THE UNIVERSITY OF
WARWICK

AUTHOR: **Mr Martin David Haigh** DEGREE: **Ph.D.**

TITLE: **Beam extrapolation and photosensor testing for the T2K experiment**

DATE OF DEPOSIT: .....

I agree that this thesis shall be available in accordance with the regulations governing the University of Warwick theses.

I agree that the summary of this thesis may be submitted for publication.

I **agree** that the thesis may be photocopied (single copies for study purposes only).

Theses with no restriction on photocopying will also be made available to the British Library for microfilming. The British Library may supply copies to individuals or libraries. subject to a statement from them that the copy is supplied for non-publishing purposes. All copies supplied by the British Library will carry the following statement:

> "Attention is drawn to the fact that the copyright of this thesis rests with its author. This copy of the thesis has been supplied on the condition that anyone who consults it is understood to recognise that its copyright rests with its author and that no quotation from the thesis and no information derived from it may be published without the author's written consent."

AUTHOR'S SIGNATURE: .....

USER'S DECLARATION

1. I undertake not to quote or make use of any information from this thesis without making acknowledgement to the author.
2. I further undertake to allow no-one else to use this thesis while it is in my care.

DATE SIGNATURE ADDRESS

.....

.....

.....

.....

.....

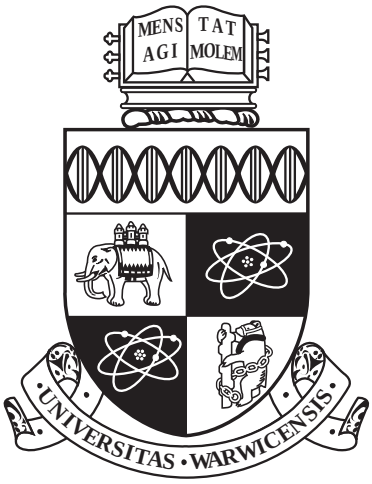

Beam extrapolation and photosensor testing for the T2K experiment

by

Mr Martin David Haigh

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Physics

May 2010

THE UNIVERSITY OF
WARWICK

Contents

List of Tables

List of Figures

Acknowledgments

As with almost all works of any merit (and that, of course, the reader must judge for herself), this thesis could not have been produced without the support of persons other than the author. I am very glad to have the opportunity to thank at least some of the people that have helped me along the way.

First of all, I would like to thank my supervisor, Gary Barker, who has kept me on the straight and narrow (more or less) for the past three-and-a-half years and especially during the production of the present work. He has made himself available for discussion with much greater frequency than could be expected of a supervisor, and has also tolerated my less than fastidious approach to workplace organisation with some forbearance. Furthermore, and to the reader's direct benefit, his thorough attention to detail has greatly reduced the number of typographical errors in this document, though the ones which inevitably remain are of course my own fault. In addition, I would like to thank Steve Boyd, who has also taken an essentially supervisory role during my time at Warwick. His long experience with neutrino experiments has proved an invaluable resource to me, and his good-natured resignation to the vicissitudes of research work has provided a welcome alternative to despair on many occasions.

I must also mention the support of Phill Litchfield, a PDRA during most of my time at Warwick and a true man of science. I am indebted to him for many valuable discussions and insights, and for his aid in my development as a physicist. I am also grateful to my contemporary Antony Carver for many interesting discussions on physics and other topics, and for the moral support and camaraderie which has been very important to me — for this I must also thank my younger colleagues, Leigh Whitehead, Andrew Bennieston and Daniel Scully.

I would like to acknowledge the confidence of Antonin Vacheret, who entrusted to me much of the management of the UK photosensor testing, which makes up a substantial part of the current work. Antonin also encouraged me to present much of the UK photosensor group's work at meetings myself. Similarly, I should also mention Fabrice Retière, who gave me the task of maintaining the MPPC simulation and asked me to present this work at the PD09 workshop.

On a more personal note, I would like to thank my parents, step-parents and siblings, who have been constantly supportive and have, as families do, provided some damping to my oscillations between under-confidence and hubris. Finally and most importantly, I would like to thank my wife Sarah, who has unselfishly put my work first throughout my Ph.D., and has managed the more worldly matters for which I am sometimes ill-equipped. I am lucky to have her, and am eternally grateful for her devotion.

Declarations

Chapter 1 of this work consists of a review of the theoretical and experimental state of the field, and is synthesised from textbooks and journal articles. Chapter 2 covers the T2K experiment and is mostly drawn from internal T2K technical notes. Chapter 6 is a short conclusion, which is again drawn from existing sources. The original material in the thesis is contained in Chapters 3–5.

Chapter 3 consists of a discussion of MPPC characterisation measurements and Monte Carlo simulation. The characterisation measurements were carried out at other institutions, but I was involved in discussions regarding the development of the analyses for this data. I was responsible for the implementation of the physics models discussed in Chapter 3 into the simulation. The saturation model used for the TFB board is entirely my own work. My development of the simulation was based on an existing framework contained in the main ND280 codebase, but I made several refinements to this framework, and produced a standalone version independent of the ND280 software. I used this version of the simulation to produce all of the MC results shown. Some of the material in this chapter was presented by me at the PD09 conference [1].

Chapter 4 contains a discussion of the MPPC QA for the ND280 Down-

stream ECal. Most of the hardware design and construction is not my own work, although I was involved in the design and testing of the LED pulser and diffuser system. I was responsible for writing the software to steer the DAQ system, and managing the QA data taking at Warwick. I was also responsible for the code to analyse the data spectra and do the high level analysis (combining the data from multiple spectra to extract MPPC properties). The former was done in conjunction with colleagues, but I was in overall charge of code development and responsible for all the analysis routines described in the thesis.

Chapter 4 also contains a discussion of MPPC calibration based on CERN testbeam data for the Downstream ECal. I was not directly involved in the test-beam data taking, though I was responsible for collating logs of the testbeam run metadata. The code to produce and apply MPPC calibration constants was written by me, though the other steps required to calibrate the MIP response of the ECal were handled by colleagues. I was responsible for the validation of the MPPC calibration software.

Chapter 5 is an evaluation of beam extrapolation techniques for T2K. The beam matrix and far/near ratio methods described in the text are similar to those used by MINOS and K2K respectively, and these experiments are referenced in the chapter. The implementation of the beam matrix calculation is my own work, though the decay physics are either drawn from standard Special Relativity or from specific sources as referenced. The systematic alteration of beam parameters and running of the Monte Carlo were done by myself, and the evaluation of ratio and matrix performance are also my work.

Abstract

Our understanding of the physics of neutrino oscillations has evolved rapidly over the past decade or so, with results from the SNO, Super-K, MINOS and CHOOZ experiments, among others, producing results favouring a three-neutrino mixing model, and significantly constraining the parameter space for this mixing. There are still several important questions to be answered however: we do not know whether θ_{13} is non-zero, or whether $(\sin^2 2\theta_{23})$ is maximal; also, we do not know the sign of the large mass splitting ΔM^2, or whether CP-violation occurs in the lepton sector. The latter is possibly the most exciting of all — leptonic CP-violation is a requirement for leptogenesis, and could therefore indicate a solution to the matter-antimatter asymmetry problem in cosmology.

The T2K long-baseline neutrino experiment is one of a new generation of neutrino projects, which will make more precise measurements of θ_{13} and θ_{23} than has been achieved by previous experiments. It uses the Super-K water Čerenkov detector at Kamioka as a far detector, and also has a suite of new near detectors. These are largely scintillator-based, but use a novel photosensor, the silicon photomultiplier (SiPM), for light readout. T2K has been leading the effort to understand and model these new sensors, and the present work will describe the current state-of-the-art in device characterisation, and also the effort to ensure the quality of the devices installed in the calorimeter of the ND280 near detector.

An important part of a long-baseline analysis is the extrapolation of the neutrino flux measured at the near detector to predict that at the far detector. Methods to do this have been developed by previous experiments; however T2K uses an innovative configuration whereby the main detectors are displaced from the neutrino beam centre, removing much of the high-energy tail in the neutrino flux to reduce background from non-quasielastic events. This thesis evaluates the effectiveness of two extrapolation techniques, used by previous experiments, for the T2K configuration.

Abbreviations

AP	Afterpulsing
APD	Avalanche Photodiode
CCQE	Charged Current Quasi-elastic
CT	Crosstalk
DCR	Dark Count Rate
MPPC	Multi-Pixel Photon Counter
NCQE	Neutral Current Quasi-elastic
PDE	Photon Detection Efficiency
p.e.	Pixel Equivalent
p.o.t.	Protons on Target
PMT	Photomultiplier Tube
TFB	Trip-t Front-end Board. These boards are used for readout of the MPPCs in ND280

Chapter 1

Overview of neutrino theory and the experimental determination of mixing parameters

1.1 History

1.1.1 Discovery of the neutrino

The neutrinos ($\nu_e,\ \nu_\mu$ and ν_τ) are the electrically neutral weak-isospin partners of the charged leptons (e, μ and τ) in the Standard Model. They interact with other particles via only the weak interaction, and have extremely small masses — the limit on the sum of neutrino masses set by cosmology is of the order a few eV, though it is rather model-dependent [2]. Until recently neutrinos were thought to be completely massless.

