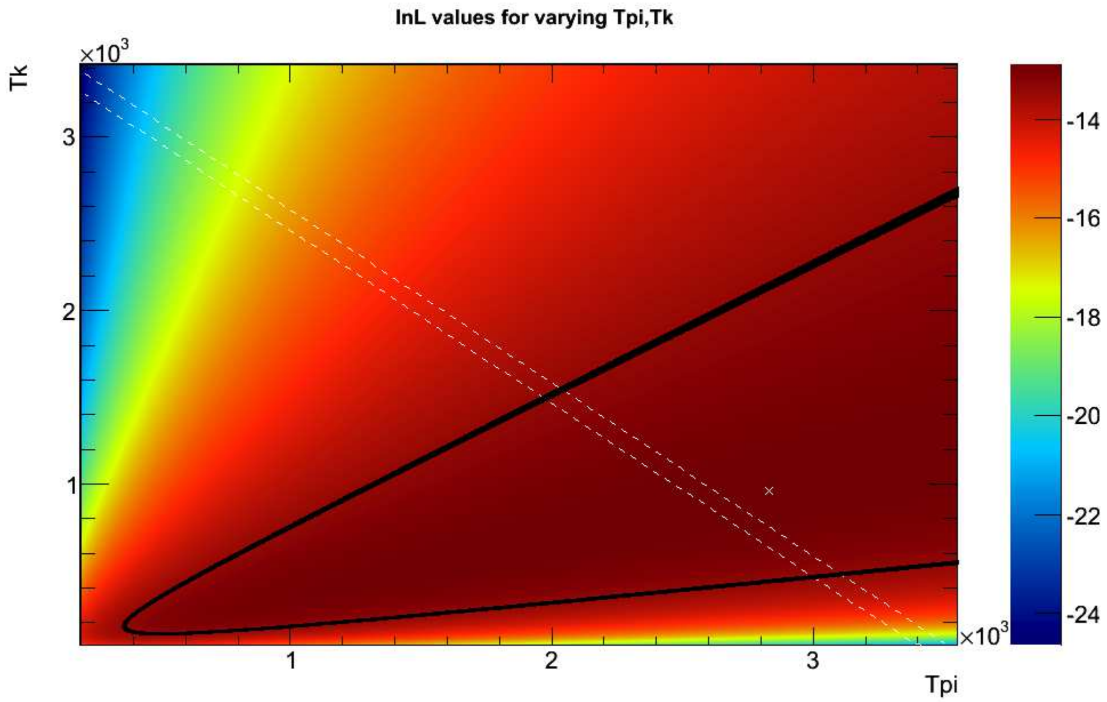

Figure 9.9: Results of the fit to the data sample using $\sigma_i = 1000$. This result corresponds to the scenario in which the cross-sections are effectively unknown, and therefore have very large uncertainties which effectively remove the penalty term from the lnL calculation. The white dashed lines mark the region in which T_π and T_K sum to 3524±59.

The best fitting totals of $\nu_\mu^{f\pi}$ and ν_μ^{fK} in the data in this case are found to be

- $\mathrm{T}_\pi = 2830^{+unknown\%}_{-87.0\%}$
- $\mathrm{T}_K = 958^{+unknown\%}_{-85.8\%}$

This demonstrates that a complete lack of knowledge of the uncertainties on the cross-section measurements would make it almost impossible to measure the relative contributions to the beam from different parents, as the 1 σ contour contains over 85% of the range of possible T values, such that almost any combination of T_π and T_K could be a plausible fit to the data. In this case there is no restriction on the c_i values from the σ values, so that the bin contents could be altered to any degree in order to exactly match the data. However they are still restricted by the property that only one correction factor value can be used per energy bin i, and must apply to both the Q1 and Q3 test distributions. Since the off-axis bins contain different bin contents due to the different spectrum shapes, this provides some constraint on the possible c_i values. This is one way in which fitting to multiple spectrum shapes is advantageous.

9.3.2 Using a limited energy range

As shown in Figure 8.1, the high energy neutrinos are produced almost exclusively by kaon decays. Therefore the data events in the high energy tail consist almost entirely of contributions from the ν_μ^{fK} spectrum. Since the spectrum contains an area where only one parent type is present, it can find the appropriate scale for the ν_μ^{fK} contributions, and therefore the likely value of T_K using these high energy bins. The optimum T_π value can then be found assuming this value of T_K.

To confirm that the fitting program can operate effectively when this is not the case, and instead contributions from both the ν_μ^{fK} and $\nu_\mu^{f\pi}$ spectra are present in all bins, we can apply the fitting program to the data sample again but this time only consider entries in lower energy bins when calculating lnL. The probability distribution functions are not changed and the optimum T values returned still correspond to the number of events in the full data spectrum, but the fitting programme will not be able to incorporate information from the high energy tail in its calculation of lnL. An energy range of 0–3.6 GeV is chosen as all energy bins in this range contain both ν_μ^{fK} and $\nu_\mu^{f\pi}$. As above, lnL is calculated according to Eq. 8.17, except that the sum over energy bins is now limited to include only bins $1 \leq i \leq 7$. This will test the program's ability to separate the ν_μ^{fK} and $\nu_\mu^{f\pi}$ in bins that contain both. The result is given in Figure 9.10.

