Study of the kaon contribution to the T2K neutrino beam using neutrino interactions in the Near Detector

Sarah Joanne Ives

High Energy Physics
Blackett Laboratory
Imperial College London

A thesis submitted to Imperial College London
for the degree of Doctor of Philosophy

October 2012

"To my wonderful husband, Wayne.
Thank you for your unwavering
support, enthusiasm and patience!"

Abstract

T2K is a long-baseline neutrino oscillation experiment. It uses an accelerator-produced neutrino beam, whereby a beam of protons impinges on a nuclear target, producing kaon and pion mesons that decay to neutrinos. The main neutrino detectors are situated at 2.5^0 off-axis from the centre of the beam. An accurate flux prediction for this off-axis beam is crucial to achieve the sensitivity required for the goals of T2K. External experiments reduce the major flux uncertainty (hadronic interactions in the target), but are inherently independent of the real and variable beamline conditions of T2K. Therefore, *in situ* measurements are required to validate the flux. This thesis uses data from the T2K near detector (ND280) to validate the flux prediction.

The normalisation of K^+-originating neutrinos at the ND280 is measured. The K^+ beam component is important since K^+ daughters dominate the high energy part of the ν_μ beam and contribute to the intrinsic ν_e contamination. As many aspects of the beam simulation affect this measurement, including the hadron production at the target and the off-axis angle, it is used to validate the entire system. The November 2010 to March 2011 data set is used, corresponding to 7.837×10^{19} protons on target. ν_μ charged-current interactions are selected (with 86.3% purity) using the ND280 tracker and binned according to the momentum and angle of the muon candidate. The Monte Carlo (MC) is fitted to the data to extract the normalisations of both K^+ and π^+ originating neutrinos, b_K and b_π respectively. The flux, cross-section and detector systematic errors are considered. The best fit point is at $b_K = 0.86$ and $b_\pi = 0.78$, consistent with the nominal MC at the 1σ level.

Additionally, results of the first time calibration of the ND280 detector, primarily of the ECal sub-detector, are presented.

Declaration

I certify that the work presented within this thesis is either my own work or has been appropriately referenced.

The original work of this thesis is presented in Chapters 5 and 6. The surrounding chapters summarise large amounts of work conducted by other people, which are presented to provide the context for the studies shown in Chapters 5 and 6. Details of each chapter and my contributions are briefly summarised below.

Chapter 1 presents an overview of neutrino oscillation physics, as well as a review of accelerator-produced neutrino beams and the difficulties of predicting the flux of such beams. This information is synthesised from a large number of published articles and textbooks and is referenced appropriately. Chapter 2 gives an overview of the experimental setup and hardware of the T2K experiment, and is largely drawn from papers published by the T2K collaboration. Chapter 3 then summarises the flux prediction of T2K, which is the output of a large amount of work by other collaborators within T2K. Chapter 4 presents the software framework used to analyse events at the T2K near detector (ND280), particularly focusing on the reconstruction algorithms. Again, this work is the product of many other people within T2K.

Chapter 5 presents the time calibration for events in the ND280. It mainly focuses on calibration of the ECal sub-detector, though some results are also shown for the P0D sub-detector. This is entirely my own work. It is worth noting that this work was conducted within the framework of the ND280 calibration working group, and as such benefited from the expertise and knowledge of people within that group. In addition, the developed methodologies were implemented within an existing ND280 calibration software framework which was developed by others.

Chapter 6 presents a measurement of the normalisation of K^+-originating neutrinos at the ND280 detector, which provides a validation of the T2K flux

prediction. This is entirely my own work. It is worth noting that this analysis depends crucially on the global reconstruction of events within the ND280 which, as described above, was developed by others. It also builds upon the outputs from the T2K working groups dedicated to evaluating the systematic uncertainties relevant for T2K analyses. These uncertainties include the flux and cross-section systematic errors, as well as the detector systematic errors for the ND280 detector. These contributions are appropriately referenced where possible.

Finally, Chapter 7 provides a conclusion for the thesis as a whole.

Sarah Joanne Ives

Acknowledgements

I would like to express my gratitude to my supervisor Dr. Yoshi Uchida for his help, guidance and encouragement throughout my PhD, especially towards the end when I was trying to finish my analysis. The detailed and often lengthy discussions about my analysis work were greatly appreciated, particularly since they sometimes took place at a moment's notice. I would also like to thank the entire T2K group at Imperial, in particular: Dr. Morgan Wascko for his useful comments and feedback at the weekly group meetings; Dr. Antonin Vacheret for his expertise and guidance during my work on the time calibration of the ECal; and Dr. Asher Kaboth for the many useful discussions in the office, and for the rigorous preparation questions for my viva!

I offer special thanks to all of my fellow T2K students at Imperial. I have thoroughly enjoyed working alongside them, and the office has been a fun (if sometimes distracting!) place to work. I would particularly like to thank Jim Dobson for the generous amount of time that he gave to discuss my analysis and to bounce around ideas. I also thank both Jim and Mark Scott for their significant help in technical and coding related issues. In addition, I wish to thank Sam Short for the many chats over coffee and in the office, which helped to keep me sane especially towards the end of my PhD!

I would also like to extend my thanks to members from the wider T2K collaboration, too many to mention here. I have learnt a great deal from working alongside such talented and dedicated physicists.

To those who I shared the experience of living in Tokai with, thank you for making it such a fun and memorable time. I have many fond memories, including (but certainly not limited to) our karaoke sessions, trips to Tokyo, Joyful Honda, the ski trip, the joys of Masago and of course "ton-katsu cheesu"!

Finally, I would like thank all of my friends and family for their encouragement and support, and for putting up with me, especially towards the end of my PhD. In particular, I would like to express my huge debt of gratitude to

my husband Wayne. Without his unwavering support and belief I would never have been able to complete my PhD. I thank him for always listening to me talk about my work, even when he had no idea what I was talking about! Wayne: this achievement is as much yours as it is mine.

Contents

List of Figures **15**

List of Tables **21**

1 Motivation **23**

1.1 Neutrino oscillation physics 23

1.1.1 Different sources of neutrinos 23

1.1.2 History of neutrino oscillation physics 24

1.1.3 Oscillation formalism 25

1.1.4 Measurements of the T2K experiment 29

1.1.5 Current knowledge 31

1.2 Accelerator-produced neutrino beams 35

1.2.1 Flux prediction of accelerator-produced neutrino beams 36

2 The T2K experiment **49**

2.1 Neutrino beam 50

2.1.1 Off-axis design 50

2.1.2 J-PARC accelerator 51

2.1.3 Neutrino beamline 52

2.1.4 Current status 54

2.2 Near detectors 55

2.2.1 Scintillator bar readout 58

2.2.2 INGRID on-axis detector 60

2.2.3 Off-axis detector 62

2.3 Far detector 71

3 Flux prediction for T2K **75**
3.1 Overview . . . 75
3.2 Tuning of hadronic interactions . . . 76
3.2.1 NA61/SHINE experiment . . . 77
3.2.2 Data sets used for hadronic interaction tuning . . . 78
3.2.3 Methodology for tuning hadronic interactions . . . 80
3.3 Uncertainties on the flux prediction . . . 87
3.3.1 Hadronic interaction uncertainties . . . 88
3.3.2 Proton beam and off-axis angle . . . 92
3.3.3 Target and horn alignment . . . 93
3.3.4 Current and magnetic field of horns . . . 94
3.3.5 Summary of the flux uncertainties . . . 94
3.4 Tuning of the flux prediction using ND280 data . . . 95
3.4.1 Simultaneous constraint of the flux and cross-section parameters using ND280 data for the T2K oscillation analysis . . . 96
3.4.2 Validation of the flux prediction via a measurement of the normalisation of K^+-originating neutrinos at ND280 . . . 97

4 ND280 offline software **101**
4.1 Overview . . . 101
4.2 Reconstruction algorithms . . . 103
4.2.1 TPC, FGD and tracker reconstruction . . . 103
4.2.2 Global reconstruction . . . 106

5 Time calibration of the ND280 **111**
5.1 Time offsets calibration for the DsECal . . . 112
5.1.1 Sources of time offsets . . . 112
5.1.2 Calibration methodology . . . 113
5.1.3 Method validation . . . 117
5.1.4 Stability of offsets . . . 120
5.1.5 Further developments . . . 124
5.2 Fibre timewalk calibration for the P0D and DsECal . . . 128
5.2.1 Methodology and results . . . 128
5.2.2 Further developments . . . 131
5.3 Time resolution of the DsECal . . . 132

6 Constraining the normalisation of K^+-originating neutrinos **137**
6.1 Importance of K^+ neutrino parents for the T2K flux prediction 137
6.2 Measurement 138
6.3 Strategy 138
6.4 Data sets 140
6.5 Modifications to the standard global reconstruction outputs 141
6.6 Event selection 142
6.6.1 Use of the global vertexing algorithm 144
6.6.2 Selection cuts 148
6.6.3 Performance of event selection 154
6.6.4 Final event sample 161
6.7 Fit methodology 165
6.7.1 Fit parameters 165
6.7.2 Likelihood definition 170
6.7.3 Detector systematic errors 174
6.7.4 Cross-section systematic errors 183
6.7.5 Flux systematic errors 194
6.8 Fit validation studies 195
6.8.1 Generation of toy MC experiments 195
6.8.2 Validation results 197
6.9 Results 208

7 Conclusion **215**

Bibliography **219**

List of Figures

1.1 Summary of experimental knowledge for the neutrino oscillation parameters 32

1.2 Schematic of the current knowledge of the neutrino squared-mass splittings 34

1.3 Kinematic distributions of particles produced by interactions of protons in the MiniBooNE target according to four different Monte Carlo hadronic interaction generators 37

1.4 Comparison of the generated ν_μ flux at the MiniBooNE detector with four different Monte Carlo generators used for hadron production in the target 38

1.5 Comparison of the ν_μ flux prediction for the K2K experiment with three different methods 42

1.6 Comparison of the observed and predicted energy distributions for ν_μ charged-current quasi-elastic events at the MiniBooNE detector 44

1.7 Comparison of the measured and predicted ν_μ charged-current energy spectra at the MINOS near detector 46

2.1 Baseline for the T2K experiment 49

2.2 ν_μ energy spectrum for different off-axis angles of the beam 51

2.3 Schematics of the T2K neutrino beamline 53

2.4 Beam operation for the period spanning the Run 1 and Run 2 physics runs 55

2.5 Relative position of the INGRID and ND280 detectors 56

2.6 An exploded view of the ND280 off-axis detector 57

2.7 Photographs of an MPPC, used to read out the ND280 scintillator based sub-detectors . . . 59

2.8 The INGRID detector viewed from upstream . . . 60

2.9 Stability of the neutrino beam centre as monitored by INGRID . . . 60

2.10 A schematic of the P0D detector . . . 62

2.11 Schematic showing the main aspects of the TPC design . . . 64

2.12 Components of an SMRD scintillator counter . . . 67

2.13 Time structure of noise signals in the MPPCs when read out by the Trip-t electronics . . . 69

2.14 Overview of the ND280 readout electronics . . . 70

2.15 Schematic of the Super-Kamiokande detector . . . 71

2.16 Example of reconstructed T2K events in Super-Kamiokande . . . 73

3.1 The layout of the NA61/SHINE experiment . . . 78

3.2 The phase space of pions and kaons contributing to the predicted neutrino flux at Super-Kamiokande, and the regions covered by NA61 measurements 79

3.3 Ratio of the T2K neutrino flux prediction with and without the hadronic interaction tuning applied . . . 84

3.4 The final T2K neutrino flux prediction at the ND280 detector . . . 85

3.5 The final T2K neutrino flux prediction at the ND280 detector, with the flux shown on a logarithmic scale . . . 86

3.6 Fractional uncertainties for the flux at ND280 . . . 96

3.7 The correlation and fractional covariance matrices between energy bins for neutrinos of different flavours and at the two off-axis detectors (ND280 and Super-Kamiokande) . . . 99

4.1 Schematic summarising the ND280 software . . . 102

4.2 Distribution of the energy loss as a function of momentum for particles traversing the TPC 105

4.3 Event display for a data event, demonstrating the capabilities of the global reconstruction 109

5.1 Simplified view of the ND280 readout electronics with the sources of time offsets 114

5.2 Time distribution for a single TFB 116

5.3 TFB number versus hit time before and after the TFB offsets are applied 118

5.4 Time distributions before and after the TFB offsets are applied 118

5.5 Plots of the TFB offsets 119

5.6 The level of alignment achieved by applying the TFB offsets to the data . 119

5.7 Alignment of the RMMs by applying the RMM offsets 120

5.8 Comparison of offsets obtained using reconstructed data and raw data . . 121

5.9 Stability of the TFB offsets over the Run 1 period 125

5.10 Variation in the RMM mean times and RMM offsets over the Run 1 period 126

5.11 Mean time relative to trigger on each RMM versus date, with RMM powercycles marked, for a different cosmic trigger 127

5.12 Fibre timewalk correction for both the P0D and ECal 131

5.13 The effect of applying the fibre timewalk correction to data 132

5.14 The effect of applying the fibre timewalk correction to data, with the mean of each charge bin in Figure 5.13 shown for easier comparison . . . 133

5.15 Plot showing the effect of the time calibration on the distribution of the time difference between bar ends 134

5.16 The time difference between bar ends as a function of the reconstructed hit charge, with the full time calibration applied 135

5.17 The time resolution as a function of charge 136

6.1 Difference between the reconstructed global vertex position and the true vertex position for both single track and multi-track vertices 146

6.2 Data-MC agreement for the difference between the global vertex position and the upstream end of the lepton track 147

6.3 The effect of the fiducial volume cuts for primary vertices in FGD1 . . . 150

6.4 The effect of the fiducial volume cuts for primary vertices in FGD2 . . . 151

6.5 Plot of vertex times 152

6.6 TPC pull distributions for all good quality TPC segments of the lepton candidate track 155

6.7 TPC particle identification likelihoods for the lepton candidate track . . 156

6.8 The number of events selected by the CC-inclusive ν_μ cuts as a function of the integrated POT 157

6.9 Purity and efficiency of the CC-inclusive ν_μ sample as a function of true kinematical quantities of the neutrino and lepton 162

6.10 The breakdown of signal and background categories in the final CC-inclusive ν_μ sample as a function of true kinematical quantities of the incoming neutrino and outgoing lepton 163

6.11 The same plots as in Figure 6.10 but with each bin normalised to 1.0 . . 164

6.12 Data-MC comparison for reconstructed quantities of the lepton candidate for the final CC-inclusive ν_μ sample, with the MC broken down by interaction type 166

6.13 Data-MC comparison for reconstructed quantities of the lepton candidate for the final CC-inclusive ν_μ sample, with the MC broken down by neutrino parent 167

6.14 p_μ - $\cos\theta_\mu$ bin occupancies for both data and MC, with the MC broken down by interaction type 168

6.15 p_μ - $\cos\theta_\mu$ bin occupancies for both data and MC, with the MC broken down by neutrino parent 169

6.16 Ratio of the number of selected events for the CC-inclusive ν_μ selection cuts versus the official selection cuts 176

6.17 Effect on the fitted parameters of varying the normalisation of Out of FGD events in the toy MC experiments 180

6.18 Effect on the fitted parameters of removing sand interactions from the toy MC experiments 181

6.19 Fractional covariance matrix for the TPC momentum distortion systematic 184

6.20 Covariance matrix for the parameters constrained by fits to single pion production MiniBooNE data 186

6.21 Fractional response of a given bin in the fit to a change in the Fermi momentum 192

6.22 Covariance matrices for several of the cross-section parameters (CC Coherent normalisation, NC Other normalisation, and W shape) 193

6.23 Flux covariance matrix for ν_μs at the ND280 195

6.24 Progressive binning used for the fit validation studies 198

6.25 Distributions of the best fit points in the b_K-b_π space for the five sets of toy MC experiments 200

6.26 Projections of the plots in Figure 6.25 onto the b_K axis 201

6.27 Projections of the plots in Figure 6.25 onto the b_π axis 202

6.28 Fractional response of a given bin in the fit to changes in the cross-section parameters 203

6.29 Distribution of $\chi^2_{\mathrm{P,min}}$ (the minimum of the χ^2_{P} surface in b_K-b_π space) for the 1500 toy experiments in toy MC set 0 204

6.30 Impact of different uncertainties on the fitted b_K distribution 206

6.31 Impact of different uncertainties on the fitted b_π distribution 207

6.32 The p_μ - $\cos\theta_\mu$ bin occupancies for the data, nominal MC and fitted MC . 208

6.33 The χ^2_{P} and b_{other} surfaces for the final fit to the data 210

6.34 Results of the fit to the data, showing the surface of $\Delta\chi^2_{\mathrm{P}}$ with the final best fit point and 1σ contour 211

List of Tables

1.1 Current best fit values and 3σ allowed ranges for the neutrino oscillation parameters 33

3.1 Decay modes of parent particles to daughter neutrinos that are simulated within JNUBEAM 76

3.2 Summary of the hadron production data used to tune the pion and kaon production in the T2K flux prediction 79

3.3 Summary of the data used to tune the total cross-sections for p, $\pi^{\pm}$ and $K^{\pm}$ interactions for the T2K flux prediction 80

5.1 TFB offsets for all TFBs on RMM 0 122

5.2 TFB offsets for all TFBs on RMM 1 123

6.1 Performance of the CC-inclusive ν_{μ} selection cuts as a function of the number of cuts applied 159

6.2 Summary of the ND280 detector systematic errors 178

6.3 Summary of the cross-section systematic errors 185

6.4 Parameters for the five sets of toy MC experiments 197

6.5 The widths of the fitted b_K and b_{π} distributions with different uncertainties included in the toy MC experiments 206

6.6 Details of the flux parameters for the official ND280 analysis 213

Chapter 1

Motivation

The first part of this chapter, Section 1.1, provides the context for the T2K experiment. It gives a brief review of the history of neutrino oscillation physics, the mathematical formalism for neutrino oscillations and the current knowledge of the oscillation parameters. T2K uses an accelerator-produced neutrino beam and the main analysis of this thesis provides a validation of the predicted flux of this beam using data from the T2K near detector. The second part of this chapter, Section 1.2, therefore focuses on the advantages of using an accelerator-produced beam compared to other neutrino sources, along with the importance and major difficulties of accurately predicting the flux of such a beam.

1.1 Neutrino oscillation physics

1.1.1 Different sources of neutrinos

It is useful to first remind the reader of the variety of sources from which neutrinos can arise. There are three known flavours of neutrino (ν_e, ν_μ and ν_τ) each with a corresponding antiparticle ($\bar{\nu}_e$, $\bar{\nu}_\mu$ and $\bar{\nu}_\tau$) [1]. These neutrinos occur from both natural and man-made sources. The sources important for the discussions in this chapter are summarised as follows:

- Solar neutrinos: ν_es are emitted in the *pp* cycle of the sun (the main reaction chain for the fusion of hydrogen to helium) [1].

- Atmospheric neutrinos: ν_e, ν_μ, $\bar{\nu}_e$ and $\bar{\nu}_\mu$ particles are produced in the hadronic showers induced by primary cosmic rays in the Earth's atmosphere [1].
- Reactor neutrinos: $\bar{\nu}_e$s are produced in the β decay of the neutron-rich fission fragments in the reactor core [2].
- Accelerator-produced neutrinos: in conventional neutrino beams, a proton beam impinges upon a nuclear target producing a beam of pion and kaon secondary mesons which in turn decay to ν_μ and ν_e neutrinos [3].[1]

As mentioned previously, T2K uses an accelerator-produced neutrino beam. It is a second generation experiment, building on the experience of the K2K and MINOS accelerator neutrino experiments, as will be discussed in Section 1.2.

1.1.2 History of neutrino oscillation physics

The standard model of particle physics assumes that neutrinos are massless. However, there is now overwhelming evidence that neutrinos oscillate from one flavour to another, which can only occur if neutrinos have non-zero mass. The discovery of neutrino oscillations was the first evidence for new physics beyond the standard model.

Neutrinos were first detected in 1956 in the experiment of Reines and Cowan through the inverse β decay ($\bar{\nu}_e p \rightarrow e^+ n$) of reactor antineutrinos [4]. Since then a host of experiments have been devised to detect and understand neutrinos using all of the different sources given in Section 1.1.1. A brief history of the most important discoveries in neutrino oscillation physics are given here, although more detailed histories can be found elsewhere (for example, [1] and [5]).

In 1968, Ray Davis and his collaborators on the Homestake experiment detected a lower flux of neutrinos from the sun than was predicted by solar models [6]. This was termed the "solar neutrino problem" and remained unresolved for nearly forty years. Before this was solved, observations consistent with neutrino oscillations were first discovered in the atmospheric neutrino sector in 1998 by the Super-Kamiokande experiment in Japan [7]. They reported a zenith angle dependent deficit of atmospheric muon neu-

[1] The proton beam must have an energy above a few GeV for significant numbers of hadrons to be produced at the target.

trinos, arising from cosmic rays interacting with the atmosphere, which was consistent with a two flavour $\nu_\mu \leftrightarrow \nu_\tau$ oscillation framework. This was the first experimental observation consistent with neutrino oscillations. In 2002, the SNO experiment in Canada provided the solution to the solar neutrino problem, providing direct evidence for solar neutrino flavour change [8]. SNO determined that, although the flux of *electron* neutrinos from the sun was less than expected, the *total* flux of neutrinos agreed with solar models, which indicated that some of the electron neutrinos had changed flavour. In 2005, the Japanese experiment KamLAND, using a reactor beam of $\bar{\nu}_e$, then showed the dependence of the $\bar{\nu}_e$ survival probability on the ratio L/E [9] (L being the distance from the source to the detector, or the baseline, and E being the energy of the neutrino). This was clear evidence for neutrino oscillations, as will be shown in Section 1.1.3. Also in 2005, the Japanese long-baseline experiment K2K observed a ν_μ deficit consistent with $\nu_\mu \leftrightarrow \nu_\tau$ oscillations using an accelerator-produced neutrino beam [10]. Following this in 2006, MINOS, another long-baseline accelerator experiment, presented results consistent with ν_μ disappearance via oscillations [11].

The combination of oscillation parameters that are probed by the different experiments depends on the energy and composition of the respective neutrino source, and the distance of the detector from the source. This will be explained in more detail in the following section.

1.1.3 Oscillation formalism

All of the neutrino oscillation observations detailed in Section 1.1.2 can be explained by a framework in which the weak eigenstates of the neutrinos ν_α ($\alpha = e, \mu, \tau$) are a superposition of the mass eigenstates ν_i ($i = 1, 2, 3$) [1], as follows:

$$\mid \nu_\alpha\rangle = \sum_i U^*_{\alpha i} \mid \nu_i\rangle. \tag{1.1}$$

The matrix $U_{\alpha i}$ is called the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing matrix and is analogous to the CKM matrix which mixes the weak and strong eigenstates of

quarks:

$$U_{\alpha i} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix}$$
$$= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{1.2}$$

where $c_{ij} \equiv \cos\theta_{ij}$, $s_{ij} \equiv \sin\theta_{ij}$, and where the angles θ_{ij} parametrise the three possible rotations between the neutrino states. Splitting the PMNS matrix into three separate matrices as above is convenient since the right hand matrix contains θ_{12}, which is the dominant angle for solar oscillations, and the left hand matrix contains θ_{23}, which is the dominant angle for atmospheric oscillations. The central matrix contains θ_{13} and δ, the CP violating phase.

Pure ν_α flavour eigenstates are emitted in weak interactions, which then propagate in time as ν_i mass eigenstates. Since the mass eigenstates ν_i have different masses they propagate at different frequencies, and therefore build up a relative phase difference. Therefore when their superposition is decomposed back into the flavour eigenstates ν_α an admixture of flavours will be present i.e. flavours which were not emitted in the weak interaction will be present. The neutrino has *oscillated*. This can be described mathematically as follows.

The massive neutrino states $\mid \nu_i\rangle$ are by definition eigenstates of the Hamiltonian with energy eigenvalues $E_i = \sqrt{p^2 + m_i^2}$ for a neutrino with momentum p. Assuming the plane wave solution to the time dependent Schrodinger equation, their evolution is given by:

$$\mid \nu_i(t)\rangle = e^{-iE_i t} \mid \nu_i\rangle. \tag{1.3}$$

Following from this and Equation 1.1, the time evolution of a pure ν_α flavour eigenstate is given by:

$$\mid \nu_\alpha(t)\rangle = \sum_i U^*_{\alpha i} e^{-iE_i t} \mid \nu_i\rangle. \tag{1.4}$$

Inverting Equation 1.1 to give $\mid \nu_i\rangle = \sum_\alpha U_{\alpha i} \mid \nu_\alpha\rangle$ and substituting into the previous equation gives:

$$\mid \nu_\alpha(t)\rangle = \sum_{\beta=e,\mu,\tau} \sum_i U^*_{\alpha i} e^{-iE_i t} U_{\beta i} \mid \nu_\beta\rangle. \tag{1.5}$$

This shows that a pure ν_α flavour eigenstate emitted at $t = 0$ in a weak interaction evolves into a superposition of all three flavour eigenstates for $t > 0$ (so long as the U mixing matrix is not diagonal); or in other words the neutrino oscillates. The probability of observing flavour state ν_β at time t is then given by:

$$P_{\nu_\alpha \to \nu_\beta} = |\langle \nu_\beta \mid \nu_\alpha(t)\rangle|^2 = \sum_{ij} U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j} e^{-i(E_i - E_j)t}. \tag{1.6}$$

Using the approximation that $p \gg m_i$ for all mass eigenstates i, the energy of each mass eigenstate can be approximated as $E_i = \sqrt{p^2 + m_i^2} \simeq p + \frac{m_i^2}{2p}$. This leads to:

$$(E_i - E_j)t = \frac{(m_i^2 - m_j^2)t}{2p} = \frac{\Delta m^2_{ij} L}{2E}, \tag{1.7}$$

in the super-relativistic approximation of the neutrino where the elapsed time t can be replaced with the distance travelled L and the momentum p can be replaced by the energy E. Finally, using the unitarity of the PMNS matrix, Equation 1.6 can be expressed in terms of its real and imaginary components to give:

$$\begin{aligned} P_{\nu_\alpha \to \nu_\beta} = \delta_{\alpha,\beta} &- 4 \sum_{i>j} \Re(U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j}) \sin^2\left(\frac{\Delta m^2_{ij} L}{4E}\right) \\ &+ 2 \sum_{i>j} \Im(U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j}) \sin\left(\frac{\Delta m^2_{ij} L}{2E}\right). \end{aligned} \tag{1.8}$$

It can be seen that the oscillation probability is a function of both the fundamental oscillation parameters and quantities that can vary (L and E). The oscillation parameters

are composed of two mass-squared splittings (the third one is not independent of the other two), Δm^2_{ij}, and the three rotation angles and δ CP violating phase of the PMNS matrix U (shown in Equation 1.2). The current knowledge of the oscillation parameters is discussed in Section 1.1.5. L and E, the distance travelled by the neutrino and the energy of the neutrino respectively, vary according to the neutrino source and, in the case of experiments using accelerator-produced beams, can be chosen.

In order to determine the dominant oscillation parameters for a given experiment, some very useful limiting cases can be employed [1]. From our current knowledge of the mass-squared splittings (as will be described in Section 1.1.5), we can assume the following relations: $|\Delta m^2_{21}| \ll |\Delta m^2_{31}| \simeq |\Delta m^2_{32}|$. For solar neutrino experiments the baseline is very long ($L \sim 10^8$ km) and the neutrino energy is very low ($E \sim 10$ MeV), meaning that $\Delta m^2_{31} L/2E \simeq \Delta m^2_{32} L/2E \gg 1$. Therefore, the oscillations due to Δm^2_{31} and Δm^2_{32} are very fast and lead to an averaged effect, whereas the oscillations due to Δm^2_{21} are slower and therefore observable. For atmospheric neutrino experiments, on the other hand, L is significantly shorter than for solar experiments and E is significantly larger ($\mathcal{O}(1$ GeV)). Therefore, in this case, $\Delta m^2_{21} L/2E \ll 1$; this means that the oscillations due to Δm^2_{21} are effectively frozen and one can consider the limit $\Delta m^2_{21} \to 0$. In addition, it is known that θ_{13} is small (as will be discussed in Section 1.1.5) and so the approximations $\sin\theta_{13} \sim 0$ and $\cos\theta_{13} \sim 1$ can be taken. These approximations and limits can be used when expanding Equation 1.8 to simplify the probability for a given oscillation. In this way, it is found that the dominant parameters for the solar ν_e oscillations are θ_{12} and Δm^2_{21}, and as such these parameters are often denoted $\theta_\odot$ and $\Delta m^2_\odot$ respectively. On the other hand, the important parameters for atmospheric neutrino oscillations, of which the dominant oscillation is $\nu_\mu \to \nu_\tau$, are θ_{23} and Δm^2_{32}; as such these parameters are often denoted θ_{atm} and Δm^2_{atm} respectively. For long-baseline experiments using accelerator-produced neutrino beams, including K2K, MINOS and T2K, L and E are chosen to probe the atmospheric parameters (θ_{23} and Δm^2_{32}) in order to make precision measurements of these parameters.

It should be noted that the oscillation probabilities derived in this section apply only to neutrinos propagating through a vacuum. When neutrinos travel through matter they experience an effective potential due to coherent scattering with the particles in the medium, altering the oscillation probabilities [12]. This leads to the MSW effect, in which a resonant enhancement of the oscillation probability is possible depending on the

density of the medium and the neutrino energy [13]. This effect is particularly important for solar ν_es, leading to a strong enhancement of their oscillation as they exit the sun. However, the effect is small for the T2K experiment and so the reader is referred to [1] for further discussion.

1.1.4 Measurements of the T2K experiment

The baseline of the T2K experiment and the peak neutrino energy are chosen in order to probe the atmospheric oscillation parameters, θ_{23} and Δm^2_{32}. As will be discussed in Section 1.2, accelerator neutrino beams have many advantages over a natural neutrino source. This enables T2K to both make precision measurements of the atmospheric parameters and also to conduct searches for the rare oscillation of $\nu_\mu \rightarrow \nu_e$.

T2K starts with an almost pure ν_μ beam (∼99% purity). Two main oscillation measurements are made. The formalisms for these measurements are given here; the results will be discussed in Section 1.1.5. The first is of ν_μ disappearance where the ν_μ oscillates into other flavours, mainly ν_τ. The second is of ν_e appearance where the ν_μ oscillates into ν_e. The oscillation probabilities can be obtained by expanding Equation 1.8. The $\left(\frac{\Delta m^2_{ij} L}{4E}\right)$ argument in Equation 1.8 can be converted from natural to experimental units leading to $\left(1.27\frac{\Delta m^2_{ij} L}{E}\right)$. The following probability is obtained for ν_μ disappearance:

$$P_{\nu_\mu \rightarrow \nu_\mu} \approx 1 - \sin^2(2\theta_{23}) \sin^2\left(1.27\frac{\Delta m^2_{32} L}{E}\right), \tag{1.9}$$

where $\Delta m^2_{ij} = m^2_i - m^2_j$ is the difference in the squares of the masses of the i'th and j'th mass eigenstates in eV^2, L is the distance from the source to the detector in km, and E is the energy of the neutrino in GeV. The full expression for ν_μ disappearance has been simplified to give Equation 1.9 using the approximations described in Section 1.1.3 (i.e., $\Delta m^2_{21} L/2E \ll 1$, $\sin\theta_{13} \sim 0$ and $\cos\theta_{13} \sim 1$). The probability for ν_e appearance is given

below:

$$\begin{aligned} P_{\nu_\mu \to \nu_e} &= P1 + P2 + P3 + P4 \\ P1 &= \sin^2(\theta_{23}) \sin^2(2\theta_{13}) \sin^2\left(1.27 \frac{\Delta m^2_{32} L}{E}\right) \\ P2 &= \cos^2(\theta_{23}) \sin^2(2\theta_{13}) \sin^2\left(1.27 \frac{\Delta m^2_{21} L}{E}\right) \\ P3 &= -J \sin(\delta) \sin\left(1.27 \frac{\Delta m^2_{32} L}{E}\right) \\ P4 &= J \cos(\delta) \cos\left(1.27 \frac{\Delta m^2_{32} L}{E}\right), \end{aligned} \tag{1.10}$$

where:

$$J = \cos(\theta_{13}) \sin(2\theta_{12}) \sin(2\theta_{13}) \sin(2\theta_{23}) \sin\left(1.27 \frac{\Delta m^2_{32} L}{E}\right) \sin\left(1.27 \frac{\Delta m^2_{21} L}{E}\right). \tag{1.11}$$

The symbols and units are the same as for Equation 1.9. No approximations are used in Equation 1.10 (unlike for Equation 1.9), which is done to show the dependence of the ν_e appearance probability on the δ CP violating phase. The $P2$ term is small enough to be safely neglected but the remaining three terms could be comparable depending on the size of δ. For ν_e appearance measurements, the θ_{13} result is usually given as a function of δ. Multiple experiments are needed to disentangle the values of θ_{13} and δ. It can be seen from Equation 1.10 that precision measurements of the θ_{23} and $|\Delta m^2_{32}|$ parameters, via the ν_μ disappearance measurement in Equation 1.9, are key to gaining further information on the values of θ_{13} and δ. It can also be seen that the sign of Δm^2_{32} impacts the ν_e appearance probability, via the sign of the $P3$ variable in Equation 1.10. Therefore, the hierarchy of neutrino masses, which is determined by the sign of Δm^2_{32} (as will be discussed further in Section 1.1.5), is another key ingredient to accurately determining the values of θ_{13} and δ.

The δ parameter can also be accessed via measurements of the CP asymmetry since the asymmetry is defined as:

$$A^{\mathrm{CP}}_{\alpha,\beta} = P_{\nu_\alpha \to \nu_\beta} - P_{\bar{\nu}_\alpha \to \bar{\nu}_\beta}, \tag{1.12}$$

which leads to [1]:

$$A^{\mathrm{CP}}_{e,\mu} = A^{\mathrm{CP}}_{\mu,\tau} = A^{\mathrm{CP}}_{\tau,e} = 4 s_{12} c_{12} s_{13} c^2_{13} s_{23} c_{23} \sin\delta \times \left[\sin\left(\frac{\Delta m^2_{21} L}{2E}\right) + \sin\left(\frac{\Delta m^2_{32} L}{2E}\right) + \sin\left(\frac{\Delta m^2_{31} L}{2E}\right) \right]. \quad (1.13)$$

As in Section 1.1.3, the oscillation phenomenology described here applies only to neutrinos propagating through a vacuum. Matter effects should be included in the full derivation of the oscillation probabilities, since the neutrinos travel through the earth on their way to the T2K far detector (Super-Kamiokande). However, as mentioned previously, these effects are small for T2K and so the reader is referred to [1] for further details.

1.1.5 Current knowledge

Figure 1.1 shows a summary of experimental knowledge for the oscillation parameters. The two different mass scales are clear, with the large atmospheric mass scale ($\Delta m^2_{32} \simeq \Delta m^2_{31} \sim 10^{-3}$ eV^2) at the top and the small solar mass scale ($\Delta m^2_{21} \sim 10^{-5}$ eV^2) further down. Super-Kamiokande, K2K and MINOS have constrained the values of θ_{23} and $|\Delta m^2_{32}|$. Solar experiments together with KamLAND have constrained the value of θ_{12} and the value with sign of Δm^2_{21}, where knowledge of the sign comes from the sensitivity of solar experiments to matter effects [1]. It should be noted that Figure 1.1 does not include the recent measurements of the θ_{13} parameter. The limit set by the CHOOZ reactor $\bar{\nu}_e$ disappearance experiment of $\sin^2 2\theta_{13} < 0.19$ (which equates to $\tan^2\theta \lesssim 0.05$) at $\Delta m^2_{32} = 2.0 \times 10^{-3}$ eV^2 [14, 15] is shown. However, recent measurements have now conclusively shown that θ_{13} is non-zero. In 2011, T2K was the first experiment to report a non-zero θ_{13}; six $\nu_\mu \rightarrow \nu_e$ candidate events were observed with an expectation of 1.5 ± 0.3 for $\theta_{13} = 0$, implying a non-zero θ_{13} at 2.5σ significance [16]. This has since been followed by more precise measurements of θ_{13} by three reactor $\bar{\nu}_e$ disappearance experiments: Double CHOOZ [17], Daya Bay [18] and RENO [19]. Daya Bay reports a 5.2σ significance for a non-zero θ_{13}. T2K has also recently reported an updated result including data taken up to May 2012 with 10 ν_e appearance candidate events which imply a non-zero θ_{13} at 3.2σ significance [20]. Table 1.1 shows the current best fit values and 3σ allowed ranges for all of the oscillation parameters.

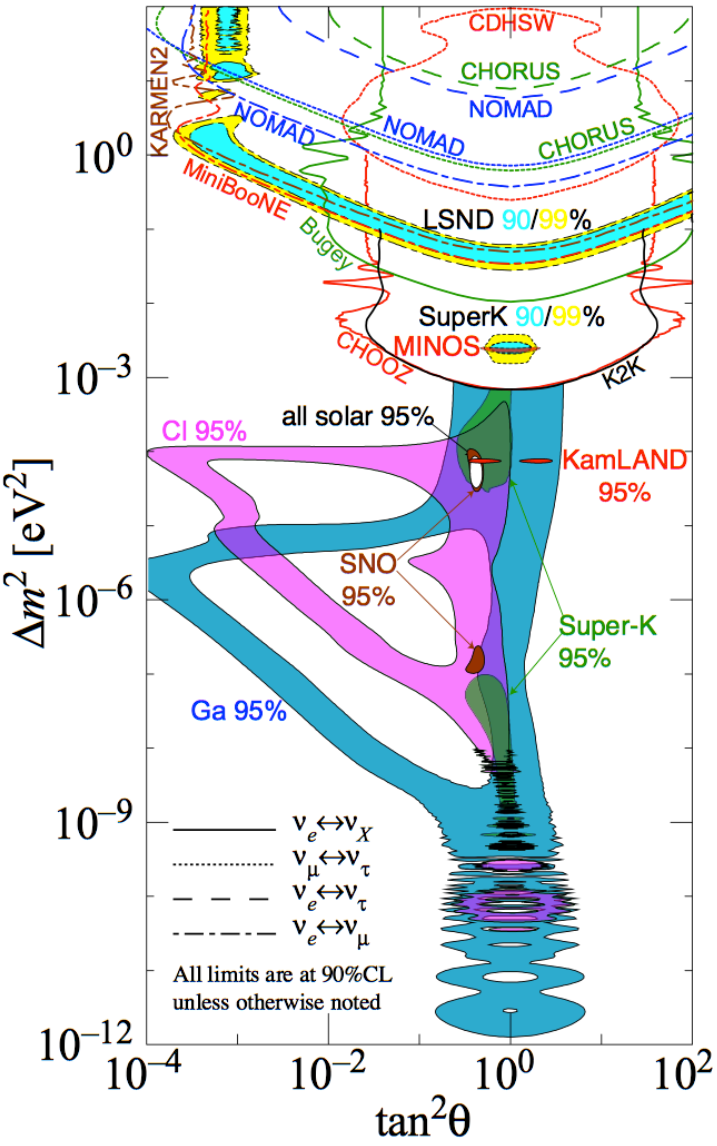


Figure 1.1: Summary of experimental knowledge for the Δm^2 and θ oscillation parameters taken from [5]. The recent T2K [16, 20], Double CHOOZ [17], Daya Bay [18] and RENO [19] results for θ_{13} are not included. Each result is based on the appropriate two-flavour neutrino oscillation approximation. Filled areas show allowed regions of the parameter space. Lines represent limits. The two different mass scales are clear, with the large atmospheric mass scale ($\Delta m^2_{32} \simeq \Delta m^2_{31} \sim 10^{-3}$ eV^2) at the top and the small solar mass scale ($\Delta m^2_{21} \sim 10^{-5}$ eV^2) further down. The benefits of using $\tan^2\theta$ as opposed to the often used $\sin^2 2\theta$ on the x-axis are discussed in [22]. In brief, this choice breaks the degeneracy between the $(0 < \theta < \pi/4)$ and $(\pi/4 < \theta < \pi/2)$ regions, which is important in showing the impact of matter effects. For pure vacuum oscillations (or experiments in which the impact of matter effects is small), the parameter contours are symmetric around $\tan^2\theta = 1$ when plotted on a log scale, where the $\tan^2\theta < 1$ and $\tan^2\theta > 1$ regions correspond to positive and negative Δm^2 respectively. However, for solar experiments in which matter effects are large and the sign of Δm^2 is known [1], this symmetry is broken.

Parameter	Best fit value ($\pm 1\sigma$)	3σ
$\Delta m^2_{21}[10^{-5}\text{eV}^2]$	$7.58^{+0.22}_{-0.26}$	6.99–8.18
$\lvert\Delta m^2_{32}\rvert(\simeq \lvert\Delta m^2_{31}\rvert)[10^{-3}\text{eV}^2]$	$2.35^{+0.12}_{-0.09}$	2.06–2.67
$\sin^2\theta_{12}$	$0.312^{+0.018}_{-0.015}$	0.265–0.364
$\sin^2\theta_{23}$	$0.42^{+0.08}_{-0.03}$	0.34–0.64
$\sin^2\theta_{13}$	$0.025^{+0.007}_{-0.008}$	0.005–0.050
$\sin^2\theta_{13}$	0.0251 ± 0.0034	0.015–0.036

Table 1.1: The current best fit values and 3σ allowed ranges for the neutrino oscillation parameters [5]. The first five rows (excluding title row) are from a global fit of all current neutrino oscillation data except that of Daya Bay and RENO. The bottom line therefore shows the average of the three recent reactor results from Daya Bay [18], RENO [19] and Double CHOOZ [17].

Figure 1.2 shows the current knowledge of the neutrino squared-mass splittings. The two possible mass hierarchies, which arise from the fact that only the magnitude and not the sign of Δm^2_{32} is known, are shown. The mass hierarchy can be probed by very long baseline neutrino experiments, such as NOvA, due to their sensitivity to matter effects. As mentioned previously in Section 1.1.4, the ν_e appearance measurement of T2K is impacted by the mass hierarchy via the sign of the $P3$ variable in Equation 1.10. Determination of the mass hierarchy is therefore key to an accurate determination of both θ_{13} and the δ CP violating phase.

T2K will soon publish new ν_μ disappearance results, using data taken up to June 2012, further constraining the values of θ_{23} and $|\Delta m^2_{32}|$. As mentioned in Section 1.1.4, these precision measurements of the atmospheric oscillation parameters are necessary for more precise measurements of θ_{13} and to set limits on the δ CP violating phase. At the time of writing, there is also discussion of running the T2K experiment in anti-neutrino mode in order to look for an asymmetry in the rate of neutrino versus anti-neutrino oscillations. As mentioned in Section 1.1.4, this can probe the δ parameter and search for CP violation in the neutrino sector.

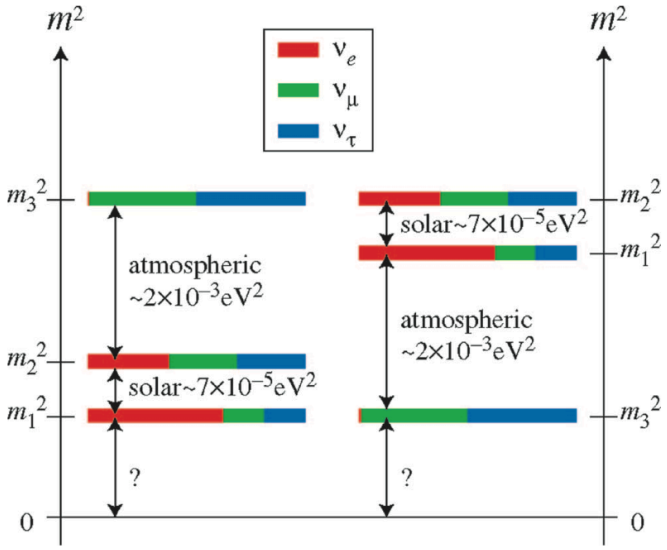


Figure 1.2: Schematic of the current knowledge of the neutrino squared-mass splittings. The value with sign of Δm^2_{21} is known, whereas only the magnitude of Δm^2_{32} is known, as described in Section 1.1.5. Therefore, both the normal (left) and inverted (right) mass hierarchies are allowed, corresponding to a positive or negative value of Δm^2_{32} respectively. The coloured bands represent the probability of finding the given weak eigenstate (ν_e, ν_μ or ν_τ) in each mass eigenstate. Figure taken from [21].

1.2 Accelerator-produced neutrino beams

Accelerator-produced neutrino beams are an important discovery tool in particle physics. A conventional neutrino beam is one that is created in the following way [3]. A proton beam is accelerated to high energy and directed at a nuclear target, thereby creating pion and kaon mesons. The charged mesons are focused by magnetic horns and then enter a decay volume in which they decay to neutrinos. Several such beams are currently in operation across the world at CERN, Fermilab and J-PARC, where the J-PARC accelerator serves the T2K experiment (as will be described in Chapter 2).

These beams have several advantages over natural neutrino sources. The peak neutrino energy (E) can be chosen because it is determined from the proton energy. In addition, the distance of the detector from the neutrino source (L), also called the baseline, can be chosen which is obviously not the case for natural neutrino sources. As explained in Section 1.1.3, by choosing the L and E values of an experiment the oscillation parameters to which that experiment is sensitive can be chosen. Furthermore, the neutrino flux of accelerator-produced beams is relatively well understood which is important for predictions of the expected events in the detectors. Finally, timing measurements of the beam can be used to greatly reduce uncorrelated backgrounds entering the detector from, for example, cosmic rays.

Long-baseline experiments using accelerator-produced neutrino beams are designed to probe the atmospheric oscillation parameters (θ_{23} and Δm^2_{23}) with baselines of several hundred kilometres and neutrino energies of a few GeV. K2K was the first such experiment followed by the MINOS experiment, as mentioned previously in Section 1.1.2. In contrast to the previous experiments, the main detectors of T2K (ND280 and Super-Kamiokande) are situated off-axis at 2.5^0. This results in a different neutrino energy spectrum at the detectors as compared to an on-axis setup, as will be described in Section 2.1.1, since the phase space of particles produced at the target whose decay neutrinos reach the detectors is different.

Accurately predicting the flux of these neutrino beams is a difficult task which has presented problems for several experiments prior to T2K. This is a particularly relevant issue for T2K since it is the first experiment to use an off-axis beam. An accurate flux prediction is crucial for T2K to achieve the sensitivity required for its physics goals.

The main analysis of this thesis, presented in Chapter 6, uses data from the T2K near detector to probe the beam system and provide a validation of the neutrino flux. The following section presents the main difficulties of an accurate flux prediction, namely the simulation of hadronic interactions in the target, along with attempts of previous experiments to constrain their flux prediction and a brief summary of the methods employed by T2K. Full details of the T2K flux prediction can be found in Chapter 3.

1.2.1 Flux prediction of accelerator-produced neutrino beams

The main uncertainty for the flux prediction of conventional neutrino beams arises from uncertainties of the hadronic interactions in the target material. The dominant uncertainty comes from the production yields of pion and kaon mesons in the interactions of the primary protons. There are also secondary effects from the reinteraction of any nucleons produced in the primary interaction and the absorption of the produced pion and kaons that can take place in thicker nuclear targets.

Figure 1.3 compares the kinematic distributions for K^+, π^+ and π^- particles exiting the target when four different Monte Carlo generators are used to model the primary interactions. These plots are for the MiniBooNE experiment (to be described later in this section), in which protons of 8.9 GeV/c impinge on a beryllium target with a length of 1.7 nuclear interaction lengths. The four generators used are: MARS [23], G4 LHEP [24, 25], G4 Bertini [26] and G4 Binary [27]. The differences between the predictions of the four generators are considerable, particularly for the kaon production. Figure 1.4 shows the resulting ν_μ flux at the MiniBooNE detector with these four different hadronic interaction generators. The discrepancies in the flux between the generators are striking, on order of $\sim$100% for some parts of the spectrum.

In view of the large discrepancies between different models of hadronic interactions, it is clearly not acceptable to simply choose one model over another, since this would result in very large flux uncertainties. Instead, it is necessary to constrain the hadronic interactions using data from external experiments. Dedicated hadron production experiments exist to make precision measurements of the yield of secondary hadrons produced in hadron-nucleus collisions. HARP [29] is one such experiment. It is a high statistics, large angular acceptance spectrometer which has made measurements of the secondary hadron yields in interactions of protons and charged pions on various nuclei (ranging

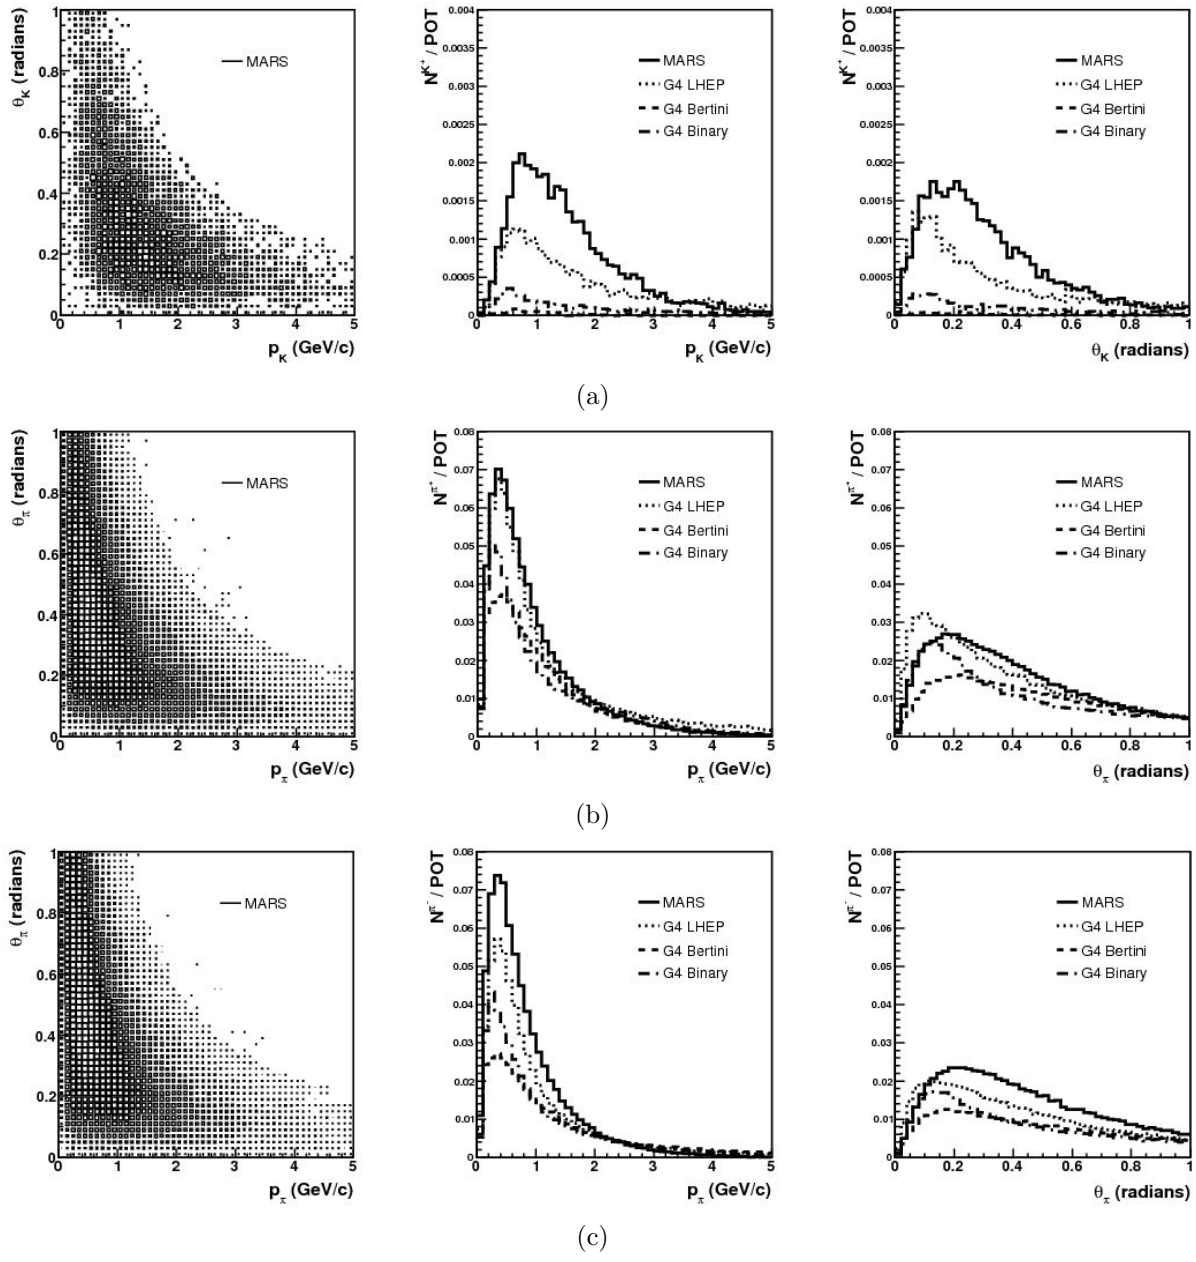


Figure 1.3: Kinematic distributions of particles produced by interactions of protons in the MiniBooNE target according to four different Monte Carlo hadronic interaction generators (see text). The following particles are shown: (a) K^+, (b) π^+, and (c) π^-. In the MiniBooNE experiment, protons of 8.9 GeV/c impinge on a beryllium target with a length of 1.7 nuclear interaction lengths. These plots are taken from [28].

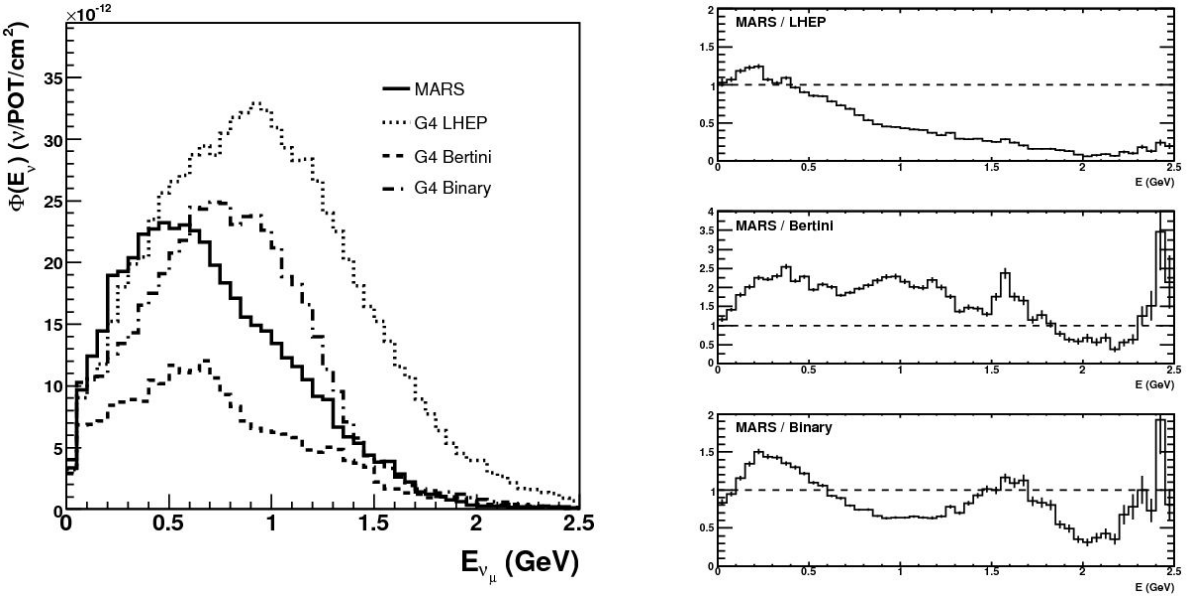


Figure 1.4: Comparison of the generated ν_μ flux (in neutrinos per proton on target per cm^2 of detector surface area) at the MiniBooNE detector with four different Monte Carlo generators (see text) used for the production of hadrons in primary interactions of 8.9 GeV/c protons on the beryllium target. The left panel shows the flux and the right panel shows the ratios of the different generators, both as a function of neutrino energy. This figure is taken from [28].

from hydrogen to lead) with various incident beam momenta. It has made dedicated measurements for both the K2K and MiniBooNE experiments, in each case using an identical incident proton beam momentum and replica target to the respective experiment (the flux predictions of both of these experiments will be discussed later in this section). NA61 is another hadron production experiment. It is a large acceptance hadron spectrometer, and will be discussed in more detail in Section 3.2.1. One of the main purposes of this experiment is to provide hadron production data to the T2K experiment to reduce the uncertainties on the T2K flux prediction, as will be discussed in Chapter 3.

The data from hadron production experiments provide a powerful tool in reducing the flux uncertainties of neutrino beam experiments from hadronic interactions in the target. However, these external measurements alone are not sufficient to constrain the flux of a neutrino experiment since important factors in the beamline of the given experiment are not accounted for. For example, some hadron production experiments use a thin target which does not take into account the reinteraction effects relevant for a thicker nuclear target of the type used in conventional neutrino beams. Also, degradation effects of the target over time will not be accounted for. In addition, the current passing through the magnetised horns of the neutrino experiment can vary with temperature, which changes the focusing of the secondary pion and kaon mesons and therefore affects the neutrino flux. Also, a component of the neutrino flux will come from interactions of the protons in material surrounding the target (for example, horns, windows and shielding) which are not measured by hadron production experiments. For this reason, *in situ* measurements are also a powerful tool in constraining the neutrino flux. A cautionary tale in this respect comes from the CERN 1967 neutrino experiment; the flux was calculated using hadron production measurements made at the IHEP accelerator, but the *in situ* measurements from a set of muon monitors in the neutrino beam suggested a factor of 2 discrepancy to the predicted flux [30]. Revised particle production measurements improved the agreement to within 15%, but this is instructive in demonstrating the importance of *in situ* measurements.

Neutrino beam experiments tend to use *in situ* measurements in combination with external hadron production measurements (if available) to constrain the flux prediction. *In situ* measurements can be made in a number of ways. The K2K experiment has a pion monitor installed after the magnetic horns to measure the momentum and direction of pions exiting the target before they enter the decay volume. The NuMI beam at

Fermilab [31], used for several neutrino experiments, has three muon monitors to measure the distribution of muons after the decay pipe. An empirical parametrisation for the meson production at the target is tuned to match the muon monitor data [32, 33].

In addition, for long-baseline experiments with both a near and far detector, measurements of the neutrino flux in the near detector can be used to improve the flux prediction at the far detector. For example, data from the near detector of the MINOS experiment taken with different configurations of the NuMI beam can be used to constrain the particle production at the target. Similarly, for the T2K experiment, data from the near detector is used to constrain the flux prediction. The flux predictions of both the MINOS and T2K experiments will be discussed later in this section.

Some details are now given for the flux predictions of some recent neutrino beam experiments to show the issues faced and the methods employed previously, particularly for constraining the hadronic interaction uncertainties in the respective neutrino beams. A summary of the methods used in the T2K flux prediction is then given for comparison to these previous experiments, with full details of the T2K flux prediction given in Chapter 3.

Flux prediction of the K2K experiment

The K2K experiment is a long-baseline neutrino oscillation experiment which uses an accelerator-produced ν_μ beam to probe the same Δm^2 and $\sin^2 2\theta$ region as that explored with atmospheric neutrinos [34]. The neutrino beam is produced in the conventional way, with 12 GeV protons impinging on an aluminium target. K2K consists of a near detector complex at KEK, approximately 300 m from the proton target, and the Super-Kamiokande far detector 250 km away, both on-axis with the neutrino beam. The near detector complex consists of a 1 kiloton water Cherenkov detector and a fine grained detector system. Super-Kamiokande is a 50 kiloton water Cherenkov detector that is also used for the T2K experiment and will be described in more detail in Section 2.3. A brief summary of the flux prediction of the K2K experiment is given here; further details should be sought in [34].

As for all experiments based on conventional neutrino beams, the flux uncertainty for K2K is dominated by the uncertainties in hadron production at the target. K2K

uses a parameterisation by J. R. Sanford and C. L. Wang [35,36] to simulate the hadron production in the target. The input parameters to this model are taken from external hadron production data. For π^+ production (the dominant parent meson for the neutrino flux), the measurements by the HARP experiment are used [37]. The HARP experiment uses the same proton beam energy and the same target material as K2K. Other data sets are used to provide the input parameters for the production of π^- and kaon mesons, as detailed in [34].

Two cross-checks of the π^+ production are performed. Firstly, the parameters obtained from the so-called "Cho-CERN compilation" are used in the Sanford-Wang model; the data used in this compilation comes mainly from the measurements of proton interactions on a beryllium target performed by Cho et al. [38]. Secondly, measurements from the pion monitor (PIMON) are used to constrain the momentum and direction of the pions exiting the target. The PIMON is a gaseous Cherenkov imaging detector and is situated just downstream of the horn magnets before the decay volume. It provides *in situ* measurements of the pions and is used to validate the flux predictions. Figure 1.5 shows the flux predictions at both the near and far detectors based on: *a*) the HARP π^+ production measurements; *b*) the Cho-CERN compilation of older (non-HARP) π^+ production data; and *c*) the PIMON measurements. The three flux predictions are consistent with each other within their uncertainties.

Flux prediction of the MiniBooNE experiment

The MiniBooNE experiment [39] was motivated by the result from the Liquid Scintillator Neutrino Detector (LSND) experiment [40] which presented evidence for $\bar{\nu}_\mu$ to $\bar{\nu}_e$ oscillations at the $\Delta m^2 \sim 1$ eV^2 scale, a scale significantly larger than those observed in both the solar and atmospheric neutrino sectors. MiniBooNE has made searches for both ν_e appearance and $\bar{\nu}_e$ appearance in ν_μ and $\bar{\nu}_\mu$ beams respectively at this Δm^2 scale, with the most recent results shown in [41]. MiniBooNE uses the Booster Neutrino Beamline (BNB) at Fermilab, a conventional neutrino beam produced by protons (8.89 GeV/c) impinging on a target (beryllium). The beam has an average energy of ∼800 MeV. The MiniBooNE detector is situated 541 m from the target, and is a spherical detector containing 800 metric tons of mineral oil. A brief summary of the flux prediction of the MiniBooNE experiment is given here; full details should be sought in [42].