The neutrino was originally proposed by Pauli to explain the continuous electron spectrum observed in beta-decay experiments. The decay

$$ {}^{A}_{Z}N \rightarrow {}_{Z+1}{}^{A}N' + e^{-} \tag{1.1} $$

is modified by the inclusion of a neutrino to

$$ {}^{A}_{Z}N \rightarrow {}_{Z+1}{}^{A}N' + e^{-} + \bar{\nu}_{e}. \tag{1.2} $$

Since the latter decay has three daughters, their energies are not uniquely determined. The shape of the electron energy spectrum, and its endpoint relative to the Q-value of the decay, suggest a massless, or nearly massless, neutrino [3] [4].

The neutrino was first directly detected by Reines and Cowan in experiments conducted between 1953 and 1956 [5], observing the interaction

$$ \bar{\nu} + p \rightarrow e^{+} + n \tag{1.3} $$

with neutrinos from a fission reactor in liquid scintillator. The neutrino signature was a coincidence between a prompt positron scintillation pulse and a delayed pulse from neutron capture on cadmium.

The neutrino accompanying a μ-lepton in interactions was shown to be distinct from the electron neutrino in an experiment carried out at Brookhaven by Lederman, Schwartz and Steinberger in 1962 [6]. They used an aluminium spark chamber, surrounded by passive shielding and anticoincidence triggers, to observe the interactions of neutrinos produced via pion decays

$$ \pi^{+/-} \rightarrow \mu^{+/-} + \nu/\bar{\nu} \tag{1.4} $$

The interactions of these neutrinos in the aluminium were observed to produce only muons, not electrons, indicating that the observed events were ν_μ charged-current interactions and no ν_e were present.

The presence of a tau neutrino was inferred from the structure of the first two lepton generations once the tau lepton was discovered — this was indirectly confirmed by the LEP measurement of the number of active neutrinos from the Z boson width [7]. The tau neutrino was directly detected in 2000 by the DONUT collaboration [8], using photographic emulsion to capture the "kinked" track of a short-lived tau lepton, produced by a neutrino interaction, decaying to a muon.

1.1.2 Neutrino mixing

Shortly before the discovery of the muon neutrino, it was suggested by Pontecorvo [9], and Maki, Nakagawa and Sakata [10], that two-flavour mixing, analogous to that observed in the K^0 system, should be seen in the neutrino system. Such mixing requires that the Hamiltonian matrix for the neutrino system mixes the ν_e and ν_μ flavour states; this is equivalent to saying that the flavour-basis eigenstates of the neutrino are not identical with the energy-basis states. For mixing to be observed, the two energy (or mass) eigenstates must be non-degenerate, so at least one must have non-zero mass. After the discovery of the tau, this theory was extended for three generations of leptons, by introducing a mixing matrix (called the PMNS matrix) analogous to the CKM matrix governing quark mixing [7].

The first experimental hint toward neutrino oscillations was the result from Ray Davis' solar neutrino observatory at Homestake mine [11]. Davis used a tank of

390,000 litres of tetrochloroethylene (C_2Cl_4, otherwise known as dry-cleaning fluid) as a target for solar neutrinos produced by ^{8}B decay. The neutrinos transmuted the chlorine in the tank via the reaction

$$^{37}\mathrm{Cl} + \nu_e \rightarrow {}^{37}\mathrm{Ar} + e^-. \tag{1.5}$$

The resultant argon (of the order 100 atoms in 390,000 litres) was flushed from the tank after about a month by purging the fluid with helium. The argon was then separated from the helium by adsorption onto charcoal and the radioactive atoms counted, via observation of their decays in a proportional counter. The results suggested a neutrino flux smaller than the prediction of the Standard Solar Model (SSM), developed by Bahcall and others [12].

A deficit in the solar neutrino flux was confirmed by later experiments including SAGE [13], GALLEX [14] and Kamiokande-II [15]. Several possible explanations existed for this discrepancy, including potential problems with the SSM. The solar model was exonerated by the SNO heavy-water Čerenkov detector in Canada [16], which was able to measure the total flux of neutrinos via neutral-current neutrino interactions on deuterium

$$\nu_\alpha + {}^2\mathrm{H} \rightarrow \nu_\alpha + p + n \qquad (\alpha = e,\ \mu,\ \tau). \tag{1.6}$$

SNO was able to combine this measurement with the charged-current and elastic-scattering event rates to confirm that the total neutrino flux agreed with the SSM. This result, backed up by the deficits reported by experiments sensitive to only ν_e, led to the conclusion that neutrinos must change flavour as they propagate through space. Other mechanisms, such as neutrino decay and decoherence, have

been proposed to explain the SNO results. However, once all oscillation data is considered, these "exotic" models are disfavoured [17] [18].

1.2 Neutrino theory

1.2.1 Mass generation

As discussed above, the explanation of neutrino flavour mixing requires massive neutrinos. Full derivations for the generation of Dirac and Majorana masses will not be given here, but this section outlines the key arguments needed — a more complete treatment can be found in, e.g., [19] [20].

In the Standard Model Lagrangian, explicit mass terms for fermions are forbidden — a Dirac mass term $m_D\bar{\psi}\psi$ can be split into its chiral components:

$$\mathcal{L}_{m_\mathrm{D}} = m_D\bar{\psi}\psi = m_D(\bar{\psi_L} + \bar{\psi_R})(\psi_L + \psi_R), \tag{1.7}$$
$$\text{where } \psi_L = P_L\psi = (1-\gamma_5)\psi, \quad \psi_R = P_R\psi = (1+\gamma_5)\psi.$$

Since $\gamma_5^\dagger = -\gamma_5$, $\bar{\psi_L} = \bar{\psi}P_R$, and we have the result:

$$\bar{\psi_L}\psi_L = \bar{\psi}P_RP_L\psi = 0; \qquad \text{similarly } \bar{\psi_R}\psi_R = 0. \tag{1.8}$$

Setting these terms to zero, (1.7) becomes

$$\mathcal{L}_{m_\mathrm{D}} = m_D(\bar{\psi_R}\psi_L + \bar{\psi_L}\psi_R). \tag{1.9}$$

However, in the Standard Model each lepton generation is described by an $SU(2)$ doublet and a singlet:

$$\psi_L \equiv \begin{pmatrix} \nu_l \\ l \end{pmatrix}_L ; l_R \qquad l = e, \mu, \tau. \tag{1.10}$$

The ψ_L doublet cannot be contracted with the singlet to give a numerical result, so terms like those in (1.9) cannot exist in the Standard Model Lagrangian.

Introducing the Standard Model Higgs doublet allows fermion masses to be generated through interactions with the Higgs vacuum expectation value (vev), which multiplies with the doublet vector ψ_L to give two terms which do give a numerical result when multiplied by the singlet. One of these terms is a mass term for the charged lepton

$$m_D \bar{l} l = m_D (\bar{l_L} l_R + \bar{l_R} l_L), \tag{1.11}$$

where m_D is equal to the product of the Higgs vev and coupling, multiplied by a factor $1/\sqrt{2}$. A term in $\bar{\nu}\nu$, giving a Dirac mass to the neutrino, can also exist if ν_R is added to the Standard Model. This does not present a problem — the lack of a right-handed neutrino in the Standard Model is a consequence of non-zero neutrino masses not being discovered until recently, rather than of theoretical prejudice.

It is also possible to generate neutrino masses by a different mechanism than for other fermions, by assuming that the neutrino is its own antiparticle. This means that the neutrino field behaves under the charge conjugation operator C as

$$C\nu_L = \nu_L^C = \pm\nu_R, \qquad C\nu_R = \nu_R^C = \pm\nu_L. \tag{1.12}$$

Such a particle is called a Majorana particle. In this case, we can include terms in the Lagrangian which directly couple the left-handed neutrino to the right-handed

antineutrino and vice versa:

$$\mathcal{L}^L = \frac{1}{2} m_L (\bar{\psi_L} \psi_R^C + \bar{\psi_R^C} \psi_L) \tag{1.13}$$

$$\mathcal{L}^R = \frac{1}{2} m_R (\bar{\psi_R} \psi_L^C + \bar{\psi_L^C} \psi_R). \tag{1.14}$$

In the case that m_D is zero, there will be two mass states $(m_L,\ m_R)$. If the Dirac mass is non-zero, there is coupling of the left and right-handed neutrino states, and the physical neutrino states (mass eigenstates) will be a superposition of pure left and right handed states. It can be shown that this gives mass eigenstates

$$m_{1,2} = \frac{1}{2} \left[(m_L + m_R) \pm \sqrt{(m_L - m_R)^2 + 4m_D^2} \right]. \tag{1.15}$$

A case which is particularly interesting is $m_L = 0,\ m_D \ll m_R$. This will give mass eigenstates $m_1 \simeq m_R,\ m_2 \simeq \frac{m_D^2}{m_R}$. The state m_1 is dominated by a ν_R component, and m_2 by ν_L. This could explain why the neutrino is observed to have a much smaller mass than other fermions — for a Dirac mass on a similar scale to the other fermions, a light left-handed neutrino, and heavy unobserved right-handed neutrino, will be produced. This is called the Seesaw Mechanism.