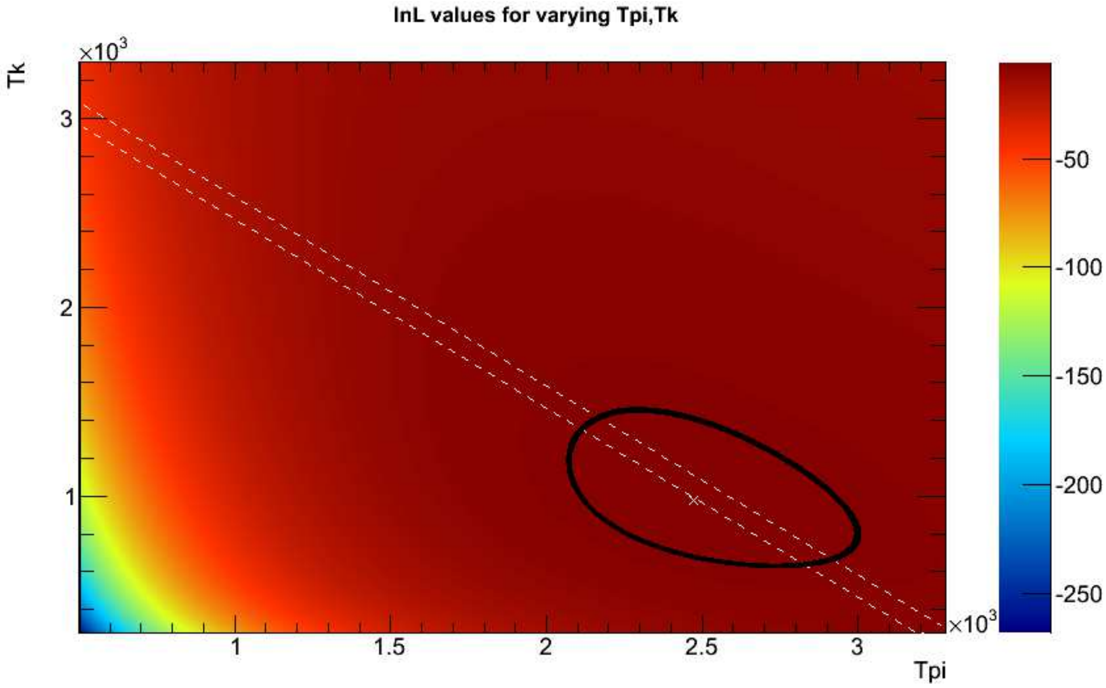

Figure 9.10: Testing the fitting procedure with a limited number of bins. Results of the fit to the data sample using only energy bins 1–7, which corresponds to 0–3.6 GeV. The white dashed lines mark the region in which T_π and T_K sum to 3524±59. If additional data samples corresponding to the same number of POT were processed, the best fit points would fall within this contour 39.4% [103] of the time.

The best fitting totals of $\nu_\mu^{f\pi}$ and ν_μ^{fK} in the data in this case are found to be

- $\mathrm{T}_{\pi} = 2474^{+21.2\%}_{-16.2\%}$

- $\mathrm{T}_{K} = 976^{+49.3\%}_{-35.3\%}$

Compared to the standard results given in Table 9.2, we see that there is a small increase in the value of T_{π}^{opt} and a small decrease in T_{K}^{opt}, however these changes are safely within the 1 σ contours, and therefore the results of this limited energy range test are equally plausible parent contributions to the data. However, reducing the quantity of data available by excluding the higher energy region of the spectrum will increase the statistical uncertainty on the results. A small number of $\nu_{\mu}^{f\pi}$ are excluded therefore the width of the 1 σ contour does not show a large increase. The uncertainties on the value of T_{K}^{opt} are significantly larger however. This is to be expected as approximately half of the ν_{μ}^{fK} spectrum is removed in this case. Therefore while the the best fit results are not significantly affected by the use of a limited spectrum range, the statistical uncertainties can be reduced by including all available data.

Tests using only limited energy ranges can in principle also confirm that the probability distribution functions used are approximately correct. Should the distribution shapes in a small energy range be significantly wrong, we would not expect the fitting procedure to return values of T which are similar to the full range results. This could be used to investigate the accuracy of the $\nu_{\mu}^{f\pi}$ and ν_{μ}^{fK} distributions once additional data and smaller cross-section uncertainties are available.

9.4 Evaluation of method and possible extensions

The results of this analysis emphasise the importance that the reconstructed spectra used for the probability distribution functions and the reconstructed data spectra match. The observed presence of a difference in reconstructed energies when the same event selection criteria and energy reconstruction methods are applied to data events and Monte Carlo events is of interest, and while the effect can be compensated for using events that are otherwise not of use to this analysis, it is important to understand the source of this difference. Some possibilities are considered.

The reconstructed energy spectra in Q1 and Q3 can each be divided into events which undergo a CCQE interaction and events which interact through a CCnonQE channel. The reconstructed energies are calculated differently for these different interaction types, as described in detail in Chapter 7. Plotting the data and Monte Carlo comparisons separately by interaction type will demonstrate if the offset is present for only one type of interaction, which would indicate that the source of the discrepancy is in the method of energy reconstruction. However, we find that the majority of the events that form the peak, and therefore are $\nu_\mu^{f\pi}$, are CCQE interactions, and that the majority of the CCnonQE interactions correspond to the higher energy ν_μ^{fK} interactions. The offset is clearly present for the CCQE events, which confirms that the source of the offset applies to the CCQE energy reconstruction formula. However, the CCnonQE spectra are broadly distributed and do not feature a prominent peak, and also contain far fewer interactions. The result is that determining if an offset is present for the CCnonQE interactions is not possible with this number of interactions. Therefore we cannot state whether the source of the offset is related to a specific method of energy reconstruction or applies to all selected events. However it may be possible to investigate this when additional data is included.

Another potential cause of the offset is a slight error in the modelled off-axis position of ND280. We have seen that the spectrum peak becomes more narrow and occurs at lower energies as the off-axis angle of the selected ν_μ increases. The Monte Carlo is generated using an off-axis angle of 2.5 degrees. If the true position of ND280 is such that the off-axis angle is slightly larger than the value used in the Monte Carlo, then the energies of the data events would be lower than the modelled events. Therefore despite the event selection criteria and energy reconstruction methods being the same, the data spectra would peak at a lower energy due to the data events all having slightly lower energies than the Monte Carlo events.

The uncertainty on the measurement of the off-axis position of ND280 with respect to the beam is measured by MUMON to be 0.21 mrad [104]. According to [100], a 16 MeV offset in peak energy corresponds to a change of approximately 1 mrad. Therefore an offset of 50 MeV would require a difference of 3.1 mrad between the actual position of ND280 and the off-axis angle used for Monte Carlo generation. Since this required error is an order of magnitude larger than the measured uncertainty in off-axis angle, we can be sure that this is not the cause of the observed discrepancy between the Monte Carlo and data spectra.

It is unfortunate that the fitting program was applied before the presence of an offset was observed and compensated for, as this prevents this analysis from being fully blind. However, additional, larger data sets are now available which can be used to repeat this analysis. The division of FGD1 used provides two surplus quadrants, Q2 and Q4, which can be used to evaluate the compatibility of the reconstructed energy spectra of the Monte Carlo and data without looking at the Q1 and Q3 events used for the analysis. Any offset present can be measured using these additional quadrants and this can be compensated for during the creation of the probability distribution functions for Q1 and Q3. These adjusted distributions can be used when fitting to the new data sets. Since new data and Monte Carlo would require processing, finer binning can be used initially, thus enabling a more precise measurement of the energy offset. This will reduce the uncertainty on the size of the offset and therefore reduce the related energy reconstruction systematic uncertainty.

The current results indicate that the pion and kaon contributions to the beam are modelled well in the Monte Carlo. Slight deviations from the expected numbers are observed, but these are comfortably within the 1 σ boundaries. This result is consistent with other, independent studies of the kaon contribution to the T2K beam, such as [95], which also finds the Monte Carlo kaon contribution to agree with data at the 1 σ level.