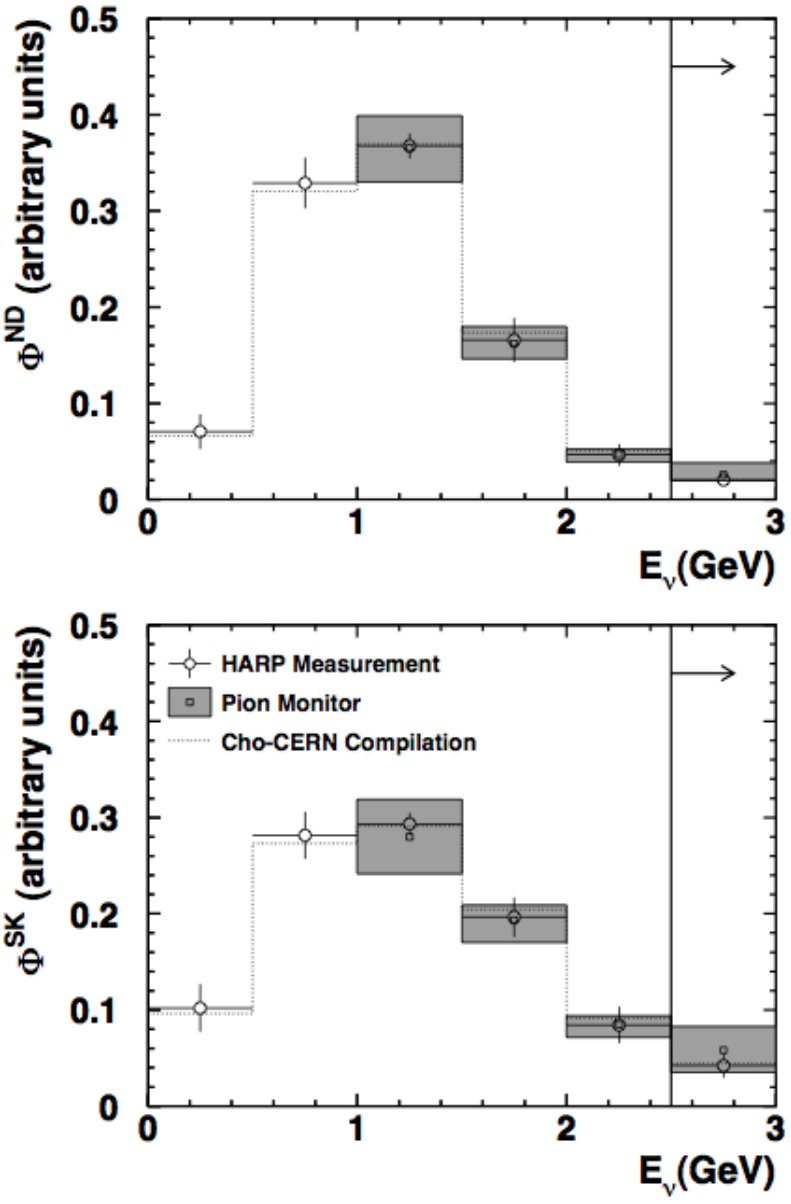


Figure 1.5: Relatively-normalised ν_μ flux predictions at the near (top) and far (bottom) detectors of the K2K experiment. The empty circles with error bars show the central values and shape-only errors based on the HARP π^+ production measurement; the empty squares with shaded error boxes show the central values and errors from the pion monitor (PIMON) measurement; and the dotted histograms show the central values from the Cho-CERN compilation of older (non-HARP) π^+ production data. The PIMON predictions are normalised such that the integrated fluxes above 1 GeV neutrino energy match those of HARP, at both the near and far detectors. The vertical line with arrow at 2.5 GeV in each plot is to illustrate that the final bin includes all events with energies above 2.5 GeV (not just those with energies in the range 2.5–3 GeV). Figure and caption (adapted) taken from [34].

As for the K2K experiment described above, and indeed all experiments based on conventional neutrino beams, the flux uncertainty for MiniBooNE is dominated by the uncertainties in hadron production at the target. External data sets are used to tune the modelling of these interactions. The $\pi^{\pm}$ production is tuned using measurements taken at HARP [43], taken at the same beam energy and with the same target material as for MiniBooNE, and measurements from the E910 experiment [44]. As for K2K, the production cross-sections are parametrised using the Sanford-Wang model. The input parameters to the model are determined from a χ^2 fit to the $\pi^{\pm}$ production data sets. For K^+ production, data from a set of other experiments is used since no data is available from the HARP or E910 experiments. These experiments provide measurements of K^+ production in proton interactions on beryllium at various proton beam momenta. The Feynman scaling hypothesis [45] (described in more detail in Section 3.2.3) is used to relate production measurements taken at different proton momenta. The production cross-section is parametrised using a model based on Feynman scaling, and a χ^2 fit is performed to the production data sets to obtain the model parameters. The production of other mesons (including K^- and K^0) is also considered, although these contributions to the flux are small and so are not detailed here.

In addition to the tuning of the production cross-sections, the hadronic cross-sections of nucleon and pion interactions on beryllium and aluminium (the material of the target and horns respectively) are also tuned, using cross-section measurements from a host of different experiments. The total cross-sections are tuned, as well as the inelastic and quasi-elastic components.

Figure 1.6 shows the observed and predicted energy spectra of ν_μ charged-current quasi-elastic events selected in the MiniBooNE data, taken from [42]. The data is observed to have a higher normalisation than the predicted spectrum by a factor of 1.21, and so the predicted spectrum is scaled by this factor in Figure 1.6 to enable comparison of the spectrum shapes. This discrepancy in normalisation could be due to errors in the flux calculation or neutrino interaction cross-sections or both, but it is useful in highlighting the difficulty of accurate flux predictions.

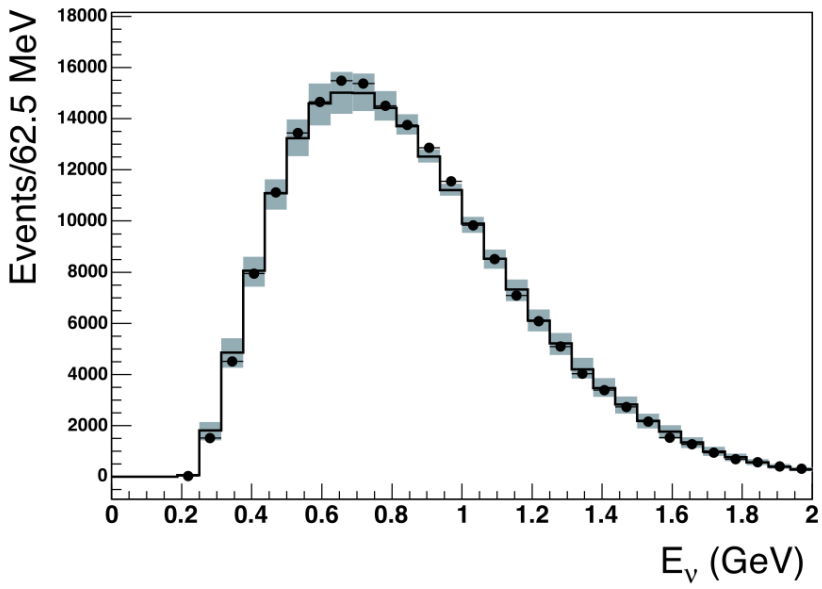


Figure 1.6: Comparison of the observed (points) and predicted (histogram) energy distributions for ν_μ charged-current quasi-elastic events selected in the MiniBooNE data. A normalisation factor of 1.21 has been applied to the predicted distribution as described in the text. The error bars on the predicted distribution are the estimated uncertainties in the shape of the spectrum once the normalisation has been fixed to match the data. Figure and caption taken from [42].

Flux prediction of the MINOS experiment

The Main Injector Neutrino Oscillation Search (MINOS) [46] is a long-baseline neutrino oscillation experiment that uses the NuMI neutrino beam at Fermilab. The NuMI beam is a conventional neutrino beam, produced by protons (120 GeV/c) impinging on a target (graphite). The beam has an average energy ranging between ∼3–8 GeV. MINOS is composed of two functionally identical detectors separated by 734 km (the near and far detectors) arranged on-axis with the beam. The detectors consist of alternating scintillator and steel planes in magnetised modules, with the scintillator planes read out by photomultipliers. The main aim of the MINOS experiment is to make precision measurements of the Δm^2 and $\sin^2 2\theta$ mixing parameters governing ν_μ disappearance at the atmospheric neutrino mass-scale, with the latest results shown in [47]. A brief summary of the flux prediction of the MINOS experiment is given here; further details should be sought in [48].

Figure 1.7 compares the measured and predicted ν_μ charged-current energy spectra at the MINOS near detector with six different configurations of the NuMI beam (with different target positions and horn currents). The points show the data and the thin line shows the FLUKA [49] Monte Carlo prediction. There are very large discrepancies of up to ∼40%. Since the magnitude and energy range of the discrepancies depend on the beam configuration, this suggests that the major source of discrepancy is from the calculation of the neutrino flux (as opposed to uncertainties in the neutrino cross-sections or detector modelling). The main uncertainty in the neutrino flux prediction arises from insufficient knowledge of hadronic production in the target. The different configurations of the NuMI beam enable the particle production to be mapped out across the hadron phase space. The six ν_μ charged-current energy spectra shown in Figure 1.7 are used to constrain the particle production. The differential pion production is represented as a parametric function which is used to tune the predicted spectra to the data in a χ^2 fit.[2] Terms are included in the fit to describe the uncertainties on beam focusing and detector systematic errors. The predicted energy spectra after the tuning of the pion production are also shown in Figure 1.7 (thick line) for each of the beam configurations; the agreement with the data is significantly improved after the tuning.

[2] K^+ production does not contribute significantly to the ν_μ flux below approximately 30 GeV.

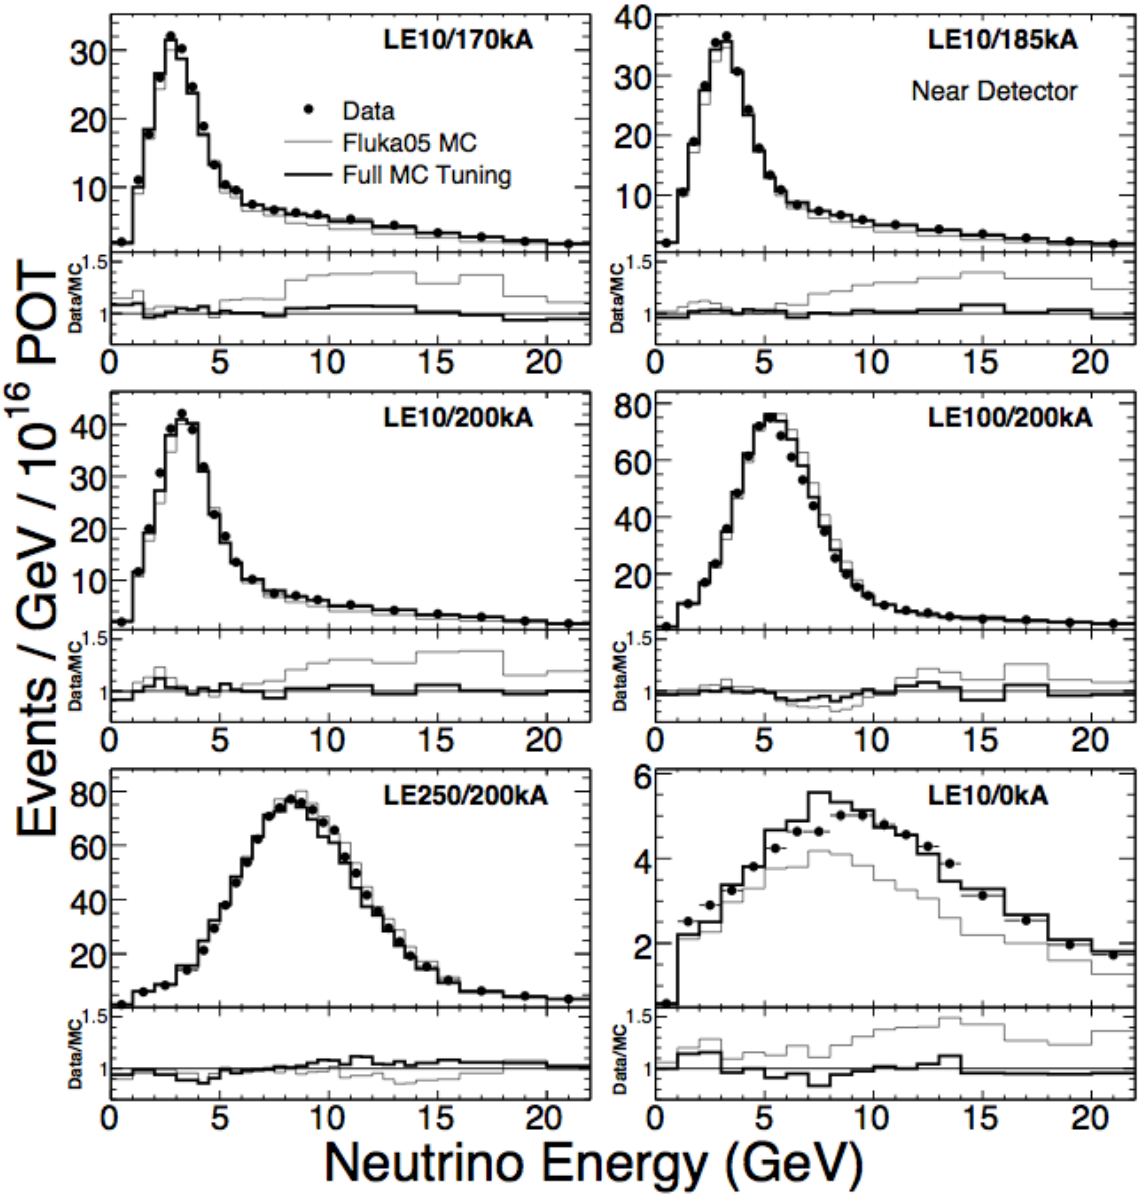


Figure 1.7: Comparison of the measured and predicted ν_μ charged-current energy spectra at the MINOS near detector for six NuMI neutrino beam configurations. Two Monte Carlo predictions are shown: one (thin line) with the *ab initio* calculation based on the FLUKA generator [49], the other (thick line) after constraining the hadron production with the procedure described in the text. The panels along the bottom of each figure show the ratio of the measured and predicted spectra. This figure is taken from [48].

Flux prediction of the T2K experiment

Full details of the T2K flux prediction are given in Chapter 3, but a summary is given here for comparison to the K2K, MiniBooNE and MINOS experiments detailed above. As for all experiments based on a conventional neutrino beam, the main source of uncertainty is the hadronic interactions in the target. The production of pion and kaon mesons is tuned using external data mainly from the NA61 hadron production experiment (which will be described in Section 3.2.1). In addition, the total cross-sections for proton, pion and kaon interactions on both carbon and aluminium (for the target and horns respectively) are tuned using external data from a large set of experiments. A variety of *in situ* monitoring devices are used to constrain the uncertainties on the flux prediction, including the primary beamline monitors and the INGRID on-axis neutrino detector. Measurements from the muon monitor are also used as a cross-check of the on-axis neutrino beam intensity and direction. All of these monitors and detectors will be detailed in Chapter 2.

In addition, measurements from the off-axis ND280 near detector can be used to further constrain the flux prediction at the off-axis Super-Kamiokande far detector, as will be discussed further in Section 3.4. For the most recent T2K oscillation results [20], measurements from the ND280 are used to simultaneously constrain the flux prediction and neutrino interaction cross-section models to improve the constraint of the oscillation parameters at Super-Kamiokande. The present measurement also uses ND280 data to validate the T2K flux prediction. A measurement is made of the normalisation of K^+-originating neutrinos at the ND280 detector (as will be described in detail in Chapter 6) which is affected by many aspects of the beam simulation, including particle production at the target and other factors such as the horn currents and the off-axis angle of the ND280. It therefore provides a probe of the beam system and a validation of the T2K flux prediction.

Chapter 2

The T2K experiment

The T2K experiment is a long-baseline neutrino oscillation experiment that uses an accelerator-produced neutrino beam. As described in Section 1.1.4, two main oscillation measurements are made. The first is of ν_μ disappearance where the ν_μ oscillates into other flavours, mainly ν_τ, and the second is of ν_e appearance where the ν_μ oscillates into ν_e. The ν_μ disappearance analysis enables precision measurements of the atmospheric oscillation parameters, θ_{23} and Δm^2_{32}. In turn, these parameters are key to making more precise measurements of θ_{13} and to setting limits on the δ CP violating phase via the ν_e appearance measurement.

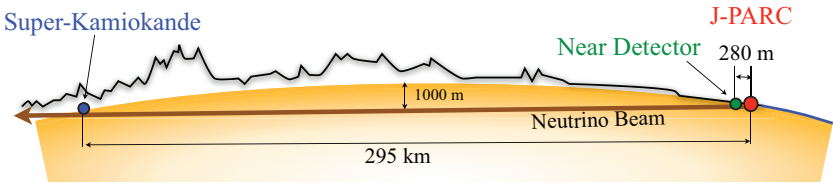


Figure 2.1: Schematic showing a neutrino's journey from the neutrino beamline at J-PARC, through the suite of near detectors situated 280 m downstream, and then underneath Japan, travelling a total distance of 295 km before entering the Super-Kamiokande detector [50].

The layout of the T2K experiment is shown in Figure 2.1. The accelerator at the J-PARC facility, on the east coast of Japan, provides a high power beam of protons to the

neutrino beamline which outputs an intense ν_μ beam (with a purity of ∼99%). At the accelerator design power of 750 kW (currently ∼200 kW) this will be the most intense neutrino beam ever built. The neutrino beam passes through the suite of near detectors, situated 280 m downstream of the beamline, and then underneath the main island of Japan, travelling a total distance of 295 km before entering the Super-Kamiokande far detector on the west coast of Japan.

T2K uses an off-axis method, whereby the neutrino beam is deliberately directed 2.5° away from the baseline connecting the neutrino production point and the Super-Kamiokande detector. This maximises the oscillation effect at 295 km and minimises the background to the ν_e appearance signal. The near detector suite comprises both an on-axis detector and a magnetised off-axis detector. The on-axis detector (INGRID) monitors the neutrino beam direction and profile. The off-axis detector (ND280) characterises the beam before oscillation, measuring the flux, flavour content and energy spectrum of the unoscillated neutrinos. Super-Kamiokande then measures the flavour composition and energy spectrum of the beam after oscillation, in order to search for $\nu_\mu \to \nu_e$ appearance and ν_μ disappearance. It is a large water Cherenkov detector, containing 50 kton of pure water and approximately 13,000 photomultiplier tubes (PMTs) to image the Cherenkov light produced by neutrino interactions in the water.

Each component of the T2K experiment is now discussed in turn. A full description of the T2K experiment is given elsewhere [50], and further details should be sought there.

2.1 Neutrino beam

2.1.1 Off-axis design

The T2K far detector, Super-Kamiokande, is 2.5° off-axis from the neutrino beam direction. The off-axis near detector, ND280, is at the same off-axis angle, located 280 m downstream of the neutrino production point in the direction of Super-Kamiokande.

As shown in Figure 2.2, the ν_μ energy spectrum narrows and moves lower in energy as the off-axis angle increases, as a consequence of the two-body decay kinematics for

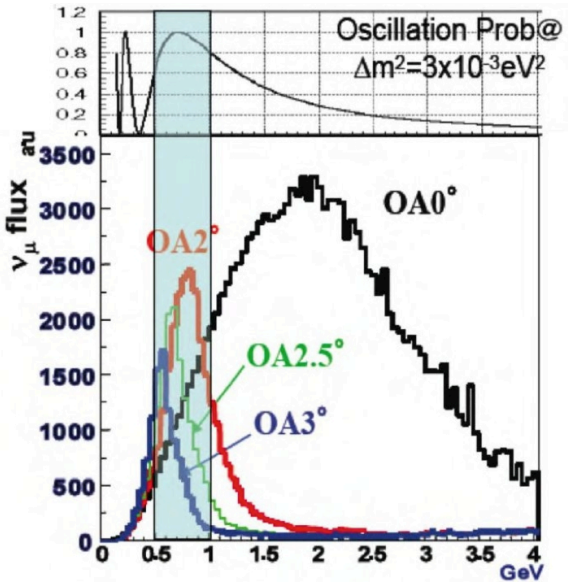


Figure 2.2: ν_μ energy spectrum (bottom) for different off-axis (OA) angles of the beam, and the corresponding ν_e appearance oscillation probability (top) with an assumed Δm^2_{32} value of 3×10^{-3} eV^2.

$\pi^+ \rightarrow \mu^+ + \nu_\mu$. By selecting an off-axis angle of 2.5°, the energy peak is aligned with that needed for maximum oscillation at Super-Kamiokande. This off-axis design has several advantages over a conventional on-axis beam. Firstly, the neutrino flux at the desired energy (the oscillation maximum) is actually higher off-axis than it is on-axis. Secondly, due to the narrow energy band, there is a large reduction in the flux of high energy neutrinos. This reduces the rate of inelastic neutrino interactions which form backgrounds to the quasi-elastic ν_e appearance signal interactions. In particular, the cross section for neutral current π^0 production, one of the main backgrounds to ν_e appearance, increases with energy. Finally, the intrinsic ν_e contamination of the beam is smaller at off-axis angles, due to the different kinematics of the ν_e producing decays.

2.1.2 J-PARC accelerator

There are three accelerators at the J-PARC experimental facility: a linear accelerator (LINAC), a rapid-cycling synchroton (RCS) and the main ring (MR) synchroton. The LINAC is designed to accelerate a beam of H^- anions up to 400 MeV. The electrons are stripped from the H^- anions at the RCS injection. The resulting proton beam is

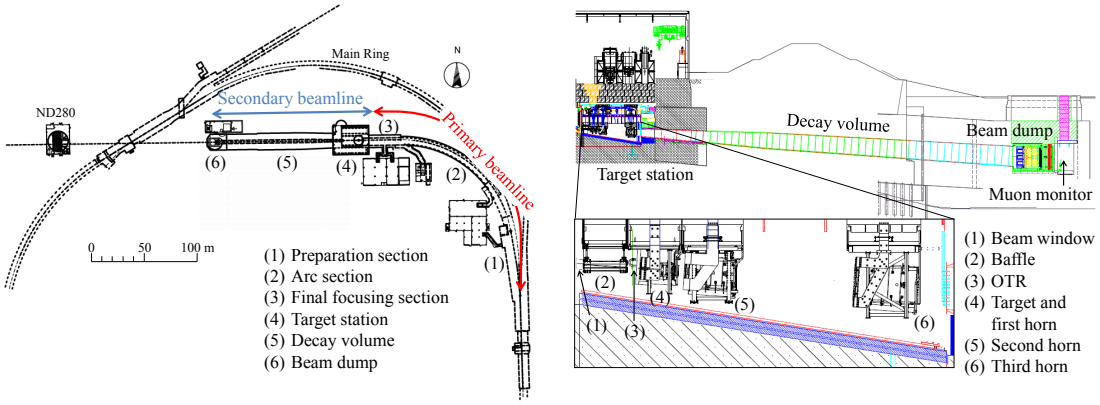
accelerated up to 3 GeV by the RCS, with a cycle rate of 25 Hz and two bunches per cycle. The RCS serves both the MR and a beamline in the Material Life Science Facility, with approximately 5% of the bunches being injected into the MR. The protons are accelerated up to 30 GeV in the MR, with eight bunches per cycle. The MR serves both the neutrino beamline and a hadron beamline. Fast extraction is used for the neutrino beamline, whereby the eight bunches are single-turn extracted in approximately 5 μs to give a beam "spill". Each beam spill consists of eight proton bunches with an inter-bunch separation of 582 ns.[1] This time structure is key to rejecting backgrounds in the various neutrino detectors, including cosmic rays.

2.1.3 Neutrino beamline

The neutrino beamline, composed of a primary and a secondary section, receives the proton beam spills and produces a neutrino beam. A schematic of the beamline is shown in Figure 2.3(a). The primary beamline directs the proton beam towards the Super-Kamiokande detector. In the secondary beamline, the proton beam hits a graphite target producing large numbers of hadrons. The positively charged hadrons are focused by magnetic horns (negatively charged hadrons could instead be focused by reversing the horn polarity) and then decay to neutrinos in the decay volume. These hadrons are dominated by π^+s, which then decay almost exclusively (99.99%) to $\mu^+ + \nu_\mu$, thereby producing a high purity ν_μ beam. However, the muons can further decay to $\bar{\nu}_\mu + e^+ + \nu_e$, which causes contamination of the ν_μ beam. A significant fraction of K^+ particles are also produced at the target which decay to ν_es as well as ν_μs, as well as other mesons which can decay to ν_es and $\bar{\nu}$s. This results in a total contamination of $\sim$1% of the ν_μ beam. The hadronic interactions in the target and the neutrino beam composition will be discussed in more detail in Chapter 3 on the T2K flux prediction.

The primary beamline consists of the preparation section (54 m long), arc section (147 m) long and final focusing section (37 m). The preparation section tunes the proton beam, using a series of normal conducting magnets, in order for it to enter the arc section. The arc section uses superconducting combined function magnets [51] to bend the beam by 80.7° to point towards the Super-Kamiokande detector, with superconducting steering

[1] There were six bunches per spill for Run 1, which (as will be described in Section 2.1.4) ended in June 2010.



(a) (b)

Figure 2.3: (a) Overview of the T2K neutrino beamline, including primary and secondary sections. (b) Side view of the secondary beamline only, where the length of the decay volume is $\sim$ 96 m. Both figures from [50].

magnets to correct the beam orbit. The final focusing section uses 10 normal conducting magnets to focus the beam onto the target and to direct the beam downwards by 3.637° in order to achieve the desired 2.5° offset at the Super-Kamiokande detector. Precise tuning of the proton beam is essential for producing a stable high intensity neutrino beam, and to this end there are approximately 100 monitoring devices around the primary beamline. The intensity, position, profile and loss of the proton beam are monitored by five current transformers (CTs), 21 electrostatic monitors (ESMs), 19 segmented secondary emission monitors (SSEMs) and 50 beam loss monitors (BLMs) respectively (details of these monitors can be found in [50]). The absolute proton beam intensity is known to within 2%, the beam position to within 450 μm, the beam width to within 200 μm and the beam loss can be measured down to 16 mW.

The secondary beamline, as shown in Figure 2.3(b), consists of the target station, decay volume, beam dump and muon monitor. The target station, expanded in the bottom part of Figure 2.3(b), contains a baffle to collimate the proton beam and therefore protect the magnetic horns, an optical transition radiation (OTR) monitor to assess the beam profile just upstream of the target, the target itself and three horns. The target is composed of a graphite rod core surrounded by a 2 mm thick graphite tube, all sealed inside a 0.3 mm thick titanium case. The graphite rod has a diameter of 2.6 cm, a length of 91.4 cm, which is equivalent to just under two interaction lengths, and a density of 1.8 gcm^{-3}. It is expected to reach temperatures of up to 700°C, and is cooled

by helium gas. The target is located inside the first magnetic horn, to maximise the number of charged mesons that are collected and then focused by the subsequent horns. Each horn consists of two coaxial aluminium alloy conductors enclosing a space between them. When a current is passed through the conductors a toroidal magnetic field is generated in this enclosed volume. When run with a current of 320 kA a magnetic field of up to 2.1 T can be produced in each horn. The running of the horns increases the neutrino flux at Super-Kamiokande by a factor of ∼16 at the spectrum peak energy (around 0.6 GeV) when compared to the horns at 0 kA. A pulse current is supplied to the horns and monitored with an uncertainty on the absolute current of less than approximately 2%. The mesons then enter the decay volume, a steel tunnel of length ∼96 m, where they decay to neutrinos. As described above, the dominant meson is π^+ which decays to $\mu^+ + \nu_\mu$ pairs. The beam dump, at the end of the decay volume, is designed to stop any undecayed pions, as well as any muons below ∼5 GeV/c. The muons that continue past the beam dump enter the muon monitor, which is composed of two detector arrays separated by approximately 1 m in the beam direction. The first array is composed of ionisation chambers and the second is composed of silicon PIN photodiodes. By measuring the distribution profile of the muons, the muon monitor determines the intensity and direction of the neutrino beam with an accuracy of better than 3% and 0.25 mrad respectively.

2.1.4 Current status

There have been three continuous physics runs: Run 1 (March to June 2010), Run 2 (November 2010 to March 2011) and Run 3 (March to June 2012). For Run 1, the Electromagnetic Calorimeter of the ND280 detector (which will be described in Section 2.2.3) had only been partially installed; the remaining modules were installed in the shutdown period between Runs 1 and 2. Run 2 was prematurely ended by the devastating March 2011 Tohoku earthquake. Intense repair work was carried out to enable the start of Run 3. This most recent run, however, is not used for the work presented in this thesis and so is not discussed further. Figure 2.4 shows the increasing number of protons per spill delivered by the MR during Runs 1 and 2, up to a value of 9×10^{13} (equivalent to a proton beam power of 145 kW). The integrated number of protons delivered to the target, or "protons on target" (POT), during these first two runs is also shown. This

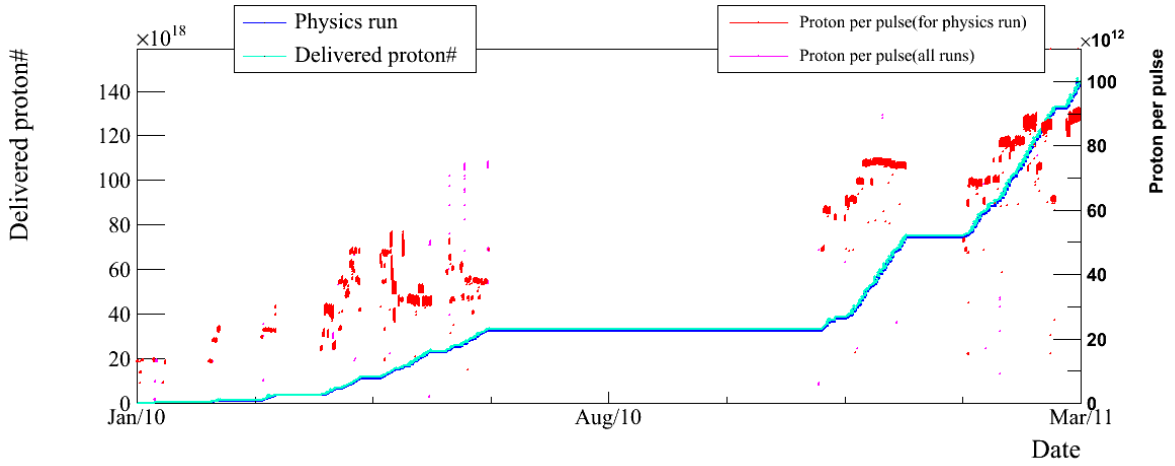


Figure 2.4: Integrated number of protons delivered by the MR to the target (POT), and number of protons per spill (pulse) for the period spanning the Run 1 and Run 2 physics runs [52].

value is directly proportional to the integrated number of neutrinos at both the near detector suite (which will be described in Section 2.2) and the far detector (which will be described in Section 2.3). The POT value is used to normalise Monte Carlo samples to the data set, to enable data to Monte Carlo comparisons of relevant quantities.

2.2 Near detectors

A suite of near detectors is situated 280 m downstream from the neutrino production point in the direction of Super-Kamiokande, housed in a pit of depth 37 m and diameter 17.5 m, and consists of the on-axis INGRID detector and the off-axis magnetised ND280 detector. The relative position of these two detectors is shown in Figure 2.5, where the cross structure of the INGRID detector is centred on the beam direction (which points into the page), and the ND280 is situated to the top right of INGRID with its centre at 2.5° off-axis from the beam direction.

The INGRID detector is composed of iron and plastic scintillator. It is designed to measure the on-axis neutrino beam direction and profile. The magnetised off-axis ND280 detector characterises the beam before oscillation, measuring the flux, flavour content and energy spectrum of the unoscillated neutrinos in the direction of Super-Kamiokande.

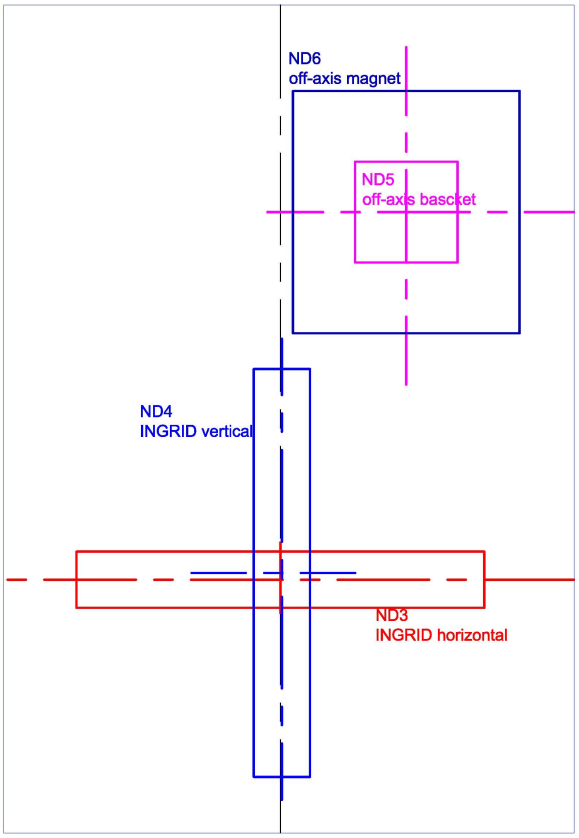


Figure 2.5: Relative position of the INGRID and ND280 detectors. The on-axis INGRID detector has a cross structure (as will be described in more detail in Section 2.2.2), with the vertical modules shown in blue and the horizontal modules shown in red. The cross is centred on the neutrino beam direction, with the neutrino beam pointing into the page. The ND280, situated to the top right of INGRID, has its centre at 2.5° off-axis from the beam direction. The ND280 (shown in more detail in Figure 2.6) includes an outer magnet and inner basket, shown in dark blue and magenta respectively.

The ND280 has several key requirements. Firstly, it must measure the inclusive ν_μ flux as a function of energy, which can then be extrapolated to Super-Kamiokande. Secondly, it must measure the ν_e contamination of the beam which constitutes an irreducible background to the ν_e appearance signal at Super-Kamiokande. Finally, it measures rates for ν_μ interactions that cause backgrounds to the ν_e appearance search at Super-Kamiokande, in particular neutral current π^0 production. To meet these requirements, the ND280 is able to reconstruct exclusive event types. These include charged current

quasi-elastic and inelastic events, and also neutral current events (particularly single π^0 production), for both ν_μ and ν_e interactions.

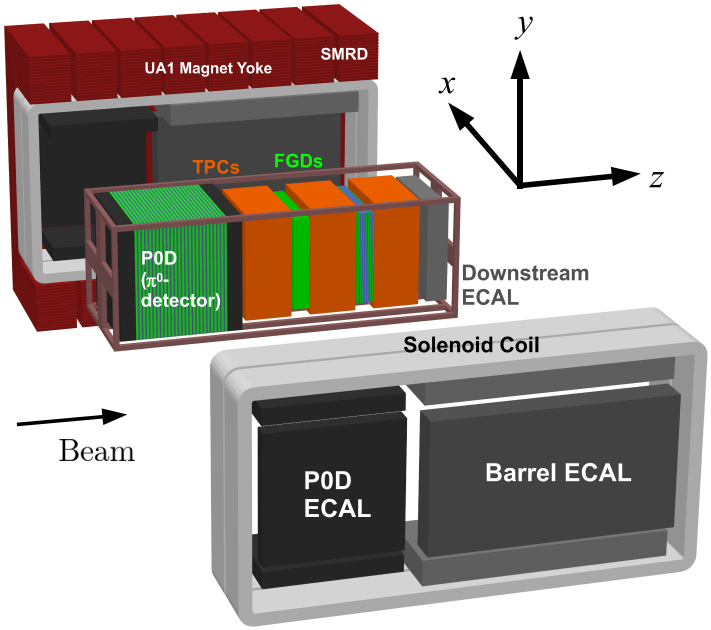


Figure 2.6: An exploded view of the ND280 off-axis detector [50].

Figure 2.6 shows an exploded view of the ND280 detector. It is composed of a number of sub-detectors, all contained inside the magnet recycled from the UA1 experiment. At the upstream end is the pi-zero detector (P0D), which consists of tracking planes of scintillator bars, interleaved with water, lead and brass sheets. The purpose of this sub-detector is to measure the rate of neutral current π^0 production. Downstream of the P0D is the so-called "tracker region", consisting of three time projection chambers (TPCs) and two fine grained detectors (FGDs). The scintillator based FGDs provide the target mass for neutrino interactions and tracking of the produced particles. The TPCs, filled with argon-based drift gas, provide tracking of the particles exiting the FGDs. They also provide momentum measurements, from the curvature of the track in the surrounding magnetic field, and particle identification based on the energy loss of the traversing particle. Surrounding the P0D and tracker region, are the electromagnetic calorimeters (ECals), made of plastic scintillator and lead, which provide near-hermetic coverage for particles exiting the inner detectors. They can provide some discrimination between particle types (based on whether the particle produces an electromagnetic shower or a

track), and are key in reconstructing π^0 particles by detecting the two decay photons. Finally, inserted between the magnet yoke elements is the scintillator based side muon range detector (SMRD). This measures the momentum of escaping muons and can also act as a veto for interactions occurring outside the ND280. Also shown in Figure 2.6 is the right-handed coordinate system used for the ND280 detector. The z axis runs along the centre of the detector at 2.5° off-axis from the beam direction. The x and y axes run along the horizontal and vertical directions of the detector respectively. This coordinate system will be used for the ND280 throughout the remainder of the thesis.

The main analysis of this thesis, presented in Chapter 6, is based upon information from the tracker region.

2.2.1 Scintillator bar readout

The scintillator based sub-detectors (INGRID, P0D, FGDs, ECals and SMRD) all use the same readout system. A wavelength-shifting (WLS) fibre running down the centre of the scintillator bar collects the scintillation light and transports it to a Multi-Pixel Photon Counter (MPPC) at the end of the bar which then converts the light into an electrical signal. Bars are instrumented either at both ends or one end only, details of which can be found in the dedicated sub-detector sections below. The WLS fibre has an absorption spectrum centred at a wavelength of 430 nm (blue) which is well matched to the peak wavelength of the scintillation light at 420 nm. The fibre then re-emits the absorbed light with a shifted wavelength centred at 476 nm (green) to better match the peak sensitivity of the MPPC. Kuraray double-clad Y-11 fibres are used with a 1 mm diameter.

MPPCs (produced by Hamamatsu) were chosen as the photosensors for T2K due to their compact design, suitability for use with WLS fibres, and ability to operate within the 0.2 T magnetic field. Photomultiplier tubes (PMTs), used in many previous neutrino experiments, were not suitable due to the degradation of their performance in the required magnetic field. Photographs of an MPPC are shown in Figure 2.7. Each MPPC consists of a square array of 667 pixels, where each pixel is an avalanche photo-diode operating in Geiger mode. The photo-diode is based on a p-n junction, which in Geiger mode, is held at a reverse bias voltage greater then the breakdown voltage. In this mode, the electric field across the diode depletion region is sufficiently large that a

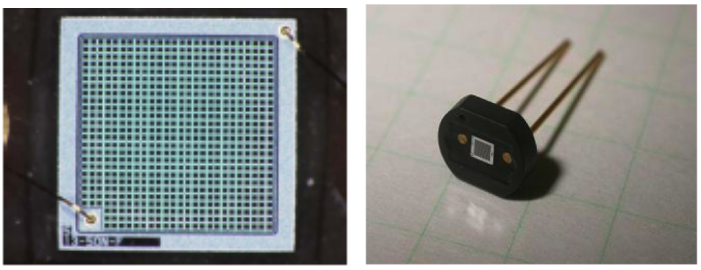


Figure 2.7: Photographs of an MPPC, used to read out the ND280 scintillator based sub-detectors: (left) magnified face view showing the 667 pixels, each with an area of 50×50 μm^2, making a total active area of 1.3×1.3 mm^2; and (right) the MPCC in its ceramic housing [50].

photo-generated free carrier produces a wave (or "avalanche") of additional free carriers through the mechanism of impact ionisation. This generates a large reverse current through the diode. The MPPC gain, G, defined as the charge produced in a single pixel avalanche, is given by $C_{\text{pixel}}(V - V_{\text{BD}})$, where C_{pixel} is the single pixel capacitance, V the operating voltage, and V_{BD} the breakdown voltage. The MPPCs are operated at about 70 V, approximately 1 V above the breakdown voltage, which with the single pixel capacitance of 90 fF gives a gain of approximately 1.0×10^6. The charge produced in an avalanche is independent of the number of free carriers generated inside the pixel by the incident light, and so the light intensity cannot be measured using the charge produced in a single pixel. Instead, the amount of light hitting the MPPC is determined by counting the number of pixels that produce avalanches. The dynamic range of the MPPC is therefore limited by the finite number of pixels.

There are three types of noise for MPPCs: uncorrelated dark noise, crosstalk and afterpulse. Uncorrelated dark noise occurs when free carriers are thermally generated inside the depletion region as opposed to being generated by an incident photon. Crosstalk is when an avalanche in a given pixel triggers an avalanche in a neighbouring pixel. Afterpulse occurs when free carriers are trapped during an avalanche, and their subsequent release triggers an avalanche in the same pixel but at a later time. Cross talk and afterpulse are referred to as correlated noise, since further avalanches are triggered by a primary avalanche. A full account of the MPPCs' operation and response is given in [53].

2.2.2 INGRID on-axis detector

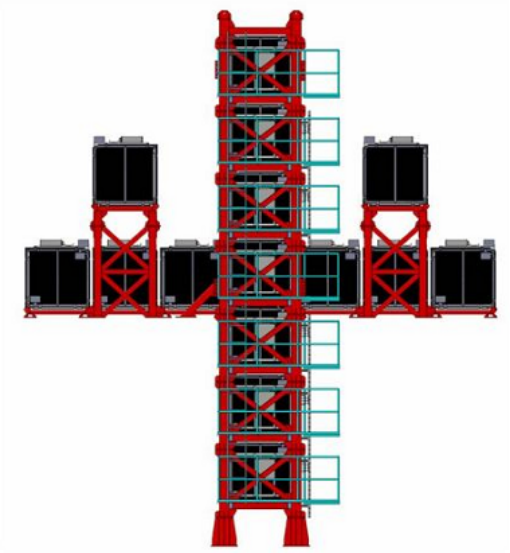


Figure 2.8: The INGRID detector viewed from upstream [50].

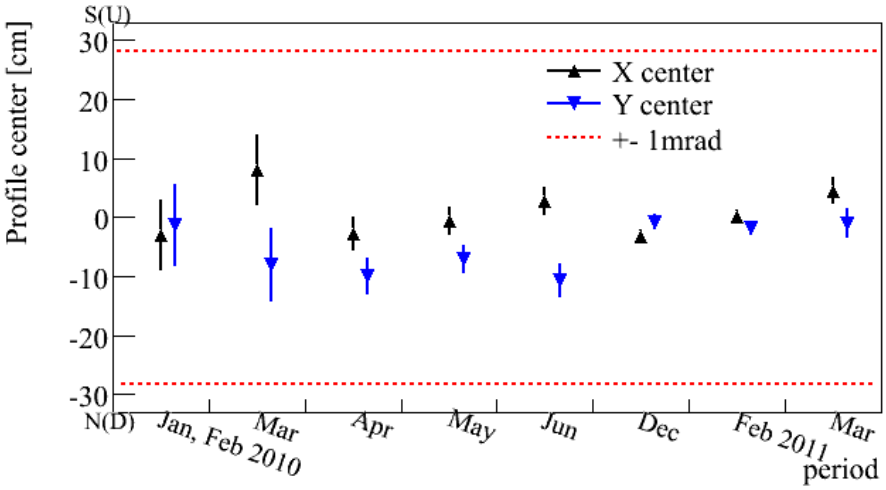


Figure 2.9: Stability of the neutrino beam centre in the horizontal (x, south-north) direction and the vertical (y, up-down) direction, as monitored by INGRID, for the period spanning the Run 1 and Run 2 physics runs [52]. The beam centre was stable to better than ± 1 mrad during this time.

The Interactive Neutrino GRID (INGRID) detector is centred on the neutrino beam axis. Its primary purpose is to monitor the direction of the beam, which is done to a

precision of better than 0.4 mrad. This measurement complements that of the muon monitor (described in Section 2.1.3), which also monitors the beam direction, by providing a measurement further downstream of the beam origin and directly using the interactions of the neutrinos themselves.

INGRID consists of 16 identical modules (with an extra non-standard module to be described later); there are seven horizontal and seven vertical modules arranged in a cross, with two additional modules mounted above the horizontal arm. The configuration is shown in Figure 2.8. Each module consists of nine iron plates interleaved between 11 tracking scintillator planes, surrounded by veto scintillator planes to reject interactions from outside the module.[2] The total mass of iron serving as a neutrino target is 7.1 tons per module. Each of the 11 tracking scintillator planes consists of 24 horizontal and 24 vertical bars, with dimensions for each bar of 1.0 cm $\times$ 5.0 cm $\times$ 120.3 cm. Each veto plane consists of 22 scintillator bars segmented in the beam direction, with dimensions for each bar of either 1.0 cm $\times$ 5.0 cm $\times$ 111.9 cm for the bottom sides or 1.0 cm $\times$ 5.0 cm $\times$ 129.9 cm for the top, left and right sides. The scintillator bars are made from polystyrene doped with scintillation fluors PPO (1%) and POPOP (0.03%). The bars are coated with a thin layer of titanium oxide to reflect escaping light back into the bulk of the bar. A wavelength shifting fibre runs through the centre of each bar and one end is connected to an MPPC, as described in Section 2.2.1. There is a 17^{th} non-standard module, called the Proton Module, designed to identify quasi-elastic neutrino interactions in INGRID. It is a tracking only detector with no iron plates, placed in the centre of the INGRID cross between the standard vertical and horizontal modules.

The INGRID measurements for the beam centre for the period spanning Run 1 and Run 2 are shown in Figure 2.9. The beam centre is stable to better than ± 1 mrad during this time.

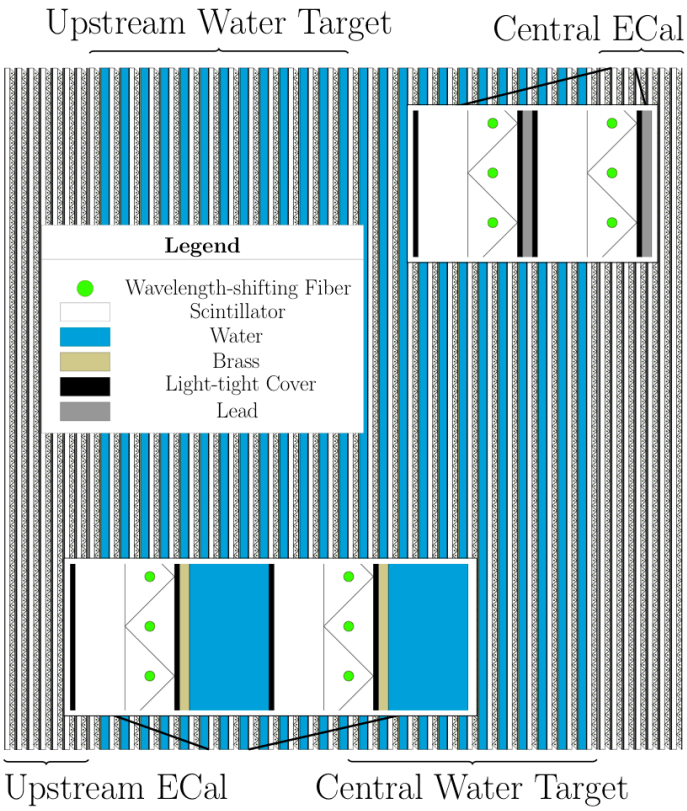


Figure 2.10: A schematic of the P0D detector [50]. The beam travels from left to right.

2.2.3 Off-axis detector

Pi-zero Detector (P0D)

The main purpose of the Pi-zero Detector (P0D) is to measure the cross-section for neutrino-induced neutral current π^0 production ($\nu_\mu + N \rightarrow \nu_\mu + \pi^0 + N'$) on a water target, since this is one of the major backgrounds to the ν_e appearance signal at Super-Kamiokande.

The main features of the P0D are shown in Figure 2.10. The P0D is made up from planes of scintillator bars, interleaved with fillable water bags and lead and brass sheets. Each scintillator plane, or P0Dule, has an array of vertical scintillator bars (134,

[2]There is no iron plate between the 10^{th} and 11^{th} tracking planes due to weight restrictions, but this does not affect the tracking performance.

each 2200 mm long) and an array of horizontal bars (126, each 2340 mm long). Each array is composed of a set of oppositely oriented triangular scintillator bars that connect in such a way that the array itself has a rectangular cross section as opposed to the triangular cross section of a single bar, as shown in Figure 2.10. Each bar is an isosceles triangle with a 33 mm base and a 17 mm height, and is made from polystyrene doped with scintillation fluors PPO (1%) and POPOP (0.03%). The bars are coated with a thin layer of titanium oxide to reflect escaping light back into the bulk of the bar. A wavelength shifting fibre runs through the centre of each bar. The fibre is mirrored at one end and is read out at the opposite end by an MPPC, as described in Section 2.2.1. The front and rear sections of the P0D, the "upstream electromagnetic calorimeter (ECal)" and "central ECal", each consist of seven P0Dules interleaved with seven stainless steel clad lead sheets (4 mm thick). The "upstream water target" and "central water target", composing the central section of the P0D, each consist of 13 P0Dules. In the upstream water target, the P0Dules are interleaved with 13 water bag layers (each 28 mm thick) and 13 brass sheets (each 1.5 mm thick). The central water target has only 12 water bag layers and 12 brass sheets. The mass of the detector with and without the water bags filled is 16.1 tons and 13.3 tons respectively.

Fine Grained Detectors (FGDs)

There are two FGDs in the ND280 detector, as shown in Figure 2.6. They provide target mass for neutrino interactions and tracking capabilities for charged particles coming from the interaction vertex. The FGDs are made from scintillator bars, which provide the target mass. Each FGD contains 1.1 tons of target material and has outer dimensions of 2300 mm $\times$ 2400 mm $\times$ 365 mm. The scintillator bars are arranged in alternating layers of horizontal and vertical bars. The upstream FGD (FGD1) consists of 30 alternating layers with 192 bars in each layer, giving a total of 5,760 bars. The downstream FGD (FGD2) consists of 14 alternating layers (again with 192 bars in each layer), with a 2.5 cm thick layer of water between each pair of horizontal and vertical layers. This gives a total of 2,688 bars and 15 cm thickness of water for FGD2. The scintillator bars are made from polystyrene doped with scintillation fluors PPO (1%) and POPOP (0.03%), and are coated with a thin layer of titanium oxide to reflect escaping light back into the bulk of the bar. A wavelength shifting fibre runs through the centre of each bar. The fibre is mirrored at one end and is read out at the opposite end by an MPPC,

as described in Section 2.2.1. The FGDs can be used to determine neutrino interaction cross-sections on both carbon (in the polystyrene bars) and on water.

Time Projection Chambers (TPCs)

There are three TPCs in the ND280 detector, as shown in Figure 2.6, which are sequentially labelled TPC1, TPC2 and TPC3, moving from the upstream to the downstream direction. They have multiple purposes. Firstly, they provide excellent three-dimensional tracking of the charged particles originating from neutrino interactions in the FGDs (the main target mass). Secondly, since they operate in a magnetic field, the curvature of the track can be used to determine the momentum of the particle. Thirdly, they provide a powerful discriminator between different types of particle, by comparing the measured energy loss through ionisation, as a function of momentum, to the expected energy loss for various particles types.

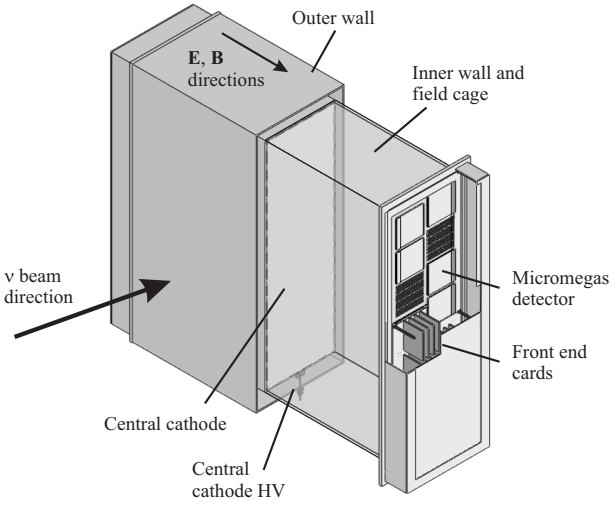


Figure 2.11: Schematic showing the main aspects of the TPC design, with the right-hand part of the outer box cut-away to show the inner sections [50].

Each TPC is constructed as shown in Figure 2.11, with outer dimensions of 2.3 m $\times$ 2.4 m $\times$ 1.0 m. An inner box, filled with argon-based drift gas, is surrounded by an outer box that holds CO_2 as an insulating gas. The inner box is divided in two by the central cathode panel, and has a readout plane on each end parallel to the cathode. Its walls are covered with a series of conducting copper strips joined by precision resistors,

which, in conjunction with the cathode, produce a uniform electric field aligned along the x direction (the same direction as the magnetic field). The inner gas is a mixture of argon, tetrafluoromethane, and isobutane (in the ratios 95:3:2 respectively). Charged particles ionise the gas as they travel through, leaving a trail of ionisation electrons in their wake. These electrons drift away from the central cathode towards one of the readout planes, where they are detected by the Micromegas detectors [54, 55]. There are twelve Micromegas modules on each readout plane with dimensions 342 mm $\times$ 359 mm, each segmented into 1728 rectangular anode pads of 7.0 mm $\times$ 9.8 mm. This totals an active surface of nearly 9 m^2 across the 72 modules of the three TPCs. A calibration system, which produces a control pattern of electrons on the central cathode, enables precise measurements of the electron drift velocity and any distortions in the electron drift due to inhomogeneous and misaligned electric and magnetic fields.

Electromagnetic Calorimeters (ECals)

The ND280 ECal is a sampling electromagnetic calorimeter, made of plastic scintillator and lead, that surrounds the inner detectors (P0D, TPCs and FGDs), as shown in Figure 2.6, and provides near-hermetic coverage for particles exiting these detectors. Its purpose is to complement the inner detectors in full event reconstruction by detecting exiting photons and charged particles. It is key to reconstructing π^0 particles produced in the tracker volume, by measuring the energy and direction of the two decay photons. It can also provide some separation between particle types (electron-muon-pion separation), according to whether the particle produces an electromagnetic shower or a track.

It consists of two main parts: the Tracker ECal, and the P0D ECal. The Tracker ECal surrounds the tracker volume and consists of six Barrel ECal modules and one Downstream ECal (DsECal) module. The Barrel ECal modules are mounted inside the UA1 magnet and surround the tracker volume on its four sides parallel to the z direction; there are two top modules and two bottom modules, since the magnet is split into two halves along the z direction, together with two side modules. The DsECal module covers the downstream exit of the tracker volume. The P0D ECal consists of six modules which surround the P0D detector on its four sides parallel to the z direction; as for the Barrel

ECal modules, there are two top modules, two bottom modules, and two side modules. This gives a total of 13 modules for the ND280 ECal.

Each ECal module is composed of layers of scintillator bars, which provide the active sampling material, interleaved with lead absorber sheets. Consecutive scintillator layers have their bars arranged perpendicular to each other, to allow three-dimensional reconstruction of electromagnetic showers and charged particle tracks. The bars have a 4.0 cm $\times$ 1.0 cm cross section, with differing lengths according to which module they are in. They are made from polystyrene doped with scintillation fluors PPO (1%) and POPOP (0.03%), and are coated with a thin layer of titanium oxide to reflect escaping light back into the bulk of the bar. A wavelength shifting fibre runs through the centre of each bar, and is read out either at one end or both ends by MPPCs, as described in Section 2.2.1.

The DsECal module consists of 34 scintillator layers and 1.75 mm thick lead sheets, providing a total of 10.6 radiation lengths. Each layer consists of 50 bars, each 2.04 m long and read out at both ends by an MPPC. Limited by the space inside the UA1 magnet, the Barrel ECal modules consist of 31 layers with the same interleaved lead sheets, providing a slightly reduced total thickness of 9.7 interaction lengths. The "long" bars, in the z direction, have a length of 3.84 m and are read out at both ends by an MPPC. The "short" bars have lengths of 1.52 m (running along the x direction) and 2.36 m (running along the y direction) for the top or bottom modules and side modules respectively, and are read out at one end by an MPPC.

The P0D ECal modules have a simpler construction, since they are not intended for full π^0 reconstruction; this is done by the dedicated P0D detector which they surround. Their main role is to provide information on energy escaping from the P0D π^0 interaction, by detecting photons that either do not convert or are only partially contained inside the P0D. Each module consists of six scintillator layers interleaved with 4 mm thick lead sheets, providing a total of 3.6 interaction lengths. The bars are 2.34 m long, all running along the z direction, and are read out at one end by an MPPC.

Magnet and Side Muon Range Detector (SMRD)

The magnet is reused from the UA1 [56] and NOMAD [57] experiments. It provides a 0.2 T magnetic field which enables the momentum and sign of charged particles to be determined by the sub-detectors which it surrounds. The magnet consists of solenoid coils, carrying a current of 2900 A, and a flux return yoke. The inner dimensions of the magnet are 7.0 m $\times$ 3.5 m $\times$ 3.6 m. The external dimensions are 7.6 m $\times$ 5.6 m $\times$ 6.1 m. The yoke has a total weight of 850 tons. The magnet is split along its length into two halves and mounted on rails, such that it can be opened to provide access to the inner sub-detectors. Each half is made up from eight C-shaped yoke elements surrounding two aluminium solenoid coils, as shown in Figure 2.6. Each C-shaped yoke element consists of 16 C-shaped steel layers in a radial arrangement (such that the inner layers are smaller in the x and y directions than the outer layers). There is an air gap of 1.7 cm between each 4.8 cm thick steel layer, giving a total of 15 air gaps for each yoke element. The field inside the magnet has an uncertainty of 2×10^{-4} T for each component (at the nominal field of 0.2 T). This level of precision reduces the systematic uncertainty of the momentum determination for charged particles.

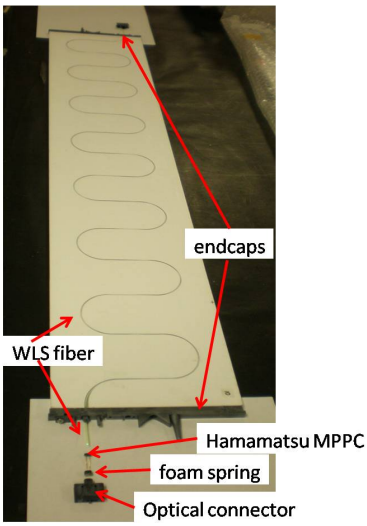


Figure 2.12: Components of an SMRD scintillator counter prior to assembly [50].

The SMRD is a system of scintillator modules inserted into the air gaps of the magnet yoke elements. The SMRD has multiple purposes. It measures the momentum of escaping muons arising from neutrino interactions in the surrounded sub-detectors. It also acts as a veto for neutrino interactions which occur outside the ND280 detector

and in the magnet yoke itself, and finally as a trigger for cosmic ray muons entering the ND280 detector. The SMRD modules populate the innermost air gaps of the yoke elements, in order to maximise the efficiency for detecting escaping particles. There are three layers of SMRD modules on the top and bottom of the yoke elements, with each module composed of four scintillation counters (875 mm $\times$ 167 mm $\times$ 7 mm). For the sides, there are three layers of SMRD modules for the five most upstream yoke elements, four layers for the sixth yoke element and six layers for the seventh and eighth yoke elements, with each SMRD module composed of five scintillation counters (875 mm $\times$ 175 mm $\times$ 7 mm). This arrangement, where the downstream yoke elements contain more SMRD modules, is due to the higher flux of particles at lower angles with respect to the beam direction. Each scintillator counter has a scintillator bar made from polystyrene and dimethylacetamide with admixtures of POPOP and para-terphenyl. The surface of the bar is coated with a white diffuse layer which acts as a reflector. An S-shaped groove is machined into the bar, along which runs a wavelength shifting fibre, as shown in Figure 2.12. Each end of the fibre is read out by an MPPC, as described in Section 2.2.1.

Readout electronics

This section gives an overview of the readout electronics of the ND280 detector. It focuses on the "Trip-t based" sub-detectors (P0D, ECals and SMRD), since this information is important for Chapter 5, which explains the timing calibration of the DsECal module and P0D. More detailed information about the electronics can be found in [50] and [58].

The P0D, ECals and SMRD sub-detectors use identical electronics to read out the MPPCs (which are described previously in Section 2.2.1), based on the Trip-t chip [59] which was originally developed for the D0 experiment. Sixteen MPPCs are connected to the each Trip-t chip, with each MPPC occupying two channels. This dual channel setup increases the dynamic range of the electronics; the output charge of the MPPC is capacitively split between a low gain channel and high gain channel (with 10 times the gain of the low gain channel), saturating at 500 and 50 photoelectrons respectively. For each channel, the Trip-t chip integrates the charge in programmable integration windows, separated by reset periods, where the integration windows are synchronised with the bunches of the neutrino beam spill. There are 23 integration windows; in

nominal running, each integration window has a duration of 480 ns, separated by 100 ns reset periods. This time structure is shown in Figure 2.13. The output of the high gain channel is routed to a discriminator. For each integration window, the discriminator will fire if the input charge exceeds a programmed threshold (which varies for different sub-detectors); the firing of the discriminator is time-stamped with an accuracy of 2.5 ns. This time-stamping process is important for the fibre timewalk effect, the calibration of which is presented in Chapter 5.

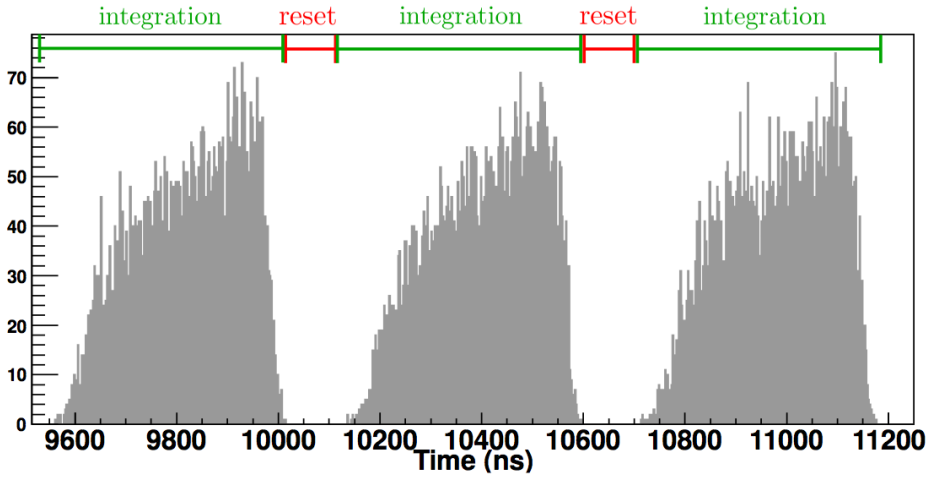


Figure 2.13: Time structure of noise signals in the MPPCs when read out by the Trip-t electronics. The Trip-t chip integrates the charge in integration windows of duration 480 ns, separated by 100 ns reset periods. The y-axis is in arbitrary units representing the accumulated charge. There are 23 integration windows for one full cycle of the readout, but only three are shown for convenience. Plot taken from [60].