1.2.2 PMNS neutrino mixing matrix

The PMNS model for neutrino mixing is based on the hypothesis that the neutrino mass eigenstates ν_i are not identical with the weak (flavour) eigenstates ν_α, but are related to them by a unitary transformation:

$$|\nu_i\rangle = \sum_{\alpha=1}^{3} U_{\alpha i} |\nu_\alpha\rangle \tag{1.16}$$

$$\text{where } U^\dagger U = I; \qquad \text{i.e. } \sum_{i=1}^{3} U_{\alpha i}^* U_{\beta i} = \delta_{\alpha\beta}. \tag{1.17}$$

Here (1.17) is just the condition satisfied by a unitary matrix. The matrix U is referred to as the *PMNS matrix*. We will show that this relationship between mass and flavour states leads to distance-dependent neutrino oscillations. The free neutrinos will be treated as plane waves — it has been shown [21] that a full wave-packet treatment gives the same result.

A neutrino produced by a weak process will be in a state of definite flavour α. In order to evolve the state to a later time t, the flavour state must be expressed in terms of mass basis states:

$$|\nu_\alpha\rangle = \sum_{\imath=1}^{3} U^*_{\alpha i}|\nu_i\rangle. \tag{1.18}$$

This expression is found by premultiplying (1.16) by $U^*_{\beta i}$, performing the summation $\sum_{i=1}^{3}$, and applying the unitarity condition (1.17). The mass states have definite energy, and therefore have a plane-wave space-time dependence $e^{i(\mathbf{p}.\mathbf{x}-Et)}$. We suppose that the momentum p of the produced neutrino has a definite value; then if $p \gg m_i$ for all i, the energies for each mass state are given by $E_i = \sqrt{p^2 + m_i^2} \simeq p + \frac{m_i^2}{2p}$. The state (1.18) evolves in time, neglecting an unobservable phase common to all terms, as

$$|\nu_\alpha(t)\rangle = \sum_{i=1}^{3} e^{-im_i^2 t/2p} U^*_{\alpha i}|\nu_i\rangle. \tag{1.19}$$

The overlap of this state with the flavour state β is given by

$$\langle\nu_\beta|\nu_\alpha(t)\rangle = \sum_{i,j}(\langle\nu_j|U_{\beta j})(e^{-im_i^2 t/2p} U^*_{\alpha i}|\nu_i\rangle). \tag{1.20}$$

Since the overlap of the basis states $\langle\nu_j|\nu_i\rangle = \delta_{ij}$, we can drop one of the summation indices. The probability to observe state β is then

$$P_{\alpha\to\beta} = |\langle\nu_\beta|\nu_\alpha(t)\rangle|^2 = \sum_{i,k} U_{\beta i} U^*_{\alpha i} U^*_{\beta k} U_{\alpha k} e^{-i(m_i^2 - m_k^2)t/2p}. \tag{1.21}$$

This expression is positive definite as expected, since the terms linked by interchanging i and k are complex conjugates. Using this fact to justify explicitly taking the real part of the expression, and also using the unitarity of the PMNS matrix, it follows that

$$P_{\alpha\to\beta} = \delta_{\alpha\beta} - 4\sum_{i>j} \Re(U_{\beta i}U^*_{\alpha i}U^*_{\beta j}U_{\alpha j}) \sin^2\left(\frac{\Delta m^2_{ij}L}{4E}\right) + 2\sum_{i>j} \Im(U_{\beta i}U^*_{\alpha i}U^*_{\beta j}U_{\alpha j}) \sin\left(\frac{\Delta m^2_{ij}L}{2E}\right) \tag{1.22}$$

where we have replaced the elapsed time t with the distance travelled L, and the momentum p with the energy E. For $p \gg m_i$ for all i, these replacements are valid. We have also introduced the symbol $\Delta m^2_{ij} = m^2_j - m^2_i$.

It can easily be shown that for ultra-relativistic neutrinos, assuming a single energy for all mass eigenstates (and therefore different momenta) gives the same result as (1.22). There is no particular significance in assuming a fixed value for one or the other quantity, unless one is known with an uncertainty on the level of the neutrino mass scale.

1.2.3 Vacuum mixing phenomenology

It is common to parameterise the PMNS matrix using an xyz Euler angle factorisation, with the inclusion of one non-trivial phase:

$$
\begin{aligned}
U_{\mathrm{PMNS}} &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \\
&= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix}
\end{aligned} \tag{1.23}
$$

where $s_{ij} = \sin\theta_{ij},\ c_{ij} = \cos\theta_{ij}$. This form of the matrix factorises mixing into three "two-flavour" cases, though of course the two leftmost matrices operate on an already-rotated basis so the interpretation is not completely elementary. Five complex phases are omitted; three of these correspond to unobservable rotations of the neutrino flavour eigenstates, the other two to relative phases of the mass eigenstates. These latter phases are unobservable for a Dirac neutrino, but affect the amplitude for neutrinoless double-beta decay [7] if the neutrino is a Majorana particle. These phases are included by postmultiplying (1.23) by the matrix $\mathrm{diag}(e^{i\alpha_1/2}, e^{i\alpha_2/2}, 1)$, with α_1, α_2 real, but the phases do not affect oscillation behaviour for either Dirac or Majorana neutrinos, so they will not be discussed further here.

For the three-flavour Dirac case, then, mixing is determined by seven parameters: three mixing angles θ_{12}, θ_{13} and θ_{23}, three mass splittings Δm^2_{12}, Δm^2_{13} and Δm^2_{23} (only two of which are independent), and one phase δ. Experimentally,

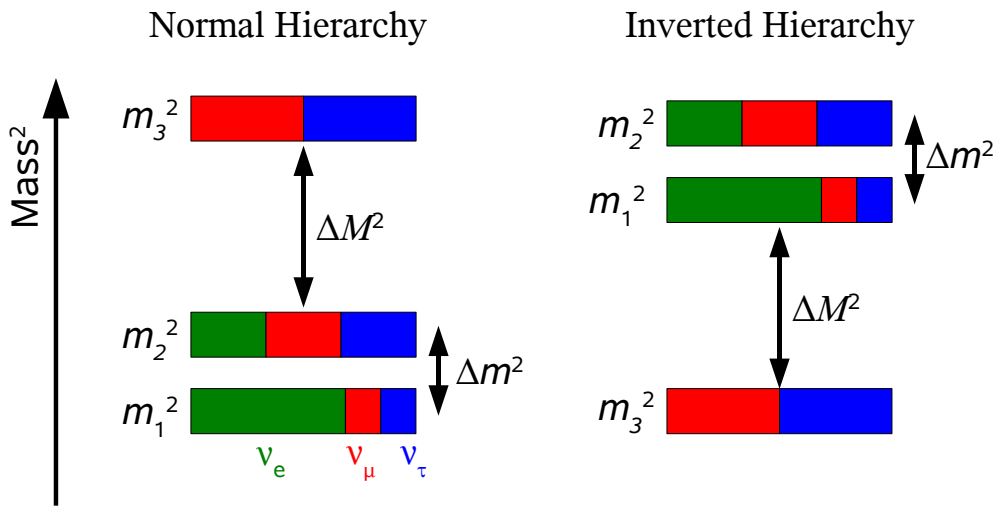

Figure 1.1: Schematic showing the mass states and their approximate flavour composition in the normal and inverted mass hierarchies. The ν_e component of ν_3 is small and may be zero. The size of Δm^2 relative to ΔM^2 is exaggerated, and the absolute mass values are unknown.

the mass difference between states 1 and 2 is much smaller than the difference between these and state 3, so the approximate parameterisation $(\Delta m^2,\ \Delta M^2)$ is often used, where $\Delta m^2 = \Delta m^2_{12}$ and $\Delta M^2 \simeq \Delta m^2_{13} \simeq \Delta m^2_{23}$. Vacuum oscillations cannot determine the sign of the mass splittings; these can be measured using matter effects (discussed in Section 1.2.4). Δm^2 is known to be positive from solar neutrino data, but the sign of ΔM^2 is presently unknown. The case where m_3 is the most massive state is referred to as the normal mass hierarchy, and the case where it is least massive as the inverted mass hierarchy, as shown in Figure 1.1.

It is expected that CP-violation is present in the lepton sector, since it has been observed in the quark sector, which exhibits otherwise analogous mixing behaviour. For antineutrino mixing, the oscillation probability is the same as that

given for neutrinos in (1.22), except that the matrix $U \to U^*$, so that CP-violation is only present if the PMNS matrix is not real ($\delta \neq 0$). δ has not yet been measured; a measurement will be challenging since it only appears in the PMNS matrix alongside s_{13}, and θ_{13}-driven oscillations are known to be small and yet to be observed. δ can be measured either by comparing ν to $\bar{\nu}$ oscillations, or by precisely measuring the shape of the oscillation probability as a function of L/E, since the last two terms in (1.22) have different L/E dependencies.