One of the current sources of uncertainty on the fit results is statistical. The more events contained in each of the data spectra, the smoother the data distributions will be, and the lower the chance that the fit will select values of T_π and T_K that agree with bin contents which have been heavily affected by statistical fluctuations. Increasing the number of events is therefore one way to reduce the 1 σ contour.

The dominant source of uncertainty on the current fit results is caused by the degree of uncertainty on the cross-section measurements. This can be seen by comparing the size of the contours in Figure 9.7 and Figure 9.8, where only the statistical uncertainty is present in the latter. Due to the presence of the penalty term in Equation 8.17, smaller uncertainties would constrain the cross-section correction factors more, reflecting the fact that the cross-section values used are more precise. Changes to the scale of cross-section corrections which are allowed would be reflected in the lnL values found for the various T_π and T_K

combinations tested, and may result in changes to the optimum T values found, while also reducing the size of the 1 σ contour.

The analysis in its current form demonstrates that a fit of this nature is an effective method of indirectly measuring the kaon and pion contributions to the neutrino beam, and produces a result which is comparable to the current level of precision for pions, and an improvement on the kaon measurement at high energies [34]. There are possible extensions that can be applied to improve the results, some of which are discussed below.

For this analysis FGD1 is divided into 4 quadrants. This provides two regions with different mean off-axis angles, and two further quadrants which can be used to check for systematic errors without affecting the events used for the analysis. This selection of off-axis angle bins is rudimentary however, and with further information about the exact geometry of the detector with relation to the beam axis, the definitions of the regions could be refined. Currently, an interaction in the top left corner of Q3 will have a very similar off-axis angle to one in the bottom-right corner of Q1 (see Figure 7.8). The result is that the peaks plotted for Q1 and Q3 will be smeared somewhat by the range of off-axis angles of the neutrinos interacting in each quadrant. One example of an improvement would be to define the two off-axis bins such that Q1 and Q3 do not touch, and the boundaries of Q1 and Q3 are arcs on circles of different radii drawn around the beam axis. This would result in the interactions contained in each off-axis bin having a smaller range of off-axis angles, and therefore sharper peaks of reconstructed energy. The gap between the regions would result in peaks with a clearer separation, and also provide event samples which would be unused by the fitting procedure, but which could be used to investigate the presence of energy reconstruction systematics before fitting to the data.

The available data can also be increased by including interactions in FGD2.

A suitable system of track selection and event reconstruction would need to be developed to select interactions occurring in FGD2. There are various possible methods of including these extra events. Detailed geometry information could be used to carefully select regions containing ranges of off-axis angles that match the Q1 and Q3 off-axis angle ranges in FGD1. Events selected in either FGD could then be combined into the same two off-axis angle bins. This would increase the number of events in the data spectra for the two samples which would increase the precision of the fit. The alternative is to consider FGD2 events separately, and calculate lnL by summing over 4 different off-axis angle bins - Q1 and Q3 in FGD1 and Q5 and Q7[8] in FGD2. This would not require careful alignment of the matching regions in the different FGDs, since the four spectra would be considered separately for fitting.

Another possible future extension is using interactions in alternative regions of ND280. Opposite sides of the barrel ECAL are further apart than opposite sides of FGD1, and therefore regions with a larger difference in off-axis angle, and energy spectra, could be defined. It is not currently possible to reconstruct interactions in the ECAL as reliably as the FGD interactions; however, as further event selections are developed for ND280 these alternative events could be used.

This analysis has assumed that the modelled distributions of $\nu_\mu^{f\pi}$ and ν_μ^{fK} used to generate the Monte Carlo are correct, and has focussed on measuring the relative contributions of these distributions. However, as discussed in §8.1, there is also a level of uncertainty on the shapes of these distributions. This method of fitting could be applied to assess how correct the $\nu_\mu^{f\pi}$ and ν_μ^{fK} distributions currently used are. The current distributions would be used as the basic forms, and then a system must be devised to parameterise alterations to the shapes of the $\nu_\mu^{f\pi}$ and ν_μ^{fK} distributions. Such parameters could include scale factors that

[8]following the quadrant naming convention for FGD1, the FGD2 quadrants could be labelled 5,6,7,8 clockwise from the top left

stretch the distributions. The fit could then operate in a similar way, except that instead of calculating lnL for different combinations of T values, the contributions from the two parents could be fixed, and instead lnL would be calculated for all combinations of allowed values of the distribution shape parameters. These lnL values could be plotted and once again the highest value of lnL would indicate the best fitting parameters, and therefore the forms of the distributions which match the data best. While it would require longer to process due to the increased number of loops, it would also be possible to allow all of T_π and T_K and several shape parameters to vary, and find the optimum combination of all parameters. This would provide the most complete understanding of the contributions made by each of the neutrino parents.

The data samples considered in this analysis, and the additional data sets suggested for inclusion, all contain the same contributions from pions and kaons. An interesting use of the fitting program would be to analyse a data set gathered with a different setting of the magnet horn currents. Changing the strength of the magnetic fields used to select mesons produced at the target may change the percentage of kaons selected to enter the decay pipe, and therefore would change the contribution to the neutrino beam from kaon decays, which would effect the number of high energy neutrinos contributing to T2K's backgrounds. To date, only a limited amount of data has been collected with different horn current applied (Run IIIb - see Appendix B for details). However, if sufficient matching data and Monte Carlo sets existed for different beam configurations, this fitting method could be used to observe any change in the kaon contribution to the beam.

Chapter 10

Conclusions

This thesis considers the two main backgrounds to the ν_μ to ν_e oscillation appearance measurement. Both background sources are dependent on the number of kaons which enter the decay pipe and decay. In order to accurately predict the rate of background events, we must have precise measurements of the kaon content in the decay pipe. Currently our knowledge of the hadron production at the target is provided by limited measurements taken at other experiments, and theoretical models.

In principle a measurement of the intrinsic high energy ν_e detected at the far detector could be used to check the accuracy of the ν_e background predictions. Measurements of high energy ν_e interaction rates at Super-K and ND280 could also be used to calculate the intrinsic ν_e fluxes at each detector (given the assumed cross-section values used by T2K). Since the near and far detectors observe the same neutrino flux but use different target materials, comparison of the calculated fluxes could reveal errors in the scintillator and water interaction cross-section values used. Applying suitable selection criteria to Super-K Monte Carlo we find the expected ν_e event rate in a nominal year of data taking at the far detector to be 10.8 ± 3.3(stat) with a background of 11.3 ± 3.4 events. Intrinsic ν_e measurements are therefore not yet useful due to the low purity and statistical limitations, as described in Chapter 4. However, with several years of data taking, additional cuts may be applied to improve the purity and an intrinsic ν_e sample may be of use.