The Trip-t chips are housed on front-end boards; each Trip-t front-end board (TFB) houses four chips. Figure 2.14 gives a schematic overview of the ND280 readout electronics, showing the TFBs and back-end electronics for the Trip-t based sub-detectors. The setup for the TPCs and FGDs, which use different front-end electronics to the Trip-t based system, is also shown but is only described very briefly here. For the Trip-t sub-detectors, the back-end electronics consist of readout merger modules (RMMs), a cosmic trigger module (CTM), slave clock modules (SCMs) and a master clock module (MCM). Each Trip-t sub-detector has a different number of TFBs and RMMs, and its own SCM. There is one Trip-t CTM (with inputs from the SMRD, P0D and DsECal detectors). There is one MCM. The TPCs and FGDs also have their own SCMs which are connected to the MCM, and the FGD has its own CTM.

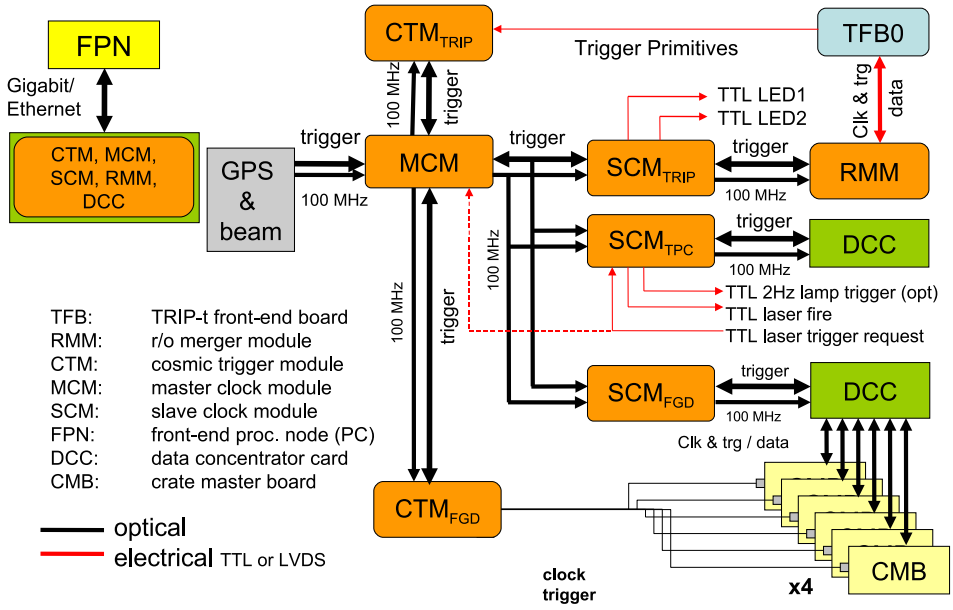


Figure 2.14: Overview of the ND280 readout electronics. The purposes of the most relevant components are explained in the text.

The MCM receives signals from the accelerator indicating the arrival time of the beam spills. It then disseminates a trigger signal to the SCMs. It also distributes clock signals to the SCMs, which are used to synchronise the electronics with UTC. The SCMs pass the trigger and clock signals to the RMMs (or equivalent for the TPCs and FGDs). Having independent SCMs is useful since it allows for stand-alone operation of each sub-detector. Each RMM is connected to up to 48 TFBs. The RMM distributes the clock and trigger signals, and receives the data after a trigger signal is received by the TFBs. The RMM then sends this data to a commercial PC to be stored and processed. The CTM determines whether there was a cosmic muon event in the detector (based on signals from up to 192 TFBs), and if necessary passes a trigger signal to the MCM. The MCM→SCM, SCM→RMM and MCM→CTM connections are all Rocket I/O driven optical links, whereas the RMM→TFB and TFB→CTM connections use electrical cables. Differences in cable length for some of the aforementioned connections and clock phase shifts due to the Rocket I/O protocol form a major part of the time calibration presented in Chapter 5.

As mentioned previously, the TPCs and FGDs use different front-end electronics to the Trip-t based system detailed above. They use custom-made “AFTER” ASIC chips. For the TPCs (FGDs), these chips sample the signals from the Micromegas pads

(MPPCs) at a period of 40 ns (20 ns) and output a waveform. Further details can be found in [50].

2.3 Far detector

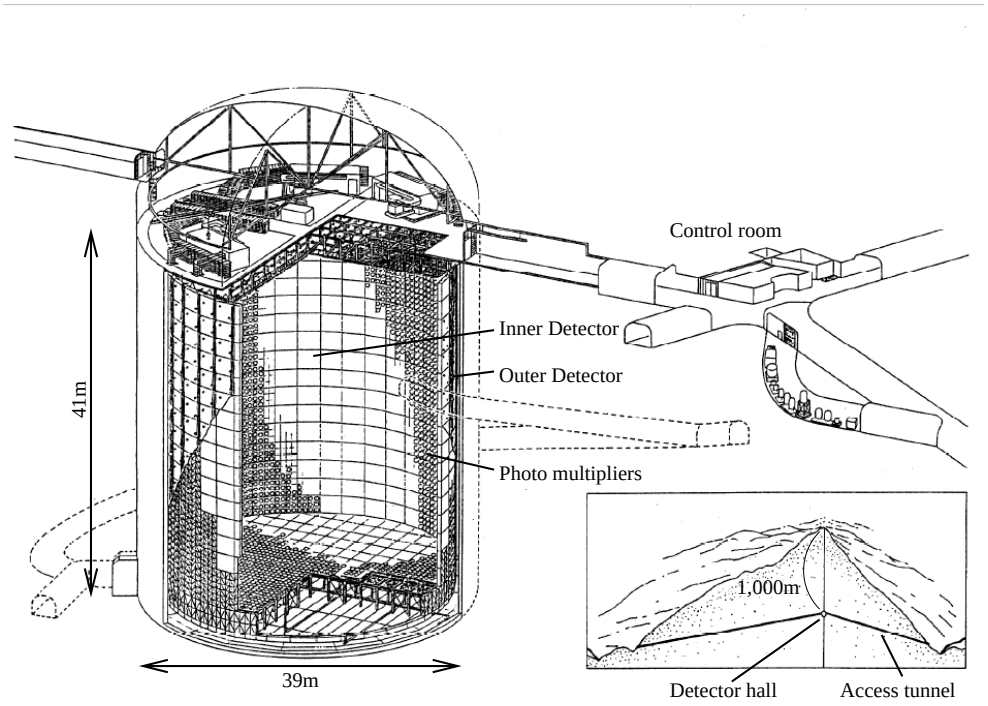


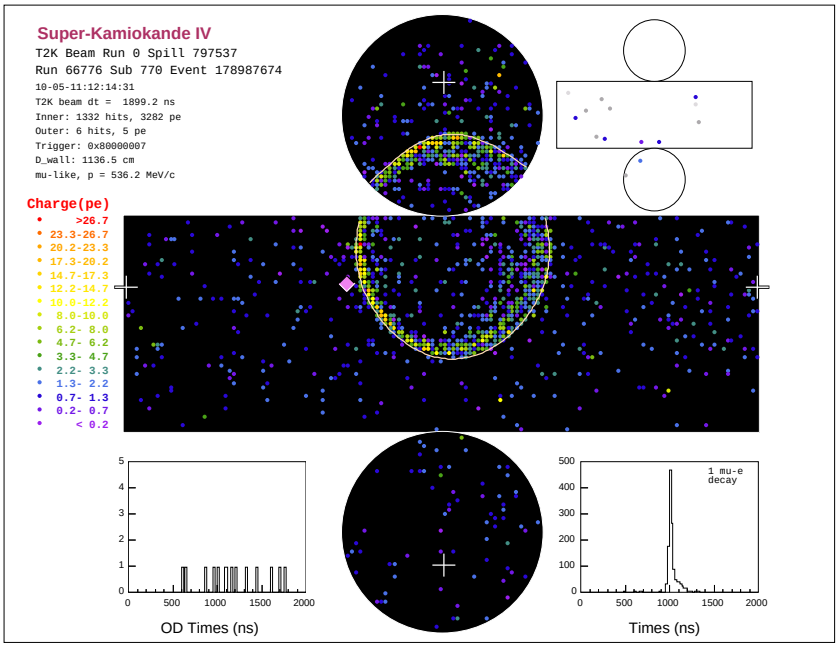
Figure 2.15: Schematic of the Super-Kamiokande detector [50].

The Super-Kamiokande detector is located 1 km deep in the Ikenoyama mountain, a cylindrical cavern filled with 50 kton of pure water. Neutrino interactions in the water will often result in the production of charged particles. As these charged particles travel through the water they will produce a cone of Cherenkov light, if they are above a certain energy threshold. This light is recorded by approximately 13,000 photomultiplier tubes (PMTs) inside the Super-Kamiokande detector, allowing details of the neutrino interaction to be reconstructed, such as the interaction vertex and the type and momenta of the product particles. Running since 1996, it is a very stable and well understood detector. It consists of an inner detector (ID) and outer detector (OD), separated by a 50 cm wide cylindrical stainless steel structure. A schematic is shown in Figure 2.15.

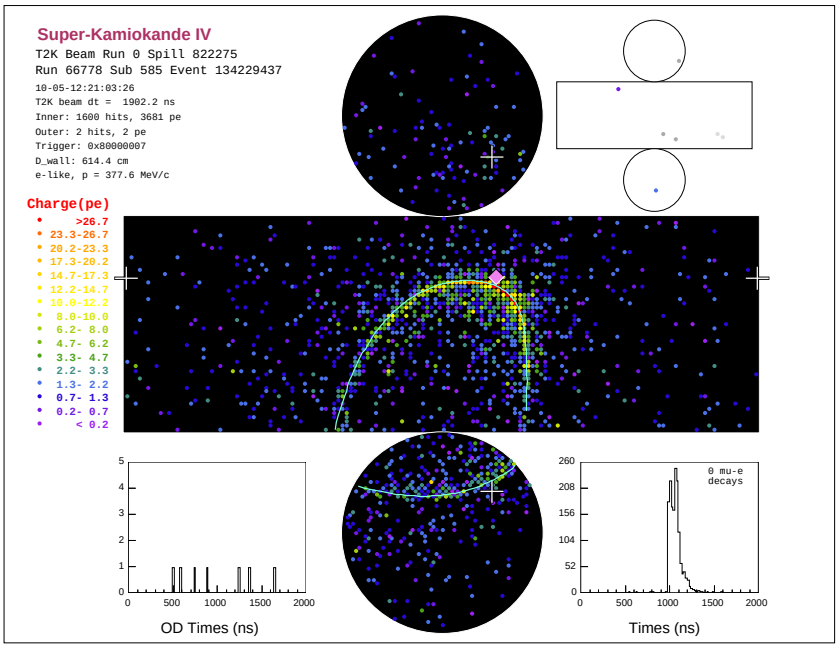
The ID cylinder has a 33.8 m diameter and 36.2 m height, with its boundary defined by the inner edge of the stainless steel structure. This wall is covered with a black sheet of plastic to minimise the scattering of photons back off the ID wall, or their passing through into the OD. On this wall, 11,129 50 cm diameter PMTs are mounted at regular intervals facing into the ID volume, which equates to a 40% surface coverage, providing sufficient spatial resolution to effectively image neutrino interactions in the ID.

Surrounding the ID is the cylindrical OD, extending approximately 2 m beyond the steel separation structure in both the radial direction and on either end. Mounted on the outer wall of the steel structure are 1,885 20 cm diameter PMTs facing outwards into the OD volume. The wall is covered with a highly reflective material to compensate for the relatively few number of PMTs in the OD. The reflective surface increases the chance of photons bouncing off the wall and ultimately entering one of the OD PMTs. Information from the OD can be used to determine whether any particles entered or left the ID. This can be used to veto cosmic ray muons, and also to determine whether an event is fully contained within the ID, which is important for determining the momentum of the event.

The PMT information can be reconstructed to determine the type of particle producing the Cherenkov ring on the detector wall. Rings are classified as "muon-like" or "electron-like". This classification is crucial for T2K in its search for ν_μ disappearance and ν_e appearance, since ν_μ and ν_e charged current interactions will produce muons and electrons respectively. Muons are unlikely to scatter, due to their relatively heavy mass, and so produce clear, sharp rings of PMT hits on the ID wall. Electrons, on the other hand, are likely to scatter due to their small mass and almost always result in an electromagnetic shower. This produces a "fuzzy" ring of PMT hits, due to the many different particles of the shower travelling in different directions. Figure 2.16 shows the PMT hit patterns for two T2K neutrino beam interactions; Figure 2.16(a) shows a "muon-like" Cherenkov ring, and Figure 2.16(b) shows an "electron-like" ring. Details of the Super-Kamiokande reconstruction can be found in [61].



(a)



(b)

Figure 2.16: Example of reconstructed T2K events in Super-Kamiokande for (a) a "muon-like" ring and (b) an "electron-like" ring. Both figures show the cylindrical detector, unrolled onto a plane. Each coloured point represents a PMT, with the colour corresponding to the amount of charge, and the reconstructed cone is shown as a white line. The second figure in the upper right corner shows the same hit map for the OD. The white crosses indicate the location of the reconstructed vertex. The diamond marks the location where a ray starting from the event vertex and heading in the direction of the beam would intersect the detector wall. Figure and caption taken from [50].

Chapter 3

Flux prediction for T2K

3.1 Overview

A 30 GeV proton beam is incident upon a graphite target, giving rise to hadrons (mainly pions and kaons) that are focused by the magnetic horns and decay to neutrinos in the decay volume, as discussed in Chapter 2. This full chain of events is simulated within the T2K neutrino beam simulation, JNUBEAM, which is based upon the GEANT3 [62] simulation tool. The modelling of hadron interactions in the simulation is done using the GCALOR model [63]. To improve the agreement with hadron interaction data from the NA61 experiment (which will be described in Section 3.2.1), the interaction of the primary protons with the target and the subsequent hadronic chains inside the target are simulated separately by the FLUKA hadron production model (2008 version) [49]. The particles exiting the target are traced through the beamline by JNUBEAM until they stop or decay. The magnetic field in the horns, described in Section 2.1.3, is calculated according to Ampère's law and varies as $1/r$, where r is the distance from the horn axis. The particles may interact with the horn material in the target station (which is modelled using GCALOR). The geometry of the beamline (including baffle, target, horns, helium vessel, decay volume, beam dump and muon monitor), described in Section 2.1.3, is replicated according to the final mechanical drawings. The decay of particles to neutrinos is simulated according to the current best knowledge of the branching ratios [64]. The decays that are simulated within JNUBEAM are given in Table 3.1.

	Parent particle				
Neutrino	π^+	K^+	K^0_L	μ^+	μ^-
ν_μ	$\mu^+\nu_\mu$ (99.9877%)	$\mu^+\nu_\mu$ (63.55%) $\pi^0\mu^+\nu_\mu$ (3.353%, $K^+_{\mu 3}$)	$\pi^-\mu^+\nu_\mu$ (27.04%, $K^0_{\mu 3}$)		$e^-\nu_\mu\bar{\nu}_e$ ($\approx 100\%$)
ν_e	$e^+\nu_e$ (1.23×10^{-2}%)	$\pi^0 e^+\nu_e$ (5.07%, K^+_{e3})	$\pi^- e^+\nu_e$ (40.55%, K^0_{e3})	$e^+\bar{\nu}_\mu\nu_e$ ($\approx 100\%$)	

Table 3.1: Decay modes of parent particles to daughter neutrinos that are simulated within JNUBEAM. The branching fractions are given in brackets, along with the name of the decay mode (where appropriate). Decay modes to $\bar{\nu}_\mu$ and $\bar{\nu}_e$ are omitted since they are simply the charge conjugates of the decay modes to ν_μ and ν_e respectively.

Once a particle decays to a neutrino, its information and history are stored along with the flavour and energy of the neutrino. The probability for that neutrino to be produced in the direction of the ND280 detector or Super-Kamiokande is also saved. Energy thresholds are used for the tracking of particles in JNUBEAM in order to limit output file size and computing time, with a threshold of 0.1 GeV for hadrons and muons and 1.0 GeV for gammas and electrons. When a particle reaches the relevant threshold, tracking of the particle is terminated and its decay is simulated (if the particle decays to a neutrino then the event history is stored).

Following this initial calculation of the flux, the simulated hadronic interactions are tuned to measurements from external data. This is very important since hadronic interactions are the main source of uncertainty in the flux prediction. A weight is applied to each simulated event to apply this tuning. Details of the hadronic interaction tuning are discussed in Section 3.2. This is followed by a discussion of the uncertainties on the flux prediction in Section 3.3. Finally, details of how measurements at the ND280 detector (described in Section 2.2.3) can be used to further constrain and validate the flux prediction are given in Section 3.4.

3.2 Tuning of hadronic interactions

As discussed previously in Section 1.2.1, hadronic interactions are the major source of uncertainty in the flux prediction. Tuning of hadronic interactions is performed using external data. There are two parts to the tuning procedure. Firstly, the differential production of $\pi^\pm$, $K^\pm$ and K^0_L particles in interactions of the beam protons on the

graphite target (and horn materials) are tuned. Secondly, the total cross-sections for p, $\pi^{\pm}$ and $K^{\pm}$ interactions are tuned, determining the attenuation rate of these hadrons along with the production rate of other hadrons that may decay to neutrinos. These tuning processes are briefly summarised in this section; full details should be sought in [65] and [66].

In this section, secondary hadrons refer to hadrons produced by the interactions of the original protons, and tertiary hadrons refer to hadrons produced by the interactions of any nucleons other than the original proton.

3.2.1 NA61/SHINE experiment

The primary source of external data for use in the hadronic interaction tuning is the NA61/SHINE experiment. This experiment was built in order to, amongst other things, provide precise hadron production measurements for the T2K experiment. It is a large acceptance hadron spectrometer located at CERN. A beam of protons, with the same energy as the T2K beam, is directed at a graphite target. A very brief description of the NA61/SHINE experiment is given here; full details can be found in [67] and the references therein.

The experimental setup is shown in Figure 3.1. A 31 GeV/c secondary hadron beam is directed at the target. A set of detectors are positioned upstream of the target in order to provide measurements of timing, particle identification and position for the incoming beam particles. Protons in the beam are identified with a purity of ∼99%. Two magnets lie downstream of the target to focus exiting particles. Tracking of the particles in the detector is done using large volume TPCs. A vertex TPC (VTPC) is located at the centre of each of the magnets. In addition, two large volume main TPCs (MTPC) are positioned downstream of the magnets. particle identification is provided both by the TPCs, via energy loss ($\mathrm{d}E/\mathrm{d}x$) measurements, and by the time-of-flight (ToF) detectors, which consist of scintillator bars with photomultiplier readout.

Two targets are used, made of the same density graphite as is used for the T2K target:

- Thin target: A 2 cm long target (∼4% of a nuclear interaction length)

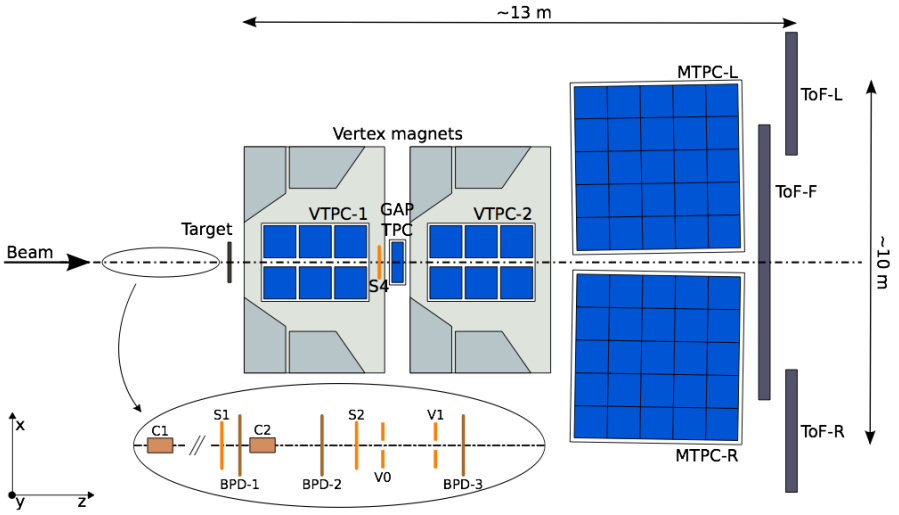


Figure 3.1: The layout of the NA61/SHINE experiment. The components are described in the main text. The figure is taken from [67].

- T2K replica target: A cylinder with a diameter of 2.6 cm and a length of 90 cm (∼1.9 of a nuclear interaction length)

The data used here is taken with the thin target; data taken with the T2K replica target will be used when available.

3.2.2 Data sets used for hadronic interaction tuning

Several data sets are used to tune the differential pion and kaon production, summarised in Table 3.2. The primary data sets are the pion production [67] and kaon production [68] results from the NA61/SHINE experiment discussed in the previous section. These data sets cover the majority of the hadron production phase space that contributes to the T2K flux, as shown in Figure 3.2. More than 90% of the pion phase space is covered, and a significant fraction of the kaon phase space is covered. Two further data sets are used to tune the kaon production in the phase space not covered by the NA61/SHINE results, described in [69] and [70]. In addition, the pion production results from the E910 experiment [44] are used to evaluate the systematic uncertainty for tertiary pion production.

Experiment	p beam momentum / GeV/c	Target	Produced hadrons
NA61/SHINE [67] [68]	31 GeV/c	C	$\pi^{\pm}$, K^{+}
Eichten *et al.* [69]	24 GeV/c	Be, Al, ...	p, $\pi^{\pm}$, $K^{\pm}$
Allaby *et al.* [70]	19.2 GeV/c	Be, Al, ...	p, $\pi^{\pm}$, $K^{\pm}$
E910 [44]	6.4-17.5 GeV/c	Be	$\pi^{\pm}$

Table 3.2: Summary of the hadron production data, taken with different proton beam momenta and target materials, used to tune the pion and kaon production in the T2K flux prediction.

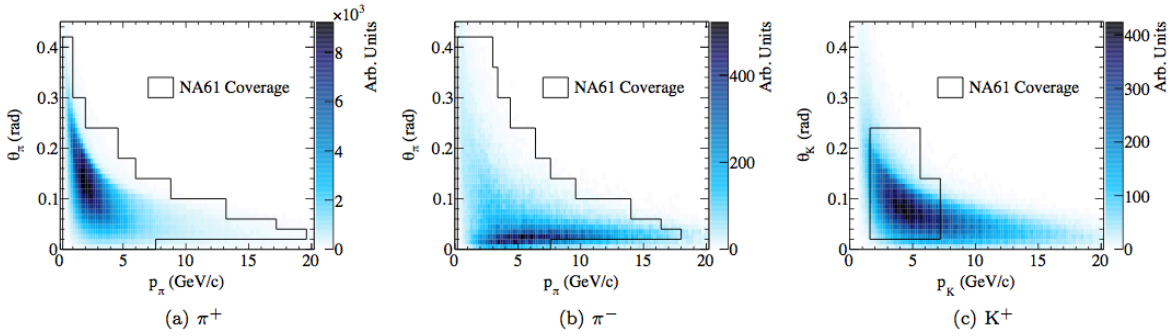


Figure 3.2: The phase space of pions and kaons contributing to the predicted neutrino flux at Super-Kamiokande, and the regions covered by NA61 measurements. The x and y axes show the momentum and angle of the hadron respectively. Figure and caption (adapted) taken from [65].

Data set	Beam particles	Target	Beam momentum / GeV/c	Measurement
Abrams et al. [71]	$K^{\pm}$	C, Cu	1-3.3	σ_{inel}
Allaby et al. [72]	π^{-}, K^{-}	C, Al, ...	20-65	σ_{inel}
Allardyce et al. [73]	$\pi^{\pm}$	C, Al, ...	0.71-2	σ_{inel}
Bellettini et al. [74]	p	C, Al, ...	19.3, 21.5	σ_{inel}
Bobchenko et al. [75]	π^{-}, p	C, Al, ...	1.75-9	σ_{inel}
Carroll et al. [76]	$\pi^{\pm}$, $K^{\pm}$, p	C, Al, ...	60-280	σ_{prod}
Cronin et al. [77]	π^{-}	C, Al, ...	0.73-1.33	σ_{inel}
Chen et al. [78]	p	C, Al, ...	1.53	σ_{inel}
Denisov et al. [79]	$\pi^{\pm}$, $K^{\pm}$, p	C, Al, ...	6-60	σ_{inel}
Longo et al. [80]	π^{+}, p	C, Al, ...	3	σ_{inel}
NA61 [67]	p	C	31	σ_{prod}
Vlasov et al. [81]	π^{-}	C, Al	2-6.7	σ_{inel}

Table 3.3: Summary of the data used to tune the total cross-sections for p, $\pi^{\pm}$ and $K^{\pm}$ interactions for the T2K flux prediction. The inelastic cross-section, σ_{inel}, and production cross-section, σ_{prod}, are defined in the main text.

Data from many different experiments is used to tune the total cross-sections for p, $\pi^{\pm}$ and $K^{\pm}$ interactions, summarised in Table 3.3. There is a subtlety in the exact cross-section that each experiment measures, which is important later when discussing the uncertainty on this tuning. Most experiments measure the inelastic cross-section, σ_{inel}, defined as the total cross-section minus the elastic cross-section. However, some experiments, including NA61/SHINE, measure the production cross-section, σ_{prod}, defined as the inelastic cross-section minus the quasi-elastic cross-section for scattering of the particle off of individual nuclei. This production cross-section represents the rate of interactions in which hadrons are produced in the final state. The measurement of each experiment is shown in the final column of Table 3.3.

3.2.3 Methodology for tuning hadronic interactions

Firstly, the tuning of the differential pion and kaon production is described. The tuning is evaluated using the differential production multiplicity of the given hadron, $\frac{dn}{dpd\theta}$, in

the momentum, p, and angle[1], θ, of the produced particle, defined as:

$$\frac{dn}{dpd\theta}(p_{\text{in}}, A) = \frac{1}{\sigma_{\text{prod}}(p_{\text{in}}, A)} \frac{d\sigma(p_{\text{in}}, A)}{dpd\theta}, \tag{3.1}$$

where σ_{prod} is the production cross-section, defined in Section 3.2.2, which is a function of the incident particle momentum, p_{in}, and the target nucleus, A. A tuning weight is calculated for each simulated interaction that produces the given hadron, and is defined as the ratio of the differential production multiplicity in data and MC:

$$w(p_{\text{in}}, A) = \frac{\left[\frac{dn}{dpd\theta}(p_{\text{in}}, A)\right]_{\text{data}}}{\left[\frac{dn}{dpd\theta}(p_{\text{in}}, A)\right]_{\text{MC}}}. \tag{3.2}$$

The differential pion and kaon production data of the NA61/SHINE experiment is provided in the form of Equation 3.1, at the same proton beam momentum as for T2K and for the same target material. Therefore, the evaluation of the tuning weights is simple for interactions in the T2K target that produce secondary $\pi^{\pm}$ or K^{+} particles covered by the NA61/SHINE data. However, the situation is more complicated for interactions outside of the target and for hadrons produced outside of the NA61/SHINE phase space. In addition, the tuning of tertiary hadron production is complicated by the fact that the secondary hadrons that produce these particles have lower momentum than the incident proton beam. In order to tackle these different situations, two scaling hypotheses are used: Feynman scaling and "A-scaling".

When analysing hadron production data from high energy inelastic collisions, Feynman observed that the invariant production cross-section ($Ed^3\sigma/d^3p$) was independent of the total centre of mass energy if described using a suitable set of parameters [45]. These parameters are the transverse momentum, p_{T}, and the Feynman scaling variable, x_{F}, defined as:

$$x_{\text{F}} = \frac{p_{\text{L}}}{p_{\text{L}}^{\text{max}}}, \tag{3.3}$$

where p_{L} is the longitudinal momentum of the produced particle and $p_{\text{L}}^{\text{max}}$ is the maximum allowed longitudinal momentum of the produced particle, both in the centre of

[1]Defined relative to the direction of the incident particle.

mass frame. This enabled data sets taken at different incident beam energies to be directly compared, by factoring out the dependence of the cross-section on the beam energy. Feynman scaling is assumed in the hadronic interaction tuning to compare the differential production rates from data taken at different incident nucleon momenta.

A-scaling is described as follows. Bonesini *et al.* [82], following work by Barton *et al.* [83] and Skubic *et al.* [84], proposed the following parametrisation to scale the differential cross-section on a given target material, A_0, to that on a different target material, A_1:

$$E\frac{d^3\sigma(A_1)}{dp^3} = \left[\frac{A_1}{A_0}\right]^{\alpha(x_\mathrm{F},p_\mathrm{T})} E\frac{d^3\sigma(A_0)}{dp^3}, \tag{3.4}$$

where:

$$\alpha(x_\mathrm{F}, p_\mathrm{T}) = (a + bx_\mathrm{F} + cx_\mathrm{F}^2)(d + ep_\mathrm{T}^2). \tag{3.5}$$

The values of the parameters a through to e in Equation 3.5 are given in [82]. However, these values are fine-tuned using the Be and Al target data sets from both [69] and [70] separately.

Secondary pion production in the horn material is tuned by scaling the NA61 pion production data on carbon to aluminium using Equation 3.4. The tuning of tertiary pion production involves extrapolating from the NA61/SHINE data down to lower incident nucleon momenta. This is done assuming Feynman scaling; tuning weights calculated from the NA61/SHINE pion production data, binned according to p-θ of the pion, are converted into the x_F-p_T phase space and applied to the tertiary production events based on the x_F and p_T of those events. In addition, for tertiary pion production in the horns, a scaling of the data from carbon to aluminium is performed according to Equation 3.4.

For the tuning of K^+ production in the phase space not covered by data from NA61/SHINE and for the tuning of K^- production, the Eichten *et al.* and Allaby *et al.* data sets shown in Table 3.2 are used. The Be target data is used. The Eichten *et al.* data is used in regions not covered by the NA61/SHINE data, and the Allaby *et al.* data is used in regions not covered by either the NA61/SHINE data or the Eichten *et al.* data. Since these data sets only measure the differential production at points that cover

a small momentum and angular range, a bi-cubic spline fit is used to interpolate between the data points. Both the target material and incident proton momentum of these data sets differ to the T2K setup. Tuning weights are derived by comparing the Be target data to the FLUKA prediction for Be at the appropriate incident proton momentum. These weights are then applied to the simulated production events assuming Feynman scaling; the calculated weights are converted to the x_{F} - p_{T} space and applied according to the x_{F} and p_{T} of the simulated events.

The tertiary kaon production is also tuned. As for the tertiary pion production, this involves extrapolating the data down to lower incident nucleon momenta. Tuning weights calculated separately from the NA61/SHINE, Eichten *et al.* and Allaby *et al.* data sets are each scaled to the appropriate incident momentum assuming Feynman scaling. A linear interpolation is then performed between a pair of weights (used in the same order as for the tuning of the secondary kaon production), according to the momentum of the secondary proton and the beam momenta of the data sets.

Kaon production from interactions outside of the target material are not currently tuned. However, an uncertainty is assigned to the flux prediction to account for this, as described later in Section 3.3.1.

In addition to tuning the pion and kaon production, the total cross-sections for p, $\pi^{\pm}$ and $K^{\pm}$ interactions on both carbon and aluminium are tuned using the data sets summarised in Table 3.3. The quantity that is tuned is the production cross-section, σ_{prod}. As described in Section 3.2.2, this is defined as the inelastic cross-section, σ_{inel}, minus the quasi-elastic component. As shown in Table 3.3, some experiments measure σ_{prod} whereas others measure σ_{inel}. It is therefore necessary to subtract the quasi-elastic cross-section from the measurements of σ_{inel}. To do this, an empirical dependence based on the work by Belletini *et al.* [74] is used to calculate the quasi-elastic cross-section of hadron h on a target material with atomic mass number A, σ_h^{qe}:

$$\sigma_h^{\mathrm{qe}} = 0.8(\sigma_{hp}^{\mathrm{el}} + \sigma_{hn}^{\mathrm{el}})A^{1/3}, \tag{3.6}$$

where $\sigma_{hp}^{\mathrm{el}}$ and $\sigma_{hn}^{\mathrm{el}}$ are the elastic cross-sections of the hadron h on the proton and neutron respectively.[2] The uncertainty assigned to this quasi-elastic subtraction will

[2]The formula derived by Belletini *et al.* is modified to include the average of the elastic cross-sections on the proton and neutron instead of on the proton only.

be discussed in Section 3.3.1. Comparisons of σ_{prod} are made between the data and MC predictions of both FLUKA and GCALOR, for interactions on both carbon and aluminium. It is observed that FLUKA is in good agreement with the data and so no tuning of the σ_{prod} values is applied to the FLUKA simulation of interactions in the target. GCALOR, on the other hand, shows significant disagreements with the data and so the σ_{prod} values used in the GCALOR simulation of interactions outside of the target are tuned to the FLUKA values.

The total effect of the hadronic interaction tuning on the neutrino flux prediction is shown in Figure 3.3. It can be seen that the pion production tuning has the largest effect at low neutrino energy, whereas the kaon production tuning is dominant at high energy. The final neutrino flux prediction at the ND280 detector, after the hadronic interaction tuning has been applied, is shown in Figure 3.4 for the four different flavours separately (ν_μ, $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$) and with each flavour broken down by neutrino parent. The same plots are shown again in Figure 3.5 but with the flux shown on a logarithmic scale.

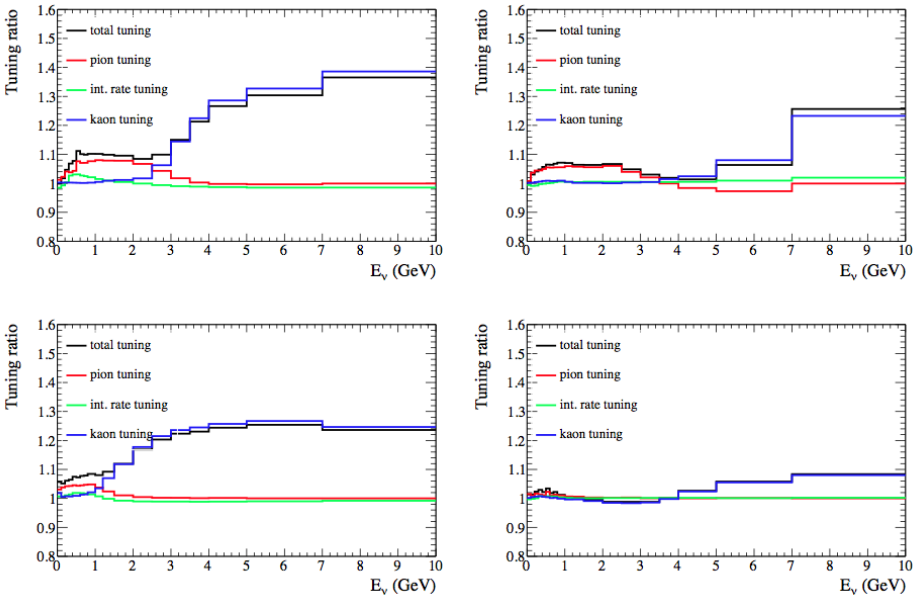


Figure 3.3: Ratio of the T2K neutrino flux prediction with and without the hadronic interaction tuning applied (described in Section 3.2) as a function of neutrino energy for: ν_μ (top left), $\bar{\nu}_\mu$ (top right), ν_e (bottom left) and $\bar{\nu}_e$ (bottom right). Plots taken from [65].

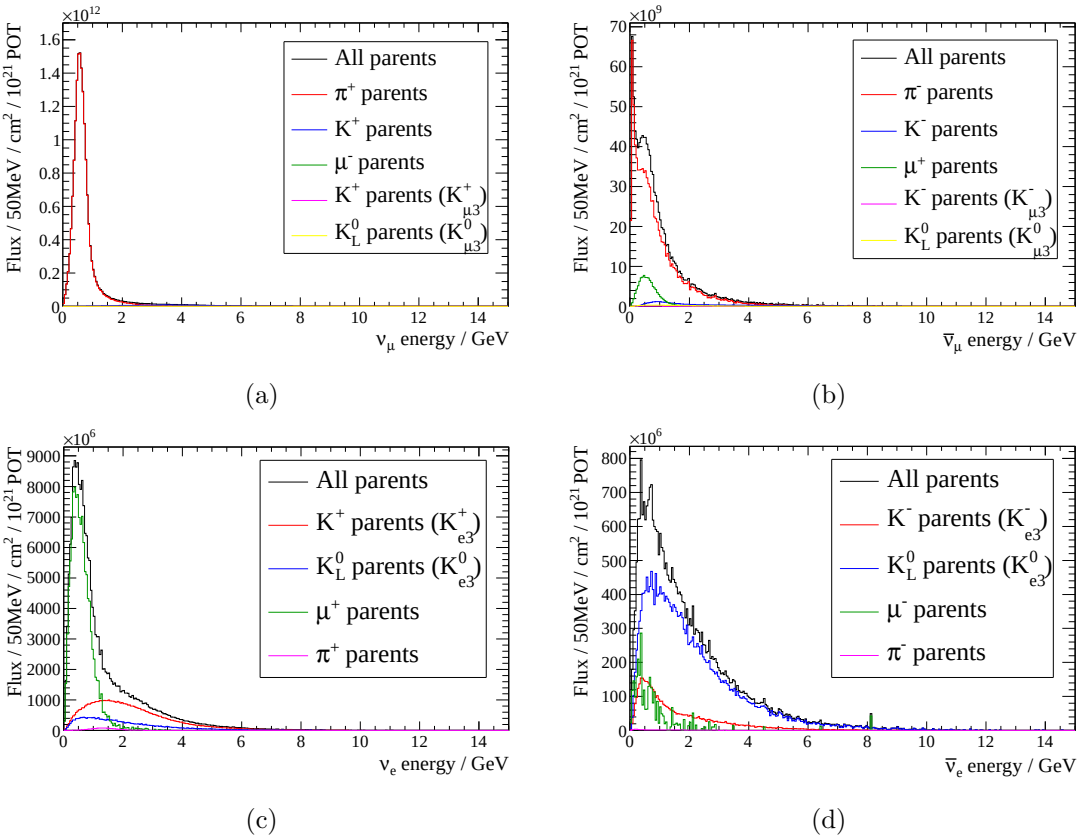


Figure 3.4: The final T2K neutrino flux prediction at the ND280 detector, after the hadronic interaction tuning has been applied, as a function of neutrino energy with contributions from: (a) ν_μ, (b) $\bar{\nu}_\mu$, (c) ν_e, (d) $\bar{\nu}_e$. The different parents for each neutrino type are shown. The equivalent plots for the flux at Super-Kamiokande can be found in [65].

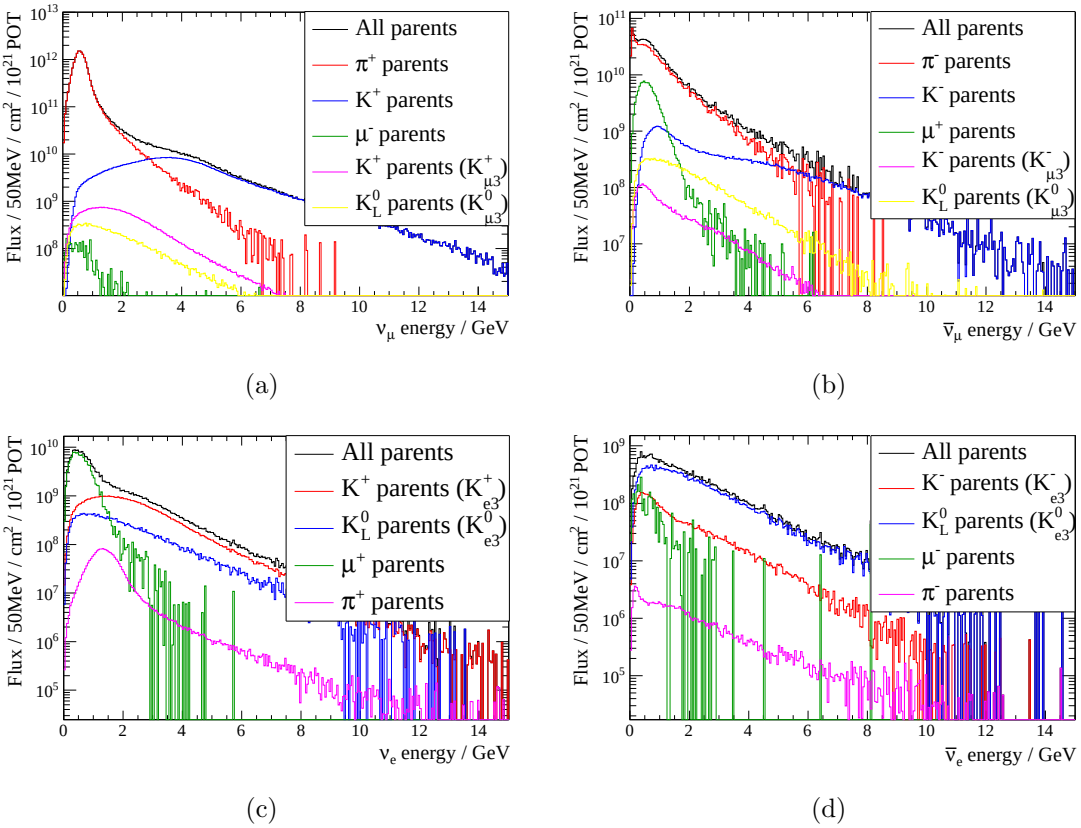


Figure 3.5: The same plots as in Figure 3.4 but with a logarithmic scale on the y axis of each plot.

3.3 Uncertainties on the flux prediction

This section describes the uncertainties on the flux prediction for T2K, which arise from the following sources:

- Errors on the hadronic interactions inside and outside of the target
- Errors on the position and direction of the proton beam, and errors on the off-axis angle of the neutrino beam
- Errors on the alignment of the target and horns
- Errors on the current and magnetic field of the horns

Each error source will be described in turn in the following sections. The details given here are summarised from [65] and [66]; full details should be sought in these references. The relative impacts of these errors on the neutrino flux will be shown in Section 3.3.5.

The dominant flux uncertainty is due to the hadronic interactions. Since there is currently no simple underlying parametrisation to describe the uncertainties on the NA61 inputs and other hadron production data, the flux uncertainty is simply parametrised in bins of true neutrino energy. A covariance matrix is evaluated for each error source between energy bins for neutrinos of different flavours and at the two off-axis detectors (ND280 and Super-Kamiokande). The final flux covariance matrix is obtained by simply summing the individual matrices and will be shown in Section 3.3.5.

To evaluate the covariance matrix for each error source, parameters within the beam simulation (for example, the hadron production model, or the proton beam profile) are varied according to their errors. The correlated changes in the neutrino energy bins are then used to calculate the covariance matrix.

For flux uncertainties represented by a set of correlated parameters, throws of the parameters are made according to their covariance. Each throw gives a new flux prediction. The elements of the fractional covariance matrix, V, can then be calculated as:

$$V_{ij} = \frac{1}{N_{\text{throws}}\phi_i^{\text{nom}}\phi_j^{\text{nom}}} \sum_{k=1}^{N_{\text{throws}}} (\phi_i^k - \phi_i^{\text{nom}})(\phi_j^k - \phi_j^{\text{nom}}), \tag{3.7}$$

where N_{throws} is the number of throws, ϕ_i^k is the flux in the ith neutrino energy, flavour and detector bin for the kth throw of the parameters, ϕ_j^k is the same but for the jth bin, and ϕ_i^{nom} and ϕ_j^{nom} are the nominal fluxes in the ith and jth bins respectively.

For flux uncertainties that are represented by a single parameter, the change in the flux under variations of the parameter of $\pm 1\sigma$ are used to calculate the covariance matrix. Equation 3.7 can be applied but with $N_{\text{throws}} = 2$, where "throws" $k = 1$ and $k = 2$ represent variations of the parameter of $+1\sigma$ and -1σ respectively. Finally, some flux uncertainties are evaluated by, for example, switching from one model to another or including an extra level of tuning. Equation 3.7 can again be applied to calculate the flux covariance, but with $N_{\text{throws}} = 1$ where the single "throw" represents the flux with this new model (or other change) implemented.

3.3.1 Hadronic interaction uncertainties

The uncertainties on the kaon production tuning will be discussed, followed by the uncertainties on the pion production tuning. The uncertainties on the production of secondary nucleons (which then interact to give tertiary kaons and pions) will also be discussed. Finally, the uncertainties on the tuning of the total cross-sections for p, $\pi^{\pm}$ and $K^{\pm}$ interactions will be described. The relative impacts of these uncertainties on the neutrino flux will be shown in Section 3.3.5.

Uncertainties on kaon production

The following sources of uncertainty on the kaon production tuning have been considered:

1. Uncertainties associated with the data

   Statistical and systematic errors on the data bins (in p-θ of the produced kaon), as well as an overall normalisation uncertainty, are considered for each of the three data sets (NA61/SHINE, Eichten *et al.* and Allaby *et al.*) used to tune the kaon production. Correlations between the bins in a given data set are considered where appropriate. A set of throws is made varying the data bins according to these

errors. Each throw gives a new set of kaon production tuning weights and therefore a new flux prediction. The flux covariance is calculated according to Equation 3.7.[3]

For the Eichten *et al.* and Allaby *et al.* data sets there is an additional error because the targets are made of beryllium rather than carbon. The uncertainty in the A-scaling, described in Section 3.2.3, is evaluated by scaling the beryllium target data to aluminium and comparing this to the aluminium target data. This is done for each experiment and meson (K^+ and K^-) separately. The uncertainties from this are included in the error of each bin when making the throws described above.

For the NA61/SHINE data, there is an additional error due to the coarse binning which is necessary due to the limited statistics of the kaon sample. A parametrised hadron production model proposed by Bonesini, Marchionni, Pietropaolo and Tabarelli de Fatis (BMPT) [82] is used to describe the shape of the data within each bin.[4] The change in the flux prediction when this shape information is considered is treated as an additional source of uncertainty.

2. Uncertainties associated with the Feynman scaling procedure

   Feynman scaling is used in the kaon production tuning both to scale the Eichten *et al.* and Allaby *et al.* data sets to the T2K proton beam momentum and to extrapolate down to lower nucleon momenta for tertiary kaon production, as described in Section 3.2.3. An alternative scaling variable, x_{R} [85], is used instead of x_{F} and the change in the neutrino flux is used to calculate the flux covariance. In addition, the Allaby *et al.* data at 19.2 GeV/c is scaled to 24 GeV/c and compared to the Eichten *et al.* data. The discrepancy is propagated through to the tuning weights to give a new flux prediction, allowing the flux covariance to be calculated.

3. Uncertainties on kaon production outside of the phase space covered by data

   The BMPT parameterisation, described above, is used to extrapolate the data to the uncovered regions of kaon phase space. The kaon production is tuned including

[3]For the NA61/SHINE data set, the data bins are treated as completely uncorrelated since the statistical uncertainties dominate. Therefore, rather than taking throws of the data bins, each bin is simply raised by 1σ.

[4]The parameters for the BMPT model are obtained by simultaneously fitting all three data sets.

this extrapolated data, and the change in the flux prediction is used to generate the flux covariance.

4. Uncertainties on kaon production outside of the target material

   As mentioned in Section 3.2.3, kaon production from interactions outside of the target, mostly in the horn aluminium, is not currently tuned. The uncertainty due to these interactions is evaluated as follows. Tuning weights are assigned to GCALOR (which is used to model interactions outside of the target) by comparing the Eichten *et al.* data on the aluminium target to the prediction. The flux covariance can then be calculated from the change in the flux prediction with and without this tuning. The process is also repeated using the NA61/SHINE data, scaled from carbon to aluminium using A-scaling.

Uncertainties on pion production

The following sources of uncertainty on the pion production tuning have been considered:

1. Uncertainties associated with the data

   The NA61/SHINE $\pi^{\pm}$ data, binned by p-θ of the produced pion, is used to tune the pion production, as described in Section 3.2.3. The errors on the data bins are dominated by systematic uncertainties, which are discussed in detail in [67]. The dominant sources of systematic error are the correction for the feed-down from strange particle decays and the particle identification. Correlations of the systematic errors between bins are considered. A set of throws is made varying the data bins according to these errors and correlations. Each throw gives a new set of pion production tuning weights and therefore a new flux prediction. The flux covariance can then be calculated according to Equation 3.7.

   This data is also used to tune the pion production in the horn material by scaling the data from carbon to aluminium, as described previously. The uncertainty on A-scaling is evaluated in the same way as for the kaon production, but using the pion instead of kaon data from Eichten *et al.* and Allaby *et al.*; the beryllium target data is scaled to aluminium and compared to the aluminium target data for each

meson (π^+ and π^-) and experiment separately. The uncertainties from this are included in the error of each scaled data bin.

2. Uncertainties associated with the tertiary pion tuning

   In order to tune tertiary pion production, the NA61/SHINE data is extrapolated down to lower incident nucleon momenta using Feynman scaling. An alternative method of tuning is performed and the change in the flux prediction is used to calculate the flux covariance. Two data sets from the E910 experiment, with proton beam momenta of 12.3 GeV/c and 17.5 GeV/c, provide an alternative source for tuning interactions of lower momentum nucleons.[5] Tuning weights are calculated for each data set, and a linear interpolation is used for tertiary production events with an incident particle momentum between the two data sets.[6] This alternative method allows for breaking of the assumed x scaling and it uses data at lower incident particle momenta to guide the scaling violation.

3. Uncertainties on pion production outside of the phase space covered by data

   The NA61/SHINE data covers most of the phase space for secondary pion production that contributes to the T2K neutrino flux. To account for the uncovered phase space, the data is fitted with the previously mentioned BMPT parameterisation which is then used to extrapolate the data into the uncovered region. The pion production is tuned including this extrapolated data and the change in the flux prediction is used to generate the flux covariance.

Other hadronic interaction uncertainties

There are two additional sources of error on the hadronic interactions: uncertainties on secondary nucleon production and uncertainties on the total cross-sections of proton, pion and kaon interactions.

Interactions of the secondary protons and neutrons inside the target contribute approximately 16% and 5% respectively to the neutrino flux. The secondary protons can

[5] A beryllium target is used in the E910 experiment and so the data is scaled to carbon using A-scaling.

[6] Prior to the linear interpolation, the tuning weights of each data set are converted from the p-θ space (of the pion) to the x_{R}-p_{T} phase space. The interpolated weight can then be applied to the tertiary production events based on the x_{R} and p_{T} of those events.

be broadly divided into high momentum protons, which are thought to be produced by quasi-elastic scattering of the primary protons with a small momentum transfer, and lower momentum protons produced by inelastic scattering of the primary protons in which hadrons are produced. For the secondary neutrons, only low momentum neutrons contribute significantly to the flux. The uncertainty of secondary proton production in the low momentum region is taken as the discrepancy between the FLUKA predictions and the Eichten *et al.* and Allaby *et al.* data sets. The same uncertainty is used for secondary neutron production, assuming isospin invariance. For the secondary proton production in the high momentum region, a 100% error is assigned due to the lack of relevant data. The secondary nucleon production is varied within these errors and the change in the predicted flux is used to generate the flux covariance.

The total cross-sections of p, $\pi^{\pm}$ and $K^{\pm}$ interactions on both carbon and aluminium are tuned using the data sets summarised in Table 3.3, as discussed previously in Section 3.2.3. Some of the experiments measure the production cross-section, σ_{prod}, defined as the inelastic cross-section minus the quasi-elastic component, whereas others directly measure the inelastic cross-section. It is therefore necessary in the tuning to subtract the quasi-elastic cross-section from the measurements of the inelastic cross-section, as described previously. Discrepancies are observed between the measurements of σ_{prod} for the proton by Denisov *et al.* [79] and those of Bellettini *et al.* [74], Carroll *et al.* [76] and NA61 [67], where the measurements of Denisov *et al.* and Bellettini *et al.* have been corrected for the quasi-elastic cross-section. This discrepancy, on the order of 10%, may be due to the difficulty of understanding which cross-section (production or inelastic) some of the experiments have measured. Therefore, a conservative approach is taken, using the size of the quasi-elastic correction as the uncertainty on the total cross-section of each particle. The total cross-sections are varied within these errors and the change in the predicted flux is used to generate the flux covariance.

3.3.2 Proton beam and off-axis angle

The position and angle of the proton beam centre at the baffle are reconstructed using measurements from the monitoring devices surrounding the primary beamline (ESMs and SSEMs) and the OTR monitor which lies just downstream of the baffle in the

secondary beamline (see Section 2.1.3 for a description of the beamline). The following sources of error are considered:

- Alignment uncertainty for the monitoring devices
- Alignment uncertainty between the primary beamline and secondary beamline (containing the target)
- Systematic errors in the measurements of the monitoring devices

Correlations between the errors on the position and angle are considered. A set of throws of the proton beam properties is made according to these errors and correlations, with each throw generating a new flux prediction. The flux covariance can then be calculated according to Equation 3.7. In addition, uncertainties in the intensity of the proton beam (measured to within 2% by the CTs, as described in Section 2.1.3) causes an error on the absolute flux normalisation.

As mentioned previously in Section 2.2.2, the INGRID detector monitors the on-axis direction of the neutrino beam at a position of 280 m from the target, which is done to a precision of better than 0.4 mrad. The positions of the ND280 and Super-Kamiokande detectors are changed in JNUBEAM by the off-axis angle uncertainty, both closer to and away from the beamline axis. The changes in the flux prediction can be used to calculate the flux covariance.

3.3.3 Target and horn alignment

The alignment of the beamline, including the target and horns, was done based on markers from GPS surveys. The effect on the neutrino flux of the systematic uncertainties in the target and horn alignments is evaluated by applying displacements and rotations in JNUBEAM. The target is rotated in both the horizontal (parallel to the beam) and vertical planes. For the horns, displacements are applied along each coordinate axis as well as rotations in both the horizontal and vertical planes. The flux covariance can be calculated from the changes in the flux prediction when these displacements and rotations are applied.

3.3.4 Current and magnetic field of horns

The electrical current of the horns is monitored using Rogowski coils and FADC electronics. Various systematic uncertainties are associated with these measurements including calibration of the coils and electronics. A total error of 2% (5 kA) is taken for the absolute horn current. The current is altered by ±5 kA in JNUBEAM, and the changes in the flux prediction are used to calculate the flux covariance.

The magnetic field in the horns is assumed to vary as $1/r$, where r is the distance from the horn axis, as described in Section 3.1. Multiple measurements have been performed to validate this model and to check for asymmetries in the magnetic field. These measurements, taken using Hall probes inserted via the instrumentation ports of the horns, find that the measured field agrees with the expected field to within 1–2%. In addition, measurements of the magnetic field were taken with the spare T2K horn, which is identical to the first of the horns along the beamline. These measurements found an unexpected field along the horn axis inside the inner conductor. Although the exact nature of this on-axis field is still under investigation, it can be modelled inside JNUBEAM by assuming asymmetries in the horn current. The observed field asymmetries, including the on-axis field, are simulated within JNUBEAM. The change in the predicted neutrino flux is used to generate the flux covariance.

3.3.5 Summary of the flux uncertainties

The fractional error for the flux at ND280 as a function of neutrino energy is shown in Figure 3.6 for all four neutrino flavours; the error is broken down into the different sources described in Section 3.3. The equivalent plots for the flux at Super-Kamiokande can be found in [65]. The dominant sources of error are the hadronic uncertainties (pion and kaon production, secondary nucleon production and the total cross-sections).[7] There is also a significant contribution from the proton beam, off-axis and alignment errors around the peak of the ν_μ spectrum.

Figure 3.7(a) shows the total correlation matrix between energy bins for neutrinos of different flavours and at the two off-axis detectors (ND280 and Super-Kamiokande).

[7]The errors on the total cross-sections (for pion, proton and kaon interactions) are referred to in the plots as the error on the hadronic interaction length.

Figure 3.7(b) shows the corresponding fractional covariance matrix. The fractional covariance between energy bins i and j, denoted as c_{ij}, is related to the error on each bin, σ_i and σ_j respectively, the flux in each bin, ϕ_i and ϕ_j respectively, and the correlation between the two bins, ρ_{ij}, as follows:

$$c_{ij} = \rho_{ij} \frac{\sigma_i}{\phi_i} \frac{\sigma_j}{\phi_j}. \tag{3.8}$$

These matrices include all of the error sources. For both matrices, there are 20 E_ν bins for each flavour and detector spaced as follows in units of GeV:

$0.0-0.1$, $0.1-0.2$, $0.2-0.3$, $0.3-0.4$, $0.4-0.5$, $0.5-0.6$, $0.6-0.7$, $0.7-0.8$, $0.8-1.0$, $1.0-1.2$, $1.2-1.5$, $1.5-2.0$, $2.0-2.5$, $2.5-3.0$, $3.0-3.5$, $3.5-4.0$, $4.0-5.0$, $5.0-7.0$, $7.0-10.0$, >10.0.

There are significant correlations between ν_μs at ND280 and Super-Kamiokande, as expected. There are also strong correlations between ν_μs at ND280 and ν_es at Super-Kamiokande (and ND280), because the μ^+ from the $\pi^+ \to \mu^+ + \nu_\mu$ then decays to give a ν_e. This means that measurements of the ν_μ flux at ND280 can be used to constrain the ν_e contamination in the beam at Super-Kamiokande (a major background to the ν_e appearance measurement of T2K).

3.4 Tuning of the flux prediction using ND280 data

The flux prediction of T2K has been described in detail in the previous sections of this chapter. It is tuned and the uncertainties are evaluated using external data and measurements from the beamline monitoring devices, as described in Sections 3.2 and 3.3. In addition to this, measurements at the ND280 detector can be used to further constrain and validate the flux prediction. Two such measurements are described below in Sections 3.4.1 and 3.4.2.

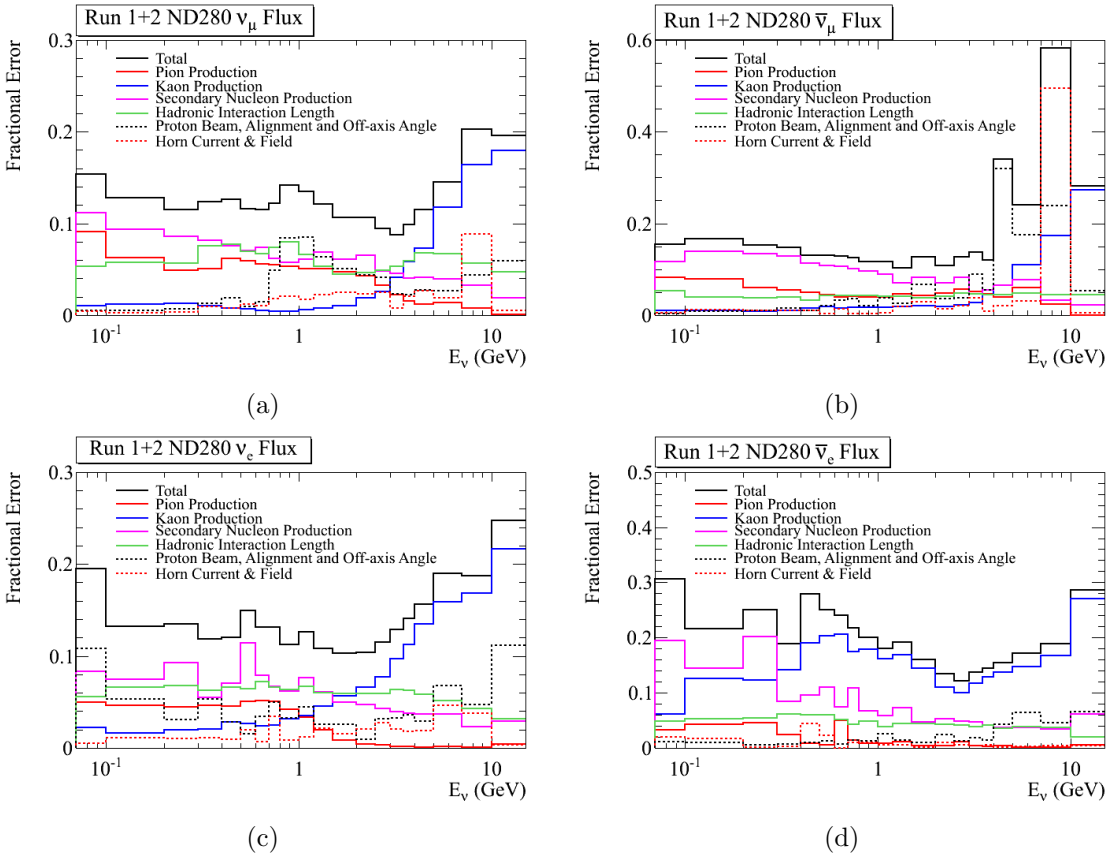


Figure 3.6: Fractional uncertainties for the flux at ND280 broken down into the different sources for: (a) ν_μ, (b) $\bar{\nu}_\mu$, (c) ν_e, and (d) $\bar{\nu}_e$ [52]. The equivalent plots for the flux at Super-Kamiokande can be found in [65].

3.4.1 Simultaneous constraint of the flux and cross-section parameters using ND280 data for the T2K oscillation analysis

For the latest T2K oscillation results [20], measurements from the off-axis ND280 near detector are used to simultaneously constrain the flux prediction and neutrino interaction cross-section models. This improves the prediction of the event rate and spectrum at the off-axis Super-Kamiokande far detector for the oscillation analysis. A sample of ν_μ charged-current events is selected from the ND280 data, and binned according to the reconstructed momentum and angle of the muon candidate (p_μ - $\cos\theta_\mu$ bins). The predicted number of events in each bin is a function of the flux and cross-section parameters, as well as parameters representing the systematic errors of the detector system. A maximum likelihood method is used to tune the Monte Carlo prediction to

the data. The likelihood includes prior constraints on the flux parameters according to the uncertainties detailed in Section 3.3. Prior constraints are also included on the cross-section parameters from cross-section measurements of external neutrino experiments such as MiniBooNE, SciBooNE and NOMAD. Full details of this methodology can be found in [20, 86, 87].