For the present work, the key result is the probability for a neutrino created with muon flavour to be detected later as a muon neutrino, or as an electron neutrino. This probability is somewhat simplified in the case that $\Delta m^2 L/2E \ll 1$ — when we remove terms in Δm^2 from (1.22) and apply the unitarity constraint $\sum_i U^*_{\beta i} U_{\alpha i} = \delta_{\alpha\beta}$, we are left with

$$P_{\alpha\to\beta} \simeq \delta_{\alpha\beta} - 4\left(\delta_{\alpha\beta}|U_{\alpha 3}|^2 - |U_{\beta 3}|^2|U_{\alpha 3}|^2\right)\sin^2\left(\frac{\Delta M^2 L}{4E}\right). \tag{1.24}$$

Note that this expression does not depend on the phases of the PMNS matrix elements, so CP violation is not included in this approximation. The results obtained from (1.24) are equally valid for neutrinos and antineutrinos.

Using the mixing-angle convention (1.23) for the PMNS matrix elements,

and converting the $\frac{\Delta M^2 L}{4E}$ from natural to experimental units, we obtain

$$P_{\mu\to\mu} \simeq 1 - 4s_{23}^2 c_{13}^2 (1 - s_{23}^2 c_{13}^2) \sin^2 \left(1.27 \frac{\Delta M^2(\mathrm{eV})^2 L(\mathrm{km})}{E(\mathrm{GeV})} \right)$$
$$\simeq 1 - \sin^2(2\theta_{23}) \sin^2 \left(1.27 \frac{\Delta M^2(\mathrm{eV})^2 L(\mathrm{km})}{E(\mathrm{GeV})} \right), \quad (1.25)$$
$$P_{\mu\to e} \simeq 4s_{23}^2 c_{13}^2 s_{13}^2 \sin^2 \left(1.27 \frac{\Delta M^2(\mathrm{eV})^2 L(\mathrm{km})}{E(\mathrm{GeV})} \right)$$
$$\simeq s_{23}^2 \sin^2(2\theta_{13}) \sin^2 \left(1.27 \frac{\Delta M^2(\mathrm{eV})^2 L(\mathrm{km})}{E(\mathrm{GeV})} \right), \quad (1.26)$$

where in equation (1.25) we have used the approximation $c_{13} \simeq 1$.These equations describe the physics relevant for the initial configuration of the T2K experiment. Similarly, we can obtain the survival probability for an electron neutrino in the same L/E regime:

$$P_{e\to e} \simeq 1 - \sin^2(2\theta_{13}) \sin^2 \left(1.27 \frac{\Delta M^2(\mathrm{eV})^2 L(\mathrm{km})}{E(\mathrm{GeV})} \right). \quad (1.27)$$

This expression is relevant to short-baseline reactor neutrino experiments.

1.2.4 Neutrino mixing in matter

It was discovered by Wolfenstein in 1978 [22] that the Hamiltonian of a neutrino in matter is subject to perturbations from coherent forward scattering with electrons and nucleons. Such a perturbation can be shown to have an effect on the mixing of neutrino flavour states. It was later shown by Mikheev and Smirnov [23] that matter effects can lead to a resonant enhancement of the mixing angle, producing large flavour transitions even where the vacuum angle is small (MSW effect). The treatment of mixing in matter is rather more complex than the vacuum case, even

for two flavours. A fairly complete discussion is given by, for example [19]; an overview of this will be presented here.

At energies much less than the weak boson masses, the Hamiltonian associated with tree-level charged-current coherent forward scattering ($\nu_e e^- \rightarrow \nu_e e^-$) can be shown to lead to an increase in energy for an electron neutrino passing through matter, of

$$V_{CC} = \sqrt{2} G_{\mathrm{F}} N_e, \tag{1.28}$$

where N_e is the electron number density. Similarly all weakly-interacting neutrinos also gain a potential due to the neutral current,

$$V_{\mathrm{NC}} = -\frac{1}{2}\sqrt{2} G_{\mathrm{F}} N_n, \tag{1.29}$$

depending only on the neutron density N_n since electron and proton contributions cancel. The neutral-current term does not affect oscillations if only active neutrino flavours exist, since the potential for each flavour state is the same. If sterile flavours exist, which do not interact with the weak bosons, mixing will be affected by the neutral-current potential of the active flavours.

The effect of the matter potential on oscillations may be calculated by adding it to the vacuum Hamiltonian H_0. The eigenstates of the vacuum Hamiltonian are the neutrino mass states $|\nu_k\rangle$:

$$H_0|\nu_k\rangle = \sqrt{p^2 + m_k^2}|\nu_k\rangle \tag{1.30}$$

and the eigenstates of the matter interaction Hamiltonian H_{I} are the flavour states $|\nu_\alpha\rangle$:

$$H_{\mathrm{I}}|\nu_\alpha\rangle = V_\alpha|\nu_\alpha\rangle \qquad V_\alpha = V_{\mathrm{CC}} \text{ for } \alpha = e,\ 0 \text{ otherwise.} \tag{1.31}$$

Since the vacuum mass eigenstates are no longer eigenstates of the full Hamiltonian, we must find the new energy-eigenstate basis by calculating the Hamiltonian matrix in the flavour basis, and finding a transformation which diagonalises it. It turns out that the flavour-basis Hamiltonian, in the two-flavour case, is given by

$$H_{\mathrm{F}} = \frac{1}{4E}\begin{pmatrix} -\Delta m^2 \cos 2\theta + A_{\mathrm{CC}} & \Delta m^2 \sin 2\theta \\ \Delta m^2 \sin 2\theta & \Delta m^2 \cos 2\theta - A_{\mathrm{CC}} \end{pmatrix}, \tag{1.32}$$

where the term due to matter, A_{CC}, is given by

$$A_{\mathrm{CC}} = 2\sqrt{2} E G_{\mathrm{F}} N_e, \tag{1.33}$$

and θ and Δm^2 are the vacuum mixing angle and mass splitting respectively. A_{CC} takes the opposite sign for antineutrinos since the potentials change sign.

(1.32) is diagonalised by the transformation $U_{\mathrm{M}}^T H_{\mathrm{F}} U_{\mathrm{M}} = H_{\mathrm{M}}$, with

$$U_{\mathrm{M}} = \begin{pmatrix} \cos\theta_{\mathrm{M}} & \sin\theta_{\mathrm{M}} \\ -\sin\theta_{\mathrm{M}} & \cos\theta_{\mathrm{M}} \end{pmatrix}, \qquad H_{\mathrm{M}} = \frac{1}{4E}\begin{pmatrix} -\Delta m^2_{\mathrm{M}} & 0 \\ 0 & \Delta m^2_{\mathrm{M}} \end{pmatrix} \tag{1.34}$$

where Δm^2_{M} and θ_{M} are the effective mass splitting and mixing angle in matter. They are related to the vacuum values by

$$\Delta m^2_{\mathrm{M}} = \sqrt{(\Delta m^2 \cos 2\theta - A_{\mathrm{CC}})^2 + (\Delta m^2 \sin 2\theta)^2} \tag{1.35}$$

$$\tan 2\theta_{\mathrm{M}} = \frac{\tan 2\theta}{1 - \frac{A_{\mathrm{CC}}}{\Delta m^2 \cos 2\theta}}. \tag{1.36}$$

It can be seen that the mixing angle in matter has a resonance for neutrinos at $A_{\mathrm{CC}} = \Delta m^2 \cos 2\theta$. Whether such a resonance exists depends on the sign of A_{CC} (i.e. whether we have a neutrino or antineutrino), and on the sign of the mass splitting. Experiments sensitive to matter effects are thus useful probes of

the neutrino mass spectrum. However, CP-violation also introduces differences in neutrino and antineutrino behaviour, and results from such experiments may need to be combined with other data in order to resolve ambiguities.

In the case of constant matter density along the neutrino path, we can simply replace the vacuum mixing angle and mass splitting with their matter counterparts to obtain an oscillation probability. In this two-flavour approximation, the relevant expression for a ν_e appearance experiment is

$$P_{\nu_\mu \to \nu_e} = \sin^2 2\theta_{\mathrm{M}} \sin^2 \left(1.27 \frac{\Delta m^2_{\mathrm{M}}(\mathrm{eV}) L(\mathrm{km})}{E(\mathrm{GeV})} \right). \tag{1.37}$$

This constant-density formula is sufficient for long-baseline experiments (though a 3-flavour analysis may be needed). Atmospheric and solar analyses, where the density profile changes along the neutrino path, require a more complicated treatment, because the effective mass basis changes with the density, leading to mixing of the effective mass states as well as the flavour states. This is a difficult problem to solve, and in general it cannot be done analytically.

1.2.5 Neutrino interactions with matter

In order to make quantitative measurements of oscillation phenomena, it is necessary to have some understanding of the interaction of neutrinos with matter. To calculate the flux and energy spectrum of neutrinos incident on a detector from the observed event spectrum, we must know the interaction cross-sections, as well as the kinematics, for the processes which can give the event topologies which will be observed. At the energies used for long-baseline experiments (a few hundred MeV to $\sim$10 GeV), the neutrino interaction cross-section is completely dominated

by interactions with nuclei, whose cross-sections cannot be calculated exactly from electroweak physics and must be determined experimentally.