Consideration of the parents of the beam ν_μ reveals that kaon decay produces the high energy ν_μ largely responsible for the π^0 background. In Chapter 5 we test the possibility of using a sample of high energy ν_μ detected at Super-K to provide a check of the modelled kaon contribution to the ν_μ flux. We conclude that this is not feasible for two reasons. Firstly the high energy tail contains too few events to provide a sample that is not severely statistically limited. Secondly, measurements of neutrino interactions reveal the product of the neutrino flux and interaction cross-section only. Therefore our ability to extract the neutrino flux is limited by the current uncertainties on the cross-sections.

A method of measuring the number of ν_μ produced directly from kaons and pions using ND280 data is developed. Existing knowledge of the kaon and pion decay branching ratios and kinematics can be used to find the combination of kaons and pions in the decay pipe that would produce the measured separate energy spectra of ν_μ produced from kaon decay, ν_μ^{fK} and ν_μ produced from pion decay, $\nu_\mu^{f\pi}$. That kaon content information could then be used to predict the number of background events. This method is therefore useful to any ν_μ beam experiment that requires detailed parent information for background calculations. Neutrino experiments in which the detector region is sufficiently wide for changes to the neutrino spectrum with differing off-axis angle to be resolved can use this technique to gain insight into the hadron production occurring at their target.

The fitting program developed in Chapter 8 provides a successful method of measuring the number of ν_μ^{fK} and $\nu_\mu^{f\pi}$ in the T2K beam. A measurement of the number of ν_μ^{fK} and $\nu_\mu^{f\pi}$ in the selected data sample is obtained, and we find that the fitting method observes 98% of the $\nu_\mu^{f\pi}$ and 106% of the ν_μ^{fK} predicted by the Monte Carlo. Full results are given in §9.2.1.

This assumes that the modelled energy spectrum shapes of ν_μ^{fK} and $\nu_\mu^{f\pi}$ that are used as the basis for the fitting method are correct, however, currently there are no experimental measurements that cover the full range of T2K energies. The version of the program tested in this thesis can be modified to find the shapes of the individual ν_μ^{fK} and $\nu_\mu^{f\pi}$ energy spectra in addition to the integrals. This possible extension and others are discussed in more detail in §9.4.

Appendix A

Meson decay kinematics

The most common pion decay mode, with a branching ratio of 99.99%, is $\pi^+ \rightarrow \mu^+ + \nu_\mu$, with an equivalent decay for the π^-. The formula for the energy of the neutrino, as a function of the pion energy, is derived here.

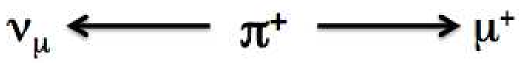

Figure A.1: Pion decay in the pion rest frame

Considering the decay in the rest frame of the pion, as shown by Figure A.1, conservation of energy and momentum gives the four vector equation

$$\boldsymbol{P_\pi} = \boldsymbol{P_\nu} + \boldsymbol{P_\mu} \tag{A.1}$$

where

$$\boldsymbol{P_\pi} = (E_\pi, 0) \tag{A.2}$$

$$\boldsymbol{P_\nu} = (E_\nu, \underline{p}_\nu) \tag{A.3}$$

$$\boldsymbol{P_\mu} = (E_\mu, \underline{p}_\mu) \tag{A.4}$$

Making $\boldsymbol{P_\mu}$ the subject of A.1, squaring, and using

$$-||\boldsymbol{P}||^2 = E^2 - |\underline{p}|^2 = m^2 \tag{A.5}$$

gives

$$m_\mu^2 = m_\pi^2 - 2\boldsymbol{P_\pi}.\boldsymbol{P_\nu} + m_\nu^2 \tag{A.6}$$

Using the approximation that the neutrino is massless, and taking the product of $\boldsymbol{P_\pi}$ and $\boldsymbol{P_\nu}$, we find the energy of the neutrino in the rest frame of the pion, E_ν^{rest}.

$$m_\mu^2 = m_\pi^2 - 2E_\pi^{rest}E_\nu^{rest} \tag{A.7}$$

$$E_\nu^{rest} = \frac{m_\pi^2 - m_\mu^2}{2E_\pi^{rest}} \tag{A.8}$$

In the rest frame of the pion, $E_\pi^{rest} = m_\pi$, therefore

$$E_\nu^{rest} = \frac{m_\pi^2 - m_\mu^2}{2m_\pi} \tag{A.9}$$

The rest frame of the pion is moving parallel to the z axis of the lab frame with speed β, and the neutrino is emitted at angle θ (in the lab frame) to the pion direction. Applying a Lorentz transformation, the energy of the neutrino in the lab frame, E_ν, is

$$E_\nu = \frac{E_\nu^{rest}}{\gamma(1 - \beta cos\theta)} \tag{A.10}$$

where β and γ are the Lorentz parameters of the pion in the lab frame.

Substituting Eq. A9, we find

$$E_\nu = \frac{m_\pi^2 - m_\mu^2}{2m_\pi\gamma(1 - \beta cos\theta)} \tag{A.11}$$

Since $p_\pi = \gamma m_\pi \beta$ and $E_\pi = \gamma m_\pi$, we find that the energy of the neutrino in the lab frame is given by

$$E_\nu = \frac{m_\pi^2 - m_\mu^2}{2(E_\pi - p_\pi cos\theta)} \tag{A.12}$$

The equivalent expression can be found for the case of 2-body K^+ decay to $\mu^+\nu_\mu$ by substituting E_π, m_π and p_π with E_K, m_K and p_K respectively.