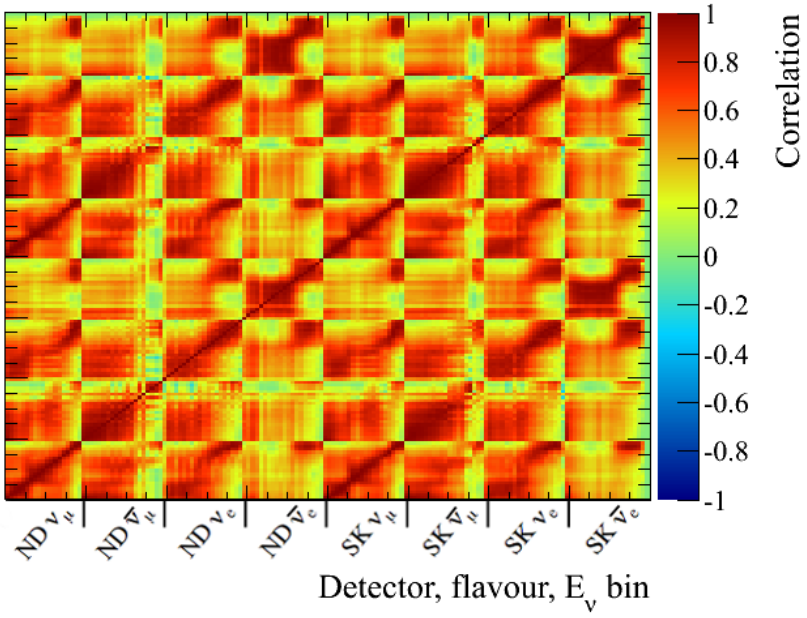
3.4.2 Validation of the flux prediction via a measurement of the normalisation of K^+-originating neutrinos at ND280

The main analysis of this thesis, which will be described in detail in Chapter 6, constrains the normalisation of K^+-originating neutrinos at the ND280 detector. This is a complementary analysis to that described in Section 3.4.1 and provides an independent validation of the T2K flux prediction. It probes a specific aspect of the beam simulation using ND280 data, namely the breakdown of neutrino parents at the ND280 detector. Since this is impacted by all of the uncertainties in the beam system that have been described in Section 3.3, as will be discussed below, this measurement provides a probe of the beam system and a validation tool for the beam simulation. The framework of the analysis described in Section 3.4.1 is used as a basis for this analysis; however, the two analyses are distinct in terms of the goals, the event selections and the fit implementations.[8] It should be noted that in order to provide an independent validation of the flux prediction, this analysis uses the flux calculation before the tuning described in Section 3.4.1 is included.

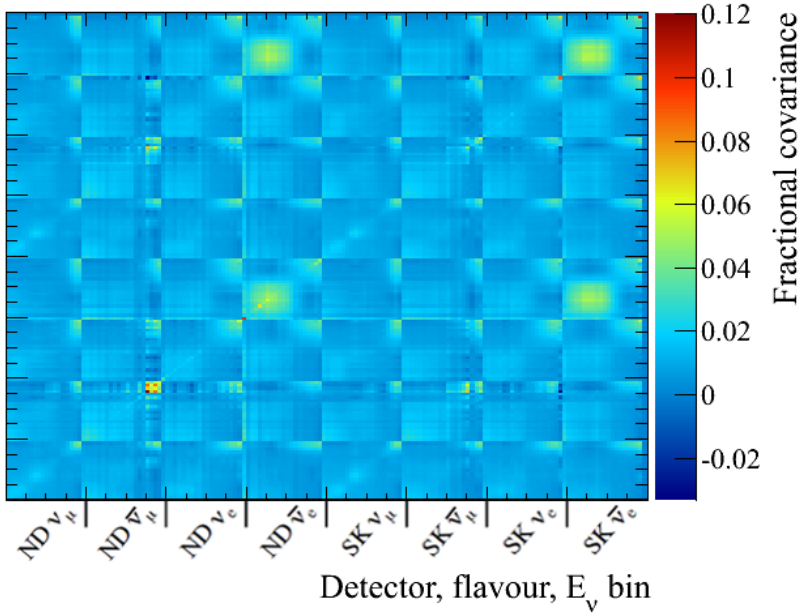
The effects of the flux uncertainties (detailed in Section 3.3) on the measurement of K^+-originating neutrinos at the ND280 detector are summarised as follows. Errors on the hadronic interactions include uncertainties on the kaon production and on the total cross-sections of proton and kaon interactions, making the impact of these errors on the measurement clear. Errors on the direction of the proton and neutrino beams directly affect the number, energy and parent breakdown of the neutrinos arriving at ND280. In addition, the alignment errors of the target and horns clearly affect the profile of

[8] Different reconstruction tools and selection cuts are used in this analysis to the select the ν_μ charged-current event sample. Also, different p_μ - $\cos\theta_\mu$ binning is used in order to probe the high energy part of the neutrino spectrum. Profile minimisation is used as opposed to global minimisation. Finally, a different treatment of the flux systematic errors is used to achieve an unconstrained measurement of the kaon normalisation. See Chapter 6 for a full description of this analysis.

the hadrons entering the decay volume, therefore affecting this measurement. Similarly, since the horns focus the charged hadrons as they leave the target, uncertainties in the current and magnetic field of the horns drive the number and energies of the K^+ particles whose daughter neutrinos arrive at the ND280. The flux covariance matrix, shown in Figure 3.7(b), will be used later to evaluate the impact of these uncertainties on this measurement.



(a)



(b)

Figure 3.7: Plots of (a) the correlation matrix, and (b) the fractional covariance matrix, between energy bins for neutrinos of different flavours and at the two off-axis detectors (ND280 and Super-Kamiokande). The fractional covariance is defined in the text. The binning on the y-axes is identical to the binning on the x-axes. For both plots, there are 20 E_ν bins within each detector and flavour group which follow the bin ordering defined in Section 3.3.5.

Chapter 4

ND280 offline software

In this chapter, a brief overview of the ND280 offline software is given. The main analysis of this thesis, presented in Chapter 6, depends crucially on the reconstruction algorithms of the ND280 detector and so particular attention is paid to describing these.

4.1 Overview

The ND280 offline software, based on the ROOT software libraries [88], encompasses the entire analysis chain; it reads in raw data and MC events, applies calibration routines and event reconstruction algorithms, and finally produces a simple summarised output ready for analysis users. A summary of the ND280 offline software suite is shown in Figure 4.1 for both data and MC events.

For data, the raw events are "unpacked" into the C++ based "oaEvent" format, which is the format used throughout the software until the final summarisation stage into simple ROOT trees. The unpacked data is then passed to the calibration packages which apply the relevant calibration routines. These calibrations are composed of both time-dependent routines, where the values of the parameters in the calibration change according to beam, detector or environmental conditions, and time-independent routines. For time-independent calibrations, the parameters can be hard-coded; for time-dependent calibrations, however, the parameters are stored in a database and accessed when the software chain is run. Part of the work presented in this thesis is based on developing calibration routines (both time-dependent and time-independent) for the

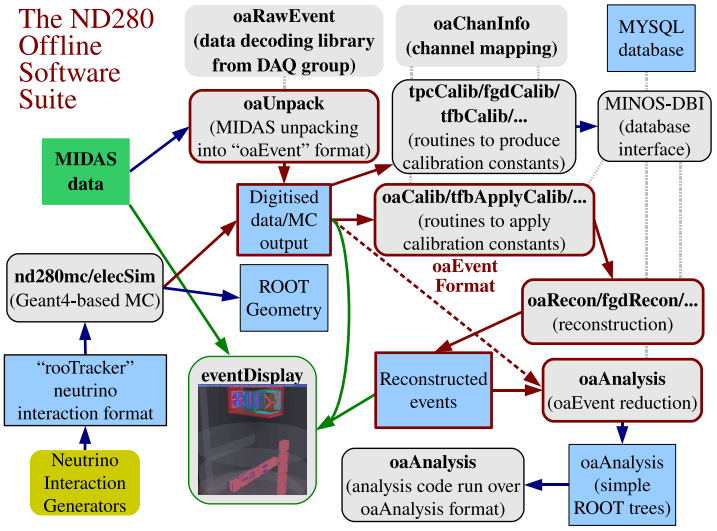


Figure 4.1: Schematic summarising the ND280 software showing the analysis chain for both raw data and MC events through to the final output analysis files. Figure taken from [50].

timestamp of events in the ND280, as will be discussed in Chapter 5. For each event, the output of the calibration stage is a list of "hits" which represent calibrated energy deposits at given positions in the detector at calibrated times. The calibrated hits are then passed onto the relevant sub-detector and global reconstruction packages. Finally, the detailed information from the reconstruction stage is summarised and distilled into ROOT-based "tree" structures ready for the end user to analyse. Also attached to this final output is extra information about the state of the beamline and detector at the time of data-taking, which can be used to determine whether the data is of analysable quality.

For MC, the input to the software chain is the simulated neutrino flux at the ND280 which, as described in Chapter 3, is generated using the JNUBEAM package. External neutrino interaction generators then simulate the interaction of the neutrinos with nuclei, using a ROOT-based description of the ND280 geometry. Both the NEUT [89] and GENIE [90] generators are fully integrated within the software framework. The output list of interactions is then input to the GEANT4-based [24, 25] detector simulation packages. The first of these packages, "nd280mc", simulates the trajectories and energy deposits of the particles produced by the neutrino interactions. The second package, "elecSim", takes the energy deposits simulated in the previous stage and models the

electronic response of the detector. At this stage, the MC events are in the same format as the "unpacked" raw data mentioned above. Therefore, from this point on, the MC events are handled in the same way as the data; they are passed through the calibration and reconstruction packages and finally to the summarisation stage ready for the end user to analyse.[1]

4.2 Reconstruction algorithms

The ND280 reconstruction begins with separate reconstruction of the hits in each sub-detector. There is a dedicated package for each sub-detector that is responsible for applying the specialised algorithms developed for that sub-detector to the calibrated hits. There is then a tracker reconstruction package responsible for combining tracks in the tracker region, defined in Section 2.2 as the three TPCs and two FGDs. Finally, the global reconstruction is performed. This is designed to combine the results from the individual and tracker reconstruction packages to form reconstructed objects that span the whole of the ND280. The analysis presented in Chapter 6 is based upon the output of the global reconstruction in the tracker region. Therefore, details of the TPC, FGD and tracker reconstruction packages are summarised in Section 4.2.1, followed by details of the global reconstruction in Section 4.2.2. Full details of these algorithms can be found in [91].

4.2.1 TPC, FGD and tracker reconstruction

The TPC reconstruction is run first. It uses separate methods for track finding and track fitting. To find tracks, TPC hits that are close in both position and time are clustered together and clusters are then joined into tracks via a pattern recognition algorithm. The track is then fitted; the likelihood of the observed charge distribution within the clusters is maximised to estimate the track parameters. In this way, the track coordinates, angles and curvature are found. In addition, the energy loss in the gas as a function of the track length is calculated. This is used for the purposes of particle identification and will be explained in more detail in the following section since the TPC particle identification is

[1] Only a subset of the calibration routines applied to data are applied to MC events, corresponding to those effects which have been simulated in the elecSim package.

central for the analysis in Chapter 6. The FGD reconstruction proceeds after the TPC reconstruction. A Kalman filter is used to incrementally match tracks from the TPC back to hits in the FGD. Any unmatched FGD hits from this stage, presumed to be from particles that start or stop in the FGD without entering a TPC, are then passed to specialised FGD-only reconstruction algorithms. The TPC-FGD matching will only create tracks that span a single TPC. The tracker reconstruction therefore joins tracks that have traversed multiple TPCs.

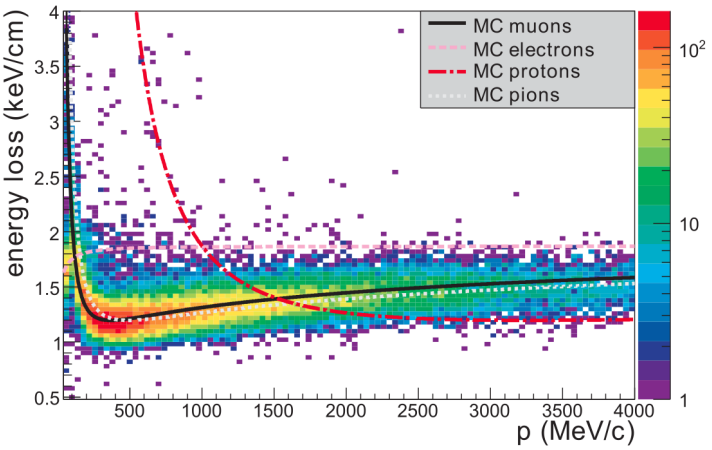
TPC particle identification

The TPC particle identification is key to the analysis in Chapter 6 and so is discussed in more detail here. The general principle is as follows. The energy loss of the particle is measured and can then be compared to the expected value based on different particle hypotheses and the measured momentum of the track.

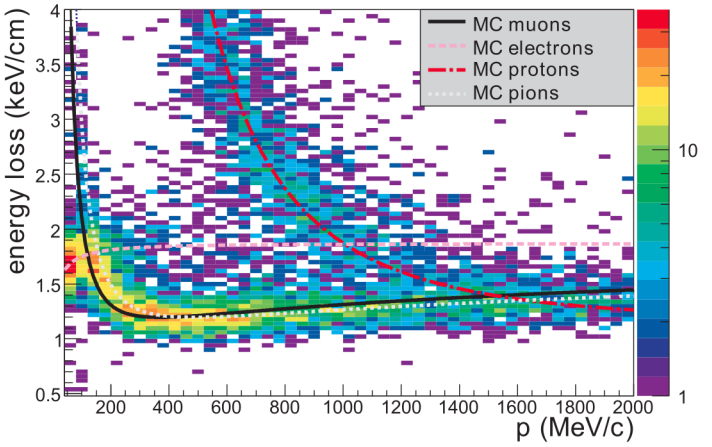
As mentioned previously, TPC hits that are close in both position and time are clustered together and then clusters are joined into a track. Each cluster has an associated total charge, obtained from the sum of charges of the individual hits. A truncated mean charge is then calculated for the TPC track; this is defined as the mean of the lowest 70% of cluster charges, an optimised approach found through Monte Carlo simulation and test beam studies. This truncated mean charge represents the energy loss of the particle in the gas. Figure 4.2 shows the measured energy loss as a function of the measured momentum of the particle. The expected curves for different particle hypotheses are also shown. For each particle hypothesis (muon, electron, proton, pion) a "pull" quantity can then be defined. The pull is a measure of the number of standard deviations that the measured energy loss, $(dE/dx)^{\text{measured}}$, is away from the expected value for a given particle type α at the observed momentum, $(dE/dx)^{\alpha}_{\text{expected}}$. It is given by:

$$\text{Pull}^{\alpha} = \frac{(dE/dx)_{\text{measured}} - (dE/dx)^{\alpha}_{\text{expected}}}{\sigma^{(dE/dx)_{\text{measured}} - (dE/dx)^{\alpha}_{\text{expected}}}}, \tag{4.1}$$

where $\alpha = e, \mu, \pi$ or p. The pull variable is used in the event selection of the analysis presented in Chapter 6. Further details of the TPC particle identification can be found in [92, 93].



(a)



(b)

Figure 4.2: Distribution of the energy loss as a function of the momentum for particles produced in neutrino interactions, compared to the expected curves for muons, electrons, protons and pions. (a) and (b) show negatively and positively charged particles respectively. Figure taken from [92].

4.2.2 Global reconstruction

The global reconstruction combines the results from the lower-level reconstruction packages to form reconstructed objects that traverse the entire ND280. It starts by attempting to match objects from the tracker reconstruction to objects in the neighbouring sub-detectors. It does this by extrapolating the tracker object to the entrance plane of a given sub-detector and searching for objects in that sub-detector near the entry point. A χ^2 fit is performed based on the positions and relative directions of the two objects and the uncertainties on these quantities. If $\chi^2 < 100$ (or $\chi^2 < 200$ for matching to P0D or SMRD objects) then the objects are matched together. The objects must also be within 300 ns of each other. A Kalman filter is then used to refit the new matched object to recalculate its position, direction and momentum, taking into account the reconstructed momentum of the two separate objects and the expected energy loss. These new larger objects are matched recursively until no more objects can be combined together. The global reconstruction then takes the so far unmatched objects in the non-tracker sub-detectors and attempts to match them together (for example, P0D objects may be matched with SMRD objects). Again this is done recursively until no more matches can be found.

Global vertexing

The global vertexing algorithm uses the output of the global reconstruction. The stages of the vertexing algorithm can be briefly summarised as follows (further details should be sought in [94]):

- A preliminary track clustering stage is performed to provide a first guess for the tracks associated with the vertex. The closest point of approach between a pair of tracks in the XZ plane (since the tracks are unperturbed by the magnetic field in this plane) is used to cluster tracks. An iterative procedure is performed to add tracks to an existing cluster.
- A Kalman filtering technique is then used to decide which of the clustered tracks are actually associated with the vertex. A cut is applied based on the effect that including each track has on the overall χ^2 of the Kalman filter.

- If the vertex from the previous stage contains more than one track, an inverse Kalman filter is used to test the effect of removing each individual track. The χ^2 contribution from each track is determined and a cut on this quantity is used to remove tracks.

- If more than one vertex is found then the vertexing algorithm attempts to identify the primary vertex. A vertex is automatically flagged as secondary if it is positioned more than 10 cm from the start or end points of any of the tracks in that vertex. From the remaining vertices, the vertex which has the highest momentum track associated with it is flagged as the primary vertex. If two or more of the vertices share the highest momentum track, then the most upstream one is selected as the primary vertex.

- In the case of a single track event or no cluster being found, the most upstream end of the highest momentum track is used as the vertex. This vertex is by definition the primary vertex in this case.

Preliminary validations of this global vertexing tool can be found in [94] and [95]. In particular, consistency checks between the performance of the algorithm in data and MC are presented in [95]. This ensures that no obvious bias is introduced into an analysis by using this tool. These checks include quantities used within the algorithm (for example, the χ^2 of the Kalman filter) as well as the efficiency of the algorithm in associating the correct tracks with a vertex.[2] In general, there is good data-MC agreement. Work is ongoing within the relevant working group to continue validation and evaluation of the systematics for this algorithm. The global vertexing tool is used in the analysis in Chapter 6; further discussion of the performance and data-MC agreement in the context of the analysis are given there.

It should be noted that the global vertexing algorithm in the software version that is used for the analysis in Chapter 6 uses the output of the tracker reconstruction rather than the global reconstruction.[3] This can mean that for a single track event where the track traverses the entire ND280, the vertex position is reconstructed at the entry

[2] An approximate vertexing efficiency is used since it is not possible to know the true efficiency for data. This compares the number of tracks associated to the vertex by the vertexing algorithm with the number of tracks within a given radius of the vertex.

[3] The tracks associated to the vertex are matched up to the corresponding global tracks at a later stage.

point to the tracker (the upstream end of the first TPC) rather than at the start of the ND280. However, this is not an issue for the analysis in Chapter 6 since the vertex must be reconstructed inside an FGD. In addition, tracks with no tracker component will not be matched to the vertex. However, since the analysis in Chapter 6 is based upon ν_μ charged-current interactions in an FGD, it is very unlikely that the emitted muon does not have an FGD or TPC component.

Figure 4.3 shows the hits for a data event along with the globally reconstructed tracks and vertex. It can be seen that the global reconstruction successfully matches tracks between different sub-detectors. In addition, the global vertexing algorithm successfully associates the two individual global tracks to a single vertex.

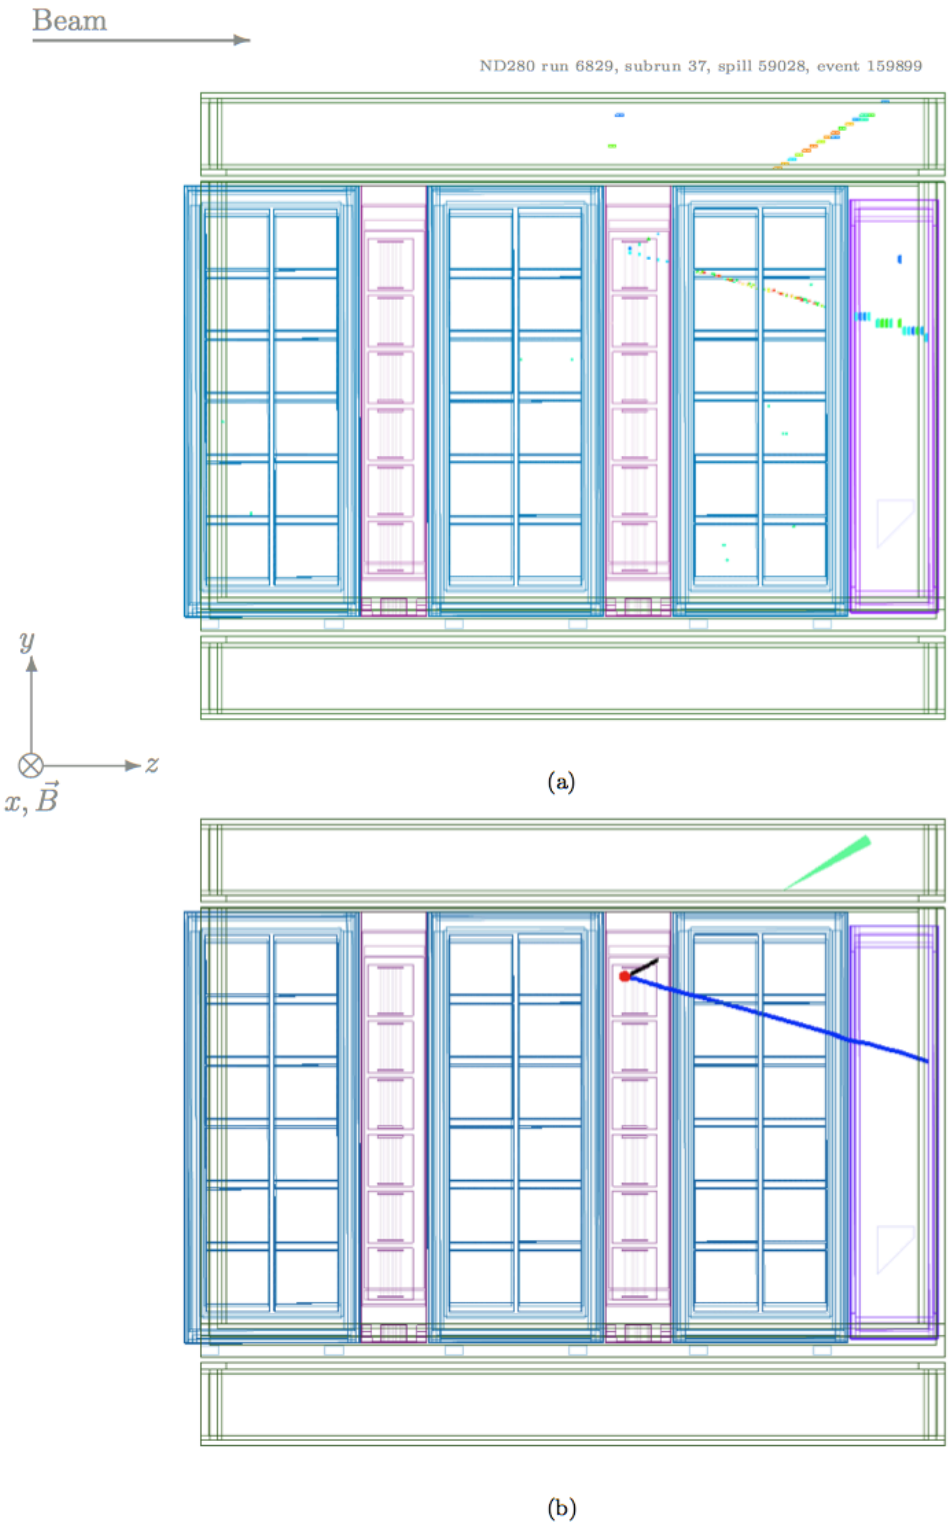


Figure 4.3: Event display for a data event, showing the tracker region and surrounding ECals. Calibrated hits are shown in (a) and the globally reconstructed vertex and tracks are shown in (b). The global reconstruction successfully matches the ECal (green) and FGD-only (black) sub-tracks, as well as the sub-detector tracks for the negative muon candidate (blue). The global vertexing algorithm associates the global tracks to a vertex in FGD2 (red circle). Figure and caption (adapted) taken from [95].

Chapter 5

Time calibration of the ND280

This chapter presents some major aspects of the first time calibration of the ND280 detector, which were performed by the author. This initial calibration was applied to the Run 1 data which, as described in Section 2.1.4, was taken from March to June 2010. The studies presented here mainly focus on the DsECal module of the ECal sub-detector. As described previously in Section 2.2.3, the ECal is composed of 13 modules. However, the DsECal was the first module to be installed and was the only one used in the Run 1 data. Some results are also shown for the P0D sub-detector.

The full time calibration chain for the Trip-t detectors, including the DsECal and P0D, includes corrections for the following effects:

- Time offsets between electronic boards
- Electronics timewalk
- Fibre timewalk

Section 5.1 describes the calibration of the time offsets between electronic boards in the DsECal due to differences in cable length using cosmic muon events. Section 5.2 presents the calibration of the fibre timewalk effect in both the DsECal and P0D. The calibration of the electronics timewalk effect is performed by people other than the author. It is mentioned here, however, because it is included when discussing, in Section 5.3, the achieved time resolution of the DsECal after all the time calibration methods are applied. In brief, the electronics timewalk effect arises because small charge deposits in the detector have a longer rise time in the electronics than large charge deposits and

therefore have an associated delay in their timestamp. Further details on the calibration of this effect, as well as details on the calibration of the other sub-detectors, can be found in [96].

The studies in this chapter are presented in the context of the knowledge that was available at the time of the work being performed. However, a brief update is given at the end of each section to describe improvements or further developments that have been made after the completion of the studies presented here.

An overview of the readout electronics has been given previously in Section 2.2.3, including the time-stamping process for energy deposits, or hits, in the detector. The main data used to develop the calibration methodologies are DsECal-only triggered cosmic runs. When running the detector in this mode, the cosmic trigger module, or CTM, only receives signals from the DsECal front-end boards, or TFBs. Based on these input signals, the CTM decides whether there was a cosmic muon event in the detector and if necessary passes a trigger signal to the master clock module, or MCM, to begin reading out the TFBs. The important time quantity for each hit in the detector is the "time relative to trigger". This quantity is used throughout this chapter and is defined as the time-stamp of the hit less the time at which the MCM issued the trigger signal. Since the trigger signal is by definition issued after the cosmic muon has passed through the detector, this quantity is always negative.

5.1 Time offsets calibration for the DsECal

This section describes the calibration of the time offsets within the DsECal module due to differences in cable lengths between the electronic boards. Cosmic muon data taken with the DsECal-only trigger configuration (described above) is used for this calibration.

5.1.1 Sources of time offsets

A simplified view of the ND280 readout electronics is shown in Figure 5.1 (see Section 2.2.3 for full details), along with the sources of time offsets. The DsECal module is read out by 56 TFBs, which are connected to 2 RMMs (28 TFBs on each), which are

in turn connected to the ECal SCM. Within this subsystem, there are several sources of offset that must be calibrated in order to align the times from each front-end channel:

- TFB offsets
- RMM offsets
- Rocket I/O shifts.

The TFB offsets are due to the different cable lengths between the RMMs and TFBs. The RMM offsets are due to the different cable lengths between the ECal SCM and the RMMs. Rocket I/O shifts occur between boards which are linked by an optical connection. The boards use the Rocket I/O protocol to transmit data along the optical link, which can give rise to a phase shift between the clocks on the boards when one or both of the boards is powercycled.[1] The clock signal is propagated from the MCM to the TFBs via two optical links, MCM→ SCM→RMM, and so the time relative to trigger of a hit can change by ±10 ns when the system is powercycled (±5 ns for each link). Since it is not possible to independently measure the Rocket I/O shifts, they are implicitly included in the measured RMM offsets. Therefore, the RMM offsets will change if the system is powercycled. In contrast, the TFB offsets are stable with time, since they arise from the electrical RMM→TFB connections. It is worth noting that these offsets only align the times internally within the DsECal module; further "cross-detector" offsets must be applied in order to align the DsECal with the other subdetectors, as described in [96].

5.1.2 Calibration methodology

The TFB and RMM offsets are calculated by lining up the time peak of cosmic ray muons across the different boards. The TFB offsets line up all TFBs on a given RMM to the average time on that RMM. For the ith TFB on the jth RMM, the TFB offset,

[1] It has later been found that Rocket I/O shifts occur not only when the boards linked by an optical connection are powercycled, but they can also occur if the optical link unexpectedly drops out and reestablishes connection, as will be discussed in Section 5.1.5.

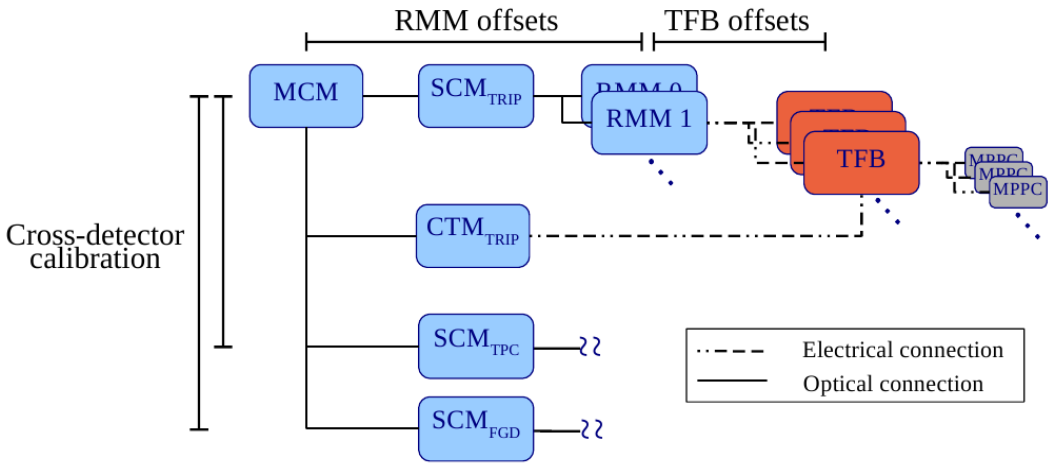


Figure 5.1: A simplified view of the ND280 readout electronics (see Figure 2.14 for the full setup). The sources of time offsets are shown. The TFB offsets are due to the different cable lengths between the RMMs and TFBs. The RMM offsets are due to the different cable lengths between the SCM and RMMs, and also include any Rocket I/O shifts due to the optical connections in the MCM→ SCM→RMM chain. The cross-detector time calibration, described in [96], aligns the times of the different sub-detectors and includes a correction for the different MCM→SCM cable lengths.

O_{ij}^{TFB}, is calculated as:

$$O_{ij}^{\mathrm{TFB}} = \mu_{ij} - \frac{\sum_{i=0}^{N_j-1} \mu_{ij}}{N_j} = \mu_{ij} - \bar{\mu}_j \tag{5.1}$$

where μ_{ij} is the mean of a Gaussian fit to the time distribution, N_j is the number of TFBs on the jth RMM and $\bar{\mu}_j$ is the average of the μ_{ij} values across all TFBs on the jth RMM. It is preferred to line up the TFBs to $\bar{\mu}_j$, i.e. the average time on the RMM, rather than to a single reference TFB, to protect against unexpected behaviour of a reference TFB. The time distribution for an example TFB is shown in Figure 5.2. The RMM offsets line up the RMMs to the average time across them. For the jth RMM, the RMM offset, O_j^{RMM}, is calculated as:

$$O_j^{\mathrm{RMM}} = \bar{\mu}_j - \frac{\sum_{j=0}^{M-1} \bar{\mu}_j}{M} \tag{5.2}$$

where M is the number of RMMs and all other symbols are as defined in Equation 5.1. The total correction for the ith TFB on the jth RMM, denoted as O_{ij}^{T}, is then simply given by:

$$O_{ij}^{\mathrm{T}} = O_{ij}^{\mathrm{TFB}} + O_j^{\mathrm{RMM}}. \tag{5.3}$$

The offsets are calculated using raw data from cosmic runs taken with the DsECal-only trigger configuration, with a charge cut of 400 ADC. The charge cut aims to minimise both the fibre timewalk effect (which will be described in Section 5.2) and the electronics timewalk effect (described at the start of the chapter), both of which are largest at low charge, and also to reject noise hits. The trigger configuration works as follows: if any one of the TFBs in the downstream half of the DsECal has more than one hit then a trigger will occur. This trigger gives good illumination across all of the TFBs. In addition, it has no significant position or directional bias that affects this calibration (the latter point is discussed in more detail in Section 5.1.3).

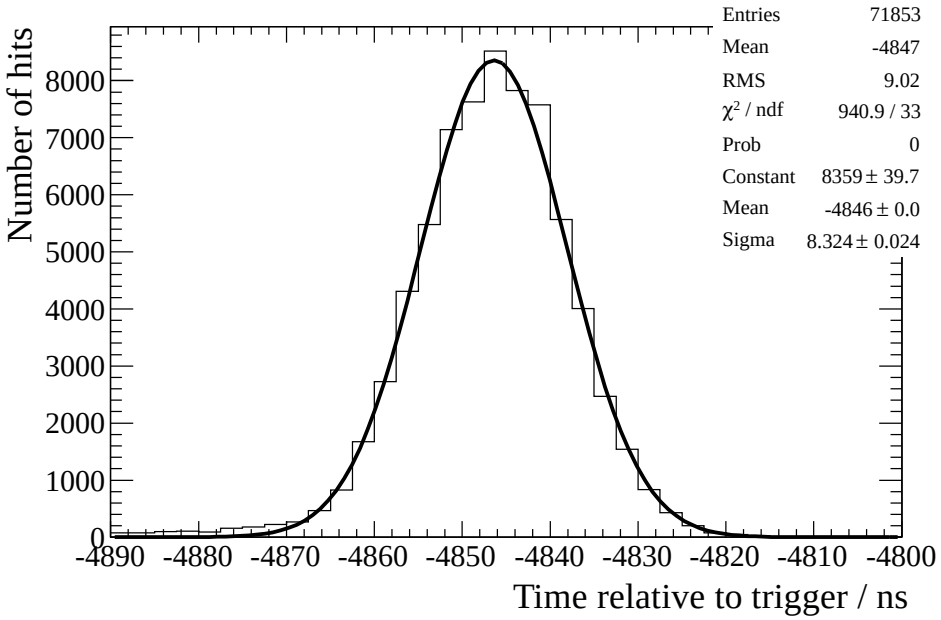


Figure 5.2: Time distribution for TFB 0 RMM 0 for ND280 run 5007 (a DsECal-only triggered cosmic run), fitted with a Gaussian. The mean of this Gaussian is μ_{00} in Equation 5.1.

The time of flight of cosmics through the DsECal is neglected. To a first approximation, the distribution of cosmics travelling through the DsECal is symmetric in the azimuthal direction. This causes a cancelling of the time of flight effect on average for the TFBs. However, since the cosmics all travel downwards (at varying angles) there remains a time bias between the TFBs at the top and bottom of the DsECal. This bias, which is on the order of 6 ns, will be built into the calculated offsets. It turns out that this bias is on the same order of magnitude as the offsets themselves, and so this is a limitation to the current methodology. However, the size of the offsets was not known *a priori*, since this methodology was the very first calibration of the DsECal. In addition, the bias has a decreasing impact the higher up the TFB is in the DsECal. In theory, the time of flight effect could be calculated and removed if reconstructed data (instead of raw data) were used. However, the reconstruction algorithms were under rapid development when this calibration was performed, and so using raw data was strongly preferred.

The TFB and RMM offsets have been calculated for all DsECal runs in the Run 1 data period, and have been entered into the calibration database to be used by the calibration software packages (as described in Chapter 4).[2]

5.1.3 Method validation

Figure 5.3 shows the effect of applying the TFB offsets to the data; the alignment of the TFBs is significantly improved, although there is still an offset between the two RMMs. Figure 5.4 shows the hit time distributions for all TFBs on RMM 0 and RMM 1 separately (the left and right hand plots respectively), before and after the TFB offsets are applied. Again, it is clear that applying the TFB offsets improves the alignment of the TFBs, causing a reduction in the RMS of the hit times. Figure 5.5 shows the TFB offset values which have been applied in Figures 5.3 and 5.4. These offsets represent the residuals between the mean time on each TFB (μ_{ij}) and the average time on the RMM ($\bar{\mu}_{ij}$), as shown in Equation 5.1. Figure 5.6 shows equivalent plots to Figure 5.5, but the TFB offsets have already been applied to the TFB time distributions. By applying the TFB offsets, the RMS of the residuals is reduced from 5.83 to 0.75 ns for RMM 0 and from 6.63 to 0.69 ns for RMM 1. This shows that the TFBs are aligned with an accuracy of 0.75 ns using this method. Figure 5.7 shows the effect of applying the RMM offsets to the data. The RMMs are aligned with an accuracy of 0.3 ns using this method.

To check that there is no significant position or directional bias in the DsECal cosmic trigger, offsets computed using both raw data and reconstructed data are compared. In contrast to the raw data, the reconstructed data allows the hit times to be corrected for the hit position along the bar, removing any bias. Figure 5.8 shows that there is close agreement between offsets calculated from raw and reconstructed data, for both the RMM offsets and TFB offsets, which gives good confidence that there is no significant bias in the DsECal cosmic trigger. The difference between the two methods has an RMS of 0.46 ns for the TFB offsets and 0.13 ns for the RMM offsets. Since these values are within the accuracy to which the TFBs and RMMs can be aligned (0.75 ns and 0.3 ns respectively), it is valid to use either the raw data or reconstructed data to calculate the offsets. The raw data is used, since this removes the need for data processing

[2]The TFB and RMM offsets are in the tables named "TFB_TIME_OFFSETS_BYRMM_TABLE" and "RMM_TIME_OFFSETS_TABLE" respectively, and are applied in the tfbApplyCalib package.

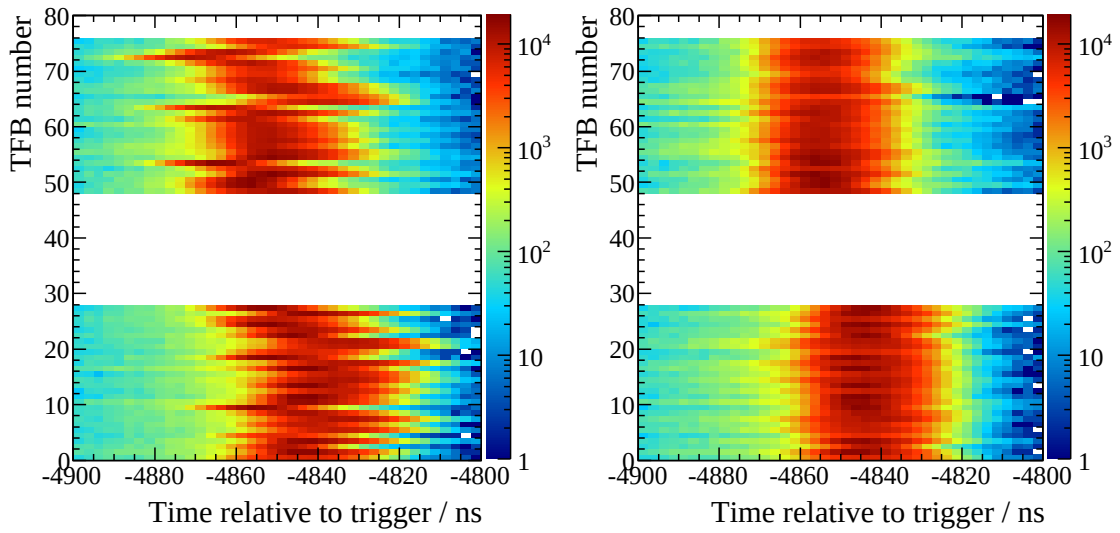


Figure 5.3: TFB number versus hit time for ND280 run 4766 (a DsECal-only triggered cosmic run). The left hand plot shows the raw time distributions. The right hand plot shows the distributions once the TFB offsets are applied. The two separate blocks of TFBs in each plot correspond to the two DsECal module RMMs. It is clear that applying the TFB offsets has improved the alignment of the TFBs, although there is still an offset between the two RMMs.

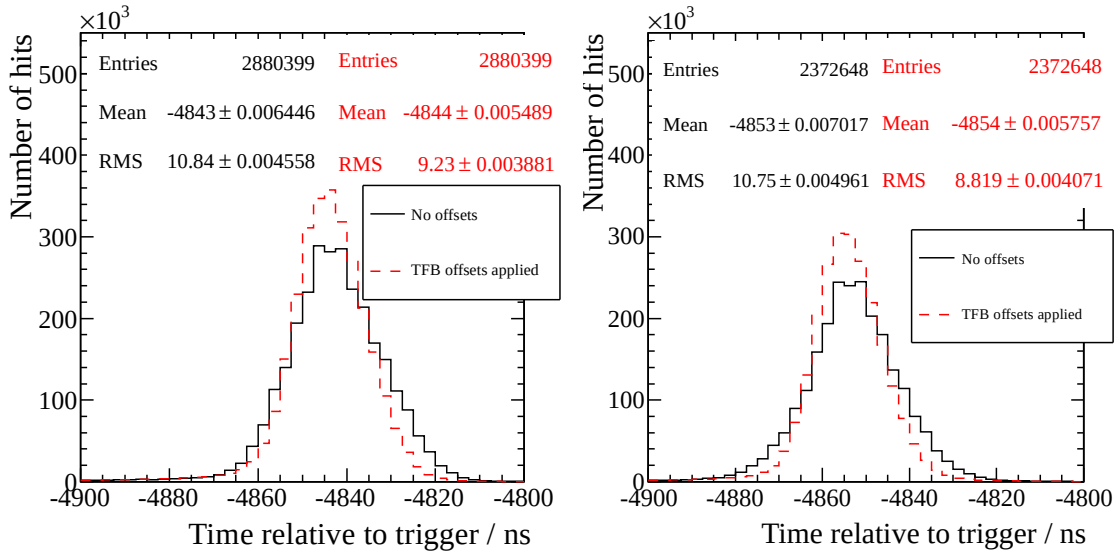


Figure 5.4: Time distributions before and after the TFB offsets are applied for ND280 run 4766. The left hand plot shows RMM 0; the right hand plot shows RMM 1. The TFB offsets have improved the alignment of the TFBs, which is shown in the reduction of the RMS, for both RMMs.

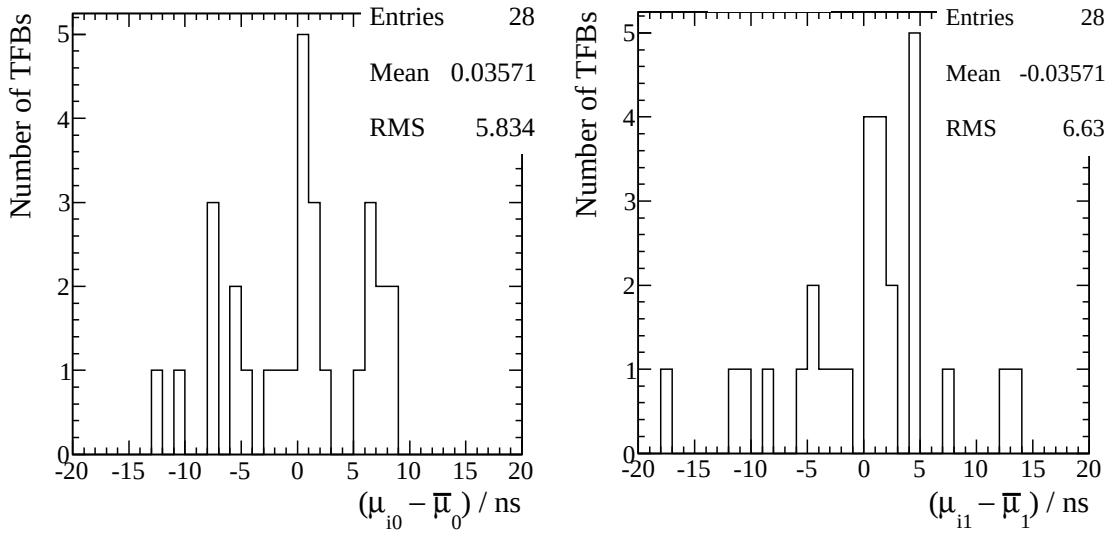


Figure 5.5: The TFB offsets calculated from ND280 run 4766. $(\mu_{ij} - \bar{\mu}_j)$ is equal to the TFB offset, as defined in Equation 5.1. The left hand plot is for RMM 0; the right hand plot is for RMM 1.

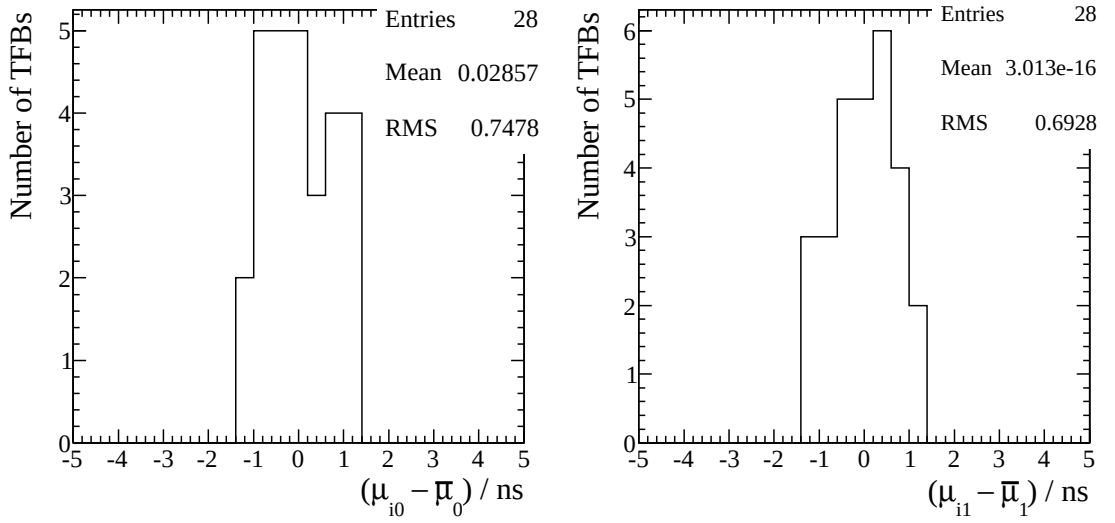


Figure 5.6: Plots showing the level of alignment that is achieved by applying the TFB offsets to the data. μ_{ij} and $\bar{\mu}_j$ are as defined in Equation 5.1. The left hand plot is for RMM 0; the right hand plot is for RMM 1. In contrast to Figure 5.5, the TFB time distributions that are fitted to give the μ_{ij} values have already had the TFB offsets applied to them. Therefore, the RMS values of the above plots show the level of alignment that is achieved between TFBs by applying the TFB offsets to the data. The TFBs are aligned to within 0.75 ns.

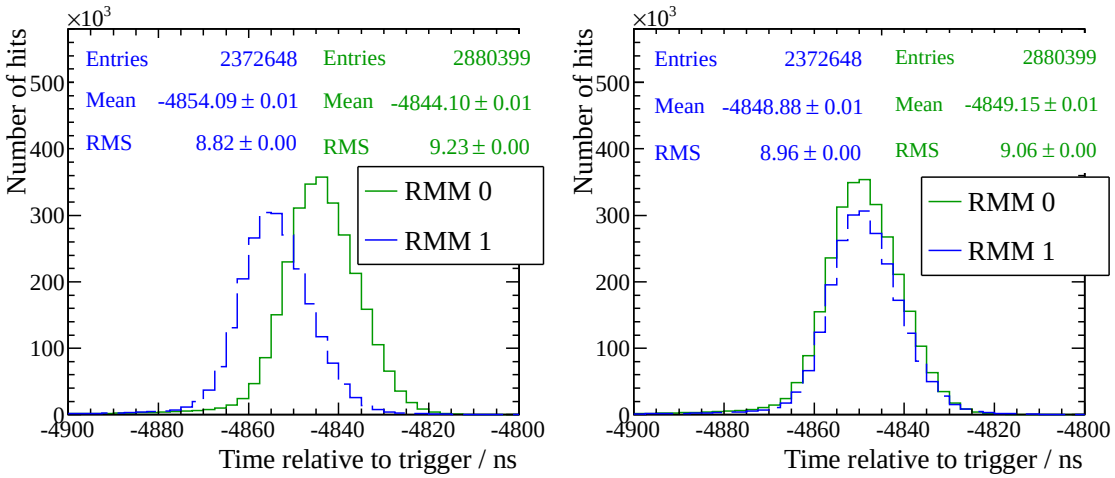


Figure 5.7: Alignment of the RMMs by applying the RMM offsets. The left hand plot shows the hit times after the TFB offsets are applied, but before the RMM offsets are applied. The right hand plot shows the hit times once both the TFB and RMM offsets are applied. By applying the RMM offsets, the RMMs are aligned to within 0.3 ns.

and any dependence on external factors, such as the reconstruction algorithms and the assumption for the speed of light in the bar.

5.1.4 Stability of offsets

Figure 5.9 shows the stability of the TFB offsets for the period spanning the Run 1 data. The Run 1 period is further subdivided into the four Main Ring (MR) runs, numbered 31 to 34, which describe the periods when the accelerator and ND280 detector configuration were kept stable. The top plot shows the TFB offsets as a function of date, for each of the 28 TFBs on RMM0; the bottom plot shows the offset residual for each TFB over the same date period. The offset residual is defined as the change of the offset from its average value over the full period. It can be seen from the bottom plot that the offsets are stable to within 2.5 ns of their average value, which is the sampling resolution of the time-stamping process (described in Section 2.2.3). Since these offsets are due purely to different cable lengths between the RMM and TFB boards, they are expected to be stable. Tables 5.1 and 5.2 show the values of the TFB offsets (these are calculated

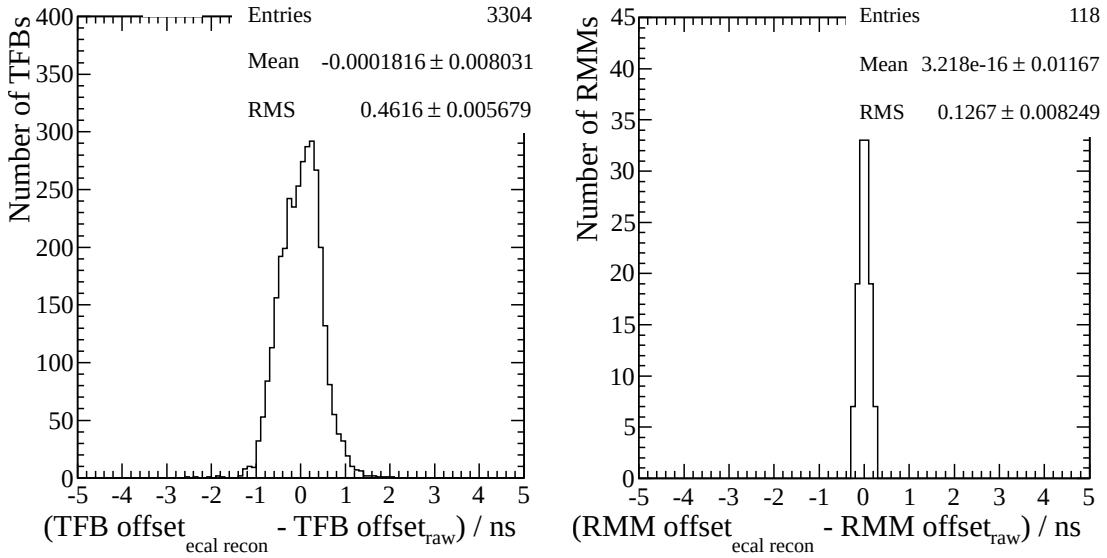


Figure 5.8: Comparison of offsets obtained using reconstructed data (where the hit times are corrected for the hit position along the bar) and raw data. The left hand plot shows the residuals for the TFB offsets, and the right hand plot for the RMM offsets. There is close agreement, to within 0.3 ns for the RMM offsets, and approximately 1 ns for the TFB offsets. All DsECal triggered runs between Jan '10 and June '10 (59 runs) are used for the comparison. There are 3304 entries in the left plot since the offset residual is calculated for each of the 56 TFBs (28 on each RMM) for each of the 59 runs. There are 118 entries on the right plot since the offset residual is calculated for each of the two RMMs for each of the 59 runs.

RMM	TFB	TFB Offset / ns
0	0	-2.75
0	1	0.36
0	2	-5.51
0	3	1.66
0	4	-5.87
0	5	6.15
0	6	0.57
0	7	6.96
0	8	-0.26
0	9	-12.03
0	10	1.01
0	11	0.99
0	12	6.20
0	13	1.08
0	14	5.84
0	15	7.20
0	16	0.27
0	17	8.24
0	18	-7.32
0	19	-1.10
0	20	8.36
0	21	7.68
0	22	-4.21
0	23	0.02
0	24	-7.65
0	25	-7.50
0	26	2.12
0	27	-10.59

Table 5.1: TFB offsets for all TFBs on RMM 0, calculated from ND280 run 4766.

from ND280 run 4766, a DsECal-only triggered cosmic run). The statistical error on the offsets is approximately 1%.

RMM	TFB	TFB Offset / ns
1	0	1.93
1	1	-4.63
1	2	-3.57
1	3	4.07
1	4	1.17
1	5	-11.33
1	6	4.06
1	7	4.94
1	8	1.17
1	9	2.66
1	10	1.00
1	11	1.30
1	12	0.66
1	13	-5.15
1	14	2.63
1	15	-10.99
1	16	12.21
1	17	13.51
1	18	4.68
1	19	4.34
1	20	-2.03
1	21	-1.04
1	22	0.45
1	23	-4.13
1	24	-17.45
1	25	-8.21
1	26	7.03
1	27	0.72

Table 5.2: TFB offsets for all TFBs on RMM 1, calculated from ND280 run 4766.

In contrast, Figure 5.10 shows that the RMM offsets drift with time during the Run 1 period (which is again subdivided into the four MR runs). The top plot shows that the mean time relative to trigger on each RMM ($\bar{\mu}_j$ in Equation 5.2) changes when the RMM boards are powercycled. This is expected, since powercycling the RMM boards will cause a Rocket I/O shift to occur between the SCM and RMM boards, as described in Section 5.1.1. However, the mean times also seem to move when there is no powercycling of the RMMs. This is not expected. To verify this behaviour, the mean times are also plotted for a different cosmic trigger configuration, shown in Figure 5.11 (the date period for this figure is a sub-period of that in Figure 5.10). As in Figure 5.10, the mean times in Figure 5.11 move even when there is no powercycling of the RMM boards. From comparison of Figures 5.10 and 5.11, it appears that there may be some discrepancies between the movements of the RMM times with the two different trigger configurations. However, the runs in these two figures were taken on different dates (due to the different triggers), and so these discrepancies could be indicative of the regularity with which the RMM times are moving. One possible explanation for these shifts is the powercyling of other boards, either the MCM or SCM. Currently it is not possible to monitor the power supplies of these boards. However, it is known that they are powercycled much less often than the RMM boards and so it does not seem likely that they can account for these regular movements. The magnitude of these unexplained movements is approximately 10 to 15 ns, which is of the same order as the TFB offsets themselves, and so it is very important that they are properly monitored and understood.

5.1.5 Further developments

Since the work presented in this section was completed, the following developments have occurred. The TFB and RMM offsets for the entire Tracker ECal have now been calculated, including the six Barrel ECal modules and the DsECal. It was necessary to develop a different calibration methodology when incorporating the Barrel ECal modules since these modules do not have their own standalone triggers, as the DsECal module does, but rather the global cosmic trigger must be used. Details of this updated methodology can be found in [96]. In addition, the unexpected changes in the RMM offsets, described in Section 5.1.4, have now been resolved. It was difficult to monitor this situation in a methodical way for the calibration methodology described in this section,

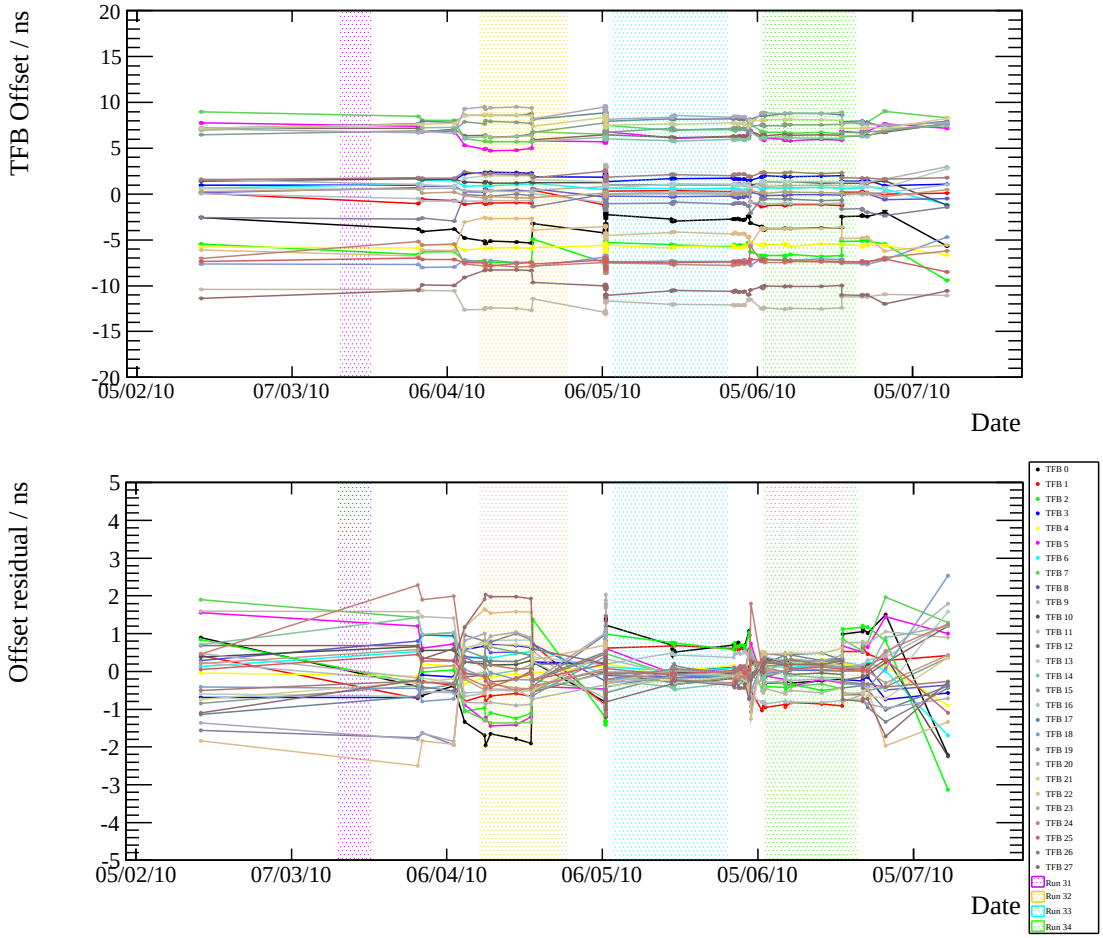


Figure 5.9: Stability of the TFB offsets over the Run 1 period, which is further subdivided into the four MR runs (Run 31 to Run 34). The top plot shows the TFB offset as a function of date for each of the 28 TFBs on RMM0. The bottom plot shows the offset residual over the same date period for each of the TFBs, where the offset residual is defined as the change of the offset from its average value over the full period. The offsets are stable to within approximately 2.5 ns, which is equal to the sampling resolution of the time-stamping process (described in Section 2.2.3). The conclusion is the same for TFBs on RMM1 (not shown here).

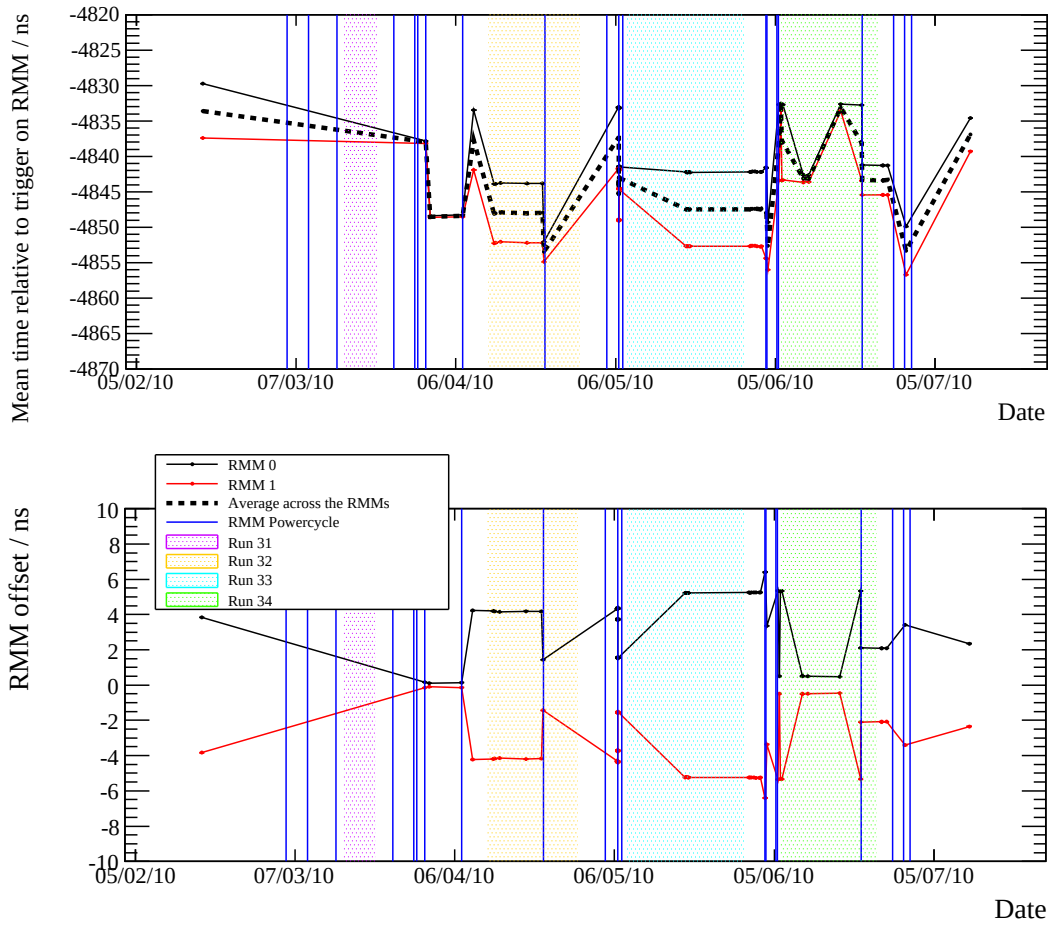


Figure 5.10: The top plot shows the variation in the RMM mean times ($\bar{\mu}_j$ in Equation 5.2) over the Run 1 period, which is further subdivided into the four MR runs (Run 31 to Run 34). The bottom plot shows the variation in the RMM offsets over the same period. RMM powercycling events are when the RMMs are switched off and on again, via the Trip-t Weiner DsECal power supply.