Several underlying interaction types have been identified as contributing to the νN cross-section. Their contributions to the total cross-section as a function of energy are shown in Figure 1.2. The most important at low neutrino energies are quasi-elastic processes:

$$\nu_l n \rightarrow l^- p, \qquad \bar{\nu}_l p \rightarrow l^+ n, \tag{1.38}$$

$$\overset{(-)}{\nu}_l\ n \rightarrow \overset{(-)}{\nu}_l\ n, \qquad \overset{(-)}{\nu}_l\ p \rightarrow \overset{(-)}{\nu}_l\ p, \tag{1.39}$$

where (1.38) are mediated by the charged current (CC) interaction, and (1.39) by the neutral current (NC).

At slightly higher energies, resonance interactions become important, since sufficient energy can be transferred to the interacting nucleon for a baryonic resonance to be formed. The sharp peaks in cross-section corresponding to the resonance masses are smeared by the Fermi momentum of the interacting nucleon, and would be difficult to measure in any case due to the spread in energy of a realistic neutrino beam. Resonance states like the $\Delta(1232)$ decay to a nucleon and pion, and π^0s produced in this way form the most important background for the ν_e-CCQE analysis in T2K. All members of the pion triplet can be produced by resonance interactions:

$$\nu_l p \rightarrow l^- p\pi^+, \qquad \nu_l n \rightarrow l^- n\pi^+, \qquad \nu_l n \rightarrow l^- p\pi^0; \qquad (\mathrm{CC}) \tag{1.40}$$

$$\left.\begin{array}{ll} \nu_l p \rightarrow \nu_l p\pi^0, & \nu_l p \rightarrow \nu_l n\pi^+, \\ \nu_l n \rightarrow \nu_l n\pi^0, & \nu_l n \rightarrow \nu_l p\pi^- \end{array}\right\} \qquad (\mathrm{NC}) \tag{1.41}$$

At still higher energies, it is necessary to consider the quark constituents of the interacting nucleon, and the cross-section is dominated by deep-inelastic scattering (DIS), where the nucleon itself is broken up. Cross-section calculations for DIS processes rely on the structure functions describing the distribution of partons as a function of their momentum fraction, which must be derived from experiment. Since the nucleon itself is not left intact in the interaction, a variety of hadronic states, including heavy-flavoured mesons, can be produced. Because the kinematics of DIS interactions are rather different to CCQE, if they are misidentified as CCQE events then the reconstructed E_ν will be somewhat lower than the true value. High-energy tails in a neutrino beam can thus produce backgrounds to CCQE measurements at lower energies.

In addition to these dominant scattering channels, other interactions are also possible. It is possible for a neutrino to interact coherently with the nucleus via the neutral current; this process produces pions at an extremely forward angle, allowing the events, at least in principle, to be distinguished from resonant production. Diffractive production, in which a W or Z particle produces a vector meson in the nuclear field, can also occur. These smaller effects are modelled in the GENIE MC, which is used for generating neutrino interactions in the T2K near detectors [24].

It should be noted that for interactions with nucleons in nuclei, further intra-nuclear processes, including charge exchange, absorption and additional particle production, can occur while the neutrino interaction products are still inside the nucleus. These effects complicate the development of models which pre-

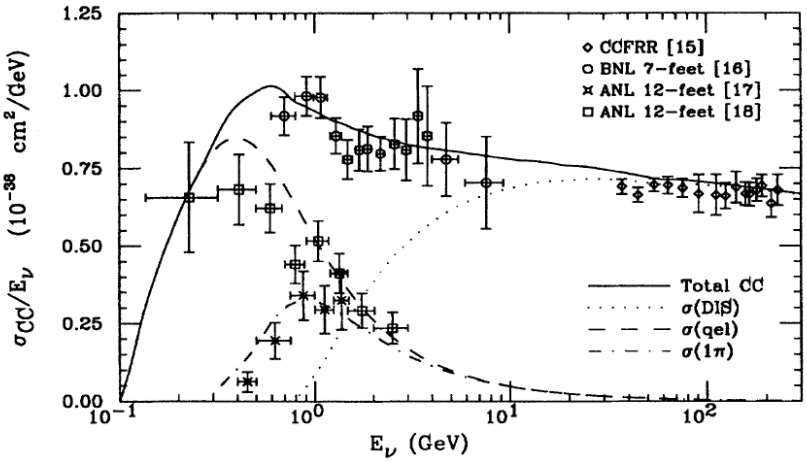

Figure 1.2: ν_μ CC cross-sections on nucleons as a function of energy [25]. The 1-π cross-section is dominated by resonance production. Note that the y-axis shows $\sigma_{\mathrm{CC}}/E_\nu$.

dict cross-sections on nuclei, since a re-interaction model must be used to relate these to the neutrino-nucleon cross-sections. Because experimental data contains only information on the final state, cleanly separating the parameterisation of the nucleon-level event and the tuning of the re-interaction model is difficult, though this can be helped by considering data from many nuclei.

In practice, experiments to measure neutrino-nucleus scattering are difficult, because production of neutrinos by conventional means, using the decays of hadrons in a secondary beamline, does not produce neutrinos of a well-defined energy, or in a tightly focused beam. This necessitates the use of detectors with large fiducial volumes, which consequently have relatively coarse granularity. Also, the vertex for an interaction must be reconstructed from the observed products, rather than being a known quantity which can itself aid reconstruction as in collider experiments. Because of these difficulties, neutrino oscillation experiments often

use configurations where systematics relating to neutrino cross-sections largely cancel — for example in T2K, CCQE event rates will be measured on water in the ND280 near detector so that the cross-section cancels when extrapolating to Super-K. However, as discussed in Chapter 5, the near and far detector fluxes are related non-trivially, and so the cancellation is not perfect. Also, the ability to discriminate between interaction types is different at the near and far detectors, and interactions on other materials are important in ND280. Because of these considerations, a precise understanding of interaction physics will still be an important contribution to the effort to reduce systematic errors.

1.3 Review of oscillation measurements

1.3.1 Present constraints on mixing parameters

A varied program of oscillation experiments has been undertaken around the world to constrain the PMNS and mass splitting parameters. The discussion here will focus on results which contribute to the current constraints on the parameter space. Solar-sector measurements will be briefly summarised, followed by a more detailed description of atmospheric-sector measurements, which are directly relevant to T2K.

1.3.2 Solar-sector oscillations

The parameters responsible for solar neutrino mixing, $\sin^2 2\theta_{12}$ and Δm^2, have been well measured. The current constraints on the parameter space come mainly

from the SNO experiment (see Section 1.1.2), Super-Kamiokande observations of solar neutrinos (Section 1.3.3) and KamLAND [26], a long-baseline reactor neutrino experiment located at Kamioka, which used a liquid scintillator-based detector to observe antineutrinos from several reactors around Japan.

About two-thirds of the neutrino flux in the spectral region observed by SNO was observed to have changed flavour from ν_e to ν_μ or ν_τ. If the MSW effect is included in the solar oscillation fit, then good agreement is seen between the solar data and the vacuum mixing-driven oscillations from the KamLAND experiment. The allowed region for solar mixing parameters is shown in Figure 1.3. This region is referred to as the "Large Mixing Angle" (LMA) solution.

1.3.3 Present constraints on ΔM^2-driven oscillations

Super-Kamiokande

Large numbers of neutrinos are produced in the Earth's upper atmosphere by the decays of particles produced by cosmic ray interactions. Evidence for the oscillation of these atmospheric neutrinos was first obtained by the Super-Kamiokande experiment in 1998 [27]. The Super-K detector, shown in Figure 1.4, was built mainly to search for proton decay, and consists of a tank holding 50 ktonne of water (22.5 ktonne fiducial mass), divided into an inner main detector and outer veto detector, with the dividing wall instrumented with inward- and outward-looking PMTs. Atmospheric neutrinos undergo charged-current interactions in the water, producing charged leptons, which are detected via the Čerenkov radiation they emit. Discrimination between electron-like and muon-like events is mainly via a

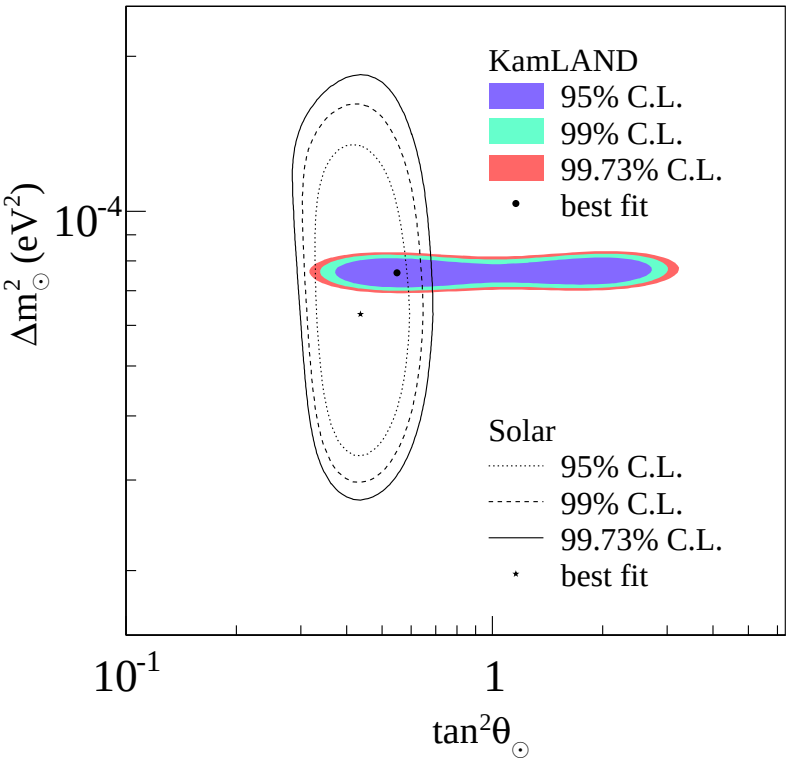

Figure 1.3: Limits on solar-sector mixing parameters from solar and KamLAND data [7].