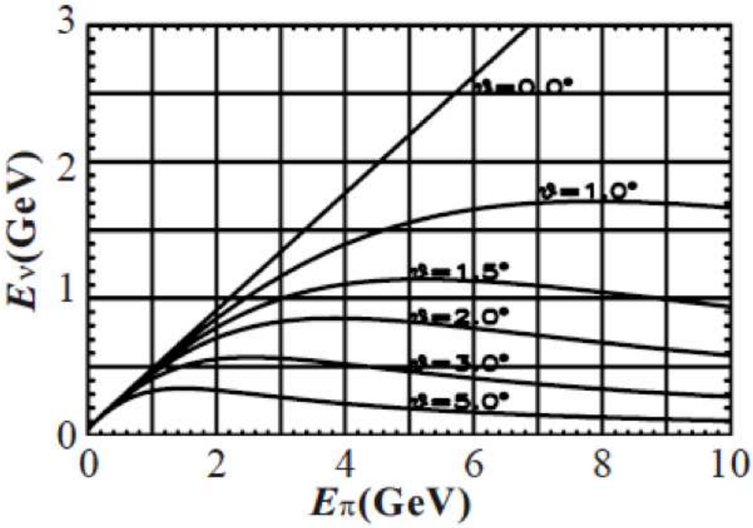

Figure A.2: Neutrino energy with respect to parent pion energy for a range of opening angles θ. Figure from [57]

This relationship is plotted for several values of cosθ in Figure A.2. When θ is large, as in the case of an off-axis sampling of a neutrino beam, the neutrino energy peaks and then becomes approximately constant with respect to the parent energy. For the special case that $\theta = 0$, which corresponds to neutrinos observed by on-axis detectors, the relationship between neutrino energy and parent meson energy is linear. Equation A11 becomes

$$E_\nu = \frac{m_\pi^2 - m_\mu^2}{2m_\pi\gamma(1-\beta)} \tag{A.13}$$

Expressing β as

$$\beta = (1 - \frac{1}{\gamma^2})^{1/2} \tag{A.14}$$

and taking the two most significant terms of the binomial expansion of this allows Eq. A11 to written in terms of γ.

$$E_\nu = \frac{m_\pi^2 - m_\mu^2}{2m_\pi}.\frac{1}{\gamma(1-(1-\frac{1}{2\gamma^2}))} = \frac{m_\pi^2 - m_\mu^2}{2m_\pi}.2\gamma \tag{A.15}$$

Using $E_\pi = \gamma m_\pi$, we find

$$E_\nu = E_\pi.\frac{m_\pi^2 - m_\mu^2}{m_\pi^2} \tag{A.16}$$

$$E_\nu \approx 0.43 E_\pi \tag{A.17}$$

Therefore E_ν is proportional to E_π when on-axis.

Appendix B

T2K data and ND280 Monte Carlo

Listed below are the beam configuration details for T2K data runs conducted up to May 2013. Run II and Run IIIc were used for the analysis presented in Ch. 9.

T2K Run	Dates	Bunches	Horn Current (kA)	Beam power (kW)	Total POT ($\times 10^{20}$)
Run I	Jan 2010 - Jun 2010	6	250	up to 50	0.32
Run II	Nov 2010 - Mar 2011	8	250	up to 150	1.11
Run IIIb	Mar 2012 -	8	205	190	0.22
Run IIIc	- Jun 2012	8	250	190	1.37
Run IV	Jan 2013 - May 2013	8	250	230	3.37

Table B.1: Summary of beam configurations for T2K data taking runs up to May 2013.

New ND280 Monte Carlo productions are generated periodically to accompany newly acquired data runs and reanalyse older data runs using updated inputs. The main goals of Production 5 were to analyse Run I, II and III data with major reconstruction improvements. Details of the three 5C Monte Carlo sets used in

Chapters 7, 8 and 9 are given in Table B.2. In each case, generated interactions within the magnet volume are used.

ND280 Run	Beam spec	P0D contents	POT/file	POT Target
2	b	water	5×10^{17}	1.2×10^{21}
2	b	air	5×10^{17}	9×10^{20}
3	c	air	5×10^{17}	3×10^{21}

Table B.2: Descriptions of the ND280 Monte Carlo sets used for the analysis presented in Chapters 7, 8 and 9. Values from [105].

Beam specifications are provided in Table B.3. In each case, the offset is 50 ns and the bunch separation is equal to 582 ns.

Beam Spec	Beam Power (kW)	Repetition T (s)	POT/Spill ($\times10^{13}$)	Bunches /Spill	Bunch Duration (ns)
a	50	3.52	3.617	6	17
b	120	3.2	7.9891	8	19
c	178	2.56	9.463	8	19

Table B.3: Descriptions of the beam specifications. Values from [105].

Appendix C

NEUT interaction codes

Listed below are the NEUT reaction codes and their corresponding neutrino interactions. Equivalent anti-neutrino interactions are represented by negative values (given in full at [93]).

NEUT code	Description	Interaction
-	CHARGED CURRENT	-
1	Elastic	ν + N $\rightarrow$ l^- + P
11	Single π from Δ resonance	ν + P $\rightarrow$ l^- + P + π^+
12		ν + P $\rightarrow$ l^- + P + π^+
13		ν + P $\rightarrow$ l^- + P + π^+
16		ν + O(16) $\rightarrow$ l^- + O(16) + π^+
21	Multi π	ν + N/P $\rightarrow$ l^- + N/P + multi π
22	Single η from Δ resonance	ν + N $\rightarrow$ l^- + P + η^0
23	Single K from Δ resonance	ν + N $\rightarrow$ l^- + Λ + K^+
26	Deep Inelastic	ν + N/P $\rightarrow$ l^- + N/P + mesons
-	NEUTRAL CURRENT	-
31	Single π from Δ resonance	ν + N $\rightarrow$ ν + N + π^0
32		ν + P $\rightarrow$ ν + P + π^0
33		ν + N $\rightarrow$ ν + P + π^-
34		ν + P $\rightarrow$ ν + N + π^+
36		ν + O(16) $\rightarrow$ ν + O(16) + π^0
41	Multi π	ν + N/P $\rightarrow$ ν + N/P + multi π
42	Single η from Δ resonance	ν + N $\rightarrow$ ν + N + η^0
43		ν + P $\rightarrow$ ν + P + η^0
44	Single K from Δ resonance	ν + N $\rightarrow$ ν + Λ + K^0
45		ν + P $\rightarrow$ ν + Λ + K^+
46	Deep Inelastic	ν + N/P $\rightarrow$ ν + N/P + mesons
51	Elastic	ν + P $\rightarrow$ ν + P
52		ν + N $\rightarrow$ ν + N

Table C.1: NEUT interaction codes, as given at [93].

Bibliography

[1] W. Pauli. Dear Radioactive Ladies and Gentlemen. *Tubingen meeting*, 1930. http://microboone-docdb.fnal.gov/cgi-bin/RetrieveFile?docid=953;filename=pauli20letter1930.pdf.

[2] E. Fermi. Versuch einer Theorie der -Strahlen. *Zeitschrift fr Physik*, 88(3-4):161–177, 1934.

[3] C. L. Cowan Jr. F. Reines. Detection of the free neutrino. *Phys. Rev.*, 92(3):830–831, 1953. http://prola.aps.org/abstract/PR/v92/i3/p830_1.