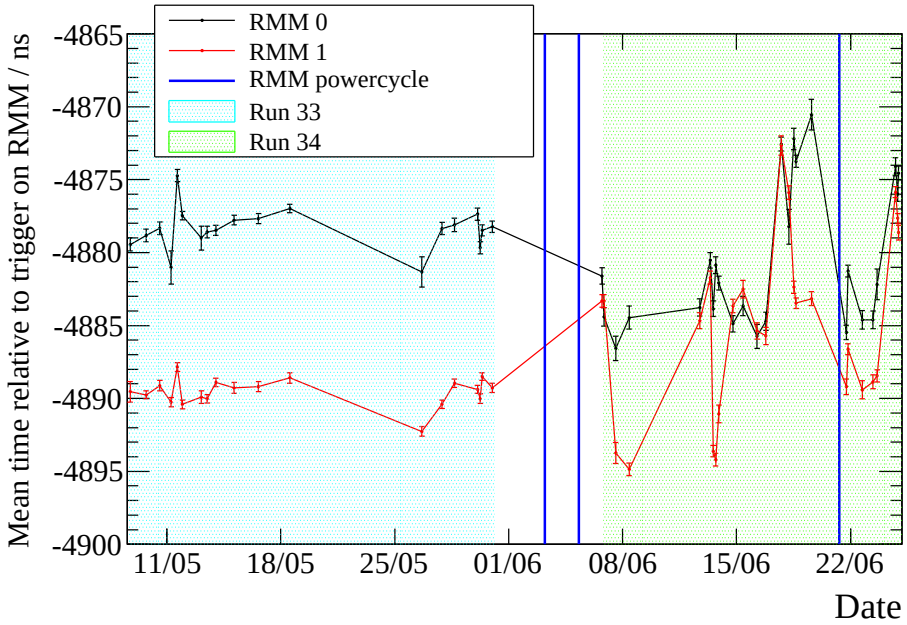


Figure 5.11: Mean time relative to trigger on each RMM versus date, with RMM powercycles marked, for a different cosmic trigger to Figure 5.10. The cosmic trigger configuration for these runs is: top 8 downstream towers of the SMRD plus the DsECal (recall that the RMM offsets (and TFB offsets) are calculated using DsECal-only triggered runs).

since the DsECal-only triggered cosmic runs necessary to calculate the offsets were only taken on an *ad hoc* basis in between periods of beam running. In contrast, the global cosmic trigger, which is now used to calculate the offsets, is interspersed between the beam spills, making it possible to regularly monitor the RMM offsets. It is now known that the behaviour described in Section 5.1.4, in which the RMM offsets change even when the boards are not powercycled, is due to the optical links between the boards temporarily "dropping out" and then reestablishing a connection. Each time this happens, a phase shift can occur between the clocks on the boards. The RMM offsets are regularly recalculated and uploaded to the calibration database.

5.2 Fibre timewalk calibration for the P0D and DsECal

5.2.1 Methodology and results

As described in Section 2.2.1, the scintillator based sub-detectors, including the P0D and DsECal, use a readout system in which a Y-11 wavelength-shifting fibre, running down the centre of the scintillator bar, collects the scintillation light and transports it to an MPPC. The MPPC then converts the light into an electrical signal. For the P0D and DsECal, which are both Trip-t based sub-detectors, the MPPCs are coupled to Trip-t chips. Each channel on a chip has a discriminator which fires and generates a time-stamp once the integrated charge crosses the discriminator threshold (as described in Section 2.2.3). The fibre timewalk effect arises as follows. The scintillation fluors in the Y-11 fibre emit light according to an exponential decay function. Due to this exponential decay, there is a time delay in the detection by the MPPC of the threshold photoelectron and therefore a time delay in the firing of the discriminator on the Trip-t chip. This time delay is the fibre timewalk. The length of the delay depends on the fibre time constant, the discriminator threshold and the total number of photoelectrons; however, it is a statistical effect, and so the magnitude of the timewalk will vary between identical events. It is worth noting that there are scintillation fluors in the scintillator bar itself (PPO and POPOP), as well as those in the wavelength shifting fibre, but these have a smaller decay time than the fibre and so are neglected. (PPO and POPOP have

time constants of 1.6 ns and 1.5 ns respectively, whereas the fibre has a time constant of approximately 7 ns.[3])

The fibre timewalk can be modelled using an analytical function. The probability that the n^{th} photoelectron out of N be detected at time t is given by:

$$P = \exp\left(\frac{-t(N-n+1)}{\tau}\right)\left(1-\exp\left(\frac{-t}{\tau}\right)\right)^{(n-1)} \tag{5.4}$$

where n is the discriminator threshold and τ is the time constant of the Y-11 fibre. The timewalk is then estimated from the first moment:

$$\langle t\rangle = \frac{\int_0^\infty tP\mathrm{d}t}{\int_0^\infty P\mathrm{d}t} \tag{5.5}$$

The P0D has a discriminator threshold of 2.5 photoelectrons (i.e. $n = 3$), giving the following equation for the fibre timewalk:

$$\langle t\rangle = \frac{\tau(2-6N+3N^2)}{N(N-1)(N-2)} \tag{5.6}$$

The ECal has a discriminator threshold of 3.5 photoelectrons (i.e. $n = 4$), giving the following equation for the fibre timewalk:

$$\langle t\rangle = \frac{2\tau(2N^3-9N^2+11N-3)}{N(N-1)(N-2)(N-3)} \tag{5.7}$$

It is worth noting that the above equations are only strictly valid when n and N are the number of photoelectrons at the MPPC arising only from photons in the fibre (excluding reflected photons). However, in the real system, there is also correlated noise in the MPPC. The discriminator sees the integrated charge *after* the MPPC response, which will include not only avalanches from the incident converting photons but will also

[3]The time constants for POP and POPOP are quoted from the compound specifications of a chemical manufacturing company, Research Products International Corp. These are intended as representative numbers only to demonstrate that they are smaller than the associated time constant of the Y-11 fibre.

include cross-talk and afterpulsing effects. Therefore, in the implementation, the charge *including* the cross-talk and afterpulsing is used as N, even though this is not strictly valid for the above formulism. The correction seems to work well (see Figures 5.13 and 5.14 described later), although the results should be compared to Monte Carlo to fully understand the different effects from the fibre, electronics and MPPC.[4]

The above issue also has an impact on the τ parameter in Equations 5.4–5.7. Strictly speaking, this is the time constant of the Y-11 fibre, which has a value of approximately 7 ns [97]. However, the afterpulsing effect also has time constants associated with it (a short time constant of approximately 17 ns, and long time constant of approximately 70 ns [53]). The time constants from both the fibre and afterpulsing effect are summed, weighted by the respective probabilities for an avalanche to originate from the fibre and an afterpulse, giving a final value of $\tau \approx 12$ ns. Figure 5.12 shows the timewalk corrections for both the P0D and ECal, given by Equations 5.6 and 5.7 respectively, with $\tau = 12$ ns.

Figure 5.13 shows the effect of applying the fibre timewalk correction to data. Each point represents a reconstructed hit. The left plot is before the hit times are corrected for the fibre timewalk effect and the right plot is after. The charge of each hit is measured in MIP-equivalent units (MEU), where MIP stands for minimum ionising particle. 1 MEU is defined as the most probable amount of charge to be deposited by a MIP-like particle which travels through the bar along the z direction of the ND280 (i.e., with a path length of 1 cm); its value is extracted from fits to cosmic muon data. Applying the fibre timewalk correction makes the time dependence on charge flatter, as intended. Figure 5.14 makes this flattening effect clearer, by plotting the mean of each charge bin for the plots in Figure 5.13. The blue points are for comparison to the red points, to show that the correction works significantly better if the time constants of the MPPC afterpulse effect, as described earlier, are included in addition to the time constant of the Y-11 fibre.

[4]The fibre timewalk correction is implemented in tfbApplyCalib for the P0D and ECal. Since it is only possible to use the fully calibrated charge in the time calibration chain, it is necessary to "de-linearise" the charge to obtain the charge including cross-talk and afterpulsing. Average linearisation factors of 0.77 for the P0D and 0.83 for the ECal are used.

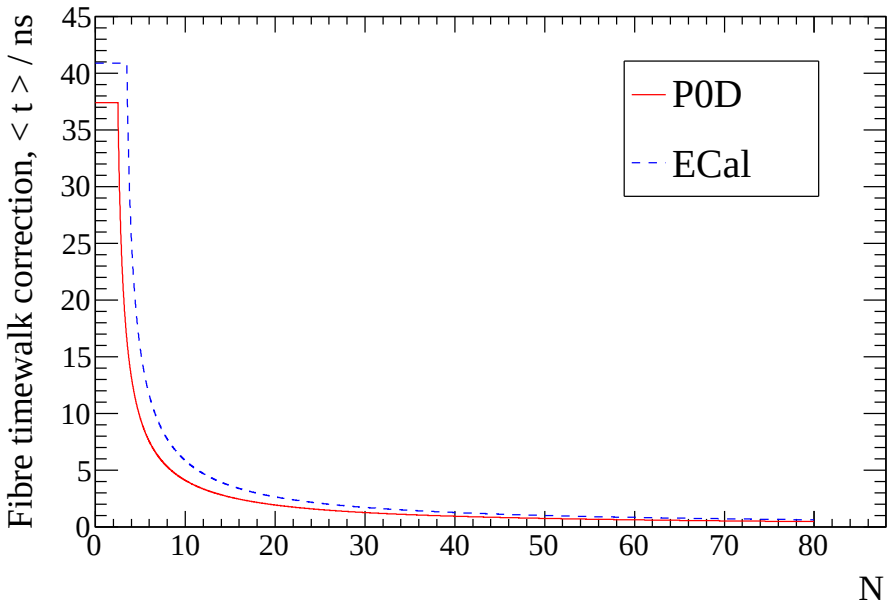


Figure 5.12: Fibre timewalk correction for both the P0D and ECal, given by Equations 5.6 and 5.7 respectively, with $\tau = 12$ ns. Note that the correction is flat below $N = 2.5$ for the P0D and $N = 3.5$ for the ECal, which is to avoid the infinities in the equations.

5.2.2 Further developments

An alternative methodology has been developed to correct the fibre timewalk effect since the completion of the work presented in this section. This new approach attempts to address the fact that there is clearly some residual timewalk even after the correction is applied in Figure 5.14. In contrast to the methodology described in this section, in which the timewalk curve is derived from first principles, the new approach uses an empirical function to parametrise the timewalk curve which is fitted to the data. At the time of writing, this new methodology has been implemented for the P0D sub-detector, and there is discussion of implementing it for the other Trip-t sub-detectors (including the ECal). Further details can be found in [98].

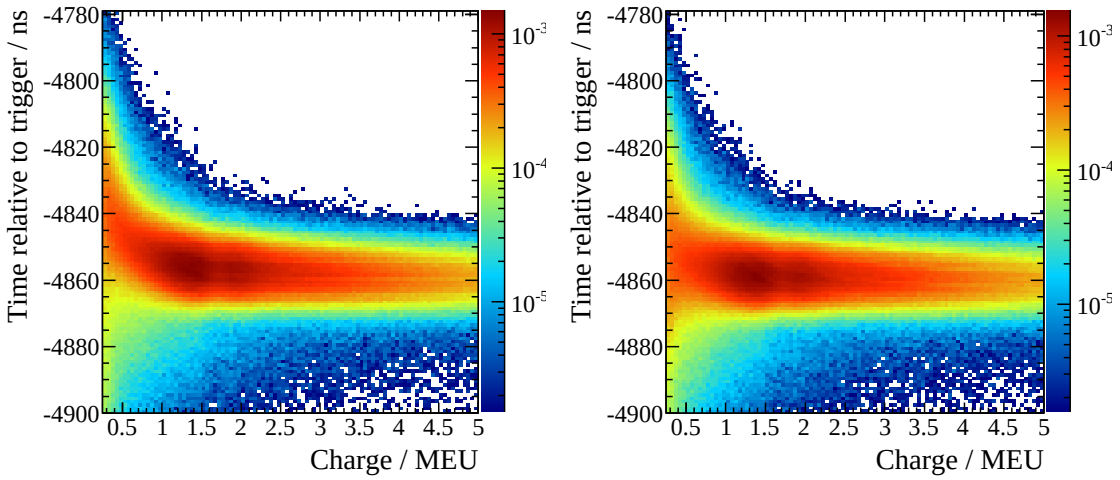


Figure 5.13: Plots show the effect of applying the fibre timewalk correction to data, with $\tau = 12$ ns. The left hand plot is before the correction is applied; the right hand plot is after. In both plots, all other time calibration stages have been applied, i.e. TFB and RMM offsets, and electronics timewalk. Each point represents a reconstructed hit, where the MEU charge unit is explained in the text. DsECal-only triggered cosmic run 4860 is used.

5.3 Time resolution of the DsECal

Achieving a good time resolution is very important, in order to, for example, effectively distinguish between separate interactions, reject noise and resolve particle direction. This section presents the achieved time resolution of the DsECal after the calibration methods are applied. As described in Section 2.2.3, the DsECal scintillator bars are read out at both ends, meaning that each bar provides two front-end channels. Therefore, to extract the time resolution for a single channel, the time resolution, σ_{res}, can be defined as:

$$\sigma_{\text{res}} = \frac{\sigma_{|t_1 - t_2|}}{\sqrt{2}} \tag{5.8}$$

where t_1 and t_2 are the hit times on either end of the scintillator bar, corrected for the light propagation time along the bar from the reconstructed hit position, and $\sigma_{|t_1-t_2|}$ is the width of a one-sided Gaussian fit to the $|t_1 - t_2|$ distribution. By using the difference

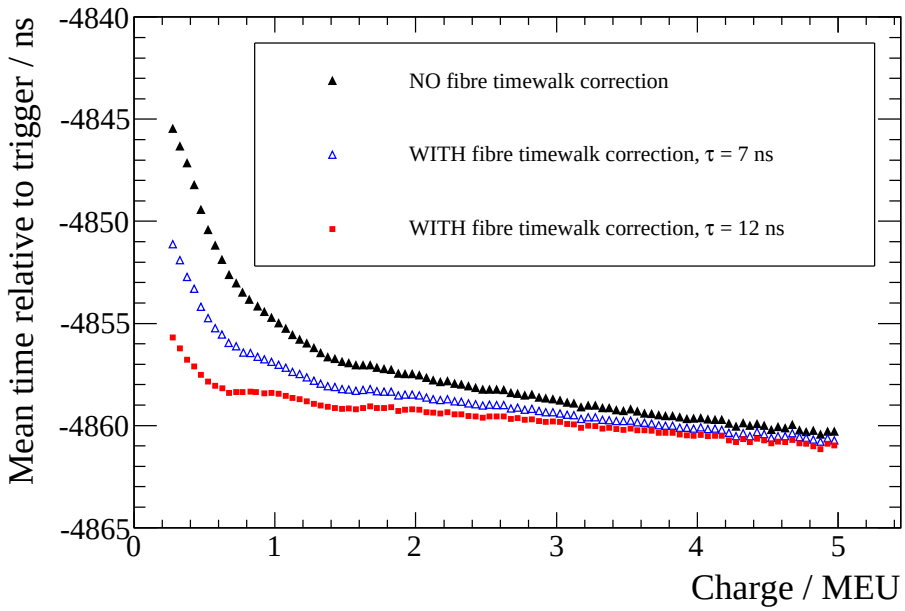


Figure 5.14: Effect of applying the fibre timewalk correction to data. The black and red points correspond to the mean of each charge bin, for the left and right hand plots of Figure 5.13 respectively. The red points are for $\tau = 12$ ns, which includes the time constants from both the fibre and MPPC afterpulsing effect; the blue points are shown for comparison with $\tau = 7$ ns, which is the time constant for the fibre only. $\tau = 12$ ns gives the better correction.

in time between the bar ends, jitters in the MCM trigger time and the time of flight of the particle can be ignored. A DsECal-only triggered cosmic run (ND280 run 4860) is used to study the time resolution, since this trigger gives good illumination across all of the DsECal channels. Single track events are selected, with hits in all 34 layers. From the hits, only double-ended hits from the middle 40 cm of the bars are considered.

As described at the start of the chapter, the full time calibration chain of the DsECal includes corrections for: *a*) RMM and TFB offsets; *b*) electronics timewalk; and *c*) fibre timewalk. Figure 5.15 shows that the distribution of time differences between bar ends gets narrower as each stage of the calibration is applied, which confirms that each stage of the calibration is working. Figure 5.16 shows the time difference between bar ends as a function of the reconstructed hit charge, with the full time calibration applied. The time difference clearly decreases as the charge increases. The charge is measured in MIP-equivalent units (MEU), as described previously in Section 5.2. Figure 5.17 shows the time resolution, σ_{res}, as a function of the reconstructed hit charge, with the full time

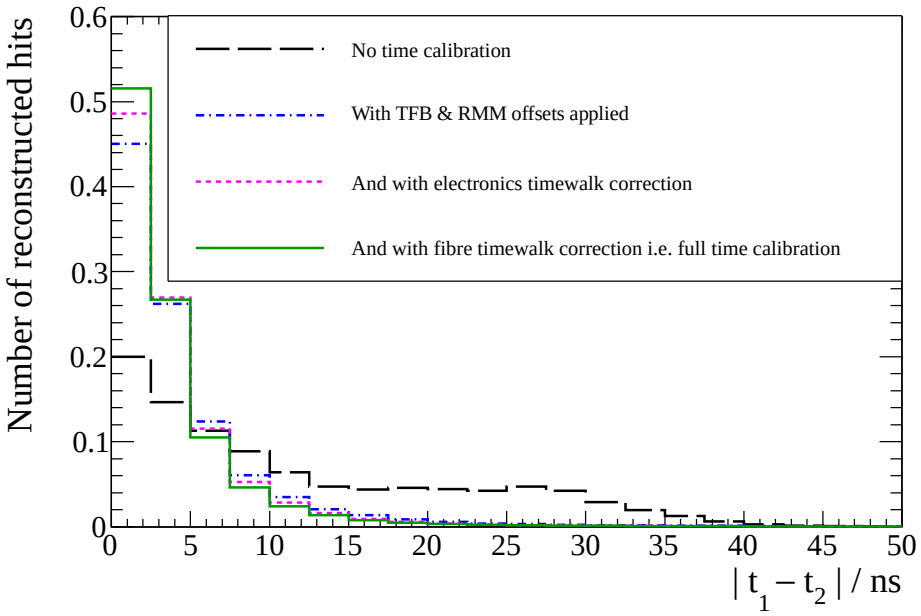


Figure 5.15: Plot showing the narrowing of the distribution of the time difference between bar ends as each time calibration stage is applied. t_1 and t_2 are the hit times on either end of the bar, corrected for the light propagation time along the bar, as defined in Equation 5.8. This demonstrates that the time calibration is working.

calibration applied, where σ_{res} is calculated from the plots in Figure 5.16 according to Equation 5.8. The time resolution clearly improves as the charge increases.[5]

For a MIP, which by definition produces hits with a charge of approximately 1 MEU, the time resolution is approximately 3.5 ns. This improves down to approximately 1.5 ns for very high charge hits. For comparison, the resolution limit for the electronics, derived from the sampling resolution of 2.5 ns for the time-stamping process, is approximately 0.7 ns. However, it is expected that the final resolution be higher than this, since there are contributions to the time resolution that cannot be removed with calibration. The two dominant examples of this are: *a*) the statistical nature of the fibre timewalk correction; and *b*) fluctuations in crossing the discriminator threshold. The fibre timewalk correction, as described in Section 5.2, is a purely statistical correction, meaning that on an event by event basis, the actual timewalk can be significantly different from the calculated correction. This leads to a remaining statistical fluctuation in the hit times, even

[5]The bump in this plot at approximately 1.7 MEU is due to an issue with the charge calibration that has since been resolved.

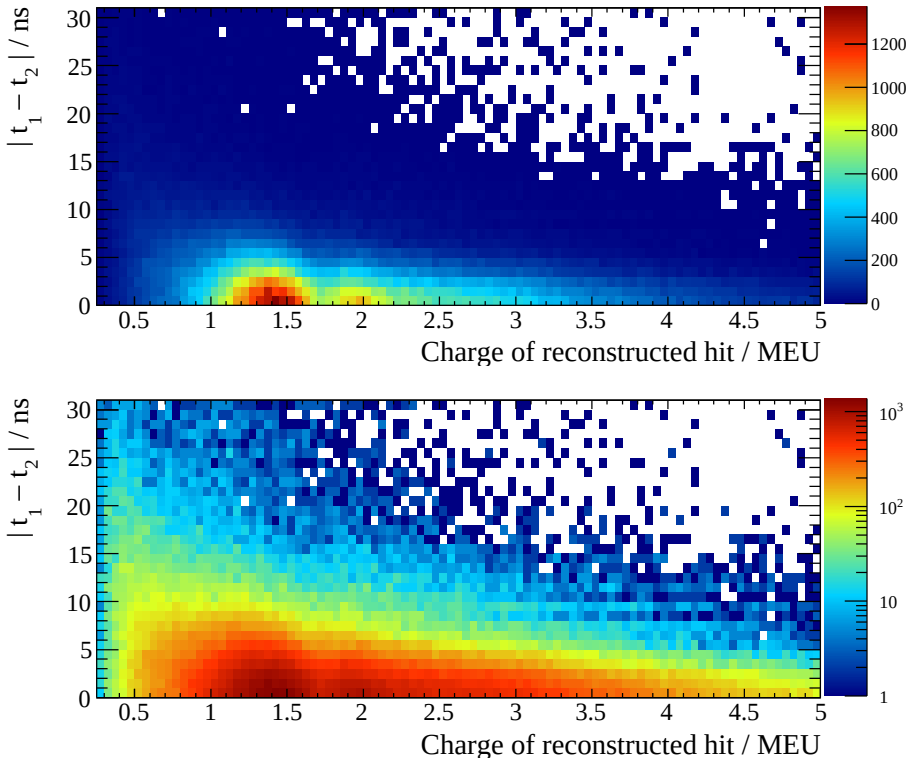


Figure 5.16: The time difference between bar ends as a function of the reconstructed hit charge, with the full time calibration applied. t_1 and t_2 are the hit times on either end of the bar, corrected for the light propagation time along the bar, as defined in Equation 5.8. The reconstructed hit charge is obtained by calibrating and combining the separate hit charges on either end of the bar. Note that the bottom plot has a logarithmic scale.

after the fibre timewalk correction is applied. The discriminator threshold contribution can be explained as follows. Smaller charges have a longer rise time in the electronics. As well as causing electronics timewalk, where there is a delay in small charges crossing the discriminator threshold, this also causes a larger fluctuation in the time at which small charges actually cross the threshold.

Using the assumption that two peaks can be resolved if they are separated by more than the average of their full widths at half maximum [99], the minimum resolvable time difference between the DsECal and an FGD can be calculated to be 5.3 ns for a MIP (where a time resolution of approximately 1 ns is assumed for the FGD). This translates into a track length of approximately 159 cm. The distance between the DsECal and FGD1 is approximately 2.5 m, compared to just over 1 m between the DsECal and

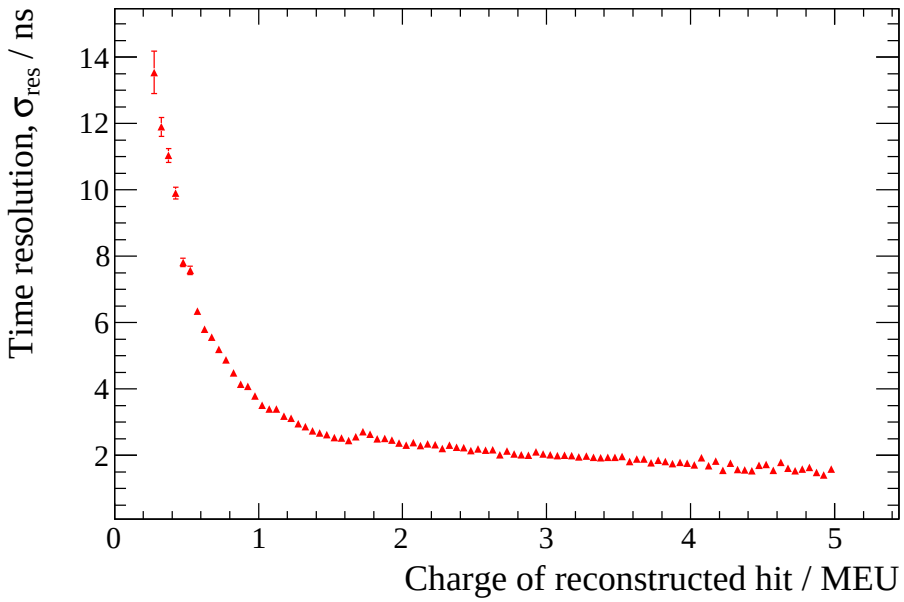


Figure 5.17: The time resolution as a function of the charge, where the time resolution is obtained by fitting a one-sided Gaussian to each bin of the plots in Figure 5.16 and dividing the width of the Gaussian by $\sqrt{2}$ (according to Equation 5.8). The time resolution clearly improves as the charge increases.

FGD2. The track direction can therefore be determined for a MIP which enters the DsECal from FGD1, but not from FGD2, unless there is a large curvature or the presence of some high charge hits. In contrast to MIPs, showering particles such as electrons tend to produce high charge hits. Assuming hits of around 5 MEU, the minimum resolvable time difference between an ECal module and an FGD can be calculated to be 2.9 ns (again assuming a time resolution of approximately 1 ns for the FGD), which translates into a distance of approximately 88 cm. Therefore, it should be possible to determine the direction of showering particles entering the DsECal module from both FGDs.

Chapter 6

Constraining the normalisation of K^+-originating neutrinos

6.1 Importance of K^+ neutrino parents for the T2K flux prediction

Both kaon and pion particles are produced when the proton beam hits the graphite target (see Section 2.1.3 for a description of the beam layout). These particles decay, giving rise to the T2K neutrino beam. The flux prediction for this beam has been discussed previously in Chapter 3; Figures 3.4 and 3.5 show the neutrino flux at the ND280, on a linear and logarithmic scale respectively, as a function of energy for the four different flavours (ν_μ, $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$) with each flavour broken down by neutrino parent. It can be seen that the muon neutrinos make up the vast majority of the flux, accounting for approximately 98% of the total. It can be seen that the low energy part of the ν_μ spectrum is dominated by π^+ daughters, whereas the high energy part is dominated by K^+ daughters. Understanding the K^+ component of the beam is therefore key to accurately predicting the T2K neutrino flux. The K^+ particles, in addition to decaying to muon neutrinos, also decay to high energy electron neutrinos, as shown in Figure 3.4(c). These electron neutrinos form a major background to the ν_e appearance measurement at Super-Kamiokande, again underlining the importance of understanding the K^+ component of the beam. The K^+ daughters can be studied using interactions of high energy neutrinos at the ND280 detector.

6.2 Measurement

A measurement is made of the normalisation of K^+-originating neutrinos at ND280. This measurement can be compared to the prediction of the nominal MC, and therefore used to validate and tune the T2K flux prediction. To be clear, this measurement does not attempt to isolate the production cross-section of K^+ particles at the graphite target. This is clearly an important factor, but the measurement is also affected by many other aspects of the beam system. For example, the off-axis angle of the ND280 detector compared to the beam direction directly affects the number, energy and parent breakdown of the neutrinos arriving at ND280, as does the current in the horns since this drives the focusing and selection of the parent particles before they decay. This measurement therefore probes all parts of the beam system that affect the normalisation of K^+-originating neutrinos at ND280. This measurement clearly only probes a specific part of the K^+ neutrino parent phase space, i.e., the K^+ parents whose neutrinos reach the ND280; it cannot provide any information on the K^+ particles whose neutrinos miss the ND280. However, this is the important part of the phase space for predictions at Super-Kamiokande, since the neutrinos at Super-Kamiokande will be a subset of those passing through the ND280.

6.3 Strategy

An enhanced sample of charged-current (CC) ν_μ interactions is selected via a series of selection cuts. There are several reasons for this choice of sample. Firstly, the high energy part of the ν_μ energy spectrum is dominated by K^+ daughters, as described previously. Secondly, it is a high statistics sample meaning that kinematical binning can be used as opposed to a simple counting experiment. Thirdly, CC ν_μ interactions have a clear, robust signal in the ND280 detector; an energetic muon is emitted from an interaction in an FGD that, in the majority of cases, then passes through the TPCs, thereby allowing superior tracking and particle identification of the track. The selected events are binned according to the reconstructed momentum of the muon candidate, p_μ, and cosine of the reconstructed muon angle, $\cos\theta_\mu$.[1] These p_μ - $\cos\theta_\mu$ bins are then passed into a fit, in which the predicted bin occupancies are fit to the observed bin occupancies

[1]The muon angle is with respect to the z-axis of the ND280 (which is shown in Figure 2.6).

by varying a set of fit parameters. These parameters include the normalisation of π^+-originating neutrinos, b_π, and the normalisation of K^+-originating neutrinos, b_K. There are also nuisance parameters which model the effects of the neutrino interaction cross-section uncertainties and systematic errors of the detector system. It is worth noting that the uncertainties on the neutrino flux prediction are not included in the fit, in order to perform an unconstrained measurement of b_K, which will be discussed in more detail in Section 6.7.5. Broadly speaking, muons from CC interactions of π^+-originating neutrinos dominate the low momentum and high angle phase space, whilst those from CC interactions of K^+-originating neutrinos dominate the high momentum and low angle phase space. By fitting b_π as well as b_K, the background due to π^+-originating neutrinos in the K^+ signal region is constrained. A profile likelihood method is used in the fit; a 2-dimensional b_K-b_π space is defined and the likelihood is maximised at each point with respect to all of the nuisance parameters. This results in a best fit point in the b_K-b_π space and a 1σ contour about this point.

This analysis was performed at approximately the same time as the official ND280 analysis for the T2K neutrino oscillation study, described previously in Section 3.4.1, which constrained the flux and cross-section parameters using ND280 data in order to propagate them to the oscillation analysis at Super-Kamiokande [20,86,87]. The analysis presented in this chapter is distinct from this official analysis in terms of the analysis goals, the event selections and the fit implementations. However, many detailed studies were performed by internal T2K working groups to provide inputs to the official analysis for both the cross-section and detector systematic errors. Where possible, these inputs are used for this analysis.

A possible extension to this analysis would be to select an enhanced sample of CC ν_e interactions and simultaneously fit this sample with the existing ν_μ sample. Some preliminary work was done to produce a CC ν_e sample. However, the number of events in this sample would be significantly smaller than in the ν_μ sample, on the order of a couple of hundred events compared to over 6000 events. Therefore, the statistical error would be much larger than for the ν_μ sample, and only very coarse binning of the events, if any, would be possible. In addition, the detector systematic errors for selecting these CC ν_e events in the ND280 were still being developed by the relevant T2K working groups at the time of this analysis being performed. For these reasons, it was decided not to include a ν_e sample in the fit.

6.4 Data sets

The Run 2 data set, as detailed in Section 2.1.4, is used in this analysis. Two data quality checks are applied. The first check ensures that the proton beam was running stably and is applied per proton beam spill, whilst the second check ensures that the ND280 detector was running stably and is applied per sub-run of data. The total POT accumulated in Run 2 is 1.055×10^{20}, dropping to 7.837×10^{19} after the data quality checks. The effective POT decrease incurred through this step is mainly due to an issue with the ND280 detector, whereby a hardware problem occurred with one of the Micromegas of TPC3 near to the beginning of the Run 2 period (in December 2010).

The Run 1 data set is not used in this analysis, since the full ECal subsystem was not installed in the ND280 detector during this period (only the DsECal module was installed and working correctly). When this analysis was originally conceived, one of the possible extensions was to simultaneously fit a ν_e sample with the ν_μ sample. Since the ECal information is very useful when trying to identify electrons from ν_e interactions, it was decided to focus on the Run 2 data set. In addition, Run 1 only contributes approximately one quarter of the combined POT from Run 1 and Run 2, and will become increasingly insignificant as the power of the proton beam is further ramped up and data is accumulated at an increasing rate.

A Monte Carlo (MC) sample, corresponding to 1.095×10^{21} POT, is used. The MC production chain is described in Section 4.1. The proton beam is simulated with a power of 100 kW, which is determined from the average beam power achieved during the Run 2 period. The NEUT neutrino interaction generator is used [89]. The full geometry of the ND280 detector is simulated, including the full ECal subsystem. However, no interactions are simulated in the cavern surrounding the ND280 detector, which is mainly composed of sand. A separate MC sample is generated to model interactions in the sand [100]. This sand MC corresponds to 7.0×10^{19} POT. Due to the simulation method of these interactions, no vertex level truth information is available for these events.

The main MC sample (excluding sand interactions) used in this analysis is tagged internally within T2K as Production 4C. This uses JNUBEAM version 11a, NEUT version v5.1.1 and v9r7p9 of the ND280 software. An additional tuning is applied to the final set of selected events, in order to take account of improvements that have been made

within the JNUBEAM neutrino flux model since the MC sample was produced; version 11b v3.1 is used.[2] The sand MC uses v9r9p1 of the ND280 software, and no additional flux tuning is applied since, as mentioned above, vertex level truth information is not available for these events. For the data sample, Production 4D is used. This also uses the newer v9r9p1 version of the ND280 software, which includes important improvements to the calibration of the ND280 sub-detectors.

From here on, unless otherwise stated, the MC sample refers to the main MC combined with the sand MC.

6.5 Modifications to the standard global reconstruction outputs

This analysis uses the global reconstruction output, described in detail in Section 4.2.2, as a starting point. However, some problems have been identified with the reconstruction tools since the production of the data and MC files. Therefore some of the outputs of the global reconstruction are modified or ignored before the analysis is performed. These modifications are as follows, and are in line with the recommendations of the internal T2K ν_μ group [101]:

- Tracks with SMRD components often have incorrect reconstructed momenta. Therefore, the momentum calculated by the tracker reconstruction, which is the stage before the global reconstruction and only includes the TPC and FGD sub-detectors (see Section 4.2), is used. The momentum from the closest TPC to the vertex is used.
- The magnetic field calculation has an issue, meaning that the reconstructed momentum is 1.59% too large for data events.[3] The momentum is rescaled to remove this effect.

[2]This tuning mainly accounts for inclusion of the NA61 kaon production data (see Chapter 3). Note that the flux tuning from the official ND280 analysis described in Section 3.4.1 is *not* included; this ensures that the analysis described here provides an independent validation of the flux prediction using the ND280 data.

[3]This value is for the Run 2 period. Different scalings are necessary for other data periods, but these are not used in this analysis.

- The energy loss correction is incorrect. It fails to account for the particle slowing down as it loses energy in the dead material and uses the wrong thickness for the FGD cover plates. A correction is applied to the track momentum to account for these effects.

6.6 Event selection

As described previously, the starting point of this analysis is to obtain a CC-inclusive ν_μ sample. These events are characterised by a muon exiting the neutrino interaction vertex. The FGDs provide the main target mass within the tracker region of the ND280 detector, and provide tracking capabilities for particles exiting the interaction vertex. The majority of particles exiting the FGDs then enter the TPCs, which provide superior tracking and particle identification capabilities, ideal for identifying this type of event. The global vertexing algorithm is used as a basis for the analysis, and will be discussed in Section 6.6.1. A set of selection cuts is then applied to the vertex and its associated tracks. A high momentum negative track must be present amongst the tracks; this is tagged as the lepton candidate which must then be identified as muon-like by the TPC particle identification. The full set of selection cuts is described in Section 6.6.2. The performance of these cuts is then discussed in Section 6.6.3. Finally, in Section 6.6.4, the sample of events passing all of the CC-inclusive ν_μ selection cuts is shown, where these events form the final input to the fit (which will be described in Section 6.7).

When discussing the selection cuts and their performance, the following signal and background categories are used. Some of these categories are self-explanatory, with details provided for the rest:

- CC ν_μ in FGD FV: CC ν_μ interactions inside the fiducial volume (FV) of one of the two FGDs (the FV boundaries will be defined in Section 6.6.2). These are the signal events.
- CC $\bar{\nu}_\mu$ in FGD FV
- CC ν_e / $\bar{\nu}_e$ in FGD FV
- NC in FGD FV: Neutral current interactions inside the FGD FV.

- Out FGD FV: Interactions inside the FGD, but outside the FV boundaries.
- Out FGD: Interactions outside of the FGD. These include interactions in other parts of the ND280, mainly in the UA1 magnet, and in the dead material around the FGD.
- Sand muons: Interactions that occur outside of the ND280 magnet, particularly in the sand upstream of the detector and the wall of the pit containing the ND280, giving rise to particles that enter the ND280. The vast majority of these particles are muons, hence "sand muons".
- Duplicated true vtx: A duplicated true vertex can happen if, for example, the tracks from a single true vertex are mis-reconstructed as coming from two separate vertices. All reconstructed vertices except the one closest to the true vertex are flagged as being duplicated true vertices.
- No matched true vtx: A reconstructed vertex not associated to any true vertex.

The CC ν_μ in FGD FV signal category can be further broken down by the neutrino interaction scattering process as follows:[4]

- QE: Quasi-elastic scattering
- RES: Resonance production
- DIS: Deep inelastic scattering
- COH: Coherent pion production

The events passing the CC-inclusive ν_μ selection cuts can also be broken down by neutrino parent type as follows:

- π^+
- K^+
- Other, including K^0_L, $\mu^\pm$, K^- and π^-

In this parent categorisation, some categories (sand muons, duplicated true vertex and no matched true vertex) are included as before due to the lack of available truth information.

[4]This categorisation is done according to the event code of the NEUT neutrino interaction generator.

These categorisations are used in the plots throughout this section.

6.6.1 Use of the global vertexing algorithm

The global vertexing algorithm is used as a basis for this analysis. Details of the algorithm are described in Section 4.2.2. The use of this algorithm over other methodologies is discussed here.

In brief, the global vertexing algorithm is composed of a basic clustering stage that groups together tracks from the same bunch into potential vertices, followed by Kalman filtering techniques that are used to decide whether these tracks are associated with the same vertex. If there is only a single track or no cluster is found then the vertex position is simply taken as the most upstream end of the highest momentum track in the bunch. The official ND280 analysis uses a methodology very similar to this single track case.[5] However, it can be seen by comparing the single track and multi-track cases that using the full vertexing algorithm outperforms simply using the upstream end of the highest momentum track. Figure 6.1 shows the vertex resolution for both single track and multi-track vertices. The x, y and z coordinates are shown separately, for true vertices in both FGD1 (left) and FGD2 (right). The resolution is better for multi-track vertices than for single track vertices, with the improvement more pronounced in FGD2 than in FGD1. This is due to the water modules in FGD2 (see Section 2.2.3). For a vertex that occurs in the water emitting a single particle, the particle will only be tracked once it reaches a scintillator bar. Simply using the start of this track as the vertex can therefore be very inaccurate. In fact, the vertex position is constrained to the centre of a bar, which gives rise to the double "hump" structure in Figure 6.1(f); the first hump is for true vertices in a bar and the second for true vertices in the water. On the other hand, the global vertexing algorithm is free to reconstruct the vertex position anywhere, including in the water, and so achieves a better vertex resolution. There is a small bias in the z position for multi-track vertices, shown in Figures 6.1(e) and 6.1(f), which shifts the vertices downstream. This is not understood and is being studied within the relevant T2K working group. However, it is small and is not expected to be a problem in this analysis.

[5] The most upstream end of the highest momentum *negative* track is used as the vertex position.

Full evaluation of the systematics associated with this algorithm are difficult and work is ongoing within the relevant working group. However, good data-MC agreement is observed for various quantities within the algorithm, as detailed in Section 4.2.2. This indicates that the global vertexing algorithm will not be introducing any significant biases into this analysis. As an additional cross-check, the data-MC agreement is studied for the difference between the global vertex position and the vertex position used by the official ND280 analysis. This is shown in Figure 6.2 for the x, y and z coordinates separately, with single track vertices on the left and multi-track vertices on the right. The data-MC agreement is good in all cases. These plots do not directly compare the performance between the global vertexing and the official ND280 vertexing (although the global vertexing is expected to perform better from the plots in Figure 6.1); they instead show that there is good data-MC agreement for the difference between the two vertex positions. This helps to ensure that the global vertexing is not introducing any significant biases into the analysis. The right hand plots of Figure 6.2, for multi-track vertices, have widths of around 10–20 mm depending on the coordinate. The left hand plots, for single track vertices, have much narrower widths of less than 1 mm. This is because, in the single track case, both algorithms simply use the upstream end of the highest momentum negative track.[6] The reason that these plots are not simply delta functions stems from a detail in the global vertexing algorithm explained previously in Section 4.2.2, namely that the algorithm only uses the output of the tracker reconstruction rather than the global reconstruction. The global reconstruction refits the entire track using a Kalman filter once the tracker output has been matched to the output from the other sub-detectors, which can cause the coordinates of the track to shift slightly. This leads to the observed non-zero widths in Figures 6.2(a), 6.2(c) and 6.2(e).[7]

[6]In fact, the global vertexing uses the upstream end of the highest momentum track (with no requirement on the charge), but this difference does not matter here since the events in Figure 6.2 have passed all of the CC-inclusive ν_μ selection cuts which include a cut on the charge of the lepton (as will be described in Section 6.6.2).

[7]The refit by the global reconstruction involves incrementally propagating the track back to planes that are essentially fixed in z, which explains why the width in Figure 6.2(e) for the z coordinate is much smaller than the widths in Figures 6.2(a) and 6.2(c) for the x and y coordinates.

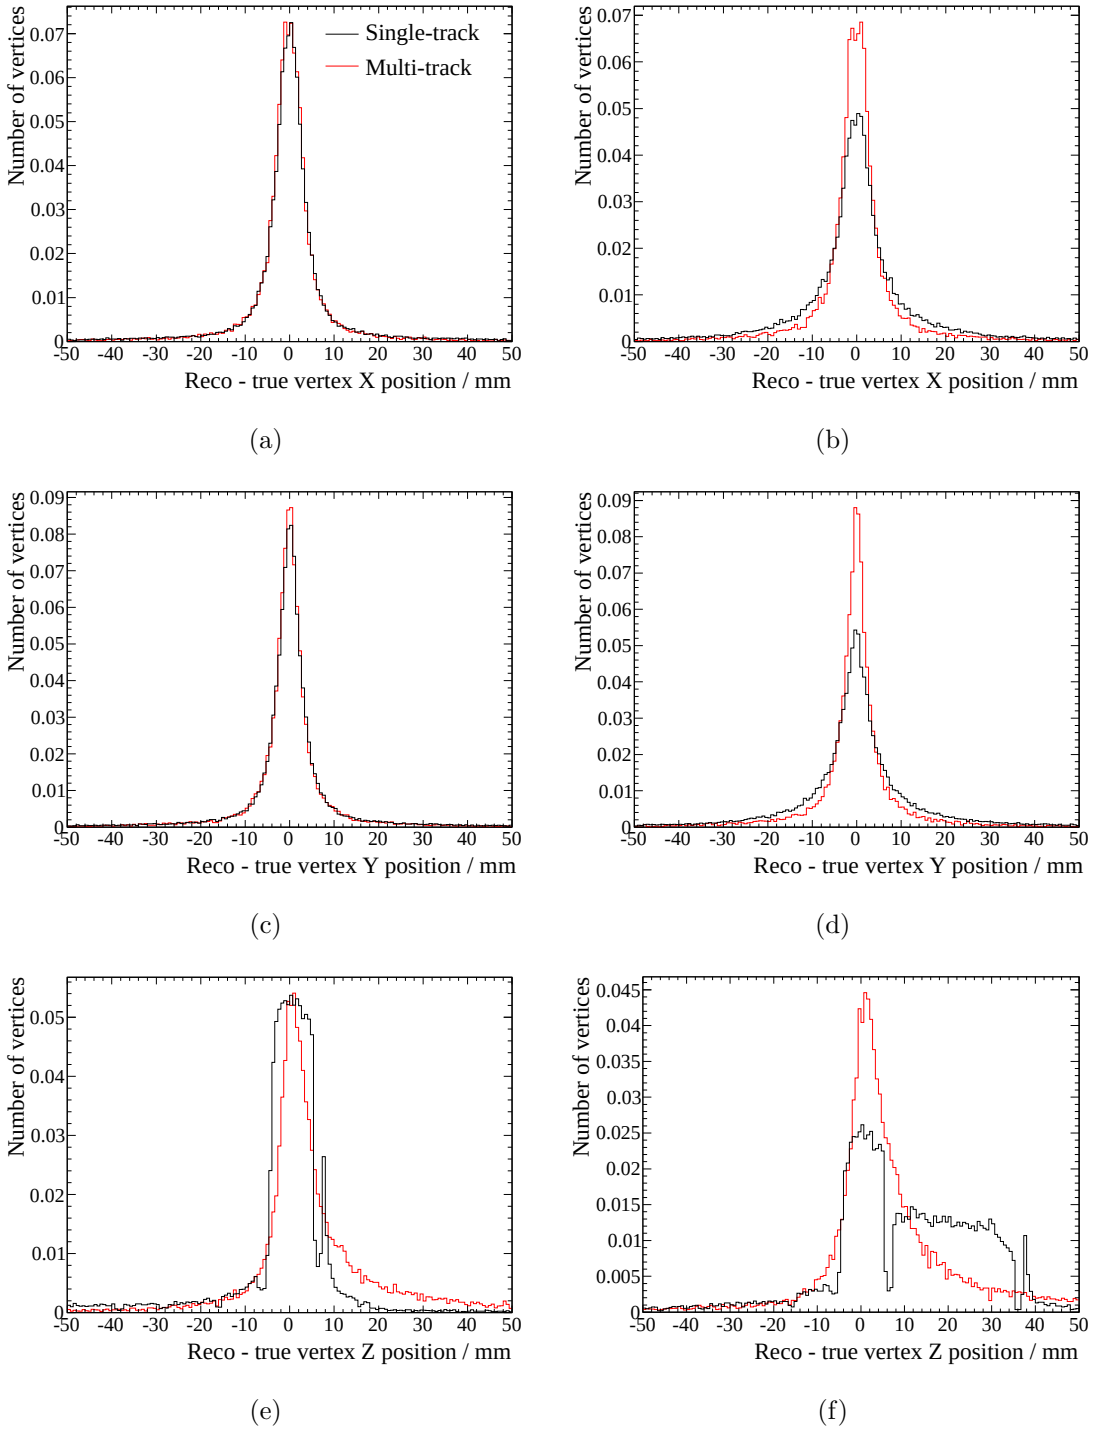


Figure 6.1: Plots showing the difference between the reconstructed global vertex position and the true vertex position for both single track and multi-track vertices. The x, y and z coordinates are shown separately. The left three plots are for true vertices in FGD1, and right for FGD2. The vertex resolution is better for multi-track vertices.

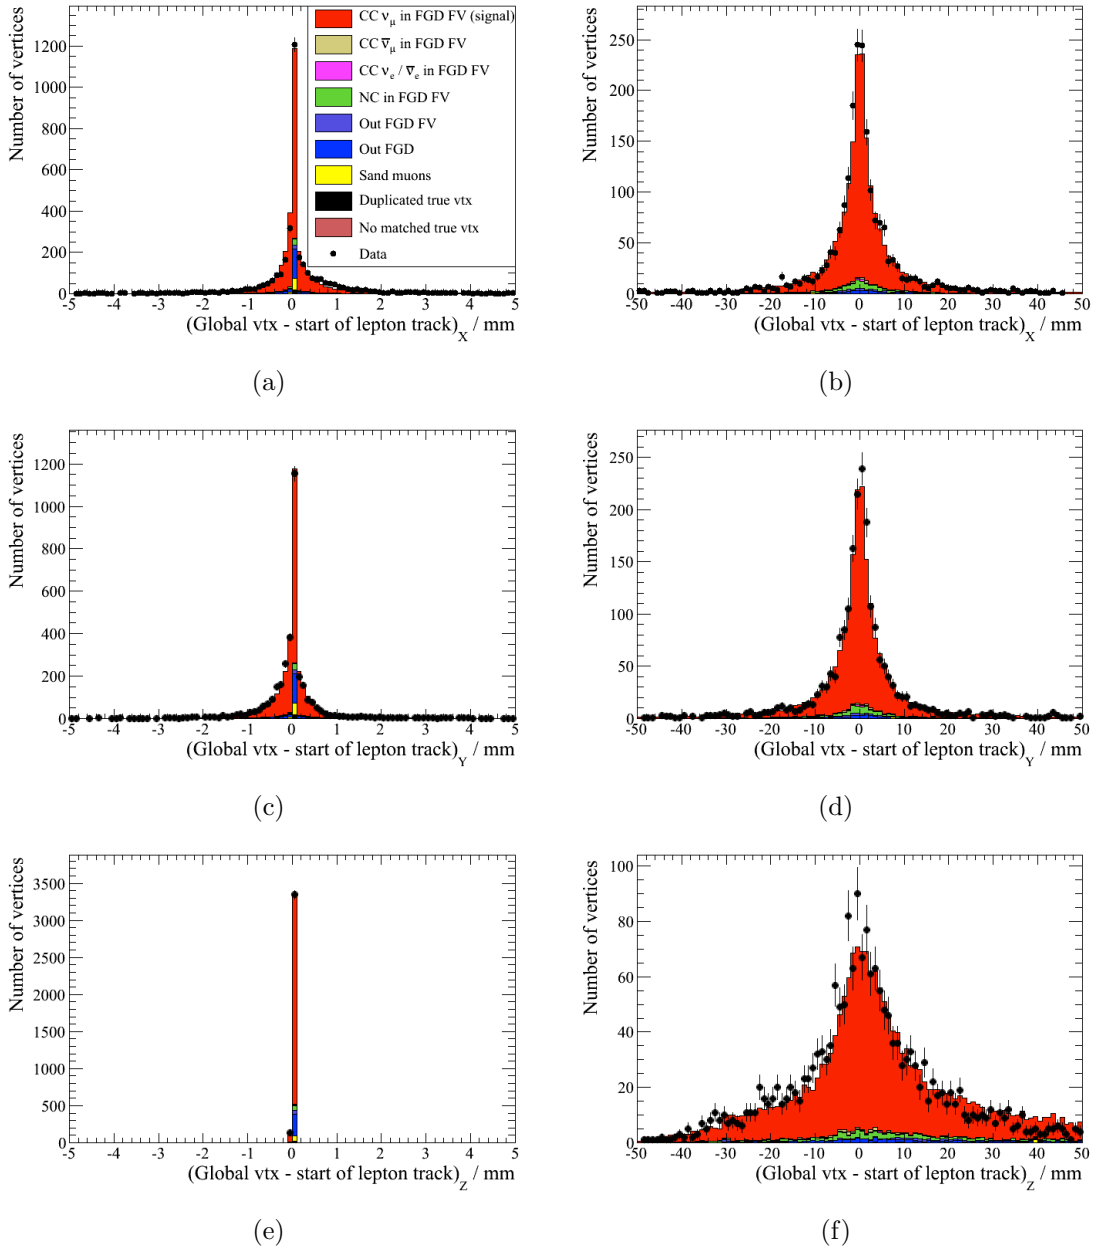


Figure 6.2: Plots showing the data-MC agreement for the difference between the global vertex position, used in this analysis, and the upstream end of the lepton track (highest momentum negative track) as used by the official ND280 analysis. The x, y and z coordinates are shown separately. Single track vertices are shown on the left, and multi-track vertices on the right. Good data-MC agreement is observed.

6.6.2 Selection cuts

The CC-inclusive ν_μ selection cuts focus on selecting CC ν_μ interactions that occur within the FGDs and whose particles then enter the TPCs. The cuts are briefly summarised in the list below, following which there is a more detailed explanation for some of the cuts as necessary:

(1) **Vertex is primary**: The global vertex must be reconstructed as a primary vertex, as described in Section 4.2.2.

(2) **Vertex in FGD fiducial volume**: The global vertex must be reconstructed with a position inside the fiducial volume of one of the two FGDs.

(3) **Vertex with time inside a beam bunch**: The global vertex must be reconstructed with a time inside one of the eight beam bunches.

(4) $\geq$ **1 track with a good quality TPC segment**: At least one of the tracks associated with the global vertex must have a good quality TPC segment (where a TPC segment is simply the part of the global track within a TPC).

(5) **Suitable lepton candidate**: Of the tracks with at least one good quality TPC segment, one must have a negative charge; this is selected as the lepton candidate. If there is more than one such track, the one with the highest momentum is selected as the lepton candidate.

(6) **Lepton starts in same FGD**: The lepton candidate track must start inside the fiducial volume of the same FGD as that within which the global vertex lies.

(7) **Lepton passes particle identification cuts**: The lepton candidate track must pass cuts based on the TPC particle identification quantities.

For selection cut (2), the FGD fiducial volumes (FV) are defined as follows, where all distances are defined in mm from the origin of the global ND280 coordinate system:

FGD1 FV: $-874.51 < x < 874.51$, $-819.51 < y < 929.51$ and $136.88 < z < 446.96$,

FGD2 FV: $-874.51 < x < 874.51$, $-819.51 < y < 929.51$ and $1500.0 < z < 1807.1$.

The x and y criteria are the same for both FGD1 and FGD2, whilst the z criteria are different for obvious reasons. As described in Section 2.2.3, the FGDs consist of

alternating layers of horizontal and vertical bars, with 192 bars in each layer. The x and y criteria are chosen to exclude 5 bars on either end of each layer. The z criteria are placed to exclude the most upstream pair of horizontal and vertical layers, but include all remaining layers.

Figures 6.3 and 6.4 show the effect of applying the FV cuts for vertices passing cut (1) in FGD1 and FGD2 respectively. For each coordinate, the vertex positions are plotted before and after the FV cut in that coordinate is applied (the FV cuts in the other two coordinates are applied). It can be seen that the Out of FGD and sand muon backgrounds are very large towards the edges of the FGDs, and that the FV cuts are placed so as to remove these events. There seems to be a data deficit within the FV. This deficit is more significant for vertices in FGD2, which could possibly indicate a lower cross-section for neutrino interactions on water than is modelled in the MC (FGD2 contains layers of water unlike FGD1). Alternatively, it could indicate a difference in particle kinematics between the data and MC which would affect the two FGDs differently due to their different angular acceptances. For both FGD1 and FGD2, there is a non-negligible fraction of vertices with a duplicated true vertex. As discussed previously, a duplicated true vertex can happen if, for example, the tracks from a single true vertex are mis-reconstructed as coming from two separate vertices. If the fraction of these events remained large, it would be necessary to study them in detail to confirm agreement between data and MC. However, it will be seen later in Table 6.1 that the fraction drops to less than 1% after all cuts are applied.

For selection cut (3), Figure 6.5 shows the times of vertices passing cuts (1) and (2). The eight bunch structure of the beam spills is clear in both data and MC. The data has a double peaked structure which is due to a shift that was introduced in the beam time during Run 2. In addition, the data is offset from the MC. Taking these features into account, it is simplest to define a loose set of bunch boundaries, as shown by the dashed lines, that are chosen according to the start time and periodicity of the beam. If a vertex does not have a time within one of these eight equally spaced bunches then it is rejected; in fact, since these eight bunches lie immediately after each other, this is equivalent to requiring that the vertex have a time inside one single accepted region which ranges from the start of the first bunch to the end of the final bunch. It is worth noting that this bunch cut could be implemented in a more sophisticated way, for example, by fitting the bunch peaks and choosing narrower cuts around the peaks. However, the current

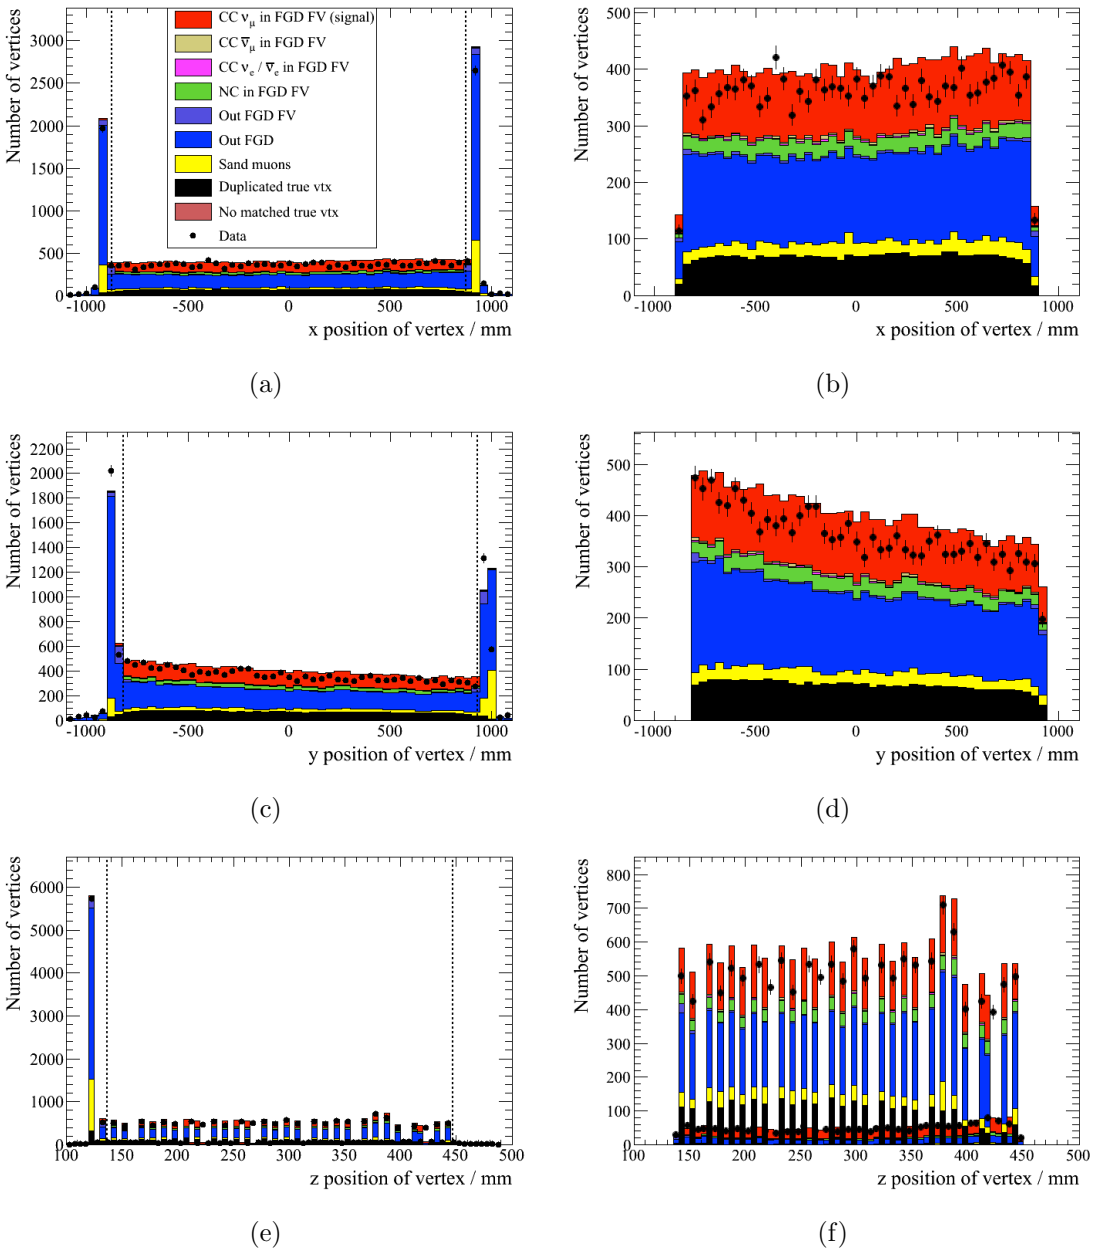


Figure 6.3: The effect of the fiducial volume (FV) cuts for primary vertices in FGD1. (a), (c) and (e) show the x, y and z positions of the vertex respectively; in each case, all FV cuts are applied except for the coordinate which is being plotted. The cut for the plotted coordinate is shown with dashed lines. A loose cut is applied for the plotted coordinate which extends a few cm beyond the edge of the FGD, to exclude vertices reconstructed in other parts of the detector where the vertexing algorithm is unvalidated. (b), (d) and (f) show the x, y and z positions of the vertex respectively once all of the FV cuts have been applied. It can be seen that by applying the FV cuts, a large fraction of the Out of FGD background is removed. The MC is normalised to the data POT.