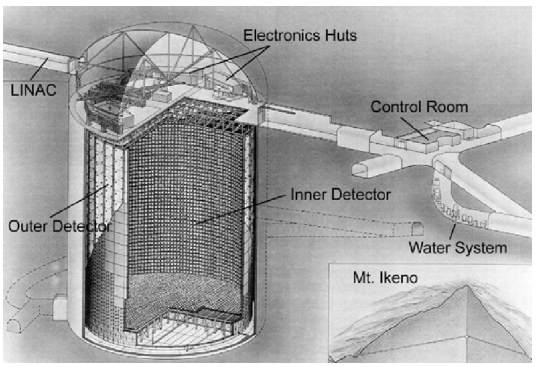

Figure 1.4: Drawing of the Super-K detector [28].

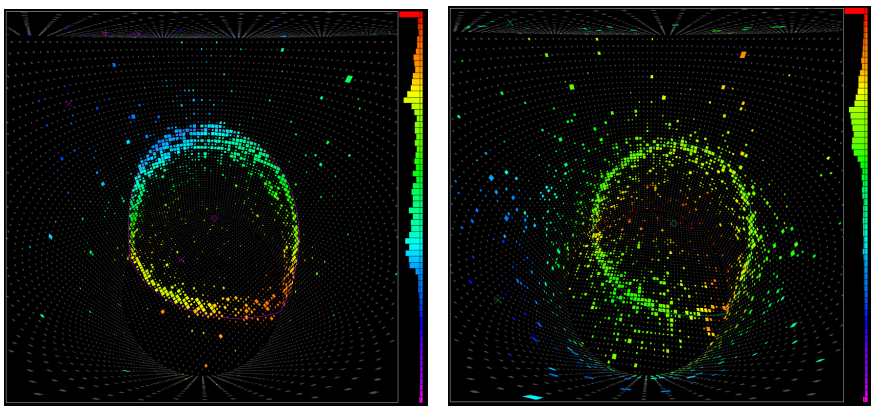

(a) (b)

Figure 1.5: Muon (a) and electron (b) events in Super-K. The colours indicate PMT hit times [29].

likelihood analysis based on the shape of the Čerenkov ring — electrons tend to produce a "fuzzy" ring because of multiple scattering. This difference is shown in Figure 1.5. The detector cannot distinguish positive from negative leptons, so observed events may have come from ν or $\bar{\nu}$.

Since Super-K is sensitive to both ν_e and ν_μ, the results enable an estimate to be made of the ratio of the $\nu_\mu + \bar{\nu}_\mu$ flux to the $\nu_e + \bar{\nu}_e$ flux. The expected value of this ratio depends on the details of cosmic ray production of neutrinos,

but it can be more precisely calculated than the absolute neutrino flux. Super-K measured a value for this ratio of about 0.65 times the expected value without oscillations.

Since the direction and energy of the incoming neutrino can be reconstructed from the kinematics of the daughter lepton, the ratio L/E can be estimated for each event. The muon neutrino flux was seen to have a deficit dependent on L/E, as expected by an oscillation hypothesis (1.25). No such structure was seen in the ν_e flux, so the results suggest that some of the ν_μs oscillate to ν_τ or some sterile flavour ν_s, and the ν_es do not participate in oscillations on this L/E scale. A detailed analysis of the Super-K dataset [30] favours $\nu_\mu \rightarrow \nu_\tau$ oscillations over $\nu_\mu \rightarrow \nu_s$; the two cases can be differentiated without direct identification of ν_τ events, by looking at the number of NC events in the detector, and by looking for distortions in the ν_μ survival probability arising from the different matter potential for ν_s compared to ν_μ and ν_τ, which couple to matter via the weak neutral current. A statistical study of the Super-K data [31], looking for hadronic τ decays, favours a ν_τ appearance signal compatible with the $\nu_\mu \rightarrow \nu_\tau$ hypothesis. The best-fit point for the Super-K data, assuming a standard 3-flavour oscillation model with matter effects and normal mass hierarchy [32], is ($\sin^2\theta_{13} = 0$, $\sin^2\theta_{23} = 0.5$, $\Delta M^2 = 2.5 \times 10^{-3}\,\mathrm{eV}^2$). The 90% C.L. intervals are $\sin^2\theta_{13} < 0.14$, $0.37 < \sin^2\theta_{23} < 0.65$.

K2K

The K2K experiment [33] was a long-baseline ν_μ disappearance experiment conducted in Japan. 12 GeV protons from the KEK-PS proton beam were used to produce pions by collision with an aluminium target, which were then focused by a set of magnetic horns and decayed to give neutrinos, in a similar manner to the T2K experiment. K2K used the Super-K detector, 250 km away, to measure the oscillated neutrino flux. A dedicated near detector was constructed and installed at KEK to measure the unoscillated flux; the expected far detector flux, without oscillations, was then calculated using a far/near ratio method (see Chapter 5). The neutrino beam flux was predicted using data from the HARP aluminium-target proton experiment at CERN [34], and the flux of parent pions was also measured in situ using a temporary pion monitor downstream of the target.

The K2K near detector comprised several different modules with complementary capabilities, arranged linearly along the beam axis. The detectors consisted of a water Čerenkov module to provide a high-statistics cross-section-independent estimate of the Super-K event rate, scintillating bar and fibre modules for precise tracking, and a downstream muon range detector. Events from the K2K beam were identified at Super-K by timing, using a GPS system. 112 events were observed which passed quality cuts, with 158.1 expected without oscillations. By fitting a 2-flavour oscillation hypothesis (1.25) to the observed energy spectrum, a best fit value of $(\sin^2(2\theta_{23}) = 1.0,\ \Delta M^2 = 2.8 \times 10^{-3}\,\mathrm{eV}^2)$ was obtained. These results are in good agreement with those obtained using Super-K atmospheric data. A search for ν_e appearance in the K2K beam was also made [35]; no

evidence for appearance was seen and a 90% C.L. limit ($\sin\theta_{13} < 0.13$) was set using the preferred value of ΔM^2.

MINOS

The MINOS experiment [36] is a long-baseline accelerator experiment built to measure ν_μ disappearance. 120 GeV protons from the Fermilab Main Injector beam are collided with a graphite target to produce neutrinos via pion decay. The neutrino beam is sampled at source by a near detector, and 735 km away by a far detector at Soudan mine in Minnesota. The near detector measures the unoscillated neutrino flux, and detailed constraints on hadron production in the target have also been obtained by moving the target with respect to the focusing system, and by changing the horn currents. The near and far detectors have different dimensions but share the same technology, so that some systematics cancel between them — both consist of layers of scintillator strips interleaved with layers of steel target material. The scintillation light is collected by wavelength-shifting fibres and read out by PMTs.

The neutrino beam is sampled on-axis by both detectors, and the neutrino energy spectrum at the detectors peaks around 3–4 GeV. In this energy region, DIS, QE and resonance processes (see section 1.2.5) all contribute significantly to the neutrino cross-section, and the inclusive charged-current interaction rate is used to determine the neutrino flux. CC and NC events are distinguished by separating the detector hits into a muon track and hadronic shower, and applying a likelihood analysis based on properties of these objects. ν_μ can be separated from events

due to $\bar{\nu}_\mu$ contamination by track curvature, since the detectors are magnetised. Neutrino energy is estimated from the muon track energy (by curvature or range), and the visible shower energy.

The measured near detector flux is extrapolated to the far detector using a beam matrix method similar to that described in Chapter 5, and the observed deviation from this prediction is used to determine the best-fit values for the mixing parameters θ_{23} and ΔM^2. The most recent published analysis [37] gives $|\Delta M^2| = (2.43 \pm 0.13) \times 10^{-3}\,\mathrm{eV}^2$ (68% C.L.), and $\sin^2(2\theta_{23}) > 0.90$ (90% C.L.).

CHOOZ

The parameters θ_{23} and ΔM^2 are now known with some precision; however the mixing angle θ_{13} is still unmeasured. The best upper limit on its value comes from the CHOOZ experiment in France [38]. CHOOZ attempted to measure electron antineutrino disappearance using neutrinos produced by two fission reactors at a power plant. The neutrinos had a mean energy around 3 MeV, and the average propagation distance was about 1 km, making the experiment most sensitive to ΔM^2 in the region of a few $\times 10^{-3}$ eV.