[4] F.B.Harrison H.W.Kruse A.D.McGuire C. L. Cowan Jr., F. Reines. Detection of the free neutrino: a confirmation. *Science*, 124(3212):103–104, 1956.

[5] G. Danby et al. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9(1):36–44, 1962. http://link.aps.org/doi/10.1103/PhysRevLett.9.36.

[6] ALEPH Collaboration et al. Precision electroweak measurements on the Z resonance. *Phys. Rept.*, 427:257–454, 2006. http://arxiv.org/abs/hep-ex/0509008.

[7] Martin Perl. Evidence for, and properties of, the new charged heavy lepton. (SLAC-PUB-1923), 1977. http://slac.stanford.edu/pubs/slacpubs/1750/slac-pub-1923.pdf.

[8] The DONUT Collaboration. Observation of tau neutrino interactions. *Phys. Lett.*, pages 218–224, 2001. http://arxiv.org/abs/hep-ex/0012035.

[9] A.W. Sunyar M Goldhaber, L Grodzins. Helicity of neutrinos. *Phys. Rev.*, 109:1015–1017, 1958. http://link.aps.org/doi/10.1103/PhysRev.109.1015.

[10] Raymond Davis Jr. Attempt to detect the antineutrinos from a nuclear reactor by the $Cl^{37}(\bar{\nu}, e^-)A^{37}$ reaction. *Phys. Rev.*, 97:766–769, 1955. `http://link.aps.org/doi/10.1103/PhysRev.97.766`.

[11] Bruce T. Cleveland et al. Measurement of the solar neutrino flux with the Homestake chlorine detector. *The Astrophysical Journal*, 496(1):505, 1998. `http://iopscience.iop.org/0004-637X/496/1/505/`.

[12] GNO Collaboration. GNO solar neutrino observations: results for GNO I. *Phys. Lett. B490 (2000) 16-26*.

[13] SAGE Collaboration. Results from SAGE (the Russian-American Gallium Solar Neutrino Experiment). *Phys. Lett. B328, 234-248 (1994)*.

[14] P. Aliani et al. The solar neutrino puzzle: present situation and future scenarios. *Results and Perspectives in Particle Physics*, 27, 2002. `http://arxiv.org/abs/hep-ph/0206308`.

[15] Super-K Collaboration. Solar ^{8}B and hep neutrino measurements from 1258 days of Super-Kamiokande data. *Phys. Rev. Lett.*, 86(25):5651–5655, 2001. `http://link.aps.org/doi/10.1103/PhysRevLett.86.5651`.

[16] R. Becker-Szendy et al. Electron- and muon-neutrino content of the atmospheric flux. *Phys. Rev. D*, 46(9):3720–3724, 1992. `http://link.aps.org/doi/10.1103/PhysRevD.46.3720`.

[17] Takaaki Kajita. Atmospheric neutrinos. *New Journal of Physics*, 6, 2004. `http://iopscience.iop.org/1367-2630/6/1/194/`.

[18] B. Pontecorvo. Mesonium and anti-mesonium. *Sov. Phys. JETP*, 6:429, 1957.

[19] S. Sakata Z. Maki, M. Nakagawa. Remarks on the unified model of elementary particles. *Prog. Theor. Phys*, 28(5):870–880, 1962. `http://ptp.oxfordjournals.org/content/28/5/870`.

[20] M Randhawa G. Ahuja, M Gupta. Constructing the CKM and PMNS matrices from mixing data. *PU Research Journal*, 57, 2007. `http://arxiv.org/abs/hep-ph/0611324`.

[21] Super-K Collaboration. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998. `http://link.aps.org/doi/10.1103/PhysRevLett.81.1562`.

[22] D.G.Michael et al. Observation of muon neutrino disappearance with the MINOS detectors in the NuMI neutrino beam. *Phys.Rev.Lett.*, 97, 2006. hep-ex/0607088.

[23] The K2K Collaboration. Evidence for muon neutrino oscillation in an accelerator-based experiment. *Phys.Rev.Lett.*, 94, 2005. hep-ex/0411038.

[24] The Super-Kamiokande Collaboration. Atmospheric neutrino oscillation analysis with sub-leading effects in Super-Kamiokande I, II and III. *Phys.Rev.D*, 81, 2010. `http://arxiv.org/abs/1002.3471`.

[25] MINOS Collaboration. Measurement of the neutrino mass splitting and flavour mixing by MINOS. *Phys. Rev. Lett.*, 106(18), 2011. `http://link.aps.org/doi/10.1103/PhysRevLett.106.181801`.

[26] T2K collaboration. First muon-neutrino disappearance study with an off-axis beam. *Phys.Rev.D*, 85, 2012. `http://www.t2k.org/docs/pub/004`.

[27] SNO Collaboration. Electron energy spectra, fluxes, and day-night asymmetries of 8B solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory. *Phys. Rev. C.*, 72(5), 2005. `p://link.aps.org/doi/10.1103/PhysRevC.72.055502`.

[28] The KamLAND Collaboration. Precision measurement of neutrino oscillation parameters with KamLAND. *Phys. Rev. Lett.*, 100, 2008. `http://arxiv.org/abs/0801.4589v3`.

[29] L Wolfenstein. Neutrino oscillations in matter. *Phys. Rev. D*, 17(9):2369–2374, 1978. `http://link.aps.org/doi/10.1103/PhysRevD.17.2369`.

[30] A. Yu. Smirnov S.P.Mikheev. Resonance amplification of oscillations in matter and spectroscopy of solar neutrinos. *Sov.J.Nucl.Phys*, 42:913–917, 1985.

[31] H.M. O'Keeffe. Low energy background in the NCD phase of the Sudbury Neutrino Observatory. pages 12–14, 2008.

[32] M. Apollonio et al. (CHOOZ). Search for neutrino oscillations on a long base-line at the CHOOZ nuclear power station. *Eur.Phys.J.*, C27:331–374, 2003. `http://arxiv.org/abs/hep-ex/0301017`.

[33] Thierry Lasserre Maury Goodman. Double Chooz, a search for the neutrino mixing angle theta-13. 2006. `http://arxiv.org/abs/hep-ex/0606025`.

[34] T2K collaboration. Indication of electron neutrino appearance from an accelerator-produced off-axis muon neutrino beam. *Phys.Rev.Lett.*, 107, 2011. `http://www.t2k.org/docs/pub/003`.

[35] The T2K collaboration. Evidence of electron neutrino appearance in a muon neutrino beam. *Phys.Rev.D*, 88, 2013. `http://link.aps.org/doi/10.1103/PhysRevD.88.032002`.