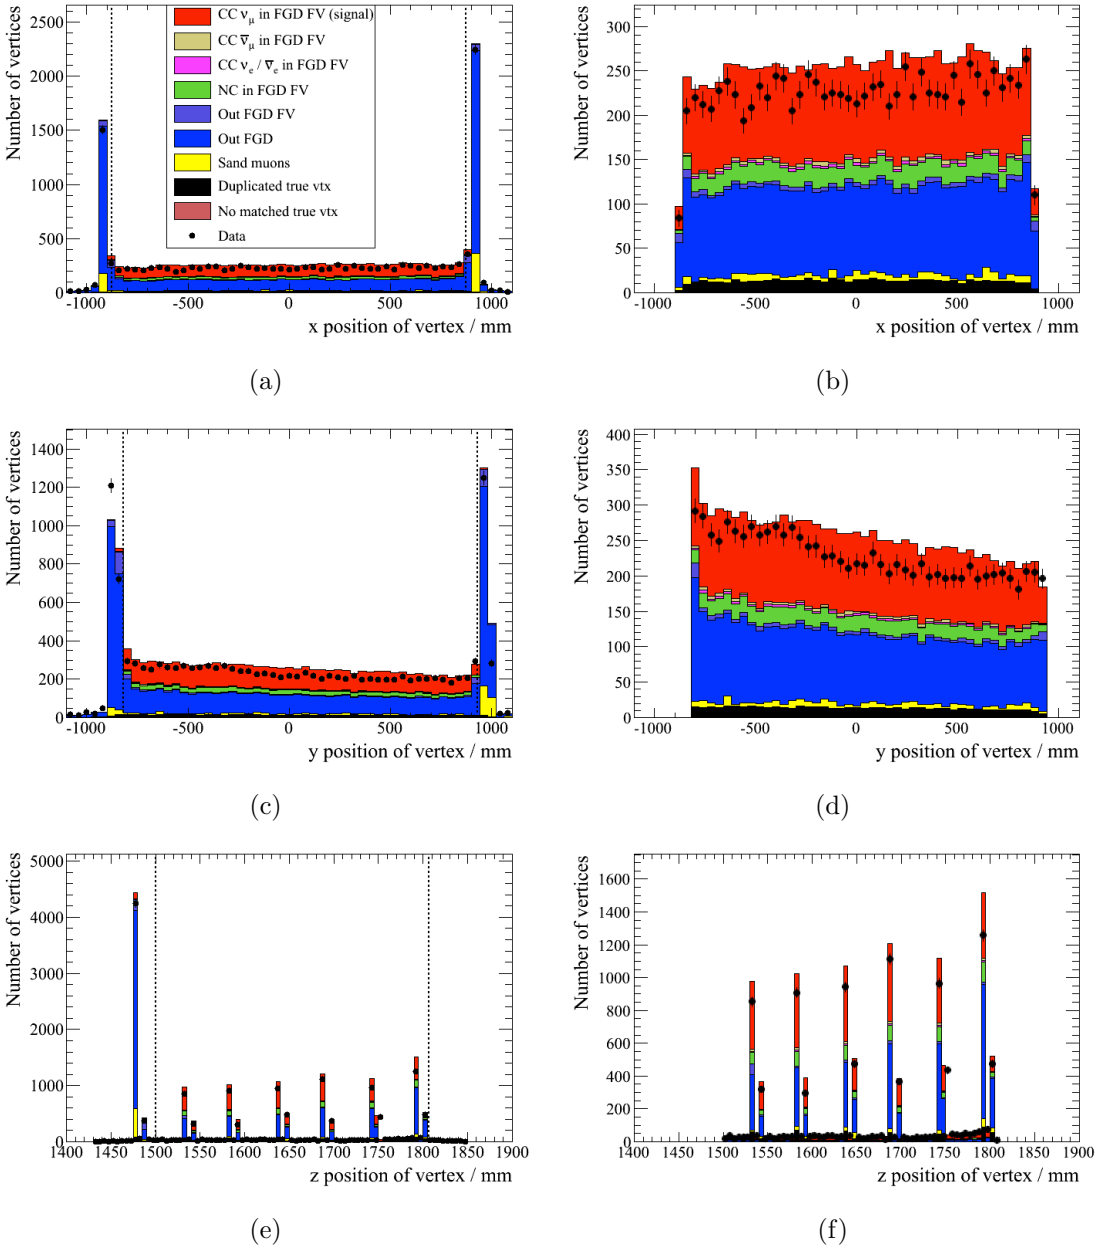


Figure 6.4: The effect of the fiducial volume (FV) cuts for primary vertices in FGD2. (a), (c) and (e) show the x, y and z positions of the vertex respectively; in each case, all FV cuts are applied except for the coordinate which is being plotted. The cut for the plotted coordinate is shown with dashed lines. A loose cut is applied for the plotted coordinate which extends a few cm beyond the edge of the FGD, to exclude vertices reconstructed in other parts of the detector where the vertexing algorithm is unvalidated. (b), (d) and (f) show the x, y and z positions of the vertex respectively once all of the FV cuts have been applied. It can be seen that by applying the FV cuts, a large fraction of the Out of FGD background is removed. The MC is normalised to the data POT.

method is simple and robust when taking into account the aforementioned features in the time distributions.

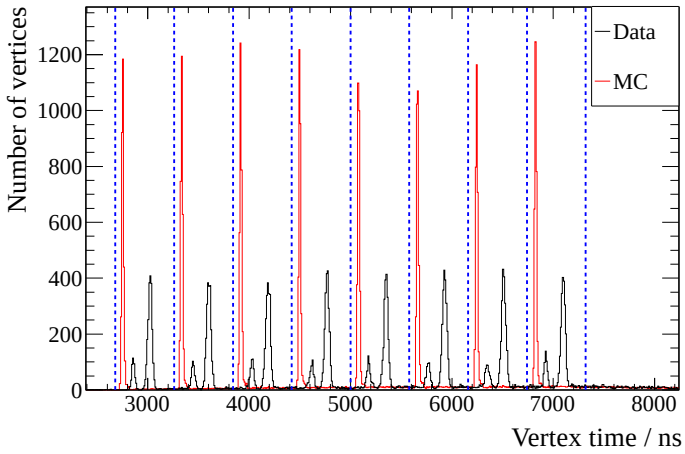


Figure 6.5: Plot of vertex times. Only vertices passing selection cuts (1) and (2) are included (i.e. primary vertices in the FGD fiducial volume). The MC is normalised to the data POT.

For selection cut (4), a good quality TPC segment is defined as one containing 18 or more hits. This cut rejects short tracks for which the reconstruction is less reliable. A study within the internal T2K ν_μ group shows that there is a significant bias in the reconstructed momentum for tracks with less than 18 hits [102].

Selection cut (6) rejects a vertex if the start position of the lepton candidate track is not reconstructed (according to the global reconstruction) inside the same FGD fiducial volume as the vertex. The global vertexing algorithm does not use the track direction information. The aim of this cut is therefore to remove vertices where a particle enters the FGD from outside but is clustered into a vertex. It will be seen later in Table 6.1 that this cut does indeed reduce the Out of FGD background from 24.2% to 20.5% of the selected events.

Selection cut (7) is based upon the particle identification (PID) quantities provided by the TPC detectors. As described previously in Section 4.2.1, the TPC provides PID discrimination using the measured values of the momentum and energy loss of the traversing particle. A pull quantity is defined which measures the number of standard deviations that the measured energy loss, $(dE/dx)^{\text{measured}}$, is away from the expected value

for a given particle hypothesis of type α at the observed momentum, $(dE/dx)^{\alpha}_{\text{expected}}$:

$$\text{Pull}^{\alpha} = \frac{(dE/dx)_{\text{measured}} - (dE/dx)^{\alpha}_{\text{expected}}}{\sigma^{(dE/dx)_{\text{measured}} - (dE/dx)^{\alpha}_{\text{expected}}}}, \tag{6.1}$$

where $\alpha = e, \mu, \pi$ or p. Figure 6.6 shows the pull distributions for all good quality TPC segments of the lepton candidate track, for the four different particle hypotheses. There is good agreement between data and MC. In each case, a value close to zero indicates that the particle hypothesis is likely to be correct, whereas a value far from zero indicates the opposite. It can be seen in Figures 6.6(b) and 6.6(c) that muons and pions are basically indistinguishable, due to their extremely similar energy loss curves, as shown previously in Figure 4.2.

As opposed to simply applying cuts to the above pull distributions, better background rejection can be achieved by constructing likelihoods for each particle hypothesis from the pulls. The likelihood of the particle being of type α, L^{α}, given the observed values of momentum and pull in the TPCs, is defined as follows:

$$\begin{aligned} L^{\alpha} &= L\left(\alpha | \{p_t, \text{Pull}^{\beta}_t\}\right) \\ &= \frac{\prod_{t=1}^{3} P_t\left(\text{Pull}^{\alpha}_t | p_t, \alpha\right)}{\sum_{\beta}\left(\prod_{t=1}^{3} P_t\left(\text{Pull}^{\beta}_t | p_t, \beta\right)\right)}, \end{aligned} \tag{6.2}$$

where t is the set of TPC indices ($t = 1, 2, 3$)[8], β is the set of particle hypotheses ($\beta = e, \mu, \pi, p$), p_t is the measured momentum of the particle in TPC t and Pull^{β}_t is the calculated pull for a particle hypothesis β, as defined in Equation 6.1, in TPC t. The variable $P_t\left(\text{Pull}^{\beta}_t | p_t, \beta\right)$ is the probability of observing a pull value Pull^{β}_t in TPC t for a particle of type β and measured momentum p_t. This distribution is assumed to be a Gaussian (centred on 0 with a width of 1), and is defined as follows:

$$P_t\left(\text{Pull}^{\beta}_t | p_t, \beta\right) = e^{-(\text{Pull}^{\beta}_t)^2/2}. \tag{6.3}$$

[8]If the particle does not traverse all three of the TPCs, then any summation over t will only sum over the TPCs that have been traversed.

A quantity, L^{MIP}, is constructed from the values of L^{α}, in order to discriminate between minimum ionising particles (i.e., muons and pions) and electrons:

$$L^{\mathrm{MIP}} = \frac{L^{\mu} + L^{\pi}}{L^{\mu} + L^{\pi} + L^{e}} = \frac{L^{\mu} + L^{\pi}}{1 - L^{p}}. \tag{6.4}$$

Two PID cuts are applied to the likelihoods of the lepton candidate track:

(1) $L^{\mu} > 0.05$,

(2) $L^{\mathrm{MIP}} > 0.8$, if $p < 500$ MeV/c.

The first cut rejects pions, protons and electrons, as shown in Figure 6.7(a). The second cut is designed to reject electrons, as shown in Figure 6.7(b). The second cut is only applied at low momentum, where the electrons are concentrated. As the momentum increases beyond 500 MeV/c, the energy loss curves of the muon and electron get closer together, resulting in an efficiency loss of true muons if the cut is applied in this region.

6.6.3 Performance of event selection

As a first basic check of stability, the number of events selected by the CC-inclusive ν_{μ} cuts is plotted as a function of the integrated POT, as shown in Figure 6.8. The Kolomogorov-Smirnov statistical test is used to determine whether this plot is compatible with a constant event rate per POT, yielding a probability of 86.9%.

Table 6.1 shows the performance of the CC-inclusive ν_{μ} cuts evaluated on MC. The number of selected events after each cut is shown, scaled to the data POT (7.837×10^{19} after data quality checks are applied). The purity and efficiency of the signal events are shown after each cut, along with the purity of each background category and its fraction of the total background. The purity of events of true type α, η^{α}, in the selected sample is given by:

$$\eta^{\alpha} = \frac{N^{\alpha}_{\mathrm{sel}}}{N_{\mathrm{totsel}}}, \tag{6.5}$$

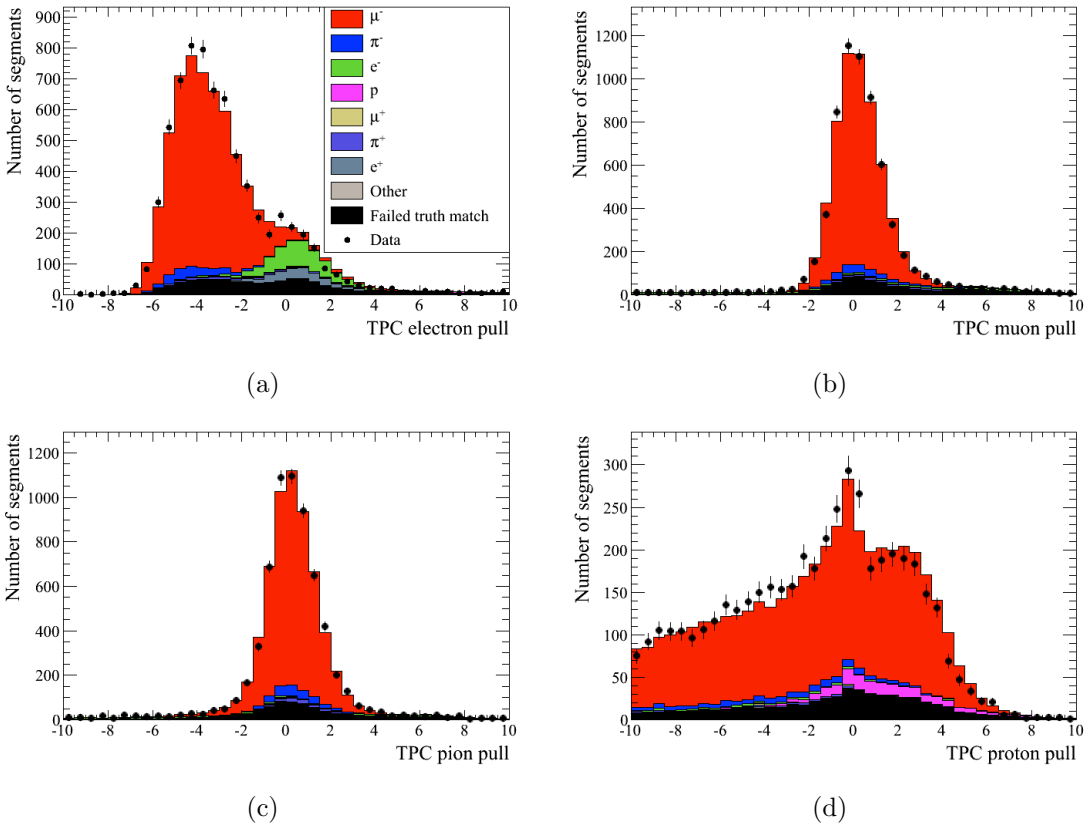
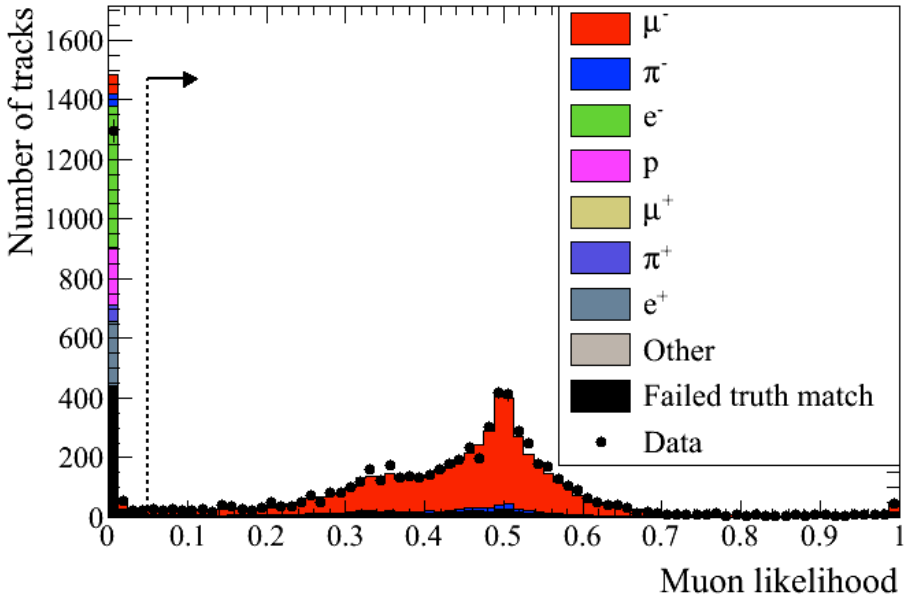
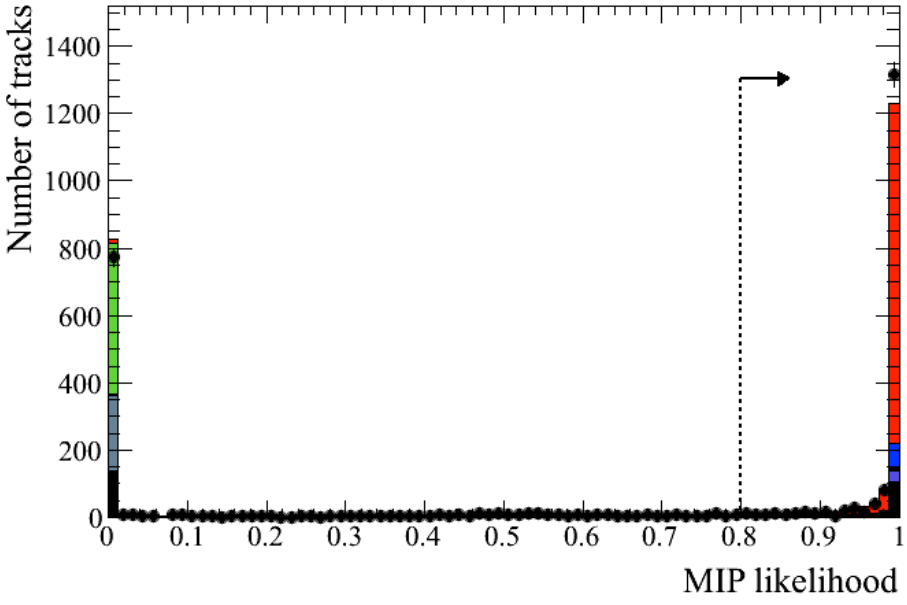


Figure 6.6: TPC pull distributions for all good quality TPC segments of the lepton candidate track: (a) electron pull, (b) muon pull, (c) pion pull and (d) proton pull. Only vertices passing selection cuts (1)–(6) are included. The MC is normalised to the data area.



(a)



(b)

Figure 6.7: TPC particle identification likelihoods for the lepton candidate track: (a) muon likelihood, and (b) MIP likelihood. For (b), only leptons with reconstructed momenta below 500 MeV/c are plotted. Only vertices passing selection cuts (1)–(6) are included. The MC is normalised to the data area.

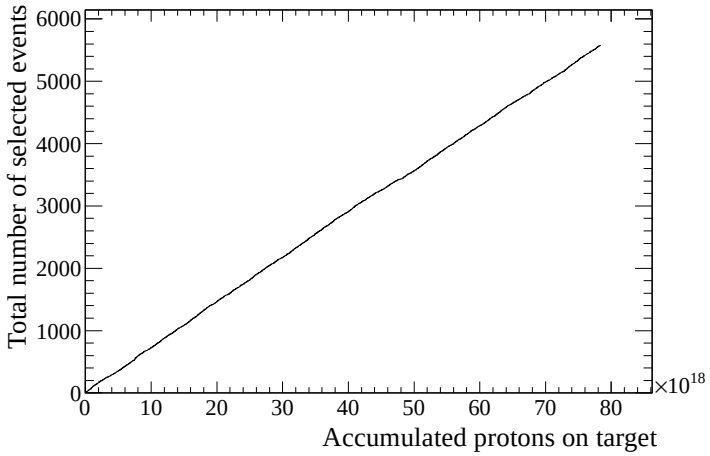


Figure 6.8: The number of events selected by the CC-inclusive ν_μ cuts, as described in Section 6.6.2, as a function of the integrated POT. Comparison of this plot with a straight line through the origin using a Kolomogorov-Smirnov statistical test gives a probability of 86.9%, indicating that this plot is compatible with a constant event rate per POT.

where N_{totsel} is the total number of selected events and N^α_{sel} is the number of those events with true type α. The efficiency of events of true type α, ϵ^α, is given by:

$$\epsilon^\alpha = \frac{N^\alpha_{\text{sel}}}{N^\alpha_{\text{tot}}}, \tag{6.6}$$

where N^α_{tot} is the total number of events of true type α before any selection cuts are applied and N^α_{sel} is as above.

The cuts perform as desired, with each cut causing an increase in the signal purity. The cuts reduce the different background categories as expected; for example, ν_e and $\bar{\nu}_e$ interactions are mostly removed by the PID cut, whilst $\bar{\nu}_\mu$ interactions are mostly removed by cut (5) since it requires a negative track. The two main background categories after all cuts are applied are: *a*) interactions which occur outside the FGD FV but are mis-reconstructed as inside (8.3% including both Out of FGD and Out of FGD FV events); and *b*) NC interactions in the FGD FV (2.7%). The NC events are most likely inelastic neutrino interactions with a nucleus where a pion or several pions are emitted. Some of these events (presumably π^+-producing interactions) are rejected by cut (5) since a negative track is required, with further events removed by the PID requirements of cut (7). The dominant background category, where interactions outside the FGD

FV are mis-reconstructed as inside, comes from several different sources. These can be neatly divided into:

- Neutral particles: If a neutral particle from an interaction outside the FGD FV enters the FGD FV it will produce no track itself but can undergo a secondary interaction and produce charged particles. These charged particles can then be reconstructed into a vertex.
- Charged particles: These are events where a neutrino interaction outside of the FGD FV produces charged particles that enter the FGD FV and are mis-reconstructed as starting inside the FGD FV.

The first of these, i.e. neutral particles, form an irreducible background which the reconstruction cannot be expected to identify. The second, on the other hand, represents a variety of different reconstruction failures. For example, a charged particle that stops in the FGD FV can be mis-reconstructed as starting inside, or a through-going charged particle which undergoes a hard scatter inside the FGD FV can be mis-reconstructed as two separate tracks with a vertex at the kink. Another quite frequent failure mode arises from interactions that occur in the dead material between an FGD and TPC producing both a backward and forward-going particle; these particles can enter the FGD and TPC respectively resulting in two tracks that are mis-reconstructed as a single track starting in the FGD.

Figure 6.9 shows the purity and efficiency of signal interactions in the final CC-inclusive ν_μ sample as a function of true neutrino energy, true lepton momentum and cosine of the true lepton angle. For the purity in Figures 6.9(b) and 6.9(c), the true lepton for the background events is either an electron or muon for CC events, or the outgoing neutrino for NC events. It should be noted that the sand muon events are not included in these plots, since no truth information is available for these events. The efficiency increases with all three of these quantities. As the neutrino energy increases, the lepton tends to have a higher momentum and be more forward-going, which leads to a higher efficiency because the lepton must travel forward from the FGD into the TPC in order to pass the selection cuts. In addition, high momentum tracks are easier to reconstruct, as are forward-going tracks. The efficiency drops almost to zero for leptons that are emitted perpendicular to the z-axis, since it is very hard to reconstruct these

		Signal		Background (BG)													
		ν_μ CC in FGD FV		$\bar{\nu}_\mu$ CC in FGD FV		$\nu_e/\bar{\nu}_e$ CC in FGD FV		NC in FGD FV		Out of FGD FV		Out of FGD		Sand muons		Duplicated true vertex	
Cuts	$N_{\text{Total}}^{\text{POT norm}}$	η	ϵ	η	% of BG	η	% of BG	η	% of BG	η	% of BG	η	% of BG	η	% of BG	η	% of BG
1) Primary vertex	216794.1	5.1%	91.9%	0.1%	0.1%	0.1%	0.1%	1.1%	1.1%	1.2%	1.3%	58.1%	61.3%	31.4%	33.0%	2.8%	3.0%
2) & in FGD FV	28981.5	33.6%	80.2%	1.0%	1.5%	0.7%	1.1%	6.9%	10.4%	1.4%	2.1%	39.3%	59.2%	4.7%	7.0%	12.5%	18.8%
3) & in time bunch	27054.1	35.8%	79.8%	1.0%	1.6%	0.8%	1.2%	7.3%	11.4%	1.5%	2.3%	40.6%	63.2%	4.1%	6.3%	8.9%	13.9%
4) & ≥ 1 track with good TPC	16592.9	48.4%	66.2%	1.5%	3.0%	1.1%	2.2%	8.1%	15.8%	1.6%	3.1%	34.0%	66.0%	2.6%	5.0%	2.6%	5.0%
5) & lepton candidate	9159.3	64.1%	48.4%	0.4%	1.0%	1.3%	3.5%	3.5%	9.8%	1.7%	4.6%	24.2%	67.4%	2.8%	7.7%	2.1%	5.9%
6) & lepton starts in same FGD	8318.3	68.4%	46.9%	0.4%	1.2%	1.3%	4.2%	3.6%	11.5%	1.5%	4.8%	20.5%	64.9%	2.3%	7.1%	2.0%	6.3%
7) & lepton passes particle-ID	6216.2	86.3%	44.2%	0.4%	2.8%	0.3%	2.3%	2.7%	19.5%	1.5%	11.0%	6.8%	49.8%	1.4%	10.4%	0.6%	4.2%

Table 6.1: Performance of the CC-inclusive ν_μ selection cuts as evaluated on MC. The leftmost column lists the cuts, as detailed in Section 6.6.2. The next column shows the total number of events selected after each cut, scaled to the data POT. The remaining columns show the breakdown of selected events after each cut into signal and background categories. Both the purity, η, and efficiency, ϵ, are shown for the signal events, along with the purity of the various background categories and their fraction of the total background. The "no matched true vertex" category is omitted here since it is negligible.

high angle tracks which essentially travel along the scintillator planes as opposed to through them.

The signal purity drops off as the neutrino energy increases, as shown in Figure 6.9(a). Figure 6.10(a) shows the breakdown of signal and background categories in the selected sample as a function of this quantity. Figure 6.11(a) shows the same as Figure 6.10(a) but with each bin normalised to 1.0, such that the fraction of a given category within each bin can be easily read off. It can be seen that an increasing fraction of the events come from both NC interactions in the FGD FV and Out of FGD events as the neutrino energy rises. For the NC background, inelastic NC interactions turn on at higher neutrino energies since a larger momentum transfer is available, giving rise to pions (or other particles) that can be (incorrectly) selected as the lepton candidate. In addition, for both of these background categories, a higher neutrino energy will tend to give rise to a higher momentum lepton candidate track. However, the charge reconstruction (which relies on the curvature of the track in the magnetic field) becomes worse as the track becomes straighter due to its high momentum, and the PID also deteriorates at high momentum as the energy loss curves of the different particles become less distinguishable.

There is a low signal purity at low true lepton momentum, as shown in Figure 6.9(b). The main backgrounds again come from NC interactions in the FGD FV and Out of FGD events, as shown in Figures 6.10(b) and 6.11(b). To be clear, in the case of NC events, the true momentum of the lepton is that of the outgoing neutrino. For these NC events, a low momentum for the outgoing neutrino means that there is more momentum available for the other outgoing particles, which increases the chance of a charged particle being successfully reconstructed in the FGD (and then selected as the lepton candidate). For the Out of FGD events, since this category is an umbrella for many different types of event there could be many different reasons why these events become more dominant at low true lepton momentum; these could include the same kinematic arguments described above for NC events if particles other than the outgoing lepton enter the FGD, or the fact that various reconstruction failures become more problematic for low momentum tracks.

Finally, there is a very large fraction of Out of FGD events for interactions where the lepton travels backwards (i.e. where the cosine of the true lepton angle is < 0), as shown in Figure 6.11(c), leading to a low purity in this region (as shown in Figure 6.9(c)). This is

consistent with backward-going tracks stopping in the FGD and being mis-reconstructed as starting in the FGD.

6.6.4 Final event sample

Having described the set of CC-inclusive ν_μ selection cuts and shown their performance, in Sections 6.6.2 and 6.6.3 respectively, the distributions of the final event sample are now shown. Figure 6.12 shows the reconstructed momentum and cosine of the reconstructed angle for the lepton candidate, with the MC broken down by true interaction type. Figure 6.13 shows the same distributions but with the MC broken down by neutrino parent. It should be noted that the data was not overlaid on the MC until both the selection cuts and fit methodology were frozen, so as not to introduce any unintentional biases into the analysis.

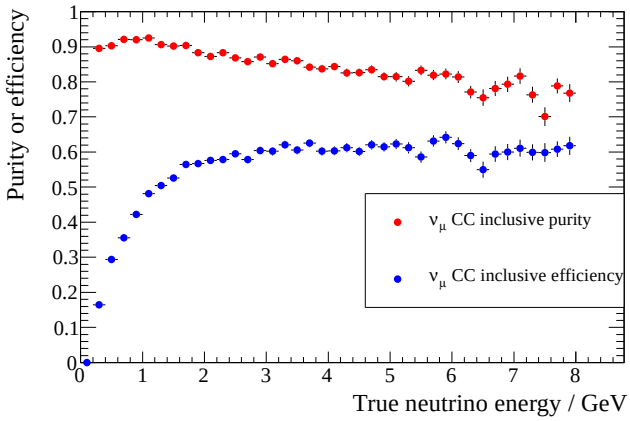
Choice of p_μ - $\cos\theta_\mu$ binning for the fit

The final event sample is binned into 32 p_μ - $\cos\theta_\mu$ bins; these bins form the input to the fit, which will be described in Section 6.7. The following bins are chosen:

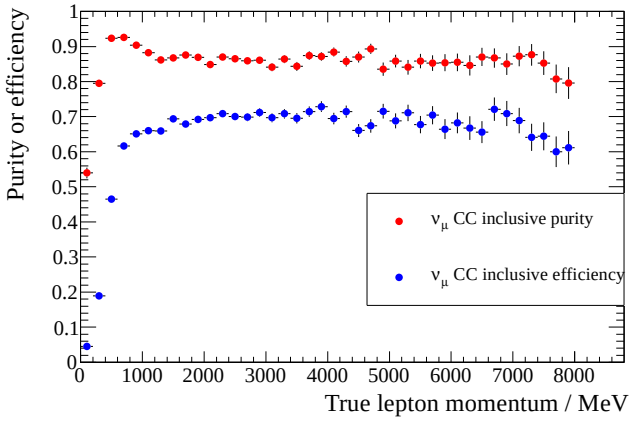
- p_μ (MeV/c): 200–400, 400–500, 500–700, 700–1000, 1000–1500, 1500–2500, 2500–4000, $>$ 4000
- $\cos\theta_\mu$: 0–0.9, 0.9–0.95, 0.95–0.98, 0.98–1.0

A bin ordering convention is chosen in which, for each p_μ bin in order, the $\cos\theta_\mu$ bins are iterated over. For example, bins 0–3 correspond to the four $\cos\theta_\mu$ bins in order with $200 < p_\mu < 400$, whilst bins 4–7 correspond to the four $\cos\theta_\mu$ bins in order with $400 < p_\mu < 500$, and so on.

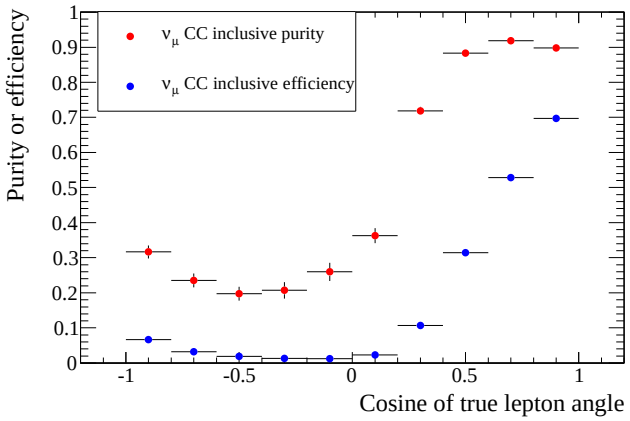
Several criteria were considered when choosing these bin boundaries. In general, the binning is made as fine as possible, without reducing any of the bin occupancies to zero. This maximises the information available to the fit. In addition, increased binning should help to break any potential degeneracies between the fit parameters, for example, between the parent normalisation parameters and the cross-section normalisation parameters (that will be described later in Section 6.7). Due to the shape of



(a)

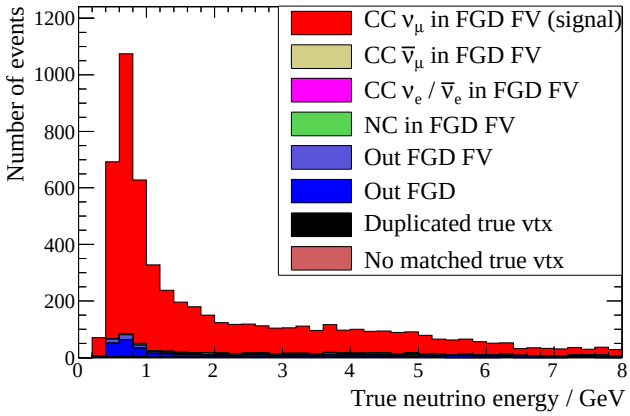


(b)

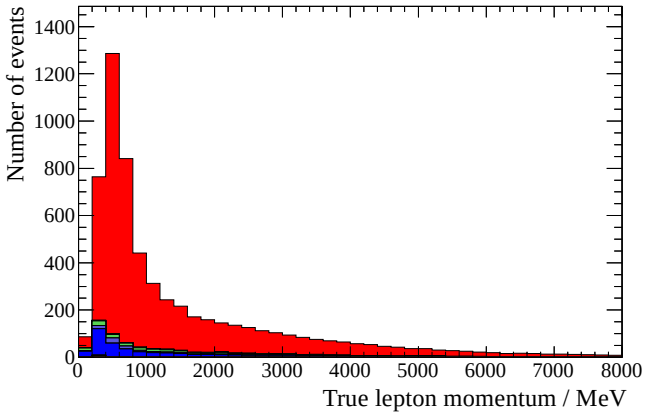


(c)

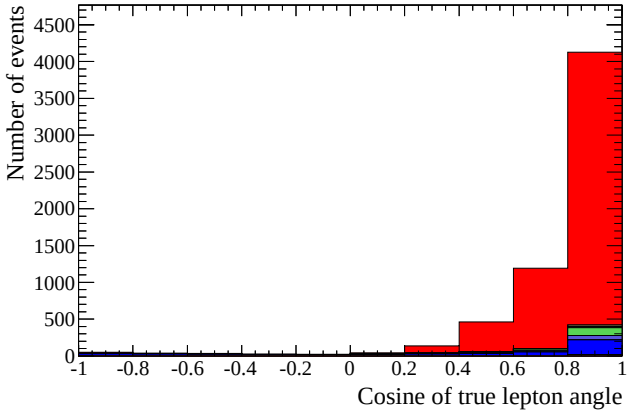
Figure 6.9: The signal purity and efficiency of the CC-inclusive ν_μ sample according to the MC, as a function of: (a) true neutrino energy, (b) true lepton momentum and (c) cosine of the true lepton angle (with respect to the z-axis of the ND280, as shown in Figure 2.6). For the calculation of purity in (b) and (c), the true lepton for the background events is either an electron or muon for CC events, or the outgoing neutrino for NC events. Note that sand muons are not included in the calculation of these plots, since the truth information is unavailable for these events.



(a)

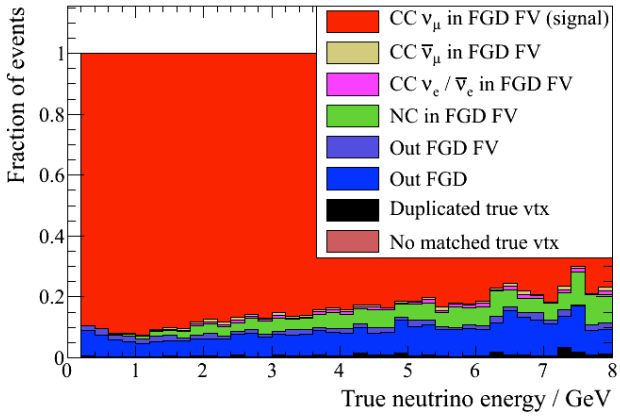


(b)

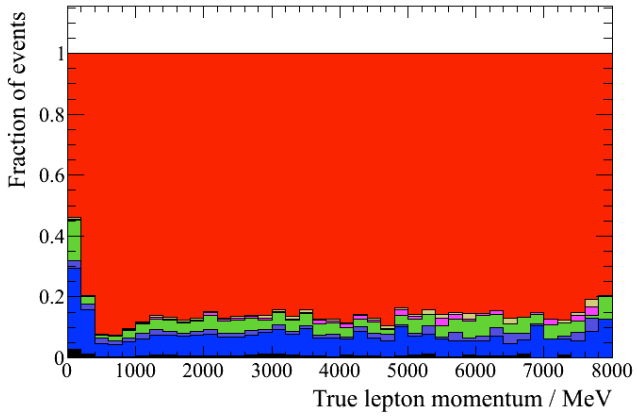


(c)

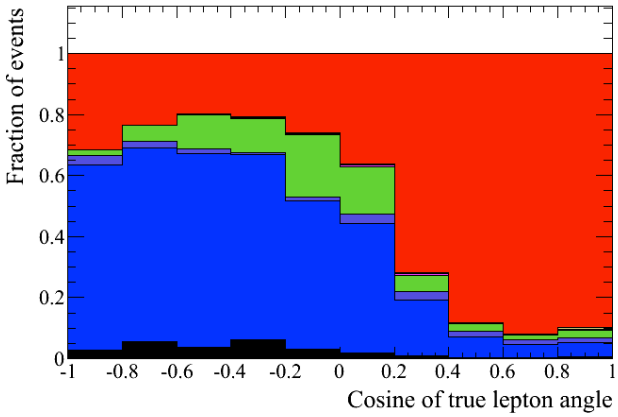
Figure 6.10: The breakdown of signal and background categories, according to the MC, in the final CC-inclusive ν_μ sample, as a function of: (a) true neutrino energy, (b) true lepton momentum and (c) cosine of the true lepton angle (with respect to the z-axis of the ND280, as shown in Figure 2.6). For the CC events, the lepton is either a muon or electron; for the NC events (both those in the FGD FV and those included in the Out of FGD categories) the lepton refers to the outgoing neutrino. Note that sand muons are not included in these plots, since the truth information is unavailable for these events.



(a)



(b)



(c)

Figure 6.11: The same plots as in Figure 6.10, but with each bin normalised to 1.0 such that the fraction of a given event category can be seen per bin.

the neutrino energy spectrum, there are significantly more events at low momentum ($p_\mu \sim 600$ MeV/c), resulting in finer binning at low momentum and coarser binning at high momentum. However, the binning in the high momentum region is kept as fine as possible, whilst avoiding empty bins, since K^+-originating neutrinos tend to have higher energies and so this is the K^+ signal region. In addition, events for which the reconstruction is unreliable are disregarded, and so backwards events ($\cos\theta_\mu < 0$) and very low momentum events ($p_\mu < 200$ MeV/c) are excluded. Figures 6.14 and 6.15 show the p_μ - $\cos\theta_\mu$ bin occupancies for both data and MC, with the MC broken down by interaction type and neutrino parent respectively. It can be seen in Figure 6.15 that π^+-originating neutrinos dominate at low momentum and high angle, whereas K^+-originating neutrinos dominate at high momentum and low angle, as would be expected due to the different energy spectra of the pion and kaon parents.

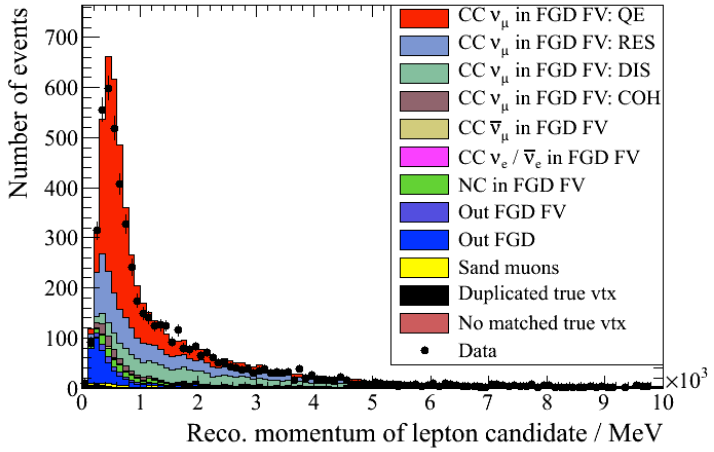
6.7 Fit methodology

6.7.1 Fit parameters

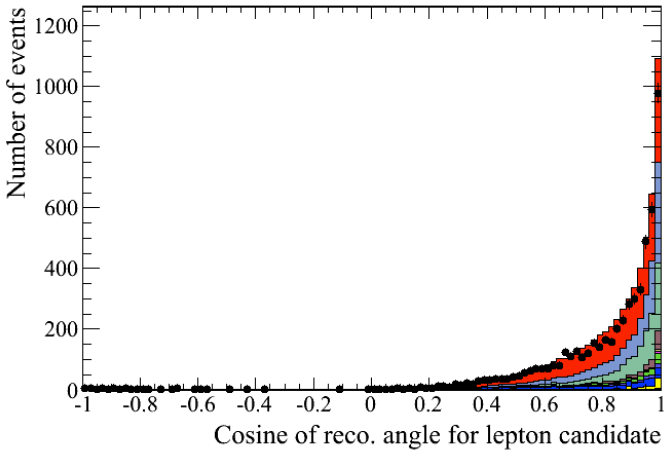
A set of fit parameters is defined. By varying these parameters, the predicted p_μ - $\cos\theta_\mu$ bin occupancies can be fit to the observed bin occupancies, as will be described in Section 6.7.2. The fit parameters are defined as:

- $\vec{b}$: Normalisation parameters for the different neutrino parent categories. There are three parameters: b_K, for K^+ parents; b_π for π^+ parents; and b_{other} including K^0_L, $\mu^\pm$, K^- and π^- parents.
- $\vec{x}$: Parameters that describe the systematic uncertainties on the underlying cross-section models
- $\vec{d}$: Parameters that model the systematic uncertainties on the detector system (the MC statistical errors are also included here, as well as some cross-section errors that are difficult to parametrise in $\vec{x}$).

The parameters of interest are b_K and b_π, whilst the b_{other}, $\vec{x}$ and $\vec{d}$ parameters are all nuisance parameters. It may be noted that there are no parameters included in the fit to describe the systematic errors on the neutrino flux prediction. This is done in order

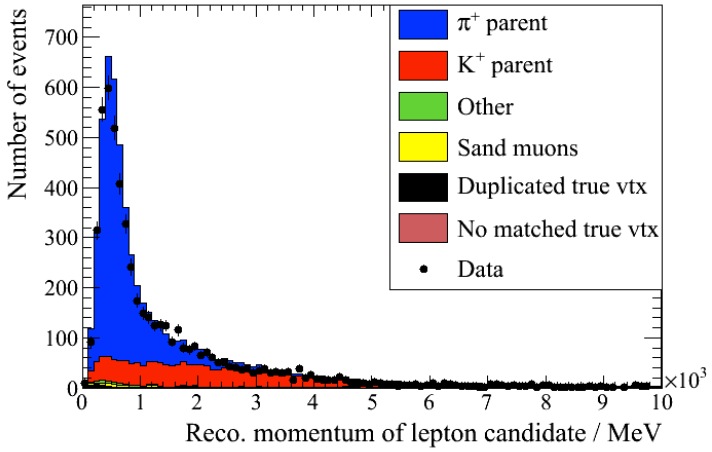


(a)

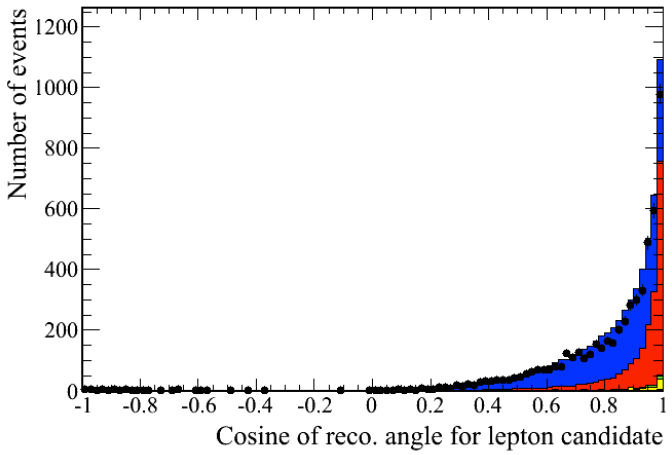


(b)

Figure 6.12: Comparison of data and MC (scaled to the data POT), with the MC broken down by interaction type, for reconstructed quantities of the lepton candidate for the final CC-inclusive ν_μ sample: (a) momentum and (b) cosine of angle (with respect to the z-axis of the ND280, as shown in Figure 2.6).



(a)



(b)

Figure 6.13: Comparison of data and MC (scaled to the data POT), with the MC broken down by neutrino parent, for reconstructed quantities of the lepton candidate for the final CC-inclusive ν_μ sample: (a) momentum and (b) cosine of angle (with respect to the z-axis of the ND280, as shown in Figure 2.6). Note that the sand muons are included as a separate category, since the neutrino parent type is not available for these events.

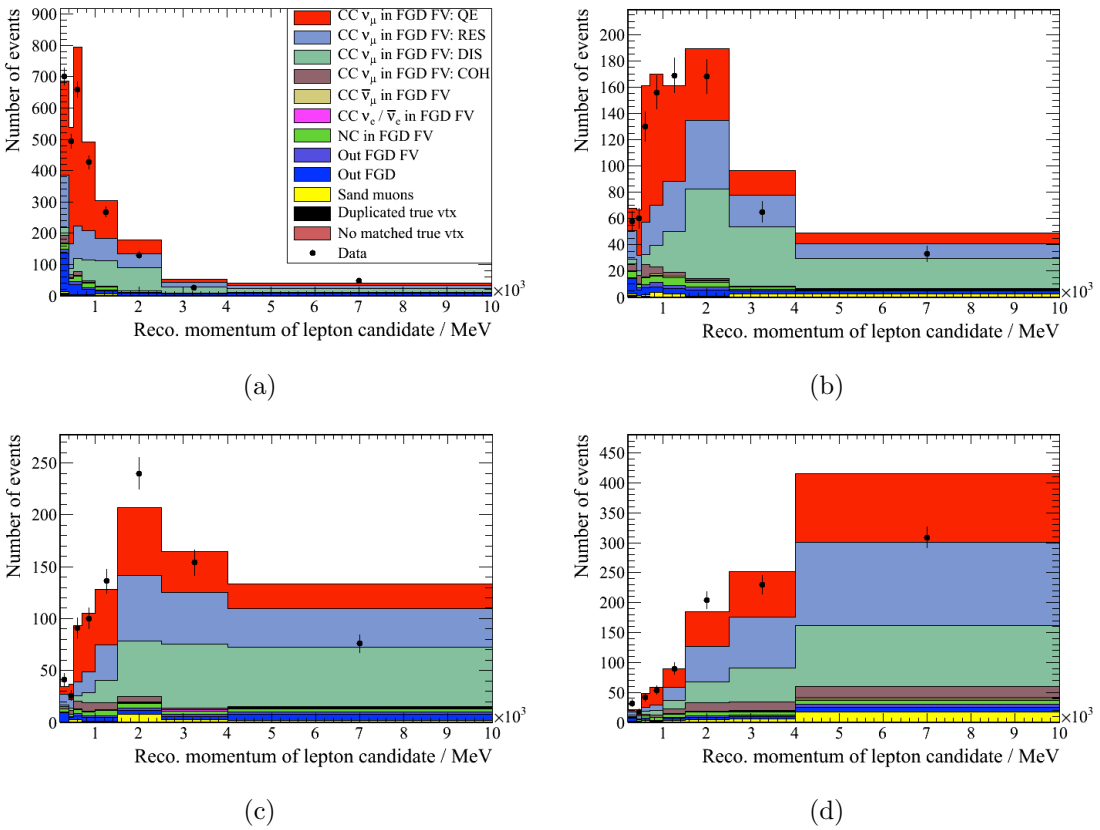


Figure 6.14: The p_μ - $\cos\theta_\mu$ bin occupancies for both data and MC (scaled to the data POT), with the MC broken down by interaction type. Figures (a)–(d) show the momentum bins for the following $\cos\theta_\mu$ bins in order: 0–0.9, 0.9–0.95, 0.95–0.98, 0.98–1.0. In each case, the final momentum bin, although only displayed up to 10 GeV/c, includes all events with $p_\mu > 4$ GeV/c.

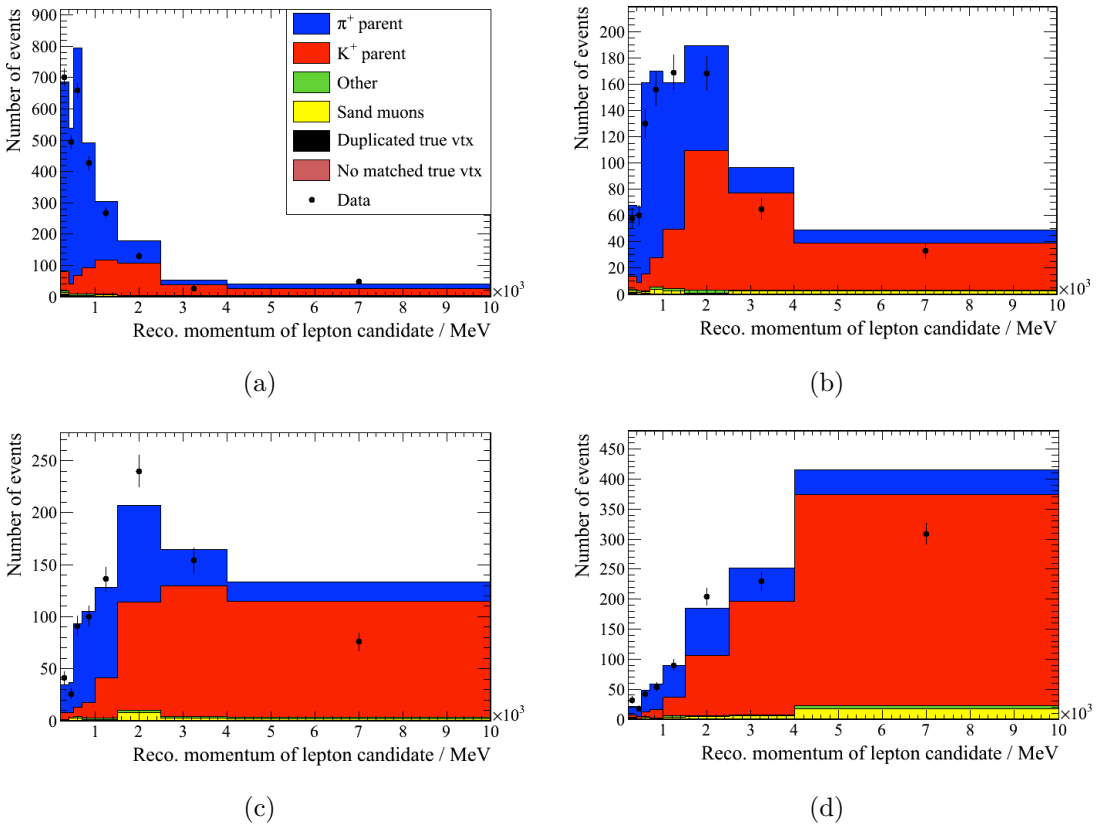


Figure 6.15: The p_μ - $\cos\theta_\mu$ bin occupancies for both data and MC (scaled to the data POT), with the MC broken down by neutrino parent. Figures (a)–(d) show the momentum bins for the following $\cos\theta_\mu$ bins in order: 0–0.9, 0.9–0.95, 0.95–0.98, 0.98–1.0. In each case, the final momentum bin, although only displayed up to 10 GeV/c, includes all events with $p_\mu > 4$ GeV/c. Note that the sand muons are included as a separate category, since the neutrino parent type is not available for these events.

to perform an unconstrained measurement of b_K, which will be discussed in more detail in Section 6.7.5.

6.7.2 Likelihood definition

A profile likelihood method is chosen for determining the values of b_K and b_π from the data. The starting point for this method is a binned maximum likelihood. The likelihood is constructed assuming N_{bins} independent bins, each with a Poisson probability of observing N_i^{d} events given an expected number of events N_i^{p} in a given p_μ - $\cos\theta_\mu$ bin i. The N_i^{p} prediction is a function of the $\vec{b}$, $\vec{x}$ and $\vec{d}$ parameters described in Section 6.7.1. A Bayesian methodology is used to include the $\vec{x}$ and $\vec{d}$ systematic uncertainties, by assuming a prior probability density function (PDF) for each of these parameter sets. The resulting likelihood is given by:

$$L(\vec{b},\vec{x},\vec{d}) = \pi_{\text{d}}(\vec{b},\vec{x},\vec{d})\pi_{\text{x}}(\vec{x}) \prod_{i}^{N_{\text{bins}}} \frac{(N_i^{\text{p}}(\vec{b},\vec{x},\vec{d}))^{N_i^{\text{d}}}}{N_i^{\text{d}}!} e^{-N_i^{\text{p}}(\vec{b},\vec{x},\vec{d})}, \tag{6.7}$$

where π_{d} and π_{x} are the prior PDFs for the detector and cross-section nuisance parameters respectively, and all other symbols are as defined previously. $\pi_{\text{x}}(\vec{x})$ is assumed to be a multivariate normal distribution, as given by:

$$\pi_{\text{x}}(\vec{x}) = \frac{1}{(2\pi)^{\frac{k}{2}}|V_{\text{x}}|^{\frac{1}{2}}} e^{-\frac{1}{2}\Delta\vec{x}(V_{\text{x}}^{-1})\Delta\vec{x}^{\text{T}}}, \tag{6.8}$$

where k is the dimension of the $\vec{x}$ parameter vector, V_{x} is a constant covariance matrix where $(V_{\text{x}})_{i,j}$ describes the covariance between the ith and jth elements of the $\vec{x}$ vector, and $\Delta\vec{x} = \vec{x} - \vec{x}_{\text{nom}}$ is the deviation of the $\vec{x}$ parameters from their nominal values. $\pi_{\text{d}}(\vec{b},\vec{x},\vec{d})$ takes the same form as $\pi_{\text{x}}(\vec{x})$:

$$\pi_{\text{d}}(\vec{b},\vec{x},\vec{d}) = \frac{1}{(2\pi)^{\frac{k}{2}}|V_{\text{d}}(\vec{b},\vec{x})|^{\frac{1}{2}}} e^{-\frac{1}{2}\Delta\vec{d}(V_{\text{d}}(\vec{b},\vec{x})^{-1})\Delta\vec{d}^{\text{T}}}, \tag{6.9}$$

where all symbols have an equivalent definition to those in Equation 6.8. However, in this case, the covariance matrix, V_{d}, is not a constant but rather a function of the $\vec{b}$ and $\vec{x}$ parameters. This is because V_{d} not only includes the detector systematic uncertainties

but also includes other errors. In particular, it includes the statistical error on the MC, as will be described later in Section 6.7.3.

In order that the likelihood function can be converted into a χ^2 statistic, as will be described later, the likelihood ratio, λ, is defined as:

$$\lambda(\vec{b},\vec{x},\vec{d}) = \frac{\pi_{\rm d}(\vec{b},\vec{x},\vec{d})\pi_{\rm x}(\vec{x})\prod_i^{N_{\rm bins}} \frac{(N_i^{\rm p}(\vec{b},\vec{x},\vec{d}))^{N_i^{\rm d}}}{N_i^{\rm d}!} e^{-N_i^{\rm p}(\vec{b},\vec{x},\vec{d})}}{\pi_{\rm d}(\vec{b}_{\rm nom},\vec{x}_{\rm nom},\vec{d}_{\rm nom})\pi_{\rm x}(\vec{x}_{\rm nom})\prod_i^{N_{\rm bins}} \frac{(N_i^{\rm d})^{N_i^{\rm d}}}{N_i^{\rm d}!} e^{-N_i^{\rm d}}}, \tag{6.10}$$

where the denominator is simply the numerator evaluated at $N_i^{\rm p} = N_i^{\rm d}$ and with the nuisance parameters equal to their nominal values. This ratio therefore evaluates the fitted likelihood (the numerator) against the maximum possible value of the likelihood (the denominator). Substituting in the expressions for $\pi_{\rm d}$ and $\pi_{\rm x}$ and simplifying, this equation becomes:

$$\begin{aligned}\lambda(\vec{b},\vec{x},\vec{d}) = {} & \frac{|V_{\rm d}(\vec{b}_{\rm nom},\vec{x}_{\rm nom})|^{\frac{1}{2}}}{|V_{\rm d}(\vec{b},\vec{x})|^{\frac{1}{2}}} e^{-\frac{1}{2}\Delta\vec{d}(V_{\rm d}(\vec{b},\vec{x})^{-1})\Delta\vec{d}^{\rm T}} e^{-\frac{1}{2}\Delta\vec{x}(V_{\rm x}^{-1})\Delta\vec{x}^{\rm T}} \\ & \times \prod_i^{N_{\rm bins}} \left(\frac{N_i^{\rm p}(\vec{b},\vec{x},\vec{d})}{N_i^{\rm d}}\right)^{N_i^{\rm d}} e^{(N_i^{\rm d}-N_i^{\rm p}(\vec{b},\vec{x},\vec{d}))}.\end{aligned} \tag{6.11}$$

It can be seen that the normalisation factors of $\pi_{\rm d}(\vec{b},\vec{x},\vec{d})$ and $\pi_{\rm d}(\vec{b}_{\rm nom},\vec{x}_{\rm nom},\vec{d}_{\rm nom})$ do not cancel, giving rise to the first term in the above equation. As mentioned previously, this is because the $V_{\rm d}$ covariance matrix includes the statistical error of the MC and is therefore a function of the $\vec{b}$ and $\vec{x}$ parameters.

According to Wilks' theorem, for a sufficiently large data set, the quantity $-2\ln\lambda$ follows a χ^2 distribution [103]. This equivalence is important since it means that this quantity can be used not only for parameter estimation but also to determine the goodness-

of-fit and for confidence level estimation. It is defined as:

$$\chi^2(\vec{b},\vec{x},\vec{d}) = -2\ln\lambda(\vec{b},\vec{x},\vec{d}) = 2\Bigg[\sum_{i}^{N_{\text{bins}}} N_i^{\text{p}}(\vec{b},\vec{x},\vec{d}) - N_i^{\text{d}} + N_i^{\text{d}}\ln(N_i^{\text{d}}/N_i^{\text{p}}(\vec{b},\vec{x},\vec{d})) + \sum_{j}^{N_{\text{xsecpars}}}\sum_{k}^{N_{\text{xsecpars}}} \Delta x_j (V_{\text{x}}^{-1})_{j,k}\Delta x_k + \sum_{l}^{N_{\text{bins}}}\sum_{m}^{N_{\text{bins}}} \Delta d_l (V_{\text{d}}(\vec{b},\vec{x})^{-1})_{l,m}\Delta d_m + \ln\left(\frac{|V_{\text{d}}(\vec{b},\vec{x})|}{|V_{\text{d}}(\vec{b}_{\text{nom}},\vec{x}_{\text{nom}})|}\right)\Bigg], \tag{6.12}$$

where N_{xsecpars} is the number of cross-section parameters (i.e., the size of the $\vec{x}$ parameter vector), and all other parameters are as defined previously.

The predicted number of events in a given p_μ - $\cos\theta_\mu$ bin i is given by:

$$N_i^{\text{p}}(\vec{b},\vec{x},\vec{d}) = d_i \sum_{\alpha}^{N_{\text{parents}}}\sum_{\beta}^{N_{\text{modes}}}\sum_{\gamma}^{N_{\text{E}_\nu\text{bins}}} b_\alpha x_\beta^{\text{norm}}(E_\gamma) w_{i\alpha\beta\gamma}(\vec{x}) T_{i\alpha\beta\gamma}, \tag{6.13}$$

where the meaning of the symbols is discussed as follows. d_i is the detector nuisance parameter for p_μ - $\cos\theta_\mu$ bin i and changes the normalisation of that observable bin according to the detector systematics (and some other contributions). The dependence of N_i^{p} on the nuisance cross-section parameters is more complicated; it is separated into energy dependent normalisations (where x_β^{norm} is the normalisation factor for interaction mode β) and parameters that are modelled with response functions (where $w_{i\alpha\beta\gamma}$ is the weighting factor for a given p_μ - $\cos\theta_\mu$, neutrino parent, interaction mode and neutrino energy bin). The detector and cross-section parametrisations are discussed in more detail in Sections 6.7.3 and 6.7.4 respectively. b_α is the normalisation factor for a given neutrino parent α; these are the parameters that are of interest in the fit. $T_{i\alpha\beta\gamma}$ is the predicted number of events for a given p_μ - $\cos\theta_\mu$, neutrino parent, interaction mode and neutrino energy bin. The N_i^{p} prediction is therefore obtained by taking the predicted number of events in the nominal MC, $T_{i\alpha\beta\gamma}$, and reweighting by the $\vec{b}$, $\vec{x}$ and $\vec{d}$ parameters which vary in the fit. It should be noted that the sand interactions in $T_{i\alpha\beta\gamma}$ are only reweighted by the $\vec{d}$ parameters, not the $\vec{b}$ or $\vec{x}$ parameters, due to the lack of truth information available for these events, as mentioned in Section 6.4.

As mentioned previously, a profile likelihood method is used for the final parameter estimation. This method is designed to effectively remove the nuisance parameters, that are not actually of interest in the final result, from the problem. These parameters include the $\vec{x}$ and $\vec{d}$, and also the b_{other} neutrino parent category. The only parameters of interest in the fit are b_K and b_π, which together are referred to as $\vec{b}_f$. This is done by defining a *profile* likelihood ratio, $\lambda_{\text{P}}(\vec{b}_f)$. This is a function only of the interesting $\vec{b}_f$ parameters and, for a given value of $\vec{b}_f$, represents the maximum value of the likelihood ratio, $\lambda(\vec{b}, \vec{x}, \vec{d})$, that can be obtained by varying the nuisance parameters. This can be expressed as:

$$\chi^2_{\text{P}}(\vec{b}_f) = -2\ln\lambda_{\text{P}}(\vec{b}_f) = \chi^2(\vec{b}_f, \hat{b}_{\text{other}}, \hat{\vec{x}}, \hat{\vec{d}}) = -2\ln\lambda(\vec{b}_f, \hat{b}_{\text{other}}, \hat{\vec{x}}, \hat{\vec{d}}), \tag{6.14}$$

where $\hat{b}_{\text{other}}$, $\hat{\vec{x}}$ and $\hat{\vec{d}}$ are the b_{other}, $\vec{x}$ and $\vec{d}$ parameters respectively that minimise the value of χ^2 (equivalent to maximising λ) for a given value of $\vec{b}_f$.

The profile minimisation is performed as follows. A 2-dimensional space of b_K versus b_π is defined. The quantity χ^2_{P} is evaluated at each point in this space by varying the nuisance parameters b_{other}, $\vec{x}$ and $\vec{d}$ and minimising χ^2 for that fixed value of $\vec{b}_f$.[9] In this way, a surface of $\chi^2_{\text{P}}(\vec{b}_f)$ is mapped out. The minimum of this surface gives the best fit value for b_K and b_π and is denoted $\chi^2_{\text{P,min}}$. The $\Delta\chi^2_{\text{P}}$ is defined as:

$$\Delta\chi^2_{\text{P}} = \chi^2_{\text{P}} - \chi^2_{\text{P,min}}, \tag{6.15}$$

where the best fit point obviously has a value of zero. The b_K-b_π space is chosen to go from 0 to 2 in both dimensions, representing a change of $\pm 100\%$ in each dimension from the value in the nominal MC (deviations beyond this seem unlikely). A granularity of 0.01 is chosen; a finer granularity is unnecessary due to the expected size of the error on the best fit values (which will be described in Section 6.8.2 on the fit validation studies).

[9] The minimisation is actually done twice, first with b_{other} fixed at 1.0 and then with it released. Since b_{other} is unconstrained in the fit, unlike $\vec{x}$ and $\vec{d}$, it can become very negative causing a negative N_i^{p} and an invalid χ^2. It is therefore initially fixed at 1.0 in order to guide the fit towards a valid minimum.