The neutrino flux from the reactors was almost entirely $\bar{\nu}_e$, and was calculated using the isotopic content of the reactor cores, their fission rates and the neutrino yield for each isotope — no near detector was employed. The detector consisted of an inner target region encased in plexiglass, a containment region around the target instrumented with inward-looking PMTs, and an outer veto region similarly instrumented. The detector was filled with scintillator-doped

paraffin, and the target region was loaded with gadolinium. The neutrinos were detected via the process $\bar{\nu}_e + p \rightarrow e^+ + n$ on the hydrogen nuclei in the paraffin, and the neutrino signature was a prompt scintillation signal from the positron, in coincidence with a delayed γ from neutron capture on the Gd. Additional spatial and energy cuts were also applied, and the neutrino energy was reconstructed from the size of the positron signal.

The neutrino flux measured by CHOOZ is compatible with a "no oscillation" hypothesis at 90% C.L. — the final limit for $\sin^2(2\theta_{13})$ is shown in Figure 1.6. The Palo Verde experiment, another reactor neutrino detector with a similar baseline, confirmed the CHOOZ result with a looser bound of $\sin^2(2\theta_{13}) < 0.17$ for large ΔM^2 [39] [40].

LSND, MiniBooNE and MicroBooNE

LSND [42] was a very-short baseline $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ accelerator experiment at Los Alamos, using a 798 MeV (kinetic energy) proton beam from the LANSCE accelerator. Pions were produced in a target (either water or a high-Z material), and stopped by a Cu block. Neutrinos were mostly produced by decay-at-rest, in the chain:

$$\pi^+ \rightarrow \mu^+ \nu_\mu;\ \mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu; \tag{1.42}$$

the charge conjugated processes were much less frequent since stopped π^- and μ^- are susceptible to nuclear capture. Producing the beam in this way virtually eliminated $\bar{\nu}_e$ contamination, limiting the irreducible background for the $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ search. Also, the neutrino spectrum was dominated by decay-at-rest events,

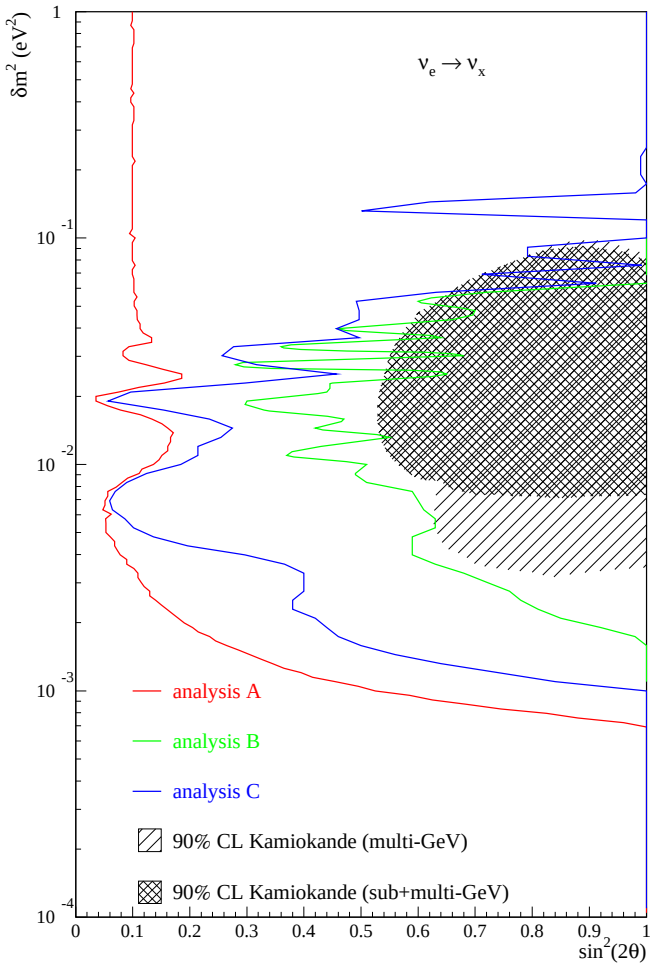

Figure 1.6: CHOOZ exclusion region for $\sin^2(2\theta_{13})$ [41]. Analysis A uses all experimental data. Analysis B uses the ratio of fluxes from the two reactors rather than the overall flux — it is therefore mostly independent of errors in the overall neutrino flux and detection efficiencies. Analysis C uses only the shapes of the neutrino spectra rather than the normalisations — the resulting limit is much less stringent.

making it easy to calculate.

The detector used was liquid scintillator instrumented with PMTs, with an outer veto tank, and placed with its centre 30 m from the neutrino source. $\bar{\nu}_e$ CC events were identified by the coincidence of an electron signal with a γ from neutron capture. The experiment observed an excess of such events, which can be attributed to $\bar{\nu_\mu} \rightarrow \bar{\nu}_e$ oscillations with probability (0.264$\pm$0.067$\pm$0.045)%. Because of the small L/E, this result favours oscillations with a mass splitting $\sim$1 eV, and is therefore incompatible with other neutrino oscillation results unless a more complex neutrino-sector model, including sterile neutrino flavours, is introduced.

The MiniBooNE experiment was built at Fermilab to directly corroborate or refute the LSND results by looking for oscillations on a similar L/E scale. MiniBooNE uses an 8 GeV proton beam, with a Be target and magnetic focusing system, to produce muon neutrinos. The detector is liquid scintillator. The first phase of MiniBooNE searched for $\nu_\mu \rightarrow \nu_e$ oscillations [43] — an excess of low-energy events was observed, but this excess is not well-accommodated by an oscillation hypothesis, and is incompatible with oscillations on the mass scale required by LSND, in the absence of CP-violation. MiniBooNE has now also taken data in antineutrino mode [44]; the initial results from this run do not show a positive oscillation signal, but do not yet exclude the LSND-allowed oscillation region.

The MicroBooNE experiment [45], currently moving into its construction phase, will use a $\sim$100 tonne liquid Argon TPC in the same beamline as MiniBooNE to study this region of the neutrino parameter space more precisely. This project

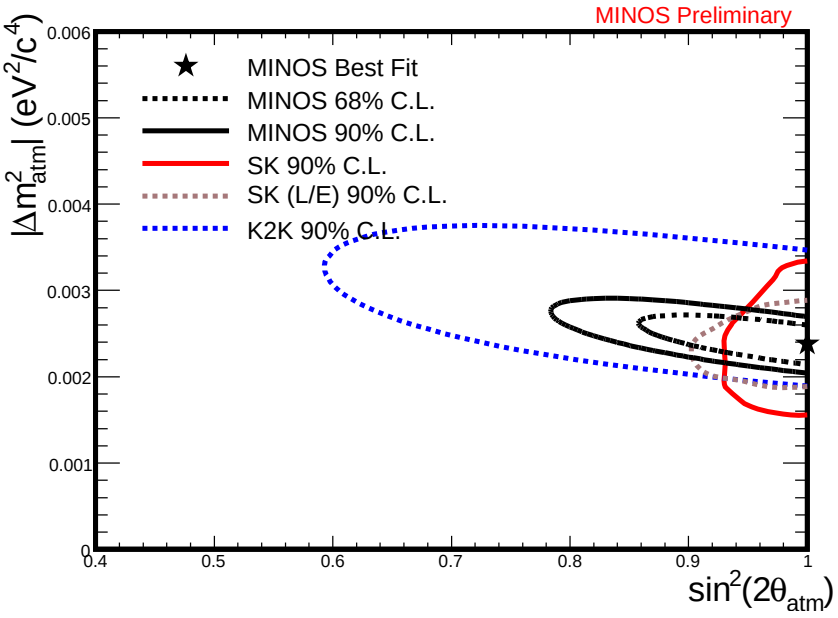

Figure 1.7: Allowed region in (ΔM^2, $\sin^2(2\theta_{23})$) space, from the MINOS, Super-K and K2K experiments [7].

also serves as an important R&D exercise for the development of much larger detectors of the same type.

Summary

The current constraints on the parameters ΔM^2 and $\sin^2(2\theta_{23})$ are shown in Figure 1.7. Maximal mixing is currently favoured by the data. MINOS has been able to put a rather better limit on the atmospheric mass splitting than Super-K, because the neutrino propagation distance is well known; for Super-K, it is a sensitive function of zenith angle, whose calculation has considerable uncertainties associated with it. As stated above, the LSND results are incompatible with this allowed region in the standard three-neutrino model.

1.3.4 Overall constraints on the PMNS matrix

Based on all the experimental results discussed, we can put limits on the individual elements of the PMNS matrix:

$$U_{\mathrm{PMNS}} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 1} & U_{\mu 1} \\ U_{\tau 1} & U_{\tau 1} & U_{\tau 1} \end{pmatrix} = \begin{pmatrix} 0.77\text{–}0.86 & 0.50\text{–}0.63 & 0.00\text{–}0.22 \\ 0.22\text{–}0.56 & 0.44\text{–}0.73 & 0.57\text{–}0.80 \\ 0.21\text{–}0.55 & 0.40\text{–}0.71 & 0.59\text{–}0.82 \end{pmatrix}, \quad (1.43)$$

where the ranges shown are 3σ limits [46].