[36] The Daya Bay Collaboration. Observation of electron-antineutrino disappearance at Daya Bay. *Phys.Rev.Lett.*, 108(17), 2012. http://link.aps.org/doi/10.1103/PhysRevLett.108.171803.

[37] The Daya Bay Collaboration. Improved measurement of electron antineutrino disappearance at Daya Bay. *Chinese Phys. C*, 37(1), 2013. http://iopscience.iop.org/1674-1137/37/1/011001/.

[38] The RENO Collaboration. Observation of reactor electron antineutrino disappearance in the RENO experiment. *Phys. Rev. Lett.*, 108, 2012. http://arxiv.org/abs/1204.0626.

[39] J. Beringer et al. (Particle Data Group). 2013 Review of Particle Physics. *Phys.Rev. D*, 86, 2012. http://pdg.lbl.gov/2013/tables/rpp2013-sum-leptons.pdf.

[40] P. Guzowski. Reconstruction of neutrino induced neutral current neutral pion events with the T2K ND280 tracker and ECAL. 2011. http://www.t2k.org/docs/thesis/013.

[41] A.D. Sakharov. Violation of CP invariance, c asymmetry and baryon asymmetry of the universe. *Sov.Phys.Esp.*, 34, 1991. http://iopscience.iop.org/0038-5670/34/5/A08/.

[42] The NOvA Collaboration. NOvA proposal to build a 30 kiloton off-axis detector to study neutrino oscillations in the Fermilab NuMI beamline. 2005. http://arxiv.org/abs/hep-ex/0503053.

[43] F. P. An et al. Observation of electron-antineutrino disappearance at Daya Bay. *Phys. Rev. Lett.*, 108, 2012. http://link.aps.org/doi/10.1103/PhysRevLett.108.171803.

[44] J. Dobson. Neutrino induced charged current π^+ production at the T2K near detector. page 22, 2012. http://www.t2k.org/docs/thesis/019.

[45] S. Hannestad. Cosmological limit on the neutrino mass. *Phys.Rev.D*, 66, 2002. http://link.aps.org/doi/10.1103/PhysRevD.66.125011.

[46] S. Pastor J. Lesgourgues. Neutrino mass from cosmology. *Adv. High Energy Phys.*, 2012. http://arxiv.org/abs/1212.6154.

[47] J.L.Rosner W.D. Arnett. Neutrino mass limits from SN1987A. *Phys.Rev.Lett.*, 58:1906–1909, 1987. http://link.aps.org/doi/10.1103/PhysRevLett.58.1906.

[48] S. Mertens C. Weinheimer G. Drexlin, V. Hannen. Current direct neutrino mass experiments. *Adv. High Energy Phys.*, (293986), 2013. http://www.hindawi.com/journals/ahep/2013/293986/.

[49] E. Majorana. Teoria simmetrica dell'elettrone e del positrone. *Nuovo Cimento*, 14, 1937. 171.

[50] The T2K Collaboration. Beyond the Standard Model. `http://t2k-experiment.org/neutrinos/beyond-the-standard-model/`.

[51] H.Ejiri et al. Spectroscopy of double-beta and inverse-beta decays from ^{100}Mo for neutrinos. *Phys.Rev.Lett*, 85, 2000. `http://ewiserver.npl.washington.edu/moon/moon-prl.pdf`.

[52] J.Hartnell. Neutrinoless double beta decay with SNO+. (TAUP2011), 2012. `http://arxiv.org/abs/1201.6169`.

[53] L.M. Widrow S. Dodelson. Sterile neutrinos as dark matter. *Phys.Rev.Lett*, 72:17–20, 1994. `http://arxiv.org/abs/hep--ph/9303287`.

[54] M.H. Shaevitz J.M. Conrad, W.C. Louis. The LSND and MiniBooNE oscillation searches at high δm^2. *Annu.Rev.Nucl.Part.Sci.*, 63, 2013. `http://arxiv.org/abs/1306.6494`.

[55] The MiniBooNE Collaboration. Event excess in the MiniBooNE search for $\bar{\nu_\mu} \rightarrow \bar{\nu_e}$ oscillations. *Phys.Rev.Lett.*, 105, 2010. `http://arxiv.org/abs/1007.1150`.

[56] M. Antonello et al. Experimental search for the LSND anomaly with the ICARUS detector in the CNGS neutrino beam. 2012. `http://arxiv.org/abs/1209.0122`.

[57] The T2K collaboration. Tokai-to-Kamioka (T2K) long baseline neutrino oscillation experiment proposal. page 14, 2006. `http://j-parc.jp/researcher/Hadron/en/pac_0606/pdf/p11-Nishikawa.pdf`.

[58] M. Apollonio et al. Search for neutrino oscillations on a long base-line at the CHOOZ nuclear power station. 2003. arXiv:hep-ex/0301017.

[59] Y. Fukuda et al. The Super-Kamiokande detector. *Nucl.Instrum.Meth.*, A501:418–462, 2003. `http://inspirehep.net/record/607144?ln=en`.

[60] The T2K Collaboration. Official documents. *Photographs and figures*. `http://www.t2k.org/docs/photos/general`.

[61] Shashank Shalgar. Neutrino oscillation toolkit. `http://sourceforge.net/projects/neutrinotoolkit/`.

[62] I. Sugai et al. Development of thick hybrid-type carbon stripper foils with high durability at 1800K for RCS of J-PARC. *Nucl. Instrum. Meth. A*, 561(1):16–23, 2006. `http://www.sciencedirect.com/science/article/pii/S0168900205026331`.

[63] The T2K Collaboration. Official documents, photographs and figures, neutrino beamline. http://www.t2k.org/docs/photos/beamline/other/Nu-Prod-Schematic/view.

[64] S. Bhadra et al. *Nucl. Instrum. Meth. A*, 703:45–58, 2013. http://arxiv.org/abs/1201.1922.

[65] The T2K collaboration. The T2K experiment. *Nucl. Instrum. Meth. A*, 659(1):106–135, 2011. http://www.sciencedirect.com/science/article/pii/S0168900211011910.

[66] Y. Oyama for T2K Collaboration. Current status of the T2K experiment. *J. Phys.:Conf. Ser.*, 347(1), 2011. http://iopscience.iop.org/1742-6596/347/1/012005/.

[67] The T2K Collaboration. Measurements of the T2K neutrino beam properties using the INGRID on-axis near detector. *Mucl. Instrum. Meth. A*, 694:211–223, 2012. http://www.sciencedirect.com/science/article/pii/S0168900212002987.