6.7.3 Detector systematic errors

Treatment of the errors in the fit

The detector systematic errors are implemented in the fit via the d parameters in Equation 6.13. The number of d parameters is equal to the number of p_μ - $\cos\theta_\mu$ bins in the fit. For a given p_μ - $\cos\theta_\mu$ bin, i, the d_i parameter scales the normalisation of that bin. These parameters are varied in the fit, with a prior constraint from the $V_{\rm d}$ detector covariance matrix. Since the d parameters are multiplicative, $V_{\rm d}$ is a fractional covariance matrix.

As mentioned previously, the $V_{\rm d}$ covariance matrix not only includes the detector systematic errors, but other contributions too. It is defined as:

$$V_{\rm d} = V_{\rm det} + V_{\rm mcstats} + V_{\rm xsec}, \tag{6.16}$$

where $V_{\rm det}$ is the matrix for the detector uncertainties exclusively, $V_{\rm mcstats}$ is the matrix for the MC statistical errors, and $V_{\rm xsec}$ is the matrix for a subset of the cross-section uncertainties which it is hard to include in the fit in alternative ways (as will be described later in Section 6.7.4). The $V_{\rm det}$ matrix is obtained by summing the covariance matrices from each individual detector systematic error separately. In actual fact, only a single detector systematic error is included in this matrix, since the rest have very small impacts on the fit, as will be explained in the following section.

The $V_{\rm mcstats}$ matrix has non-zero diagonal elements only. These elements are calculated as follows from the statistical error, $\epsilon_{i\alpha\beta\gamma}$, on each $T_{i\alpha\beta\gamma}$ bin prediction in Equation 6.13:

$$(V_{\rm mcstats})_{i,i}(\vec{b},\vec{x}) = \sqrt{\sum_{\alpha}^{N_{\rm parents}} \sum_{\beta}^{N_{\rm modes}} \sum_{\gamma}^{N_{\rm E_\nu bins}} \left[b_\alpha x_\beta^{\rm norm}(E_\gamma) w_{i\alpha\beta\gamma}(\vec{x}) \epsilon_{i\alpha\beta\gamma}\right]^2} \tag{6.17}$$

Since $\vec{b}$ and $\vec{x}$ vary in the fit, $V_{\rm mcstats}$ (and therefore $V_{\rm d}$) is recalculated after each iteration of the fit parameters.

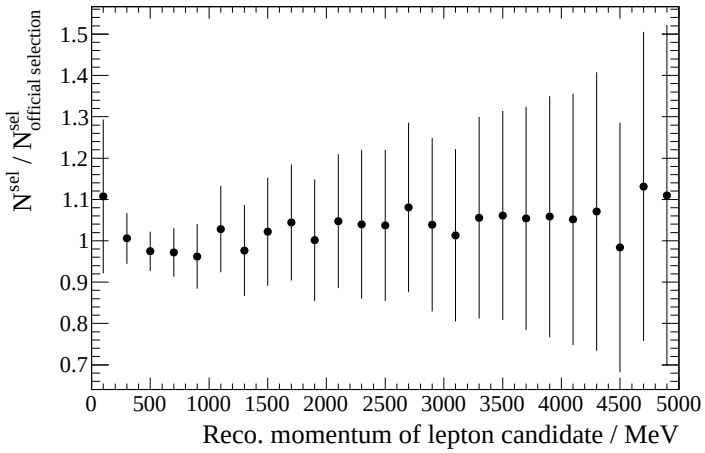
Evaluation of the errors

The detector systematic errors used in this analysis build where possible upon the outputs of the internal T2K ν_μ group. This group has performed extensive studies to evaluate the detector systematic errors at ND280 to pass as inputs to the official ND280 analysis, as mentioned in Section 6.3. A covariance matrix in p_μ - $\cos\theta_\mu$ bins has been produced for each systematic error. Ideally, these could be used to obtain V_{det} in Equation 6.16; however, this is not possible because different p_μ - $\cos\theta_\mu$ binning is used in this analysis compared to the official one. In addition, there are differences between the selection cuts used here (described in Section 6.6.2) and the "official" CC-inclusive ν_μ selection cuts. However, the outputs of this group can be used to determine the *relative importance* of the different systematic errors, meaning that only the significant ones are evaluated for this analysis. The rationale for this will be given below.

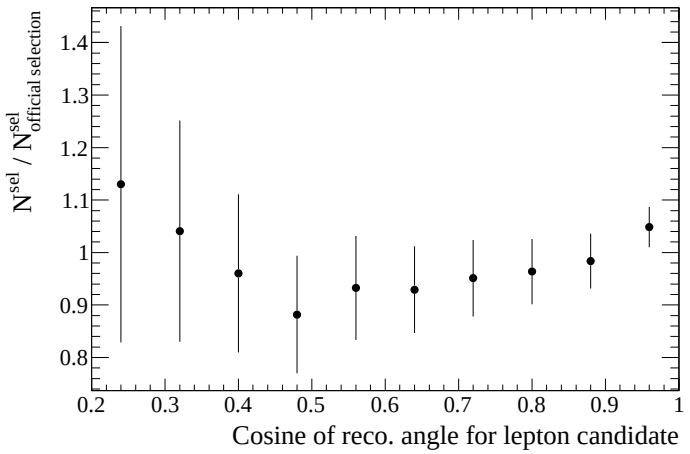
The main ways in which the selection cuts used here differ from the official ones are: *a*) the use of the global vertexing algorithm (Section 6.6.1) to determine the vertex position as opposed to simply using the start of the highest momentum negative track; *b*) the use of FGD2 events as well as FGD1 events; and *c*) no explicit veto against through-going or backward-going tracks. Figure 6.16 shows the ratio of selected events for the two sets of selection cuts as a function of both the reconstructed momentum and cosine of the reconstructed angle for the lepton candidate (the two quantities used in the fit). It can be seen that the ratio remains within approximately 5–10% of 1 as a function of both quantities, and indeed is consistent with unit value within the statistical errors. The selections are therefore deemed sufficiently similar such that the detector systematic errors evaluated for the official selection can be used to determine the relative sizes of the systematic errors for this analysis. In this way, it can be decided which systematic errors should be evaluated and which are small enough to neglect.

Table 6.2 summarises the detector systematic errors evaluated for the official selection cuts.[10] The first three columns show the source of systematic error, the sample used to evaluate it and the size of the error respectively. For some sources of systematic, there is no single error available; for example, the Out of FGD events are made up of many different failure modes all of which are evaluated separately. The final column shows the

[10]The official analysis splits the CC-inclusive ν_μ events into two samples (CCQE and the remainder), and so there are some extra systematic errors to the ones shown here that cause migrations between these two samples. However, these systematic errors are not relevant for this analysis.



(a)



(b)

Figure 6.16: Ratio of the number of selected events for the CC-inclusive ν_μ selection cuts described in Section 6.6.2 versus the official selection cuts. The ratio is shown as a function of the reconstructed quantities used in the fit: (a) momentum and (b) cosine of angle (with respect to the z-axis of the ND280, as shown in Figure 2.6), for the lepton candidate. The similarity in output of the two selections means that the detector systematic errors derived from the official selection can be used to determine the relative sizes of the errors in this analysis. Decisions can therefore be made as to which systematic errors should be studied in detail and which are small enough to be neglected.

impact on the selected events, specifically the maximum error on a p_μ - $\cos\theta_\mu$ bin when the respective systematic error is applied to the official CC-inclusive ν_μ sample (this corresponds to the maximum diagonal component of the detector covariance matrix for that systematic error). The p_μ - $\cos\theta_\mu$ bins referred to here are the ones used in the official analysis, which are slightly different to the ones used in this analysis. However, the relative sizes of these errors are still instructive in deciding which systematics to study or neglect.

It can be seen that the two dominant detector systematic errors are the TPC momentum distortion and the Out of FGD events (with maximum p_μ - $\cos\theta_\mu$ errors of 5.7% and 9% respectively.) These systematic errors are evaluated fully for the selection cuts and p_μ - $\cos\theta_\mu$ binning used in the analysis, and are described below. The systematic due to sand muons is also evaluated. Although it does not have a large impact for the official selection (1.3%), the lack of veto for through-going tracks for this analysis means that it should be evaluated. The remaining systematic errors are small, with maximum p_μ - $\cos\theta_\mu$ errors up to $\sim$2%. Errors of this size are insignificant when compared to the very large cross-section errors which will be discussed in Section 6.7.4; the relative impacts of the different systematic errors will be shown in Section 6.8.2. In addition, due to the distinct peaks in neutrino energy for kaon-originating and pion-originating neutrinos (shown in Figure 3.4(a)) combined with the relatively coarse p_μ - $\cos\theta_\mu$ binning used in this analysis (described in Section 6.6.4), only very broad spectral information is required in order to constrain the kaon peak, meaning that systematics such as the TPC momentum resolution and TPC momentum scale are suppressed.

Out of FGD events

As described above, one of the dominant systematic errors for the official selection is on the Out of FGD events. These events also form one of the main background categories for this analysis, as shown in Table 6.1, and so the impact of this systematic error must be evaluated. The Out of FGD events encompass a range of different scenarios and failure modes, described previously in Section 6.6.3. A host of detailed studies has been performed within the ν_μ group to evaluate the uncertainties on each of these sub-categories, finally producing an overall covariance matrix in p_μ - $\cos\theta_\mu$ bins [104]. Unfortunately this matrix cannot be used due to the different binning in this analysis, and to replicate this work is a prohibitively

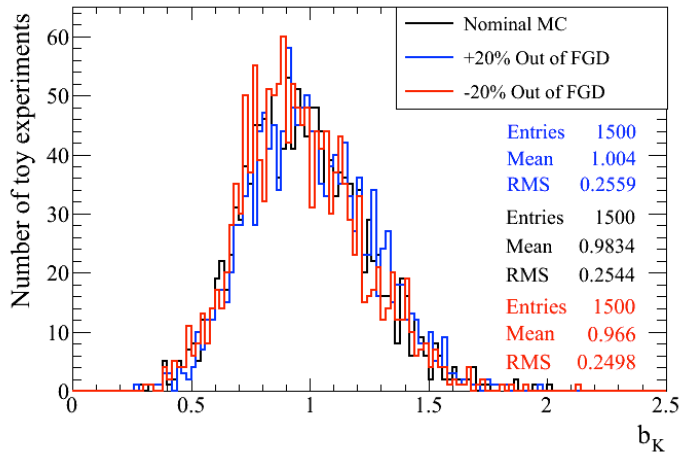
Source of systematic	Data sample	Error size / %	Max. error on p_μ - $\cos\theta_\mu$ bin / %
Track quality cut	Beam data / MC	0.1	0.02
TPC single track efficiency	Beam data / MC	0.5	0.8
TPC double track efficiency	Beam data / MC	0.6	2.0
TPC particle identification	Beam data / MC	0.1	0.6
TPC momentum scale	External measurements	0.51	2.2
TPC momentum distortion	Special MC	-	5.7
TPC momentum resolution	Beam data / MC	-	2.5
TPC-FGD matching efficiency	Sand interactions + cosmics	< 1	< 1
Fiducial mass	External measurements	0.67	0.67
Charge misidentification	Beam data / MC	-	1.1
Cosmic rays	Special MC	Negligible	Negligible
Sand muons	Special MC	15	1.3
Out of FGD events	Several samples	-	9
Pileup	Beam data / MC	0.24	0.24

Table 6.2: Summary of the ND280 detector systematic errors evaluated by the internal T2K ν_μ group for the official CC-inclusive ν_μ selection cuts [101]. The first three columns show the source of systematic error, the data sample used to evaluate the systematic and the size of the error (where appropriate) respectively. The final column shows the maximum fractional change in the number of events in a p_μ - $\cos\theta_\mu$ bin when this systematic is applied to the official CC-inclusive ν_μ sample. Note that the bins referred to here are those used in the official analysis, which are slightly different to the ones used in this analysis. However, the relative sizes of these numbers are still useful in determining the relative importance of the systematic errors.

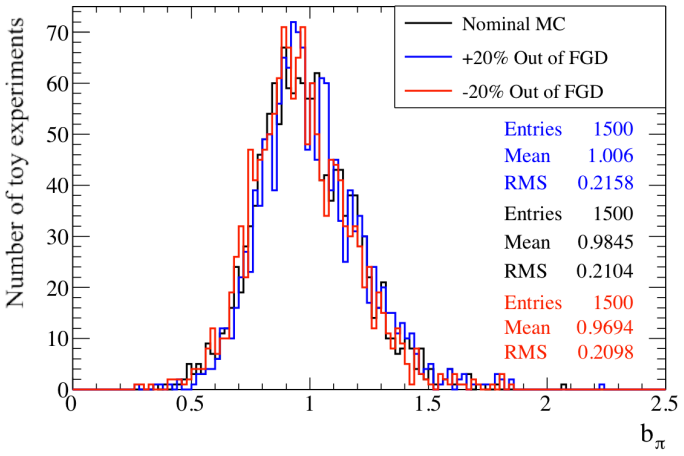
large task. However, it is found that the impact of the Out of FGD events on the final fitted parameters, b_K and b_π, is small enough that this systematic need not be included in the V_{det} matrix in Equation 6.16. This is justified as follows. An overall normalisation error of 20% conservatively covers all of the different errors evaluated by the ν_μ group. Three sets of 1500 toy MC experiments are therefore produced, with normalisations of 0.8, 1.0 and 1.2 for the Out of FGD events (the process for producing these toy MC experiments will be described in detail in Section 6.8.1). These toy MC experiments are then fitted, using the profile minimisation method described in Section 6.7.2. The resulting b_K and b_π distributions are shown in Figure 6.17. It can be seen that the mean changes by approximately $\pm 1.7\%$ and $\pm 1.5\%$ for b_K and b_π respectively when the Out of FGD normalisation is varied by $\pm 20\%$. This is small when compared to the effect of the cross-section errors on b_K and b_π, which will be discussed in Section 6.8.2, and so it is reasonable to neglect this systematic error.

Sand muons

As described previously, sand interactions refer to neutrino interactions which occur outside of the ND280 magnet, particularly in the sand upstream of the detector and the wall of the pit containing the ND280. This can give rise to a muon travelling into the ND280 detector. A special MC sample is available for these events since they are not modelled in the default MC. Unlike for the official selection, there is no explicit veto in this analysis on through-going tracks and so the impact of these sand interactions is evaluated. These interactions account for 1.7% of events passing the CC-inclusive ν_μ selection cuts, and their contribution to the p_μ - $\cos\theta_\mu$ bins is shown in Figure 6.15. To assess the impact of these events on the fitted b_K and b_π parameters, two sets of 1500 toy MC experiments are produced, one with the sand interactions included and one without. These toy MC experiments include all systematic and statistical errors, and their production will be described in Section 6.8.1. These toy MC experiments are then fitted, using the profile minimisation method described in Section 6.7. The resulting b_K and b_π distributions are shown in Figure 6.18. By removing the sand interactions, which is equivalent to assuming a 100% error on their normalisation, the mean b_K and b_π drop by approximately 4% and 1% respectively. The ν_μ group assigns an error of 15% on



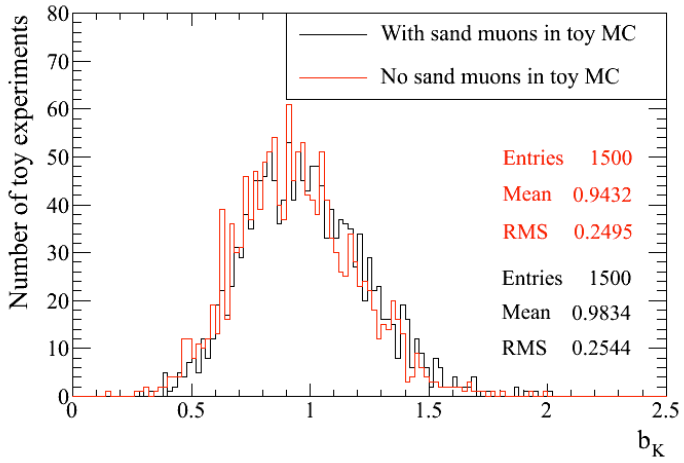
(a)



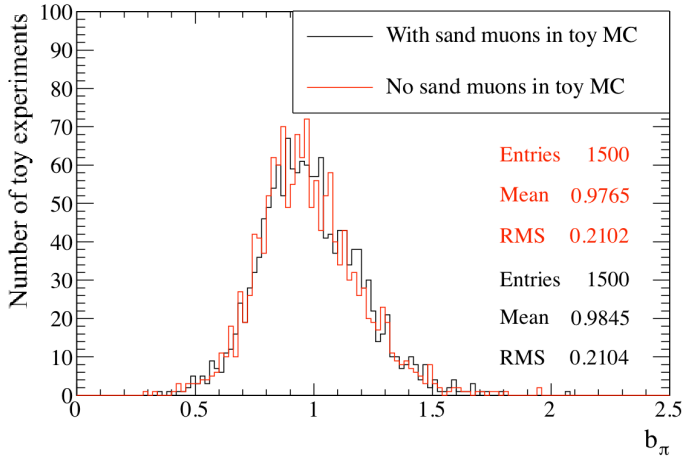
(b)

Figure 6.17: Effect of varying the normalisation of Out of FGD events in toy MC experiments on the fitted parameters: (a) b_K, and (b) b_π. Three sets of 1500 toy MC experiments are produced with Out of FGD normalisations of 0.8 (red), 1.0 (black) and 1.2 (blue). These toy MC experiments include all systematic and statistical errors, and their production will be described in Section 6.8.1. The variation in the mean is approximately $\pm 1.7\%$ and $\pm 1.5\%$ for b_K and b_π respectively. These changes are small when compared to the effect of the cross-section errors, as will be discussed in Section 6.8.2.

the normalisation of sand interactions, for the official selection. Using this as a guide, the effect on b_K and b_π will be at the sub-% level, and so can be neglected.



(a)



(b)

Figure 6.18: Effect of removing sand interactions from toy MC experiments on the fitted parameters: (a) b_K, and (b) b_π. Two sets of 1500 toy MC experiments are produced, with and without the sand interactions. These toy MC experiments include all systematic and statistical errors, and their production will be described in detail in Section 6.8.1. Removing the sand interactions is equivalent to assuming a 100% error on their normalisation.

TPC momentum distortion

The magnetic field inside the ND280 (generated by the UA1 magnet, described in Section 2.2.3) is not perfectly uniform but has small inhomogeneities. These

inhomogeneities cause deflections of the drift electrons in the TPCs and result in distortions of the reconstructed tracks. Events in both FGD1 and FGD2 are used in this analysis, meaning that both TPC2 and TPC3 tracks are used. Since TPC3 sits closer to the end of the ND280 detector (in the z direction) it suffers from larger distortions than TPC2. The MC assumes a perfect magnetic field. In order to correct this to the real field, a mapping survey of the ND280 field was conducted in 2009 using a custom built device containing Hall probes. The MC field is then corrected to the measured field map, a process referred to as the "field correction".[11] Despite these corrections, there are remaining biases between the reconstructed momentum, p_{reco}, and the true momentum, p_{true}, of particles traversing the TPCs. Studies of these biases have been performed within the ν_μ group for tracks in TPC2 and TPC3. It is actually the bias in the *inverse* momentum $(1/p_{\text{reco}} - 1/p_{\text{true}})$ that is studied since this has a Gaussian distribution. The bias is studied with and without the field correction applied in the reconstruction of events, and the difference is taken as the systematic error on the inverse momentum due to TPC momentum distortions. In other words, an error of 100% is assumed on the field correction, and this is propagated through to give the systematic error on the inverse momentum. This error is found to be 1% for tracks in TPC2 and 4.8% for tracks in TPC3.

The impact of the TPC momentum distortion on the CC-inclusive ν_μ sample is evaluated by scaling the momentum of each track in the MC according to the appropriate error, depending on which TPC the track traverses.[12] The momentum scaling is applied as follows:

$$p^{\text{reco}}_{\text{scaled}} = \frac{p^{\text{reco}}_{\text{nom}}}{1 + p^{\text{reco}}_{\text{nom}}(\epsilon/1\text{GeV})} \tag{6.18}$$

where $p^{\text{reco}}_{\text{scaled}}$ is the scaled reconstructed momentum, $p^{\text{reco}}_{\text{nom}}$ is the nominal reconstructed momentum and ϵ is the systematic error due to the distortions quoted for a 1 GeV track (with a value of 1% and 4.8% for tracks in TPC2 and TPC3 respectively).

[11] There is a further stage of correction to take into account remaining data-MC differences in the field inferred using calibration targets, which are assumed to be due to differences between the real field and the measured field map, but these corrections are not relevant here.

[12] For a track which traverses both TPC2 and TPC3, the error is applied according to which TPC was used to estimate the momentum of the track (see Section 6.5).

The covariance matrix is calculated using the same methodology as for the flux covariance matrices (as described in Section 3.3); the error is applied and the correlated bin changes are used to calculate the covariance. The elements of the fractional covariance matrix, in p_μ - $\cos\theta_\mu$ bins, are given by:

$$V_{ij} = \frac{(N_i^{\text{scale}} - N_i^{\text{nom}})(N_j^{\text{scale}} - N_j^{\text{nom}})}{N_i^{\text{nom}} N_j^{\text{nom}}} \tag{6.19}$$

where the symbols are defined as follows. N_i^{scale} and N_i^{nom} are the occupancies of the ith p_μ - $\cos\theta_\mu$ bin with the momentum scaling applied and for the nominal sample respectively. N_j^{scale} and N_j^{nom} are the same but for the jth bin.

The covariance matrix for the TPC momentum distortion is shown in Figure 6.19. Since this is the only detector systematic error for which a covariance matrix is generated, this matrix is the V_{det} matrix in Equation 6.16. To be clear, ideally the V_{det} matrix would also include a covariance matrix for the systematic error on the Out of FGD events, since this systematic error actually has a larger impact on the official CC-inclusive ν_μ sample than the TPC momentum distortion, as shown previously in Table 6.2. However, as discussed earlier, generating a covariance matrix for this umbrella category of systematic errors is a prohibitively large task. The decision not to include the Out of FGD covariance matrix in V_{det}, though ideally it would be included, has been justified previously by showing that the impact of the Out of FGD events on the final fitted parameters, b_K and b_π, is sufficiently small.

6.7.4 Cross-section systematic errors

Parametrisation and errors

The cross-section systematic errors used in this analysis are those recommended by the internal T2K cross-section working group [105] and are the same as those used for the official ND280 analysis, as mentioned in Section 6.3. Detailed studies have been performed within this working group in order to provide a cross-section parametrisation that is suitable for fits to spectral ND280 data (as opposed to simple counting experiments). Since the ND280 data does not yet cover the full phase space and kinematics

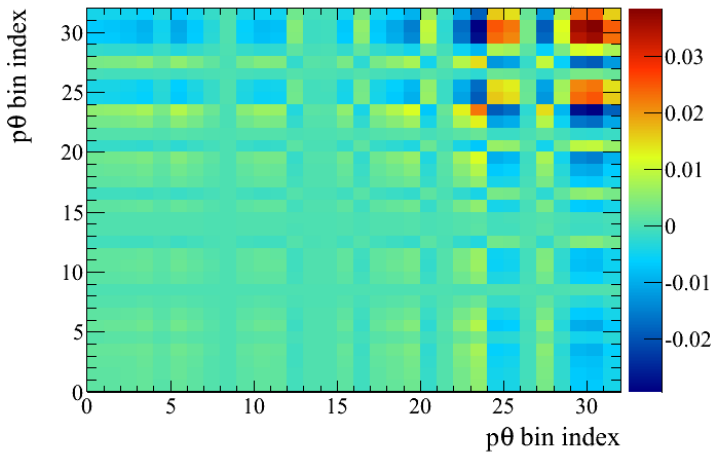


Figure 6.19: Fractional covariance matrix for the TPC momentum distortion systematic. The x and y axes follow the p_μ - $\cos\theta_\mu$ bin ordering given in Section 6.6.4.

of final state particles produced by neutrino interactions, it is necessary to use external neutrino scattering data to evaluate these systematic errors. Data from the MiniBooNE experiment (discussed previously in Section 1.2.1) is fitted to extract values and uncertainties for the cross-section model parameters to which the data is sensitive. There are plans to include data from other experiments (for example, SciBooNE, K2K and NOMAD) in these fits but this has not yet been done. For the remaining parameters, results from other experiments are used to deduce the uncertainties (but the data is not fitted). Table 6.3 summarises all of the cross-section parameters, showing their nominal values and uncertainties which are input to the fit, along with the impact of each parameter on the total event rate of the CC-inclusive ν_μ sample. All of the cross-section parameters are treated as uncorrelated when input to the fit, except for the three parameters constrained by fits to single pion production data (M_A^{RES}, CC1π E1 and NC1π^0). The covariance matrix for these parameters is shown in Figure 6.20. A brief description of each cross-section parameter and how the values in Table 6.3 were evaluated is now given. Full details should be sought in [105] and [106]. It should be noted that final state interactions (FSI) are not considered here, unlike in the official analysis. These are interactions within the nucleus of the particles emitted from the primary neutrino vertex that affect the number and type of particles observed exiting the nucleus. However, since FSI variations mainly change the number and kinematics of secondary (non-lepton) tracks, they are neglected for the CC-inclusive selection in this analysis.

Parameter	E_ν range / GeV	Nominal value	Error / %	Change in event rate / %	Category
CCQE E1	$0 < E_\nu < 1.5$	1.0	0.11	3.9	1
CCQE E2	$1.5 < E_\nu < 3.5$	1.0	0.3	1.9	1
CCQE E3	$E_\nu > 3.5$	1.0	0.3	1.0	1
CC1π E1	$0 < E_\nu < 2.5$	1.63	0.43	5.3	1
CC1π E2	$E_\nu > 2.5$	1.0	0.40	3.7	1
NC1π^0	all	1.19	0.43	0.1	1
M_A^{QE}	all	1.21 GeV/c^2	0.45	8.8	2
M_A^{RES}	all	1.16 GeV/c^2	0.11	2.7	2
p_F ^{12}C	all	217 MeV/c	30	0.7	2
SF ^{12}C	all	0 (off)	1 (on)	0.1	2
CC Other Shape	all	0.0	0.40	1.8	2
CC Coherent	all	1.0	1.0	5.2	3
NC Other	all	1.0	0.30	0.8	3
W shape	all	87.7 MeV/c^2	45.3	1.5	3

Table 6.3: Summary of the cross-section systematic errors. For each cross-section parameter, the energy range, nominal value and error are shown, along with the impact of changing the parameter by its error value on the total CC-inclusive ν_μ event rate. The final column gives a categorisation for how the parameters are implemented in the fit, as will be described in Section 6.7.4. The first four columns of the table are adapted from [105].

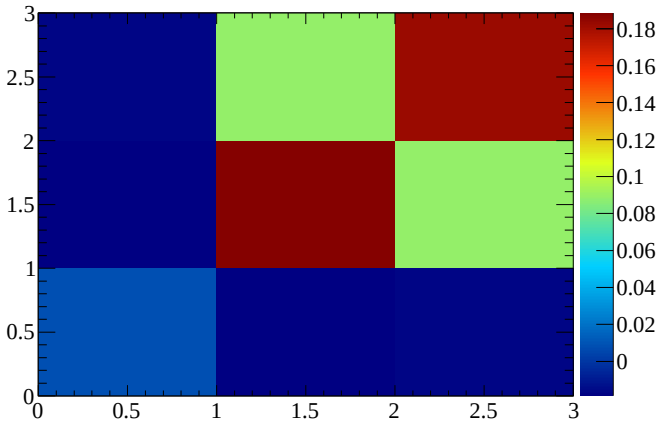


Figure 6.20: Covariance matrix for the parameters constrained by fits to single pion production MiniBooNE data. Parameters 0, 1 and 2 on the x and y axes are M_A^{RES}, CC1π E1 and NC1π^0 respectively.

M_A^{QE} and low energy CCQE normalisation (CCQE E1)

M_A^{QE} is the axial mass in the axial form factor of the nucleon for charged-current quasi-elastic (CCQE) interactions. CCQE E1 is a simple normalisation factor for the CCQE events. The MiniBooNE CCQE sample is used as a starting point to obtain values and errors for these parameters [107]. This sample has a correction applied to remove contamination due to FSI interactions.[13] The double differential cross-section, as a function of the outgoing lepton kinetic energy and the angle between the incoming and outgoing leptons, is fit to a NEUT MC sample containing only true CCQE interactions, using a χ^2 minimisation. The values of M_A^{QE} and the low energy CCQE normalisation, CCQE E1, are varied in the fit. The latter is included due to the 10.7% uncertainty on the MiniBooNE flux (the MiniBooNE flux is primarily below 1.5 GeV/c^2). The fitted M_A^{QE} (1.64 GeV/c^2) is larger than other published results [107, 108], which is thought to be due to differences in treatment of the low Q^2 (momentum transfer to the hadronic system) region. Bin covariances are not available for this data set and so it is not possible to accurately determine the error on the fitted M_A^{QE}. Therefore, the NEUT nominal value of M_A^{QE} is used (1.21 GeV/c^2), with the error defined as the difference between the

[13] Final state interactions cause the original CCQE sample to be contaminated with other cross-section modes, primarily CC1π^+. A subtraction of this mode is therefore performed in an attempt to isolate primary (nucleon level) CCQE interactions.

fitted and nominal values plus the fit error ($1.64 - 1.21 + 0.03 \sim 0.45$ GeV/c^2). For CCQE E1, the nominal value is used with an error of 11% taken directly from the error on the MiniBooNE flux.

Fermi momentum of the nucleus, p_F

For the cross-section models, NEUT uses a relativistic Fermi gas model to describe the target nucleus. One of the parameters in this model is the Fermi momentum of the nucleus, p_F, which (along with other parameters) characterises the nuclear potential. The value of this parameter is nucleus dependent. The value and uncertainty for carbon nuclei, shown in Table 6.3, is taken directly from electron scattering data [109].

Spectral function (SF)

In contrast to the Fermi gas model used by NEUT to describe the target nucleus, a more complicated model exists called the spectral function (SF) [110]. Comparisons with electron scattering data suggest that this is the more accurate representation, and so the difference between the Fermi gas model and the SF is assigned as a cross-section systematic. This systematic can be represented by a discrete parameter, where 0 represents the SF model switched off (i.e., using the default NEUT Fermi gas model) and 1 represents it switched on. As will be described in the next section, it is necessary to define a 1-sigma error for each parameter so that a response function can be built describing the impact of the change in the parameter on the bins in the fit. In contrast to the other cross-section parameters, which already have assigned errors, this is not a natural quantity for this discrete parameter, but it is decided to use the difference in the cross-section between the SF-off and SF-on states as the 1-sigma error. The NuWro generator is used to generate events with the SF model switched on (since this functionality is not available in NEUT), and the fractional difference in cross-section between this and the nominal NEUT MC is evaluated as a function of true neutrino energy and true momentum and angle of the outgoing lepton. The SF parameter can then be treated in the fit as continuous between 0 and 1, and the fitted value of the SF parameter can be interpreted as how "SF-like" the data is.

Parameters constrained by single pion production fits

By performing a simultaneous fit to three different MiniBooNE data sets, the values of several cross-section parameters governing single pion production can be

constrained. These parameters are: M_A^{RES}, W shape, the low energy CC1π normalisation (CC1π E1) and the NC1π^0 normalisation. M_A^{RES} is the axial mass in the axial form factor of the nucleon for resonant interactions. The W shape parameter is an empirical parameter that allows the shape of the NC1π^0 $|\vec{p}_{\pi^0}|$ spectrum to be modified in order to better match the data, as described in [106]. The other two parameters are simple normalisations. The MiniBooNE data sets used are the CC1π^0 [111], CC1π^+ [112] and NC1π^0 [113], where in each case the signal is defined by the particles leaving the nucleus and they are not corrected for FSI variations. A simultaneous fit is performed using a χ^2 minimisation. Additional parameters to those mentioned above are included in the fit that affect the MC prediction but that the fit has minimal power to constrain; penalty terms are therefore included in the χ^2 for these parameters. The χ^2 neglects the correlations between the different data sets, and in addition the bin-by-bin covariances are not available for all three data sets. For this reason, an *ad hoc* procedure is applied to scale the errors from the fit to match the size of the MiniBooNE flux-integrated cross-section errors. Two additional sources of error on the fit parameters are also evaluated: FSI uncertainties[14] and pionless delta decay[15]. The fitted values of M_A^{RES}, CC1π E1 and the NC1π^0 normalisation and their errors (after the scaling procedure and the inclusion of FSI and pionless delta decay) are shown in Table 6.3. The covariance between these three parameters is shown in Figure 6.20. For the W shape parameter, there is a large difference between the fitted and nominal values, suggesting that it is probably acting as a proxy for some general model change. Therefore, the nominal NEUT value (87.7 MeV/c 2) is input to the fit, and the error is taken as the difference between this value and the fitted value.

Normalisation of CC coherent pion production (CC Coherent)

External experiments [114, 115] show results consistent with no coherent pion production at neutrino energies of $\mathcal{O}$(1 GeV). In addition, the normalisation of CC coherent pion production is one of the additional parameters in the single pion production fits detailed above and the best fit value is consistent with no coherent pion production at the 2σ level. For these reasons, a nominal value of 1.0 with an error of 100% is input to the fit for this normalisation parameter.

[14] The pions can undergo absorption, scattering or charge exchange in the nucleus

[15] In NEUT, the Δ resonance can decay without emitting any pions, therefore reducing the number of observed CC1π events in the data samples.

High energy normalisations

The high energy CCQE normalisations (CCQE E2 and CCQE E3) and the high energy CC1π normalisation (CC1π E2) are not directly constrained using the MiniBooNE data, unlike for the low energy normalisations as detailed above, due to the lack of neutrino flux in the high energy region. A discrepancy of around 30% has been observed in the CCQE cross-section between the MiniBooNE results [107] (with an average neutrino energy of $\sim$0.8 GeV) and those of the higher energy NOMAD experiment [116] (operating with neutrino energies of $3 < E_\nu < 100$ GeV). An error of 30% is therefore assigned to CCQE E2 and CCQE E3. The error on CC1π E2 is assigned by examining the discrepancy between the MiniBooNE CC1π^+ data set and the nominal NEUT MC at high neutrino energy ($\sim$2 GeV) and extrapolating above that; an error of 40% is decided upon.

Normalisation of other NC modes (NC Other)

The NC Other category is defined as all NEUT NC interaction modes except NC1π^0 (which has its own normalisation constrained in the single pion fits described above). This includes NC elastic, multi-π and deep inelastic interactions, as well as resonant interactions that produce 1$\pi^\pm$, a γ, η or K. These modes contribute a negligible amount to the MiniBooNE CC1π^0, CC1π^+ and NC1π^0 data sets and so cannot be constrained in the single pion fits detailed above. Studies are performed varying different parameters within the NEUT model by their 1-sigma errors, and seeing the effect on the different interaction modes, detailed further in [117]. There are limitations to this method (the parameter variations do not account for all of the observed discrepancies between MC and external cross-section data) but using these results as a guideline, an error of 30% is assigned to this cross-section category.

Shape uncertainty for other CC modes (CC Other Shape)

The CC Other category is defined as all NEUT CC interaction modes except CCQE and CC1π (which have their own normalisations as described above). This includes CC multi-π and deep inelastic interactions, as well as resonant interactions that produce a γ, η or K. From the MINOS results for the total CC-inclusive cross-section [118], the error at the lower end of their neutrino energy range ($\sim$4 GeV) is seen to be of $\mathcal{O}(10\%)$. Using this as a reference point, and the fact that the error on the cross-section decreases with neutrino energy, E_ν, the shape error for the CC

Other category is defined as:

$$\sigma_{\mathrm{CCOther}} = \frac{0.4}{E_\nu} \tag{6.20}$$

Since the modes in question have a threshold energy of ∼0.6 GeV, the division by E_ν is not a concern.

Treatment of the cross-section systematic errors in the fit

The cross-section systematic errors are divided into three different categories, which define how the systematic is implemented in the fit. The category of each systematic is shown in the final column of Table 6.3. These categories are:

1. Uncertainties on the overall normalisations of specific interaction modes. These uncertainties are energy dependent, and are implemented in the fit via the x_β^{norm} parameters in Equation 6.13.

2. Uncertainties that can change the cross-section in a non-trivial way. These uncertainties are modelled by response functions ($w_{i\alpha\beta\gamma}$ in Equation 6.13) which are discussed further below.

3. Uncertainties that are modelled in the same way as the detector systematics. These uncertainties are included in the d_i parameters in Equation 6.13, which represent the uncertainty in normalisation of a given observable p_μ - $\cos\theta_\mu$ bin. There are several reasons why parameters are included in this category, which is discussed below.

The response function, $w_{i\alpha\beta\gamma}(\vec{x})$ in Equation 6.13, represents the fractional change in the number of events for a given p_μ - $\cos\theta_\mu$, neutrino parent, interaction mode and neutrino energy bin under a new set of cross-section parameters, $\vec{x}$. The response of the bin is calculated for each parameter separately, and then these responses are combined into a total response for the set of new cross-section parameters. A reweighting program called T2KReWeight is used to calculate the responses; instead of having to regenerate the MC for each new parameter value (which is very expensive computationally), a weight is calculated for each event given the new cross-section value. For a change in parameter j from its nominal value to x^j_{new}, the weight for a given event is given by the

ratio of the new cross-section, $\sigma(x^j_{\text{new}})$, to the nominal, σ_{nom}:

$$r(x^j_{\text{new}}) = \frac{\sigma(x^j_{\text{new}})}{\sigma_{\text{nom}}}. \tag{6.21}$$

The fractional response of a given bin for the change in parameter j, $W_{i\alpha\beta\gamma}(x^j_{\text{new}})$, is then calculated by summing the weights of all N events in that bin and dividing by the total number of predicted events in the bin, $T_{i\alpha\beta\gamma}$ (as defined in Equation 6.13):

$$W_{i\alpha\beta\gamma}(x^j_{\text{new}}) = \sum_{e}^{N} r_e(x^j_{\text{new}})/T_{i\alpha\beta\gamma}. \tag{6.22}$$

The responses for each of the N_{xsecpars} parameters separately are then combined into an overall response of the bin for the set of new cross-section parameters, $\vec{x}_{\text{new}}$, by:

$$w_{i\alpha\beta\gamma}(\vec{x}_{\text{new}}) = \prod_{j}^{N_{\text{xsecpars}}} W_{i\alpha\beta\gamma}(x^j_{\text{new}}). \tag{6.23}$$

The values of x^j_{new} are varied as the fit is run. In theory, $W_{i\alpha\beta\gamma}$ can be calculated for any value of x^j_{new}. However, in order to speed up the fit, a lookup table is produced for a series of new values of each parameter. The parameters are varied by $\pm1\sigma$, $\pm2\sigma$ and $\pm3\sigma$, where the error values are given in Table 6.3. Interpolation is used for x^j_{new} values between these points. A graphical representation of this lookup table is shown in Figure 6.21 for an example bin and cross-section parameter. The treatment of the SF parameter is slightly different because, in contrast to the other parameters, it is a discrete parameter where 0 represents the SF model switched off (i.e., using the default NEUT Fermi gas model) and 1 represents it switched on. As described previously, it is decided to use the difference in the cross-section between the SF-off and SF-on states as the 1-sigma error. The parameter is then treated in the fit as continuous between 0 and 1. The lookup table therefore consists of only two points (nominal and $+1\sigma$), and the value of the SF parameter can be interpreted as how "SF-like" the data is.

As described above, uncertainties in Category 3 are modelled in the same way as the detector systematic errors. The cross-section uncertainties are translated into uncertainties in the number of events in the observable p_μ - $\cos\theta_\mu$ bins, and included in the V_{d} detector covariance matrix in Equation 6.12. The parameters included in this cate-

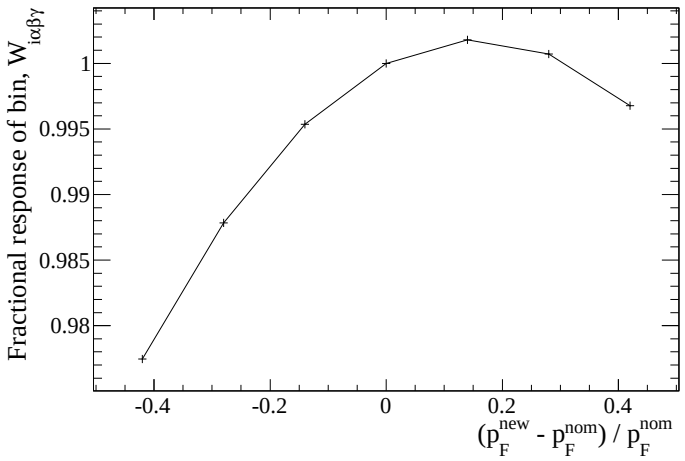
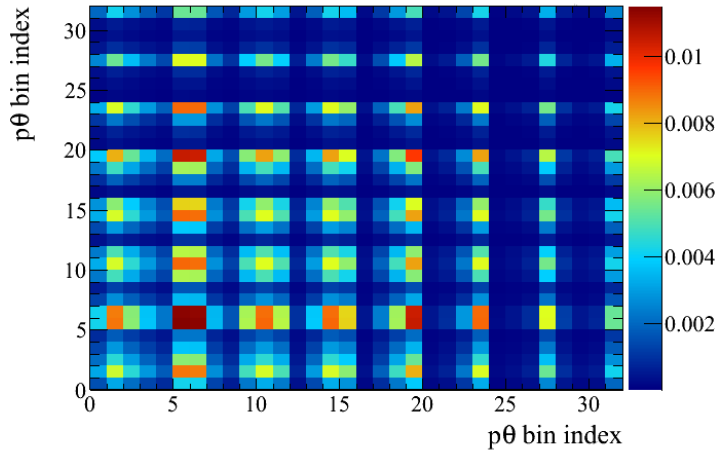


Figure 6.21: Fractional response of a given bin, $W_{i\alpha\beta\gamma}$, to a change in the Fermi momentum, p_F. The 7 values on the x axis correspond to -3σ, -2σ, -1σ, nominal, $+1\sigma$, $+2\sigma$ and $+3\sigma$. The bin shown here is $500 < p_\mu < 700$ MeV/c, $0 < \cos\theta_\mu < 0.9$, π^+ neutrino parents, CCQE interaction mode and $0 < E_\nu < 1.5$ GeV.

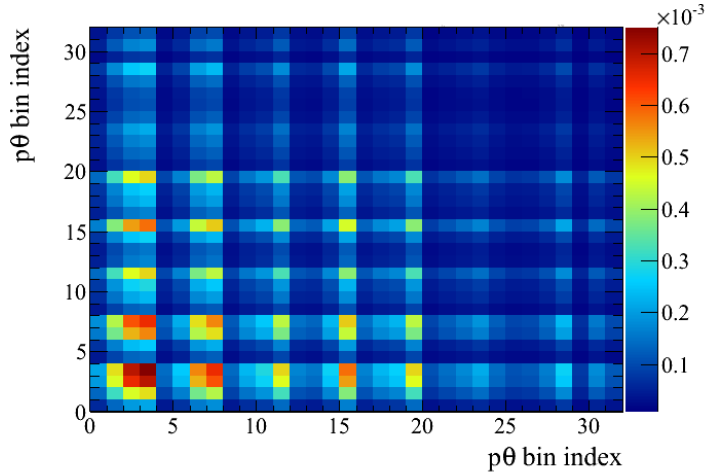
gory are the CC Coherent normalisation, NC Other normalisation and W shape, with covariances V_{CCCoh}, V_{NCOther} and V_{Wshape} respectively. These covariance matrices are calculated in a similar way as the detector matrices (see Section 6.7.3); the parameter is shifted by $\pm 1\sigma$ and the correlated p_μ - $\cos\theta_\mu$ bin changes are observed. The matrices are shown in Figure 6.22. The V_{d} covariance matrix from Equation 6.16 becomes:

$$V_{\mathrm{d}} = V_{\mathrm{det}} + V_{\mathrm{mcstats}} + V_{\mathrm{CCCoh}} + V_{\mathrm{NCOther}} + V_{\mathrm{Wshape}}. \tag{6.24}$$

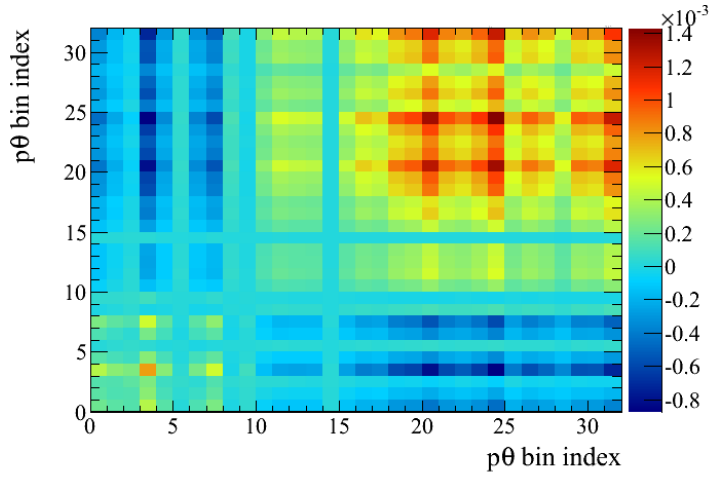
The parameters are included in the fit in this way due to their very large errors (for example, 100% for the CC Coherent normalisation). The size of the errors means that the parameters could approach the physical boundaries when varied within the fit. On the other hand, the impact of these parameters on the overall p_μ - $\cos\theta_\mu$ bin normalisations is relatively small, and so by including these parameters in the V_{d} matrix this boundary issue is avoided.



(a)



(b)



(c)

Figure 6.22: Covariance matrices for: (a) CC Coherent normalisation, (b) NC Other normalisation, and (c) W shape. In each case, the x and y axes follow the p_μ - $\cos\theta_\mu$ bin ordering given in Section 6.6.4, and the z axis shows the covariance, V_{ij}, between the ith and jth bins.

6.7.5 Flux systematic errors

The estimation of the flux systematic errors has been described previously in Chapter 3. The dominant uncertainty on the flux prediction comes from the hadron production at the target. Currently, there is no simple underlying parametrisation to describe the uncertainties on the NA61 inputs and other hadron production data. Therefore, the flux uncertainty is simply parametrised in bins of true neutrino energy. Figure 3.7(b) shows the total covariance matrix between energy bins for neutrinos of different flavours and at the two off-axis detectors (ND280 and Super-Kamiokande).

The flux systematic errors are not included in the fit (as mentioned in Section 6.3). This is done in order to perform an *unconstrained* measurement of the b_K parameter. As discussed previously in Section 3.4.2, this measurement is affected by all of the major flux uncertainties, including, for example, the kaon production at the target, the horn current and the off-axis angle of the ND280. Therefore including the flux systematic errors in the fit would directly constrain the b_K parameter, which is not desired. Ideally, the flux uncertainties would be included on the background pion-originating neutrinos (fitted with b_π), but this would involve factorising the existing flux errors into separate pion and kaon contributions which is not readily available. However, the full set of flux uncertainties *is* included when generating the toy MC experiments, as will be described in Section 6.8.1. This is done to give the most realistic approximation to the actual data, and it enables the robustness of the fitter to be evaluated including all of the shape and normalisation flux uncertainties. It will be seen that the fit performs well with the full set of flux errors included in the toy MC experiments, demonstrating that the existing fit methodology is robust.

In actual fact, only the errors for ν_μs (at the ND280) are used when generating the toy MC experiments since the contribution to the CC-inclusive ν_μ sample from the other flavours is very small ($\sim$1%). The sub-section of the total covariance matrix which corresponds to ν_μs at the ND280 is shown again in Figure 6.23 for convenience. The 20 true neutrino energy (E_ν) bins are spaced as follows in GeV:

0.0–0.1, 0.1–0.2, 0.2–0.3, 0.3–0.4, 0.4–0.5, 0.5–0.6, 0.6–0.7, 0.7–0.8, 0.8–1.0, 1.0–1.2, 1.2–1.5, 1.5–2.0, 2.0–2.5, 2.5–3.0, 3.0–3.5, 3.5–4.0, 4.0–5.0, 5.0–7.0, 7.0–10.0, $>$ 10.0.

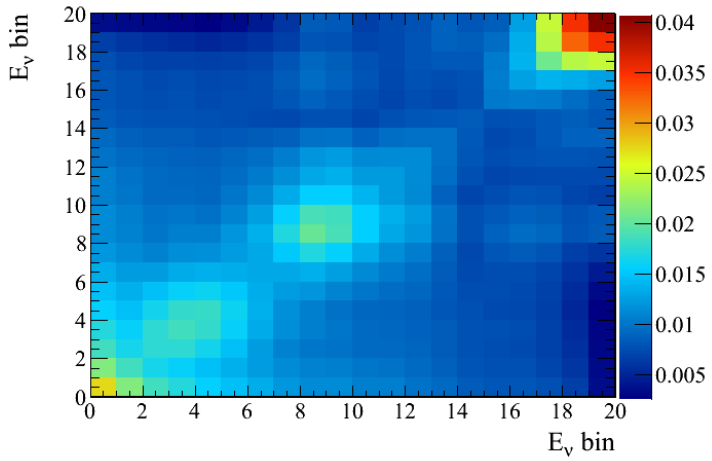


Figure 6.23: Flux covariance matrix for ν_μs at the ND280, where the x and y axes follow the E_ν binning described in Section 6.7.5. This is a sub-section of the full flux covariance matrix shown in Figure 3.7(b).

6.8 Fit validation studies

6.8.1 Generation of toy MC experiments

Toy MC experiments are produced with variations that include all systematic errors (detector, cross-section and flux) and also statistical errors. By producing many toy MC experiments and fitting them, the robustness of the fit can be validated under realistic changes in the normalisation and shape of the fitted distributions according to the systematic errors and also statistical fluctuations in the data.

The inclusion of the systematic errors in the toy MC experiments is described first. The cross-section and detector errors are included via parameters $\vec{x}$ and $\vec{d}$ respectively, as defined previously in Section 6.7.1. Additionally, the flux systematic errors are included, even though these errors are not included in the fit, which is done to give the most realistic approximation to the real data. As discussed previously in Section 6.7.5, the flux errors are parametrised as normalisation factors of true neutrino energy bins; these parameters are called $\vec{f}$. The number of events for a toy MC experiment in the ith p_μ - $\cos\theta_\mu$ bin is a function of the $\vec{f}$, $\vec{x}$ and $\vec{d}$ systematic error parameters, as well as the

neutrino parent parameters, $\vec{b}$, and is given by:

$$N_i^{\text{toy}}(\vec{f},\vec{b},\vec{x},\vec{d}) = d_i \sum_{\alpha}^{N_{\text{parents}}} \sum_{\beta}^{N_{\text{modes}}} \sum_{\gamma}^{N_{\text{E}_\nu\text{bins}}} b_\alpha f_\gamma x_\beta^{\text{norm}}(E_\gamma) w_{i\alpha\beta\gamma}(\vec{x}) T_{i\alpha\beta\gamma}, \tag{6.25}$$

where the meaning of the symbols is discussed as follows. f_γ is the flux normalisation in neutrino energy bin γ. All other parameters are as defined in Equation 6.13, but are repeated here for convenience: d_i is the detector nuisance parameter for p_μ - $\cos\theta_\mu$ bin i; b_α is the normalisation for neutrino parent type α; x_β^{norm} is the normalisation factor for interaction mode β; $w_{i\alpha\beta\gamma}(\vec{x})$ is the cross-section response function for p_μ - $\cos\theta_\mu$, neutrino parent, interaction mode and neutrino energy bin $i\alpha\beta\gamma$, and is calculated as described in Equation 6.23; and, finally, $T_{i\alpha\beta\gamma}$ is the predicted number of events for bin $i\alpha\beta\gamma$ in the nominal MC. As before for Equation 6.13, sand interactions in the $T_{i\alpha\beta\gamma}$ prediction are only reweighted by the $\vec{d}$ parameters, due to the lack of truth information available for these events.

Throws of the $\vec{f}$, $\vec{x}$ and $\vec{d}$ parameters are made according to the prior values and covariance matrices described in Sections 6.7.5, 6.7.4 and 6.7.3 respectively.[16,17] One throw of $\vec{f}$, $\vec{x}$ and $\vec{d}$ will define one toy MC experiment. The number of events in each p_μ - $\cos\theta_\mu$ bin for that toy experiment can be calculated using Equation 6.25. In addition to the systematic errors described above, statistical uncertainties are also included in the toy MC experiments. There are statistical uncertainties on both $T_{i\alpha\beta\gamma}$ and N_i^{toy} in Equation 6.25. Statistical throws of these values are taken according to a Poisson distribution.

The $\vec{b}$ neutrino parent parameters are not thrown in the same way as for the $\vec{f}$, $\vec{x}$ and $\vec{d}$ parameters, since there is no covariance matrix for these parameters. However, it is desirable to vary the $\vec{b}$ parameters and test the response of the fit to this. Therefore, five sets of 1500 toy MC experiments are produced according to the above procedure, each with b_K and b_π fixed at different values.[18] The values are shown in Table 6.4.

[16]The Cholesky decomposition method is used to make throws of the parameters according to their covariance matrices.

[17]In this case, the detector covariance matrix, V_{d}, defined in Equation 6.16, does not include V_{mcstats}. It is necessary to include the MC statistical error in the detector matrix in the fit because the error changes with each iteration of the fit parameters, as described in Section 6.7.3. However, for a given toy MC experiment, the $\vec{b}$ and $\vec{x}$ parameters are fixed and so the MC statistical error can be included in a simpler way, according to a Poisson distribution, as described in the main text.

[18]b_{other} is fixed at 1.0.

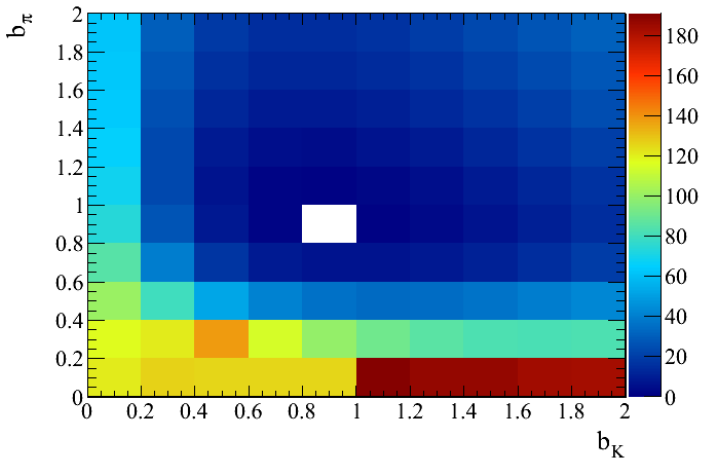
Toy MC set	b_K	b_π
0	1.0	1.0
1	1.5	1.0
2	0.5	1.0
3	1.0	1.5
4	0.5	0.5

Table 6.4: Five sets of 1500 toy MC experiments are produced according to the procedure outlined in Section 6.8.1, each with different values of b_K and b_π.

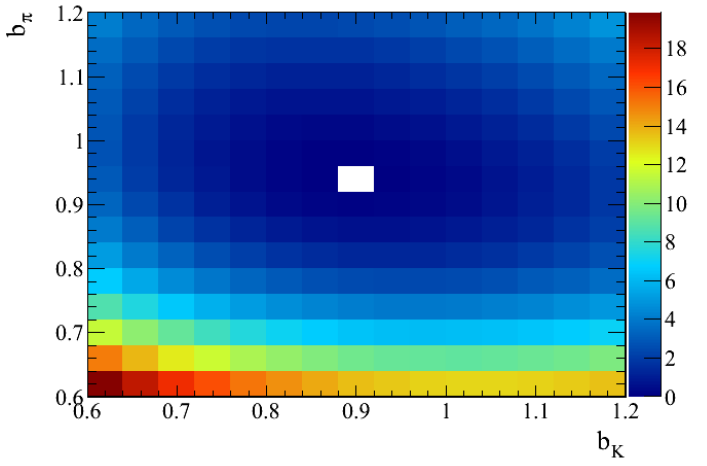
6.8.2 Validation results

As described in Section 6.7.2, a 2-dimensional space of b_K versus b_π is defined, and the quantity χ^2_{P} is evaluated at each point by minimising with respect to the nuisance parameters. In this way, a surface of χ^2_{P} is mapped out, with the minimum point in the surface giving the best fit values of b_K and b_π. The b_K-b_π space goes from 0 to 2 in both dimensions with a granularity of 0.01, resulting in 40,000 minimisation points. Scanning this entire grid is prohibitively slow for testing and validating the fitter. Therefore, progressive binning is used. An initial coarsely binned grid is used, divided into 10 bins from 0 to 2 in both dimensions. The bin with the lowest χ^2_{P} is found. This bin and the eight surrounding bins define the next grid, which is split into 15 bins of 0.04 in both dimensions. The process is repeated and a final grid is defined, split into 12 bins of 0.01 in both dimensions. The bin with the lowest χ^2_{P} in this final grid defines the best fit values for b_K and b_π. An example of this progressive binning is shown in Figure 6.24, where the best fit values of b_K and b_π are 0.895 and 0.945 respectively (taken as the centre of the bins).

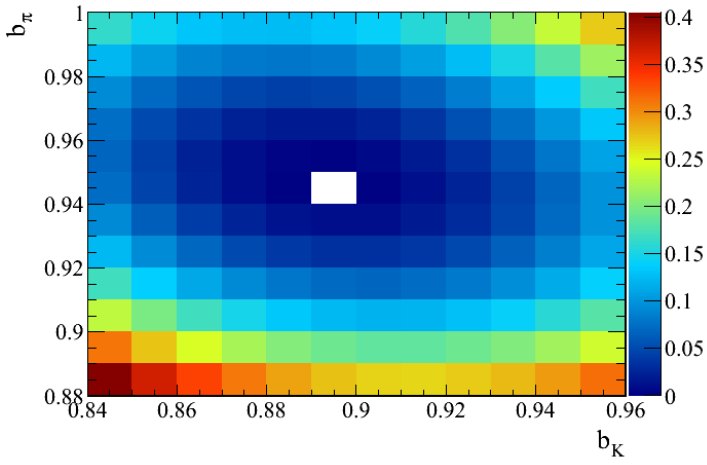
The toy MC experiments described in Section 6.8.1 can be used to validate several different aspects of the fit. The response of the fit to changes in b_K and b_π is examined. The goodness of fit and 1σ contour are also validated. Finally, the relative impact of the different sets of systematic errors (flux, cross-section and detector) upon b_K and b_π is studied. These validation studies are described in the following sections.



(a)



(b)



(c)

Figure 6.24: Progressive binning used for fit validation studies. Shown for an example toy MC experiment. In each case, the z-axis is $\Delta\chi^2_{\mathrm{P}}$ (see Equation 6.15) with the best fit bin shown in white with a value of zero. (a) shows the initial coarsely binned grid with a granularity of 0.2. The grids in (b) and (c) have granularities of 0.04 and 0.01 respectively, and are defined by the best fit bin and eight surrounding bins of (a) and (b) respectively.

Response of fit to changes in b_K and b_π

Figure 6.25 shows the distribution of best fit points in the b_K-b_π space for each of the five sets of toy MC experiments shown in Table 6.4. Each toy MC set has different fixed values for b_K and b_π, and consists of 1500 toy experiments with variations representing all systematic errors (flux, cross-section and detector) and statistical variations. Figures 6.26 and 6.27 show the projection of the best fit points in Figure 6.25 onto the b_K and b_π axes respectively, for the same five toy MC sets. It can be seen that the means of the fitted b_K and b_π parameters agree well with the input values, and the deviations of the means from the true values are well within the widths of the distributions. This shows that the fit is working well.

It can be seen in Figures 6.26 and 6.27 that the b_K and b_π distributions are slightly non-Gaussian. This can be explained as follows. The dependence of the p_μ - $\cos\theta_\mu$ bins on some of the cross-section parameters is observed to be non-linear. The response functions of some example bins are shown in Figure 6.28 for M_A^{QE} and M_A^{RES}. In each case, it is clear that the bin occupancy decreases by a greater fraction than it increases for a change in the parameter of the same magnitude but opposite sign. This could lead to b_K and b_π distributions skewed in the negative direction, as is observed. However, the observed skew is small and since the means of the fitted b_K and b_π distributions track changes in the input b_K and b_π values very well, as discussed previously, this is not a problem.

Goodness of fit

The goodness of fit, given by $\chi^2_{\mathrm{P,min}}$, is expected to follow a χ^2 distribution with the number of degrees of freedom (NDF) equal to the number of bins, N_{bins}, minus the number of fitted parameters, N_{params}:

$$\mathrm{NDF} = N_{\mathrm{bins}} - N_{\mathrm{params}}. \tag{6.26}$$

The χ^2 from the fit to data can be compared with this distribution to determine whether the proposed model is in agreement with the data.