[68] Haigh M. Beam extrapolation and photosensor testing for the T2K experiment. *PhD Thesis*, 2010. http://www.t2k.org/docs/thesis/006.

[69] E Lorenz D Renker. Advances in solid state photon detectors. *JINST*, 4, 2009. http://iopscience.iop.org/1748-0221/4/04/P04004.

[70] M. Yokoyama et al. Development of Multi-Pixel Photon Counters. *Conference proceedings*, 2006. http://inspirehep.net/record/717989?ln=en.

[71] T2K Collaboration Tatsuya Kikawa. Development of the new T2K on-axis neutrino detector INGRID proton module. http://www.t2k.org/docs/poster/003, 2010.

[72] S. Assylbekov et al. The T2K ND280 off-axis pi-zero detector. *Nucl. Instrum. Meth. A*, 686:48–63, 2012. http://arxiv.org/abs/1111.5030.

[73] The T2K Collaboration. Time projection chambers for the T2K near detectors. *Nucl. Instrum. Meth. A*, 637, 2011. http://arxiv.org/abs/1012.0865.

[74] P. Baron P. Bene V. Berardi et al. N. Abgrall, B. Andrieu. Time Projection Chambers for the T2K near detectors. *T2K-TN-005*, 2010. http://arxiv.org/abs/arXiv:1012.0865.

[75] P. A. Amaudruz. The T2K fine-grained detectors. *Nucl. Instrum. Meth. A*, 696:1–31, 2012. http://arxiv.org/abs/1204.3666.

[76] D. Allan et al. The electromagnetic calorimeter for the T2K near detector ND280. *JINST*, 8, 2013. `http://iopscience.iop.org/1748-0221/8/10/P10019/`.

[77] S. Aoki et al. The T2K Side Muon Range Detector (SMRD). *Nucl. Instrum. Meth. A*, 698:135–146, 2013. `http://arxiv.org/abs/1206.3553`.

[78] Y. Itow et al. The JHF-Kamioka neutrino project. (C01-12-05.2):239–248, 2001. `http://arxiv.org/abs/hep-ex/0106019`.

[79] J. Albert. Indication of electron neutrino appearance in an accelerator-produced muon neutrino beam. *Thesis*, 2012. `http://www.t2k.org/docs/thesis/016`.

[80] S. Mine. New POLfit. *T2K-TN-005*. `http://www.t2k.org/docs/technotes/005`.

[81] D. Scully. Neutrino induced coherent pion production. 2013. `http://www.t2k.org/docs/thesis/038`.

[82] C. Giganti. Particle identification in the T2K TPCs and study of the electron neutrino component in the T2K neutrino beam. 2010. T2K Thesis.

[83] T2K Collaboration. The T2K neutrino flux prediction. *Phys.Rev.D*, 87(1), 2013. `http://prd.aps.org/abstract/PRD/v87/i1/e012001`.

[84] K. Ueno. T2KSK MC and reconstruction. `http://www.t2k.org/meet/collab/archive/t2k201008/talks/satam/t2kskmc`.

[85] N. Abgrall et al. Neutrino flux prediction. 2011. T2K internal note 38.

[86] A. Fasso et al. The physics models of FLUKA: status and recent developments. *CHEP 03, La Jolia, Cal.*, 2003. `http://arxiv.org/abs/hep-ph/0306267`.

[87] NA61/SHINE Collaboration. Pion emission from the T2K replica target: method, results and application. *Nucl. Instrum. Meth. A*, 701:99–114, 2012. `http://inspirehep.net/record/1121706`.

[88] Y. Hayato. A neutrino interaction simulation program library NEUT. *Acta Phys.Polon. B*, 40:2477–2489, 2009. `http://inspirehep.net/record/844435?ln=en`.

[89] R. Terri. Queen Mary, University of London, r.terri@qmul.ac.uk.

[90] T. Dealtry et al. T2K 6.5700×10^{20} POT 3-flavour muon-neutrino disappearance analysis (VaLOR). *T2K Tech Note*, (183). `http://www.t2k.org/docs/technotes/183/valor2013run1234disapv1/view`.

[91] Particle Data Group. Lepton summary table. http://pdg.lbl.gov/2011/tables/rpp2011-sum-leptons.pdf, 2011.

[92] The T2K Collaboration. Measurement of neutrino oscillation parameters from muon neutrino disappearance with an off-axis beam. *Phys. Rev. Lett.*, 111, 2013. http://journals.aps.org/prl/abstract/10.1103/PhysRevLett.111.211803.

[93] NEUT reaction codes. http://nngroup.physics.sunysb.edu/~mcgrew/t2k/nd280mc/v4r14/dox/classND280NeutKinematicsGenerator.html.

[94] The T2K Collaboration. Observation of electron neutrino appearance in a muon neutrino beam. *Phys. Rev. Lett.*, 112, 2013. http://journals.aps.org/prl/abstract/10.1103/PhysRevLett.112.061802.

[95] S. Ives. Study of the kaon contribution to the T2K neutrino beam using neutrino interactions in the near detector. https://workspace.imperial.ac.uk/highenergyphysics/Public/theses/Ives.pdf.

[96] C. Bojechko et al. CCQE-like and CC-non-QE-like ν_μ event selections in the ND280 tracker using Run 1+2 data. Technical Report 093, 2012. http://www.t2k.org/docs/technotes/093/tn093v04/view.

[97] M. Zito C. Giganti. Particle identification with the T2K TPC. *T2K Technical Note*, (001).

[98] H.M. O'Keeffe. Department of Particle Physics, Lancaster University, h.okeeffe@lancaster.ac.uk.

[99] A. Dabrowska. Numu CC analysis in the Ecal and Ecal induced background in FGD CC analysis. 2013.

[100] K. Suzuki. http://www-he.scphys.kyoto-u.ac.jp/member/k.suzuki/temp/t2kseminar_beam_v3.pdf.

[101] ROOT Development Team. TMath Function Members. http://root.cern.ch/root/html/TMath.html#TMath:KolmogorovTest.

[102] C. Andreopoulos. RAL, STFC, costas.andreopoulos@stfc.ac.uk.

[103] PDG G. Cowan. Statistics. *Particle Data Group 2013*, (Chapter 37). http://pdg.lbl.gov/2013/reviews/rpp2013-rev-statistics.pdf.

[104] Kento Suzuki. Reduction of beam direction uncertainty by MUMON. http://www.t2k.org/meet/collab/archive/201109/talks/thuram/mumon_plenary_2011Sep_v3.pdf/view, Sep 2011.

[105] T2K collaboration. Production 5 summary. http://www.t2k.org/nd280/datacomp/production005/mcp/prod5mcSummary.