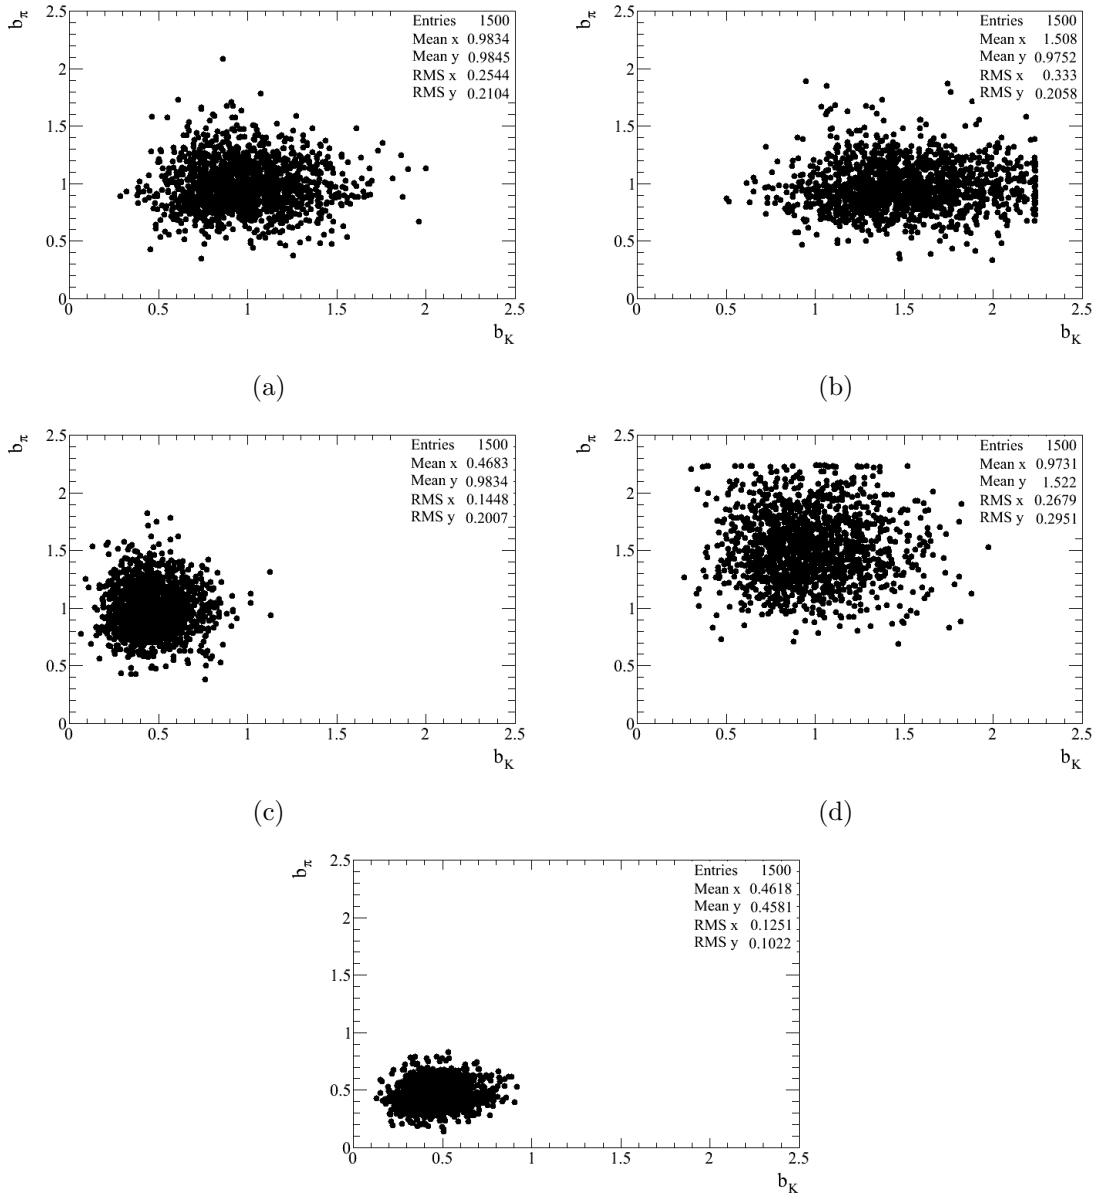


Figure 6.25: Distributions of the best fit points in the b_K-b_π space for the five sets of toy MC experiments shown in Table 6.4. Plots (a)-(e) are for sets 0–4 in order. Each toy MC set consists of 1500 toy experiments with variations representing all systematic errors (flux, cross-section and detector) and statistical variations, as described in Section 6.8.1.

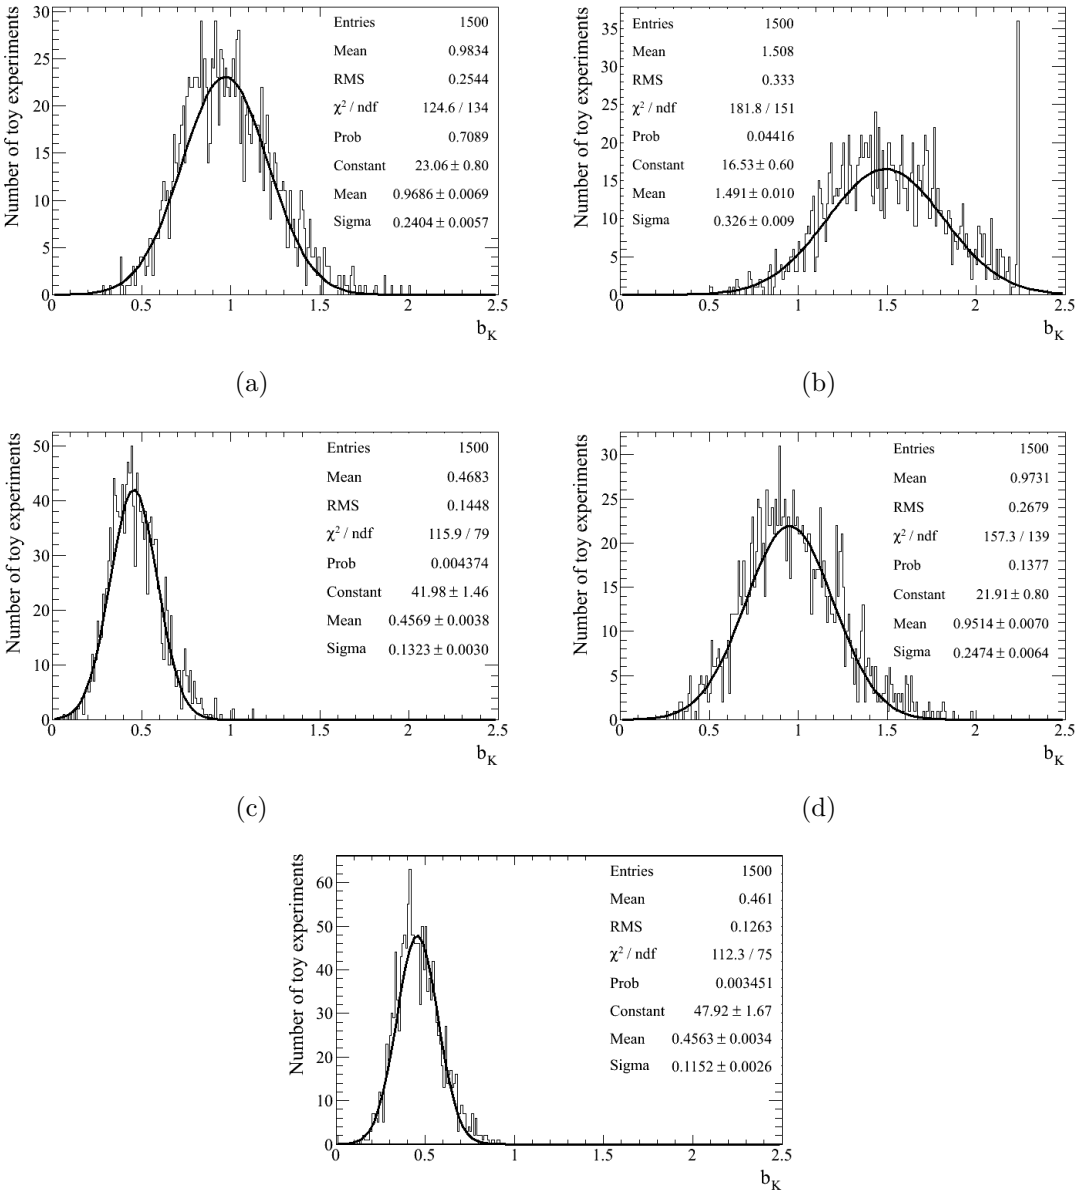


Figure 6.26: Projections of the plots in Figure 6.25 onto the b_K axis.

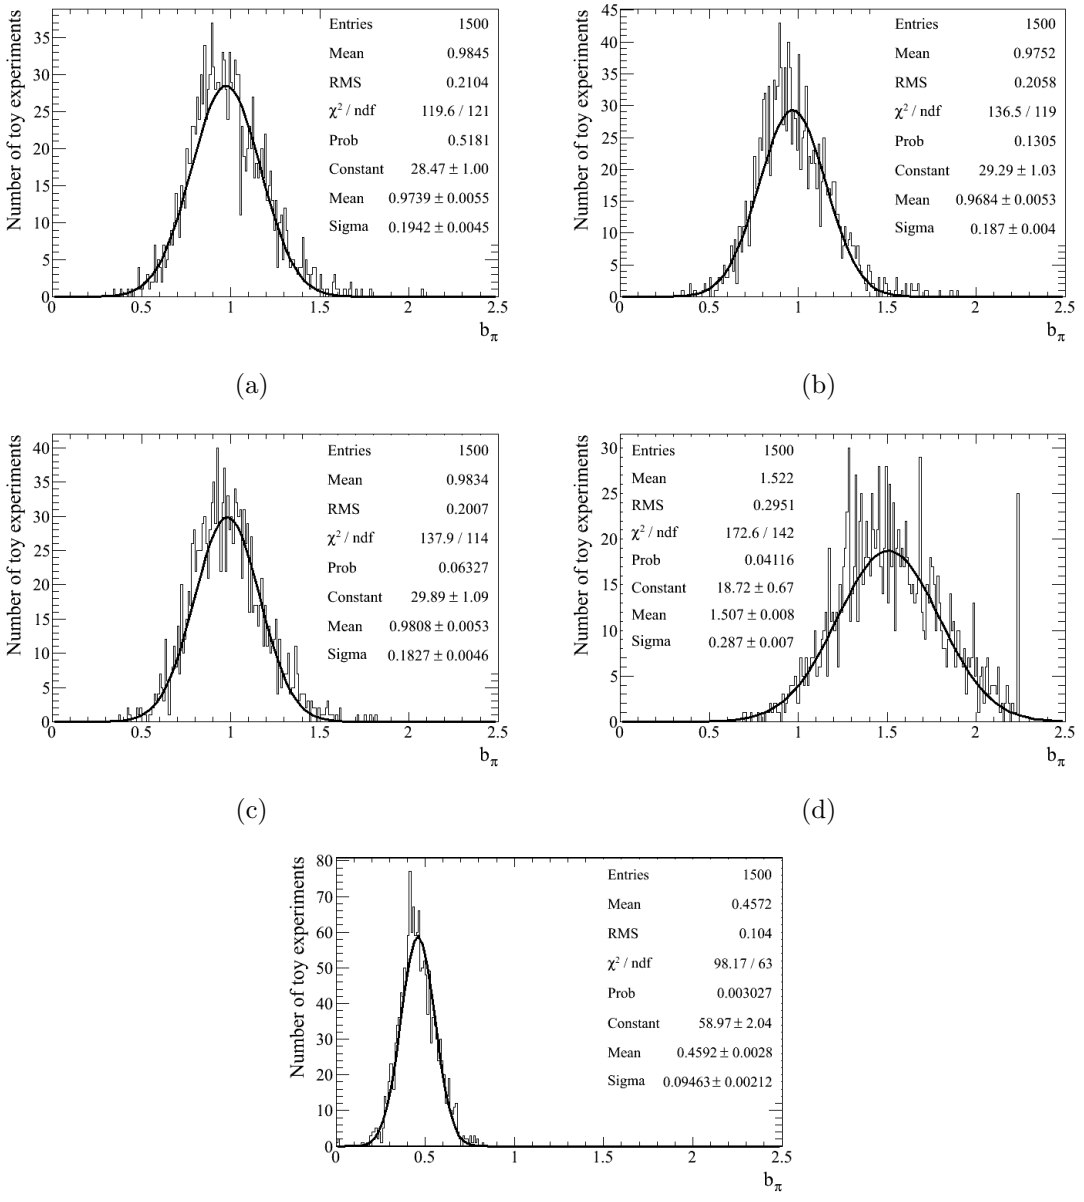
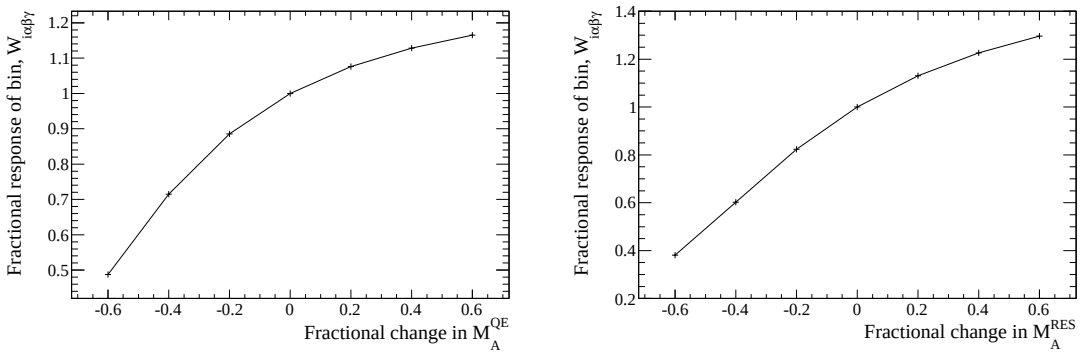


Figure 6.27: Projections of the plots in Figure 6.25 onto the b_π axis.



(a) (b)

Figure 6.28: Fractional response of a given bin, $W_{i\alpha\beta\gamma}$ (defined in Equation 6.22), to changes in a given cross-section parameter. (a) is for parameter M_A^{QE} and bin $400 < p_\mu < 500$ MeV/c, $0 < \cos\theta_\mu < 0.9$, π^+ neutrino parents, CCQE mode and $0 < E_\nu < 1.5$ GeV. (b) is for parameter M_A^{RES} and bin $1500 < p_\mu < 2500$ MeV/c, $0.98 < \cos\theta_\mu < 1.0$, K^+ neutrino parents, CC resonance mode and $1.5 < E_\nu < 3.5$ GeV. In both cases, the response of the bin is non-linear.

Figure 6.29 shows the distribution of $\chi^2_{\mathrm{P,min}}$ for the 1500 toy experiments in toy MC set 0 in Table 6.4 (i.e., with nominal values for b_K and b_π). The distributions for the other toy MC sets are very similar. It is fitted with a χ^2 function with the NDF allowed to float. The returned NDF is approximately 32.5, which is compared to the expected NDF as follows. There are 32 p_μ - $\cos\theta_\mu$ bins and, at each point in the b_K-b_π space, a total of 46 fitted parameters (3 parent normalisation parameters, $\vec{b}$, 32 detector parameters, $\vec{d}$, and 11 cross-section parameters, $\vec{x}$). If all of these parameters were free, then the NDF would simply be given by Equation 6.26.[19] However, all except the 3 parent normalisation parameters have prior constraints upon them (encoded in the V_d and V_x matrices in Equation 6.12), which are calculated from external data. Therefore, the NDF is given by:

$$\mathrm{NDF} = N_{\mathrm{bins}} + N_{\mathrm{prior}} - N_{\mathrm{params}}, \tag{6.27}$$

where N_{prior} is the number of parameters with prior constraints. This gives an expected NDF of 29. The best fit NDF from Figure 6.29 of 32.5 is therefore slightly higher than expected. However, the discrepancy is not large and could be assigned to a statistical fluctuation. In addition, there are possible degeneracies between some of the fit parame-

[19]The NDF would actually be negative in this case.

ters, for example, the b_K parameter is probably correlated with some of the high energy cross-section normalisation parameters. Since this would effectively reduce N_{params}, this would lead to an inflated effective NDF.

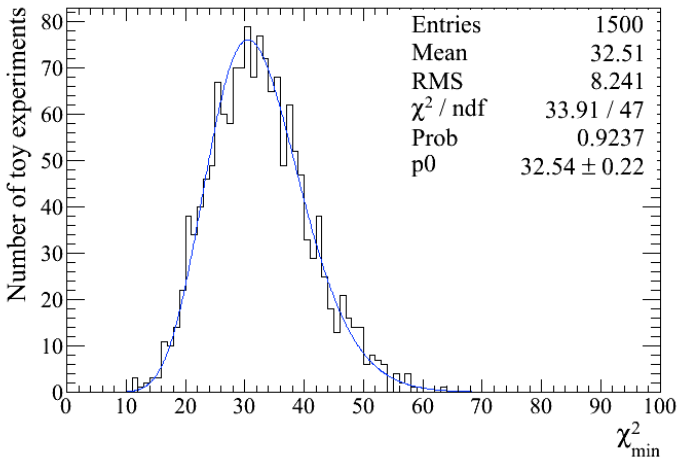


Figure 6.29: Distribution of $\chi^2_{\text{P,min}}$, the minimum of the χ^2_{P} surface in b_K-b_π space (see Equation 6.14), for the 1500 toy experiments in toy MC set 0 in Table 6.4. This is fitted with a χ^2 function, with the number of degrees of freedom allowed to float. The best fit NDF is ∼32.5.

Validation of 1σ contour

As well as finding the best fit point in the b_K-b_π space (where $\Delta\chi^2_{\text{P}} = 0$), it is necessary to define the 1σ contour around this point. In general, for the joint estimation of two parameters, the condition for a coverage probability of 68.27% (i.e., the 1σ contour) is $\Delta\chi^2_{\text{P}} = 2.30$ [103]. The coverage obtained using this condition is tested for the 1500 toy MC experiments in toy MC set 0 in Table 6.4 (i.e., with nominal values for b_K and b_π). The coverage is calculated as the fraction of toy MC experiments for which the true b_K-b_π point ($b_K = 1.0$ and $b_\pi = 1.0$) lies within the 1σ contour ($\Delta\chi^2_{\text{P}} < 2.30$). It is found to be 43.1%, meaning that the contour is undercovering. This is actually expected for the following reason. The 1σ contour will only be correct if all nuisance parameters are included in the fit. As explained in Section 6.8.1, these toy MC experiments include variations from all systematics: flux, cross-section and detector. The fit however only includes nuisance parameters for the cross-section and detector systematic errors, *not* the flux uncertainties. The flux uncertainties are deliberately omitted from the fit, in

order to perform an unconstrained measurement of b_K, as described in Section 6.7.5. The contour is therefore not expected to be correct. As a further test, the toy MC experiments are reproduced *without* the flux uncertainties included, and the coverage calculation is repeated. A coverage of 60.1% is obtained, which is closer to 68.27% but still not correct. Strictly speaking, this condition of $\Delta\chi^2_{\mathrm{P}} = 2.30$ is only valid for cases where the fitted parameter follows a Gaussian PDF and where the data sample is infinite [103]. However, the data sample used here is finite and, as discussed previously, the b_K and b_π distributions obtained from many toy experiments are not perfectly Gaussian. This could explain the observed undercoverage even when the flux uncertainties are excluded from the toy MC experiments.

In order to obtain the correct 1σ contour, the condition upon $\Delta\chi^2_{\mathrm{P}}$ is varied until a coverage of 68.27% is obtained. For the toy MC experiments *including* the flux uncertainties (and all other systematic errors and statistical uncertainties), the appropriate condition is found to be $\Delta\chi^2_{\mathrm{P}} = 4.61$ (this condition increases the coverage from 43.1% up to the desired 68.27%). This is used to draw the contour for the final fit to data, as will be shown in the results section (Section 6.9).

Impact of different systematic errors upon the fit

It is useful to understand the relative impacts of the different systematic errors on the fitted parameters, b_K and b_π. These results are used when deciding which detector systematic errors it is important to evaluate and which ones can be neglected, as discussed previously in Section 6.7.3. Figures 6.30 and 6.31 show the b_K and b_π distributions respectively with different uncertainties included in the toy MC experiments. The results are summarised in Table 6.5. It can be seen that the widths of the distributions increase dramatically when the cross-section errors are included, whereas the flux and detector errors have much smaller impacts. The cross-section systematic errors are therefore the dominant source of uncertainty on the b_K and b_π parameters.

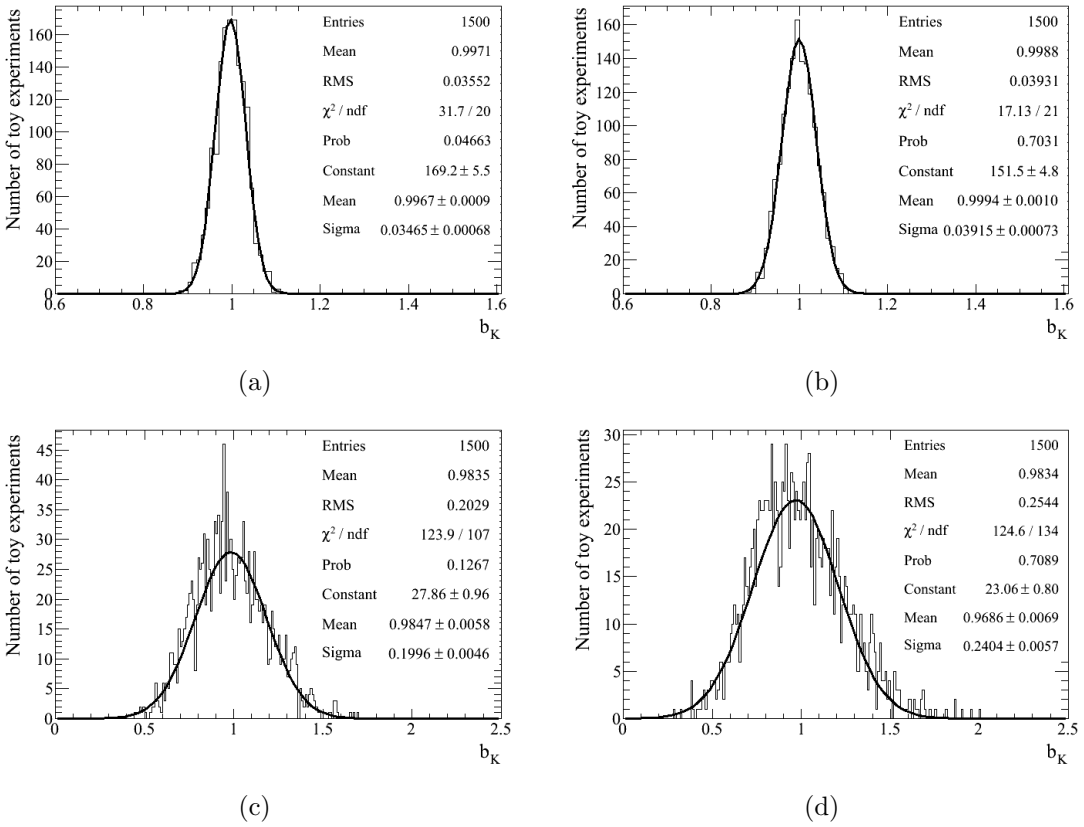


Figure 6.30: Impact of different uncertainties on the fitted b_K distribution. In each case, 1500 toy MC experiments are fitted, with the following uncertainties included in the toy experiments: (a) statistical errors only; (b) statistical and detector errors; (c) statistical, detector and cross-section errors; and (d) statistical, detector, cross-section and flux errors. The toy experiments are produced with true b_K = 1.0 and b_π =1.0.

Uncertainties included in toy MC	Width of b_K %	Width of b_π %
Statistical only	3.5	2.1
Statistical & detector	3.9	2.3
Statistical & detector & cross-section	20.0	14.9
Statistical & detector & cross-section & flux	24.0	19.4

Table 6.5: Summary of results shown in Figures 6.30 and 6.31. The widths of the b_K and b_π distributions are shown with different uncertainties included in the toy MC experiments. The widths are quoted in % assuming a mean of 1.0 (even though the means differ slightly from this in some cases).

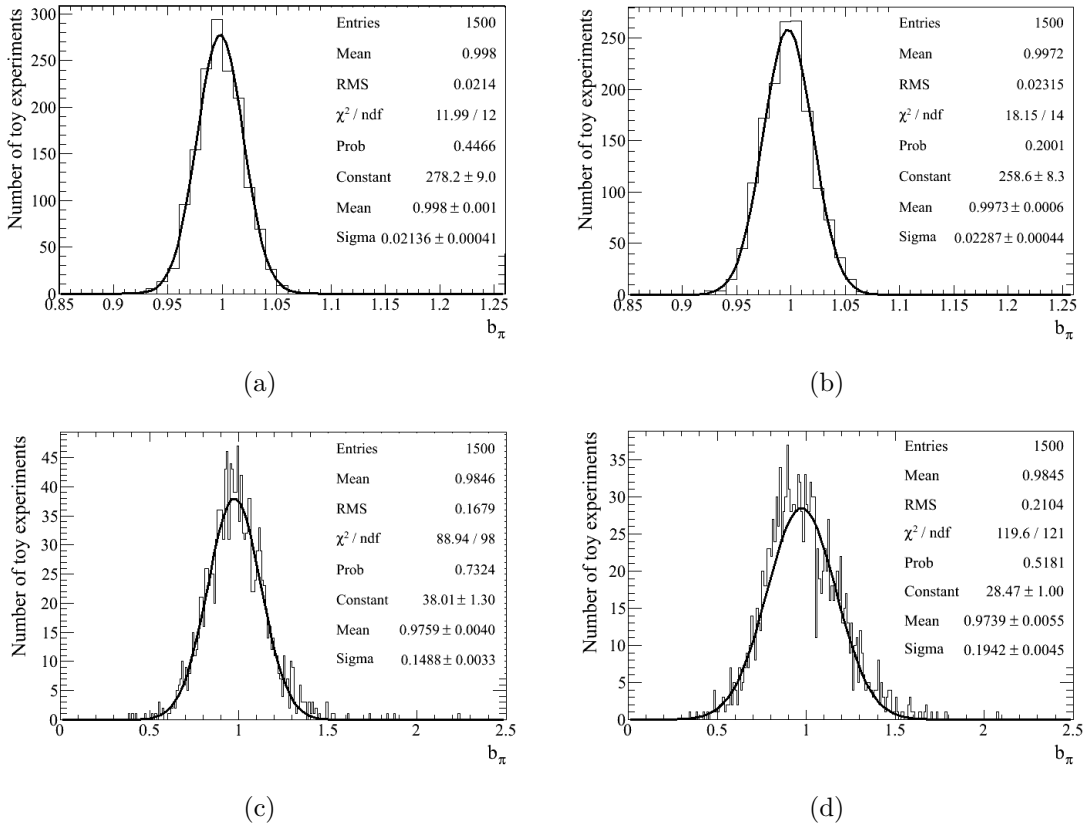


Figure 6.31: Impact of different uncertainties on the fitted b_π distribution. In each case, 1500 toy MC experiments are fitted, with the following uncertainties included in the toy experiments: (a) statistical errors only; (b) statistical and detector errors; (c) statistical, detector and cross-section errors; and (d) statistical, detector, cross-section and flux errors. The toy experiments are produced with true b_K = 1.0 and b_π =1.0.

6.9 Results

The results of the fit to the final CC-inclusive ν_μ data sample are discussed in this section. Figure 6.32 shows the p_μ - $\cos\theta_\mu$ bins for the data and nominal MC, which together form the inputs to the fit. The fitted MC is also shown, where the bin occupancies are calculated according to Equation 6.13 with the values of $\vec{b}$, $\vec{x}$ and $\vec{d}$ at the best fit point. The agreement between data and MC is clearly improved after the fit has been performed.

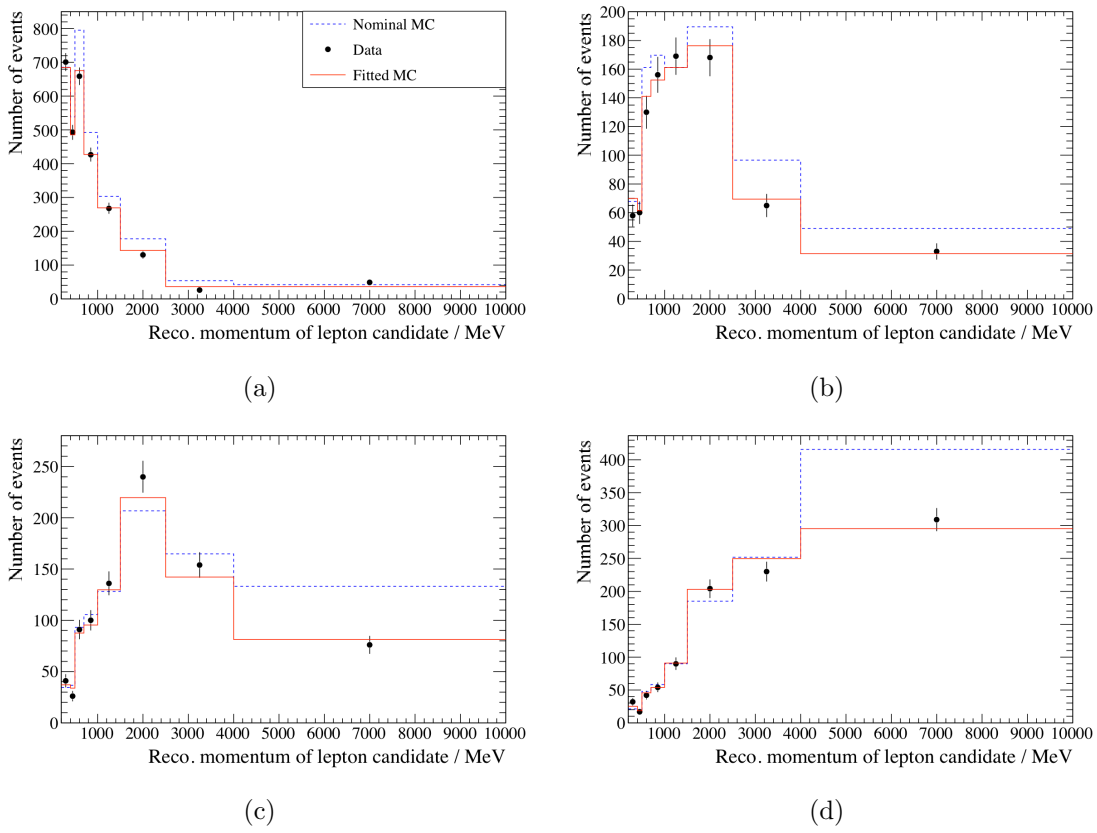


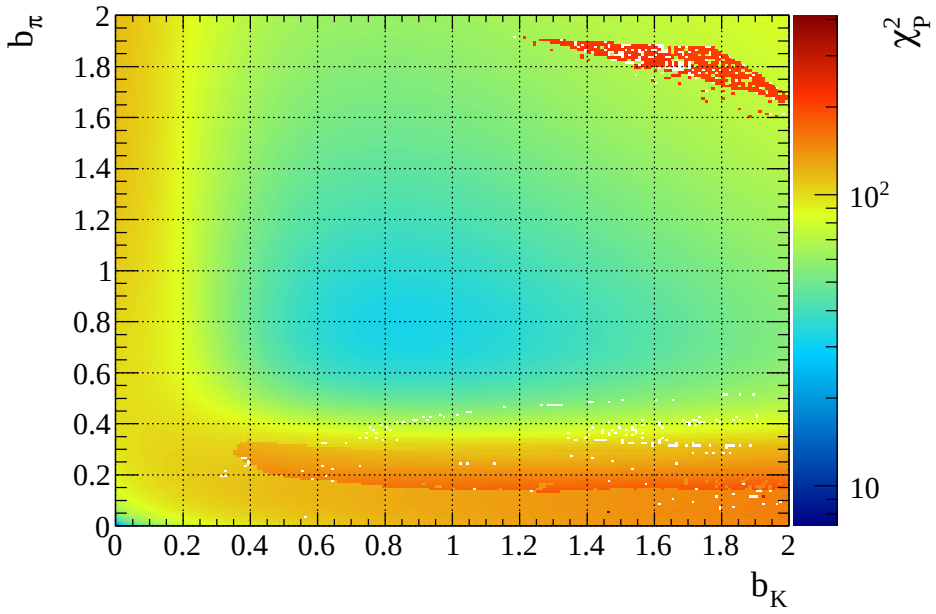
Figure 6.32: Plots showing the p_μ - $\cos\theta_\mu$ bin occupancies for the data (black), nominal MC (blue) and fitted MC (red). Figures (a)–(d) show the momentum bins for the following $\cos\theta_\mu$ bins in order: 0–0.9, 0.9–0.95, 0.95–0.98, 0.98–1.0 (as described in Section 6.6.4). Events with $p_\mu > 10$ GeV/c are included in the final momentum bin (4–10 GeV/c) for ease of display.

Figure 6.33(a) shows the χ^2_{P} surface (defined in Equation 6.14) in the b_K-b_π space. The minimum of this space defines the best fit values of b_K and b_π. It can be seen that there are actually two minima in this surface: the first at $b_K = 0.005$ and $b_\pi = 0.005$ with

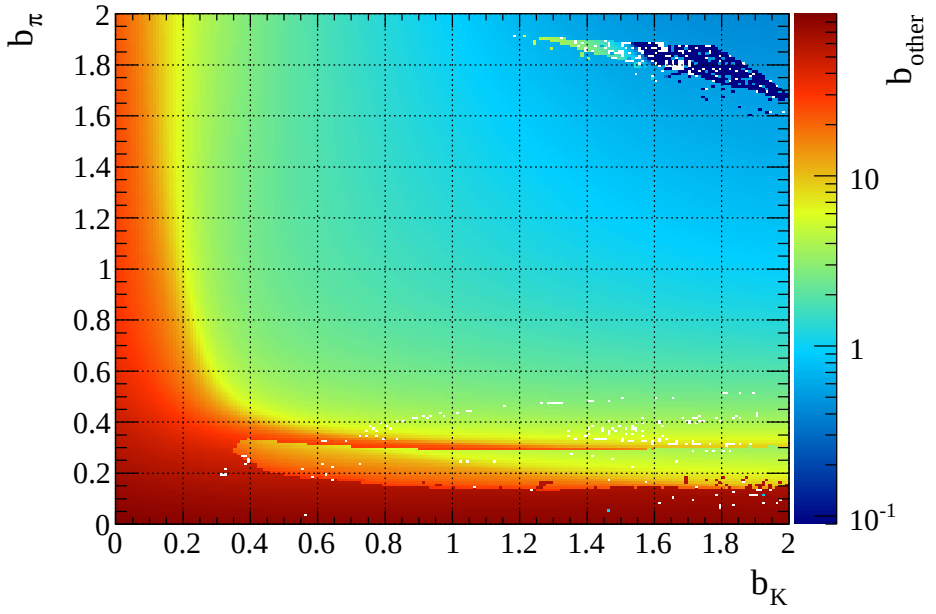
$\chi^2_{\mathrm{P}} = 21.2$; and the second at $b_K = 0.855$ and $b_\pi = 0.775$ with $\chi^2_{\mathrm{P}} = 32.1$. Since the first has a lower χ^2_{P} it is the primary minimum. Figure 6.33(b) shows the value of b_{other} in the b_K-b_π space. This is one of the nuisance parameters in the fit and is the normalisation for neutrinos with parents other than K^+ and π^+ (as described in Section 6.7.1). It can be seen that the value of b_{other} at the primary minimum is very high ($b_{\mathrm{other}} = 89.8$), whereas it is much closer to its nominal value of 1.0 at the secondary minimum ($b_{\mathrm{other}} = 2.0$). As described previously in Section 6.7.5, an unconstrained measurement of b_K and b_π is desired, and so no prior constraints are imposed upon the $\vec{b}$ parameters. The b_{other} parameter is therefore free to move far from the nominal value with no penalty. However, it is highly unlikely that the true value of b_{other}, and indeed the true values of b_K and b_π, are so far from their nominal values and so, for this reason, the primary minimum is disregarded. The white points in Figure 6.33(a) show points where the fit minimisation failed. These occur when b_π is very low or when b_K and b_π are very high, compared to their nominal values of 1.0. There are also some very high values of χ^2_{P} in these regions. It is to be expected that, for values of b_K or b_π far from the nominal, the fit will struggle to reconcile the data and MC, and so these points are not a problem. The χ^2_{P} value of 32.1 at $b_K = 0.855$ and $b_\pi = 0.775$ can be compared to the distribution in Figure 6.29, and a p-value of 46% is obtained. This shows that the data is in good agreement with the model.

Figure 6.34 shows the $\Delta\chi^2_{\mathrm{P}}$ surface (defined in Equation 6.15) in the b_K-b_π space. The primary minimum in Figure 6.33(a) has been disregarded as discussed above. Figure 6.34(b) is a magnified version of Figure 6.34(a). The best fit point (where $\Delta\chi^2_{\mathrm{P}} = 0$) is shown by a red star and is at $b_K = 0.855$ and $b_\pi = 0.775$. The black line is drawn at $\Delta\chi^2_{\mathrm{P}} = 4.61$ which, as discussed in Section 6.8.2, defines the 1σ contour around this point. It can be seen that the nominal point at $b_K = 1.0$ and $b_\pi = 1.0$ is just inside the 1σ contour. The contour is close to a circle which shows that the b_K and b_π parameters are essentially uncorrelated. This is reasonable since these parameters control two distinct parts of the neutrino energy spectrum and are free to move independently within the fit.

These results can be compared to the results of the "official" ND280 analysis, described previously in Section 3.4.1, which simultaneously constrains the flux and cross-section parameters in a fit to ND280 data in order to propagate them to the oscillation analysis at Super-Kamiokande. A sample of ν_μ charged-current events is selected from

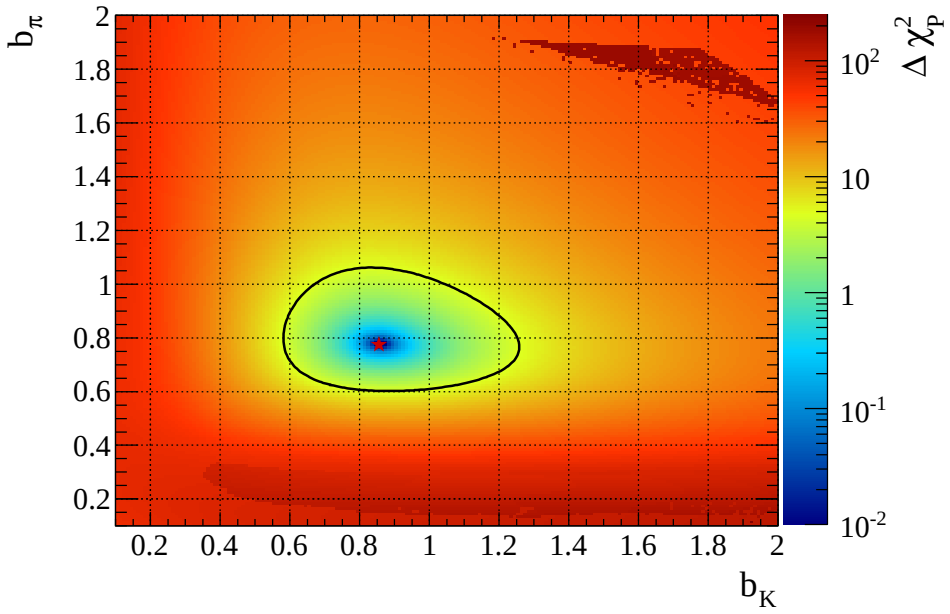


(a)

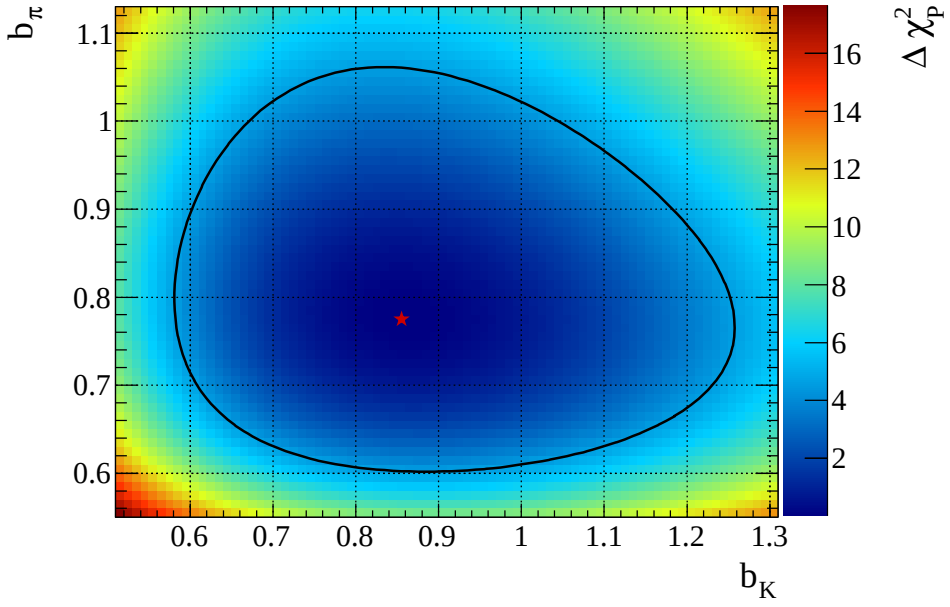


(b)

Figure 6.33: Results of fit to data. (a) shows the χ^2_{P} surface (defined in Equation 6.14) in the b_K-b_π space; there are two minima in this surface. (b) shows the value of b_{other} in this same b_K-b_π space. The white points in both plots show fits where the minimisation failed; these occur when b_π is very low or when b_K and b_π are very high.



(a)



(b)

Figure 6.34: Results of fit to data. Surface of $\Delta\chi^2_{\mathrm{P}}$ (defined in Equation 6.15) in b_K-b_π space, with the primary minimum in Figure 6.33(a) (at b_K = 0.005 and b_π = 0.005) disregarded. (b) is a magnified version of (a). The best fit point, where $\Delta\chi^2_{\mathrm{P}} = 0$, is shown by a red star. The black line shows the contour of $\Delta\chi^2_{\mathrm{P}} = 4.61$, representing the 1σ contour, as described in Section 6.8.2.

the ND280 data and binned according to the reconstructed momentum and angle of the muon candidate. A maximum likelihood method is then used to tune the MC prediction, which is a function of the flux and cross-section parameters, to the data. As explained earlier in Section 6.3, the analysis presented in this chapter is distinct from the official analysis in terms of the analysis goals, the event selections and the fit implementations. However, it is interesting to compare the results from the two analyses as far as possible. A full description of the official analysis and results can be found in [87], but the most relevant results are discussed here. There are 11 parameters controlling the flux prediction at ND280 in the official analysis, with each one being a normalisation factor for a bin of true neutrino energy for the ν_μ events at ND280. The energy ranges, prior values and fitted values of these flux parameters are shown in Table 6.6. It is difficult to directly compare these parameters to b_K and b_π, since they control different (if overlapping) aspects of the model. However, it can be seen from Table 6.6 that they all have fitted values less than 1.0, which is consistent with the best fit results of $b_K = 0.855$ and $b_\pi = 0.775$, in as far as they are also less than 1.0. The flux parameters of the official analysis appear to be slightly higher than the b_K and b_π parameters, but given the large errors on both sets of parameters, they are consistent.

In summary, the results of this analysis show that the best fit normalisations of K^+-originating and π^+-originating neutrinos are both lower than in the nominal MC, with values of 0.855 and 0.775 respectively. As shown in Figure 6.34, these results are consistent with the nominal MC at the 1σ level. However, the 1σ contour around this best fit point is very wide. The main source of uncertainty is from the large cross-section systematic errors, as shown in Section 6.8.2. As these cross-section errors are further constrained by ND280, the sensitivity of this method will increase.

Flux parameter	Energy range / GeV	Prior value	Fitted value
E1	0–0.4	1.0	0.959 ± 0.094
E2	0.4–0.5	1.0	0.961 ± 0.098
E3	0.5–0.6	1.0	0.950 ± 0.088
E4	0.6–0.7	1.0	0.931 ± 0.085
E5	0.7–1.0	1.0	0.893 ± 0.094
E6	1.0–1.5	1.0	0.889 ± 0.093
E7	1.5–2.5	1.0	0.919 ± 0.080
E8	2.5–3.5	1.0	0.953 ± 0.069
E9	3.5–5.0	1.0	0.962 ± 0.079
E10	5.0–7.0	1.0	0.973 ± 0.103
E11	7.0–30.0	1.0	0.980 ± 0.146

Table 6.6: Details of the 11 parameters controlling the ND280 flux for the official ND280 analysis. Each parameter is a normalisation factor for a given bin of true neutrino energy for the ν_μ events at ND280. The range in true neutrino energy, the prior value and the fitted value are shown for each flux parameter. Further details of this analysis and the full set of results can be found in [87].

Chapter 7

Conclusion

The T2K long-baseline neutrino oscillation experiment uses an accelerator-produced neutrino beam. To produce this beam, high energy protons impinge on a nuclear target to produce kaon and pion mesons which are then focused by magnetic horns and decay to neutrinos. It is a second generation experiment, building on the experience of the previous K2K and MINOS experiments. However, in contrast to these experiments, the main detectors (ND280 and Super-Kamiokande) are situated *off-axis* by 2.5^0.

Accurately predicting the flux of accelerator neutrino beams is a difficult task which has presented problems for several experiments prior to T2K. This is a particularly relevant issue for T2K since it uses the first off-axis beam. An accurate flux prediction is crucial for T2K to achieve the sensitivity required for its physics goals. The main flux uncertainty comes from hadronic interactions in the target material. An external hadron production experiment, called NA61, is key to reducing this uncertainty since it provides measurements of the hadronic production in the target at the same proton beam energy as T2K and with the same target material. However, these external measurements alone are not sufficient to constrain the flux since important factors in the T2K beamline are not accounted for, including target degradation over time, temperature related effects in the horn focusing, and interactions downstream of the target. Therefore, *in situ* measurements of the flux are central to a full validation. The work presented in this thesis uses data from the off-axis near detector (ND280) to validate the flux prediction.

A measurement is made of the normalisation of K^+ originating neutrinos at the ND280. The K^+ component of the beam is important since K^+ daughters dominate the high energy part of the ν_μ beam and contribute to the intrinsic ν_e contamination.

This measurement is affected by many aspects of the beam simulation, including the hadron production inside and outside of the target, the horn current and off-axis angle. It therefore provides an *in situ* probe of the entire beam system. The Run 2 data set is used, corresponding to 7.837×10^{19} protons on target. A sample of ν_μ charged-current interactions are selected using the ND280 tracker. This is a high statistics sample with a clear, robust signal in the detector. The events are binned according to the momentum and angle of the muon candidate. A fit methodology is developed to extract from the data the normalisation of both K^+ and π^+ originating neutrinos, b_K and b_π respectively. Extensive validation of this methodology is performed before applying the fit to the data. A profile likelihood method is used, in which the likelihood is maximised at each point in a 2-dimensional b_K-b_π space, resulting in a best fit point in this space and a 1σ contour. Nuisance parameters are included in the fit to model the effects of the neutrino interaction cross-section uncertainties and systematic errors of the detector system. The uncertainties on the flux prediction are not included in the fit in order to perform an unconstrained measurement of b_K, but are included in the validation studies. The best fit point is at $b_K = 0.86$ and $b_\pi = 0.78$ which is consistent with the nominal MC at the 1σ level. This measurement therefore shows that the ND280 data is consistent with the existing flux prediction within the current statistical and systematic uncertainties. At present, the 1σ contour around the best fit point is rather wide (with a radius on the order of 30%) which is mainly driven by the large cross-section uncertainties. As more data is taken at the ND280, the cross-section uncertainties will be reduced which in turn will increase the sensitivity of this measurement.

In addition, the results of the first time calibration of the ND280 detector, primarily the ECal sub-detector, are presented. The ND280 is a new complex detector that requires robust calibration. Good time resolution is important for effective noise rejection, event reconstruction and particle direction determination. Time offsets between electronic boards due to differences in cable length are calibrated using cosmic muon data. The fibre timewalk effect, which arises due to properties of the wavelength shifting fibres, is also calibrated. The resulting time resolution of the ECal after these calibrations are applied is presented.

Bibliography

[1] E. K. Akhmedov, Neutrino physics, in *Proceedings of ICTP Summer School in Particle Physics*, pp. 103–164, 1999, arXiv:hep-ph/0001264.

[2] G. Gratta, Nucl. Phys. Proc. Suppl. **85**, 72 (2000), arXiv:hep-ex/9905011.

[3] S. E. Kopp, Phys. Rept. **439**, 101 (2007), arXiv:physics/0609129.

[4] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, Science **124**, 103 (1956), http://www.sciencemag.org/content/124/3212/103.full.pdf.

[5] Particle Data Group, J. Beringer *et al.*, Phys. Rev. **D86**, 010001 (2012), Section 13 (Neutrino Mass, Mixing, and Oscillations).

[6] R. Davis, D. S. Harmer, and K. C. Hoffman, Phys. Rev. Lett. **20**, 1205 (1968).

[7] Super-Kamiokande Collaboration, Y. Fukuda *et al.*, Phys. Rev. Lett. **81**, 1562 (1998), arXiv:hep-ex/9807003.

[8] SNO Collaboration, Q. Ahmad *et al.*, Phys. Rev. Lett. **89**, 011301 (2002), arXiv:nucl-ex/0204008.

[9] KamLAND Collaboration, T. Araki *et al.*, Phys. Rev. Lett. **94**, 081801 (2005), arXiv:hep-ex/0406035.

[10] K2K Collaboration, E. Aliu *et al.*, Phys. Rev. Lett. **94**, 081802 (2005), arXiv:hep-ex/0411038.

[11] MINOS Collaboration, D. Michael *et al.*, Phys. Rev. Lett. **97**, 191801 (2006), arXiv:hep-ex/0607088.

[12] L. Wolfenstein, Phys. Rev. **D17**, 2369 (1978).

[13] S. P. Mikheev and A. Y. Smirnov, Sov. J. Nucl. Phys. **42**, 913 (1985).

[14] CHOOZ Collaboration, M. Apollonio *et al.*, Phys. Lett. **B466**, 415 (1999), arXiv:hep-ex/9907037.

[15] CHOOZ Collaboration, M. Apollonio *et al.*, Eur. Phys. J. **C27**, 331 (2003), arXiv:hep-ex/0301017.

[16] T2K Collaboration, K. Abe *et al.*, Phys. Rev. Lett. **107**, 041801 (2011), arXiv:hep-ex/1106.2822.

[17] DOUBLE-CHOOZ Collaboration, Y. Abe *et al.*, Phys. Rev. Lett. **108**, 131801 (2012), arXiv:hep-ex/1112.6353.

[18] DAYA-BAY Collaboration, F. An *et al.*, Phys. Rev. Lett. **108**, 171803 (2012), arXiv:hep-ex/1203.1669.

[19] RENO collaboration, J. Ahn *et al.*, Phys. Rev. Lett. **108**, 191802 (2012), arXiv:hep-ex/1204.0626.

[20] T. Nakaya, New Results from T2K, presented at NEUTRINO 2012, 2012.

[21] S. F. King, Contemporary Physics **48**, 195 (2007), http://www.tandfonline.com/doi/pdf/10.1080/00107510701770539.

[22] A. de Gouvea, A. Friedland, and H. Murayama, Phys. Lett. **B490**, 125 (2000), arXiv:hep-ph/0002064.

[23] N. V. Mokhov and S. I. Striganov, MARS15 overview, in *AIP Conference Proceedings 896*, p. 50, 2007.

[24] GEANT4 Collaboration, S. Agostinelli *et al.*, Nucl. Instrum. Meth. **A506**, 250 (2003).

[25] http://geant4.web.cern.ch/geant4/, accessed online 10$^{\text{th}}$ September 2012.

[26] A. Heikkinen, N. Stepanov, and J. P. Wellisch, eConf **C0303241**, MOMT008 (2003), arXiv:nucl-th/0306008.

[27] G. Folger, V. N. Ivanchenko, and J. P. Wellisch, Eur. Phys. J **A21**, 407 (2004).

[28] David W. Schmitz, *A measurement of hadron production cross-sections for the*

simulation of accelerator neutrino beams and a search for muon neutrino to electron neutrino oscillations in the $\Delta m^2 \sim 1\ eV^2$, PhD thesis, Columbia University, 2008.

[29] HARP Collaboration, M. Catanesi *et al.*, Nucl. Instrum. Meth. **A571**, 527 (2007).

[30] D. Bloess *et al.*, Nucl. Instrum. Meth **91**, 605 (1971).

[31] S. E. Kopp, The NuMI neutrino beam at Fermilab, in *Particle Accelerator Conference (PAC 05) Proceedings*, 2005, arXiv:physics/0508001.

[32] L. Loiacono, Neutrino Beam Flux Systematics, presented at NuFact09, 2009.

[33] M. Jerkins, Measuring the NuMI Beam Flux for MINERvA, presented at NUINT - 2011, 2011.

[34] K2K Collaboration, M. Ahn *et al.*, Phys. Rev. **D74**, 072003 (2006), arXiv:hep-ex/0606032.

[35] J. R. Sanford and C. L. Wang, AGS Internal Report, 1967.

[36] C. L. Wang, Phys. Rev. Lett. **25**, 1068 (1970).

[37] HARP Collaboration, M. Catanesi *et al.*, Nucl. Phys. **B732**, 1 (2006), arXiv:hep-ex/0510039.

[38] Y. Cho *et al.*, Phys. Rev. **D4**, 1967 (1971).

[39] MiniBooNE Collaboration, A. A. Aguilar-Arevalo *et al.*, Nucl. Instrum. Meth. **A599**, 28 (2009), arXiv:hep-ex/0806.4201.

[40] LSND Collaboration, A. Aguilar *et al.*, Phys. Rev. **D64**, 112007 (2001), arXiv:hep-ex/0104049.

[41] MiniBooNE Collaboration, A. A. Aguilar-Arevalo *et al.*, (2012), arXiv:hep-ex/1207.4809.

[42] MiniBooNE Collaboration, A. A. Aguilar-Arevalo *et al.*, Phys. Rev. **D79**, 072002 (2009), arXiv:hep-ex/0806.1449.

[43] HARP Collaboration, M. Catanesi *et al.*, Eur. Phys. J. **C52**, 29 (2007), arXiv:hep-ex/0702024.

[44] E910 Collaboration, I. Chemakin *et al.*, Phys. Rev. **C77**, 015209 (2008), arXiv:nucl-ex/0707.2375.

[45] R. P. Feynman, Phys. Rev. Lett. **23**, 1415 (1969).

[46] MINOS Collaboration, D. Michael *et al.*, Nucl. Instrum. Meth. **A596**, 190 (2008), arXiv:0805.3170.

[47] MINOS Collaboration, P. Adamson *et al.*, Phys. Rev. Lett. **106**, 181801 (2011), arXiv:hep-ex/1103.0340.

[48] MINOS Collaboration, P. Adamson *et al.*, Phys. Rev. **D77**, 072002 (2008), arXiv:hep-ex/0711.0769.

[49] A. Ferrari, P. R. Sala, A. Fasso, and J. Ranft, FLUKA: A Multi-Particle Transport Code, in *CERN-2005-010*, 2005.

[50] T2K Collaboration, K. Abe *et al.*, Nucl. Instrum. Meth. **A659**, 106 (2011), arXiv:1106.1238.

[51] T. Nakamoto *et al.*, IEEE Transactions on Applied Superconductivity **14**, 616 (2004).

[52] Official plots provided by the T2K collaboration.

[53] A. Vacheret *et al.*, Nucl. Instrum. Meth. **A656**, 69 (2011), arXiv:1101.1996.

[54] Y. Giomataris, P. Rebourgeard, J. Robert, and G. Charpak, Nucl. Instrum. Meth. **A376**, 29 (1996).

[55] I. Giomataris *et al.*, Nucl. Instrum. Meth. **A560**, 405 (2006), arXiv:physics/0501003.

[56] UA1 Collaboration, G. Arnison *et al.*, Phys. Lett. **B126**, 398 (1983).

[57] J. Altegoer *et al.*, Nucl. Instrum. Meth. **A404**, 96 (1998).

[58] A. Vacheret, S. Greenwood, M. Noy, M. Raymond, and A. Weber, The Front End Readout System for the T2K-ND280 Detectors, in *Nuclear Science Symposium Conference Record*, 2007.

[59] J. Estrada, C. Garcia, B. Hoeneisen, and P. Rubinov, MCM II and the Trip chip,

in *FERMILAB-TM-2226 (D0 note 4009)*, 2002.

[60] Pawel Guzowski, *Reconstruction of neutrino induced neutral current neutral pion events with the T2K ND280 Tracker and ECAL*, PhD thesis, Imperial College London, 2011.

[61] Super-Kamiokande Collaboration, Y. Ashie *et al.*, Phys. Rev. **D71**, 112005 (2005), arXiv:hep-ex/0501064.

[62] R. Brun, F. Carminati, and S. Giani, GEANT Detector Description and Simulation Tool, in *W5013*, CERN, 1994.

[63] C. Zeitnitz and T. A. Gabriel, The GEANT-CALOR interface, in *International Conference on Calorimetry in High Energy Physics*, p. 376, 1993.

[64] Particle Data Group, J. Beringer *et al.*, Phys. Rev. **D86**, 010001 (2012), Summary tables (for both leptons and mesons).

[65] T2K Collaboration, The T2K Neutrino Flux Prediction, paper in preparation.

[66] N. Abgrall *et al.*, Flux Prediction and Uncertainties for the 2012a Oscillation Analysis, T2K internal technical note (T2K-TN-099), 2012.

[67] NA61/SHINE Collaboration, N. Abgrall *et al.*, Phys. Rev. **C84**, 034604 (2011), arXiv:hep-ex/1102.0983.

[68] NA61/SHINE Collaboration, N. Abgrall *et al.*, Phys. Rev. **C85**, 035210 (2012), arXiv:hep-ex/1112.0150.

[69] T. Eichten *et al.*, Nuclear Physics B **44**, 333 (1972).

[70] J. V. Allaby *et al.*, High-energy particle spectra from proton interactions at 19.2 GeV/c, in *Technical Report 70-12*, CERN, 1970.

[71] R. J. Abrams *et al.*, Phys. Rev. **D1**, 1917 (1970).

[72] J. V. Allaby *et al.*, Yad. Fiz. **12**, 538 (1970).

[73] B. W. Allardyce *et al.*, Nucl. Phys. **A209**, 1 (1973).

[74] G. Bellettini *et al.*, Nuclear Physics **79**, 609 (1966).

[75] B. M. Bobchenko *et al.*, Sov. J. Nucl. Phys. **30**, 805 (1979).

[76] A. S. Carroll *et al.*, Phys. Lett. **B80**, 319 (1979).

[77] J. W. Cronin, R. Cool, and A. Abashian, Phys. Rev. **107**, 1121 (1957).

[78] F. F. Chen *et al.*, Phys. Rev. **99**, 857 (1955).

[79] S. P. Denisov *et al.*, Nucl. Phys. **B61**, 62 (1973).

[80] M. J. Longo and B. J. Moyer, Phys. Rev. **125**, 701 (1962).

[81] A. V. Vlasov *et al.*, Sov. J. Nucl. Phys. **27**, 222 (1978).

[82] M. Bonesini, A. Marchionni, F. Pietropaolo, and T. Tabarelli de Fatis, Eur. Phys. J. **C20**, 13 (2001), arXiv:hep-ph/0101163.

[83] D. Barton *et al.*, Phys. Rev. **D27**, 2580 (1983).

[84] P. Skubic *et al.*, Phys. Rev. **D18**, 3115 (1978).

[85] F. E. Taylor *et al.*, Phys. Rev. **D14**, 1217 (1976).

[86] K. Mahn, Constraining systematic errors using the T2K near detector, presented at NUFACT 2012, 2012.

[87] P. de Perio, M. Hartz, K. Mahn, and S. Oser, Constraining the flux and cross section models with data from the ND280 detector for the 2012a oscillation analysis, T2K internal technical note (T2K-TN-106), 2012.

[88] R. Brun and F. Rademakers, Nucl. Instrum. Meth. **A389**, 81 (1997).

[89] Y. Hayato, Acta Physica Polonica B **40**, 2477 (2009).

[90] C. Andreopoulos *et al.*, Nucl. Instrum. Meth. **A614**, 87 (2010), arXiv:hep-ph/0905.2517.

[91] A. Hillairet *et al.*, ND280 Reconstruction, T2K internal technical note (T2K-TN-072), 2011.

[92] T2K ND280 TPC Collaboration, N. Abgrall *et al.*, Nucl. Instrum. Meth. **A637**, 25 (2011), arXiv:1012.0865.

[93] C. Giganti and M. Zito, Particle Identification with the T2K TPC, T2K internal technical note (T2K-TN-001), 2009.

[94] G. Wikstrom, Global Kalman vertexing in ND280, T2K internal technical note (T2K-TN-046), 2011.

[95] James E. Y. Dobson, *Neutrino Induced Charged Current π^+ Production at the T2K Near Detector*, PhD thesis, Imperial College London, 2012.

[96] The ND280 calibration group, Calibration of the ND280 scintillator detectors, T2K internal technical note (T2K-TN-037), 2012.

[97] P. Masliah and A. Vacheret, Characterisation of the MPPC linearity response with TRIPt electronics and Y-11 fibre, T2K internal technical note (T2K-TN-047), 2011.

[98] The ND280 calibration group, Calibration of the ND280 scintillator detectors (Prod 5), upcoming T2K internal technical note (T2K-TN-130).

[99] G. Knoll, *Radiation Detection and Measurement* (John Wiley & Sons, 1999), p. 115.

[100] K. Kowalik and J. Lagoda, The simulation of beam neutrinos interactions outside the ND280 detector, T2K internal technical note (T2K-TN-077), 2012.

[101] C. Bojechko *et al.*, CCQE-like and CC-non-QE-like ν_μ event selections in the ND280 tracker using Run 1+2 data, T2K internal technical note (T2K-TN-093), 2012.

[102] A. Cervera, PID and Quality cut optimisation, Internal T2K ν_μ meeting, 19^{th} July 2011.

[103] Particle Data Group, J. Beringer *et al.*, Phys. Rev. **D86**, 010001 (2012), Section 36 (Statistics).

[104] F. Dufour, T. Lindner, and S. Oser, Systematics on Out-of-Fiducial-Volume Backgrounds in the ND280 Tracker, T2K internal technical note (T2K-TN-098), 2011.

[105] P. de Perio *et al.*, Cross section parameters for the 2012a oscillation analysis, T2K internal technical note (T2K-TN-108), 2012.

[106] P. de Perio *et al.*, Implementation of the NIWG Cross Section Parametrization, T2K internal technical note (T2K-TN-113), 2012.

[107] MiniBooNE Collaboration, A. A. Aguilar-Arevalo *et al.*, Phys. Rev. **D81**, 092005 (2010), arXiv:hep-ex/1002.2680.

[108] C. Juszczak, J. T. Sobczyk, and J. Żmuda, Phys. Rev. **C82**, 045502 (2010).

[109] E. J. Moniz *et al.*, Phys. Rev. Lett. **26**, 445 (1971).

[110] A. M. Ankowski and J. T. Sobczyk, Phys. Rev. **C77**, 044311 (2008).

[111] MiniBooNE Collaboration, A. A. Aguilar-Arevalo *et al.*, Phys. Rev. **D83**, 052009 (2011).

[112] MiniBooNE Collaboration, A. A. Aguilar-Arevalo *et al.*, Phys. Rev. D **83**, 052007 (2011).

[113] MiniBooNE Collaboration, A. A. Aguilar-Arevalo *et al.*, Phys. Rev. **D81**, 013005 (2010).

[114] K2K Collaboration, M. Hasegawa *et al.*, Phys. Rev. Lett. **95**, 252301 (2005), arXiv:hep-ex/0506008.

[115] SciBooNE Collaboration, K. Hiraide *et al.*, Phys. Rev. **D78**, 112004 (2008).

[116] NOMAD Collaboration, V. Lyubushkin *et al.*, Eur. Phys. J. **C63**, 355 (2009), arXiv:hep-ex/0812.4543.

[117] P. de Perio *et al.*, NEUT Systematic Studies for 2010a Analysis, T2K internal technical note (T2K-TN-032), 2011.

[118] MINOS Collaboration, P. Adamson *et al.*, Phys. Rev. **D81**, 072002 (2010).