1.4 Projects under construction

1.4.1 NOνA

NOνA [47] is a long-baseline neutrino oscillation experiment under construction in the United States, utilising the existing NuMI beamline used by the MINOS experiment, but with a new set of near and far detectors. Its main goals are to measure θ_{13} and determine the mass hierarchy (sign of ΔM^2). It uses a higher-energy configuration than MINOS (by altering the target position and horn currents), but the detectors are placed off the beam axis, giving a peak neutrino energy at the far detector of around 5 GeV. The off-axis configuration serves to substantially reduce the high-energy tail of the neutrino spectrum, which is a major source of background events. The baseline of the experiment is 810 km.

The NOνA near and far detectors are very similar, in order to minimise systematic uncertainties in extrapolating measurements between them. The far detector is a completely active design, made from plastic cells 15.5 m$\times$3.9 cm$\times$6.00 cm

doped with TiO_2 and filled with mineral oil-based liquid scintillator. Wavelength-shifting fibres are strung through the cells and read out using APDs. The cells are arranged in a crossed layer structure, and the total mass is 15 ktonne. The near detector also consists mostly of fully active cells, with an iron-liquid scintillator segmented region downstream for muon catching.

1.4.2 Daya Bay and Double CHOOZ

Daya Bay [48] is a short-baseline (i.e. sensitive to only ΔM^2-driven mixing) reactor $\bar{\nu}_e$ disappearance experiment in China. This experiment will use antineutrinos generated by six fission reactors, with a total power of 17.4 GW, to make a precision measurement of $\sin^2(2\theta_{13})$. In order to achieve the goal of sensitivity to $\sin^2(2\theta_{13}) > 0.01$, Daya Bay must achieve considerably lower systematic errors than the previous generation of reactor projects. In particular, this requires a measurement of the $\bar{\nu}_e$ flux more precise than that given by reactor power measurements, so both near and far detectors must be used.

All detector modules used by the experiment will be identical. The modules have a similar design to the CHOOZ detector, with an inner Gd-doped liquid scintillator region serving as the fiducial volume, surrounded by an undoped γ-catching region, and an outer veto. Modules will be surrounded by a PMT-instrumented water chamber with resistive plate chambers above, to act as a veto for muons and other radiation. The positions of the detector modules are chosen to be suitable distances from the reactors, while maintaining sufficient overburden to reduce cosmic radiation levels. There are two main blocks of reactors, with

two near detector modules situated close to each; the far detector will consist of four modules, approximately 2 km from each set of reactors. Since the modules are identical, results of similarly-situated modules can be compared to give consistency checks, and it is also planned to swap modules between locations.

Another reactor experiment, Double CHOOZ [49], is being prepared on the site of the CHOOZ project. The experiment has a similar far detector to CHOOZ, but with a larger mass and additional outer veto detectors to reduce cosmic background. A near detector will be added in order to reduce systematic errors from the neutrino flux. Double CHOOZ aims to achieve a sensitivity to $\sin^2\theta_{13} > 0.03$.

1.4.3 Summary

The estimated $\sin^2(2\theta_{13})$ sensitivities and timescales for the new generation of neutrino experiments are compared to that of the T2K project in Figure 1.8. The T2K proposal corresponds to 5 years of running at a beam power of 0.75 MW, though in practice this will take somewhat longer since the proton beam will take some time to ramp up to full intensity (see Section 2.2). This is reflected in its sensitivity curve.

Although all the experiments discussed are designed to measure θ_{13}, they are highly complementary because of their different configurations. Because Daya Bay and Double CHOOZ are $\bar{\nu}_e$ disappearance searches, they are insensitive to CP-violation, and to the size of θ_{23}. They are also insensitive to matter effects because of the low neutrino energy and short baseline. This means that a reactor

search alone cannot measure all the unknown neutrino-sector parameters, but it also means that it is able to make an unambiguous determination of θ_{13} which may be needed to resolve the degeneracies in the results of long-baseline searches which are sensitive to more effects.

T2K and NOνA will also sample different parts of the parameter space, since the neutrino energy peak is much higher at NOνA. This means, from (1.33), that matter effects are much more significant for NOνA. NOνA will attempt to determine the neutrino mass hierarchy by running with both neutrinos and antineutrinos, but the measurement will contain some ambiguity, since both CP-violation and matter effects produce an asymmetry between neutrino and antineutrino oscillations. In this case, it may not be possible to make a separate determination of the parameters using NOνA data alone, but a combination of the measurements with T2K, which is sensitive to CP-violation but has limited sensitivity to matter effects, may allow the ambiguities to be removed.

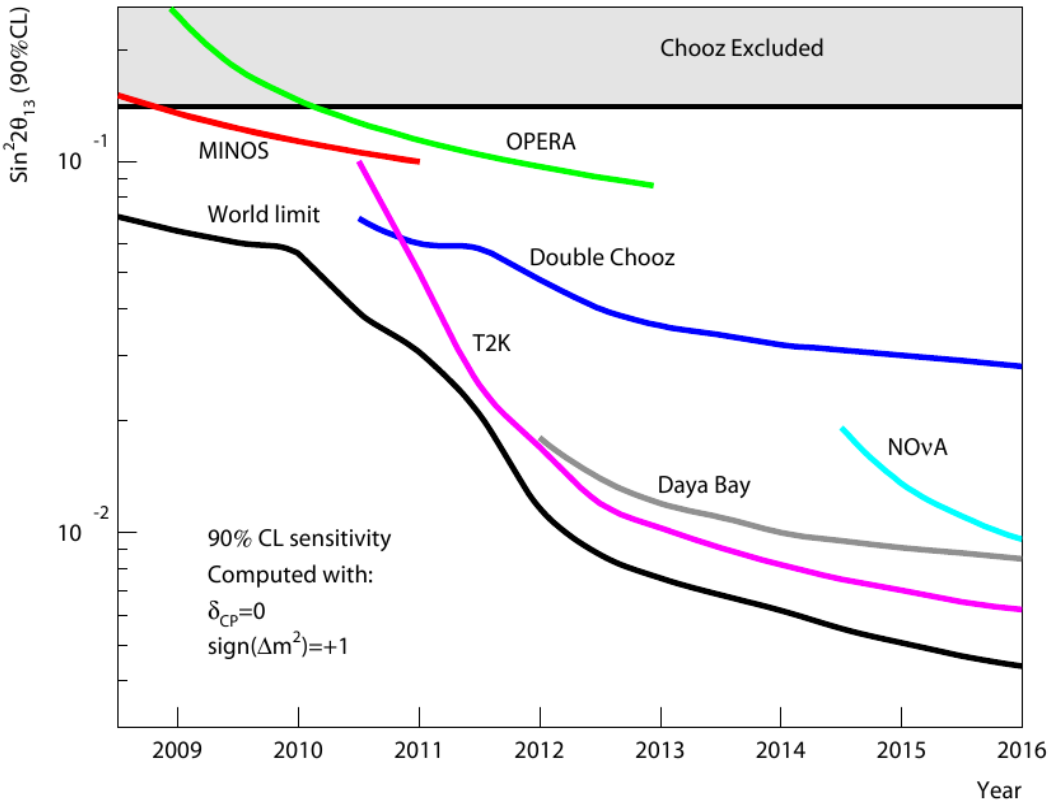

Figure 1.8: Comparison of projected $\sin^2 2\theta_{13}$ sensitivities and timelines, for T2K and its contemporaries. For T2K, a beam power of 0.1 MW in the first year, 0.75 MW from year 3, and a linear transition in between, are assumed. For NOνA, 6.5×10^{20} protons on target/yr is assumed [50].

Chapter 2

The T2K experiment

The T2K project [51] is a long-baseline neutrino oscillation experiment, searching for ν_μ disappearance and ν_e appearance in a beam of mainly ν_μ. The experiment is based in Japan, and uses beam from the 30 GeV proton synchrotron at the J-PARC facility in Tokai. The Super-Kamiokande water Čerenkov (see Section 1.3.3) is used as the far detector, and is complemented by a suite of new near detectors at Tokai. The experiment started taking data in November 2009 with a low-power beam and incomplete near detector, and commissioning of the full near detector is due to finish in September 2010.

The present chapter will begin with a discussion of the motivation and physics goals of the T2K experiment. The beamline and near detectors will be briefly described, along with a discussion of the modifications made to Super-K to accommodate T2K data taking. Finally, the physics reach of the experiment over its planned duration will be presented. The context of T2K within the broader field of neutrino physics, and possible extensions and successors to the experiment,

will be discussed in Chapter 6.

2.1 Motivation and physics goals

The primary goal of T2K is to measure the neutrino mixing parameter θ_{13}. As discussed in Chapter 1, previous measurements have only been able to place an upper limit on this parameter. Because all non-real terms in the PMNS matrix (1.23) contain a factor of $\sin\theta_{13}$, it must be non-zero for CP-violation in the lepton sector to be observed in flavour oscillations. Since leptonic CP-violation is a requirement of a leptogenesis-based explanation of the matter-antimatter asymmetry in the universe [52], a measurement of θ_{13} is important for cosmology as well as particle physics. The value of θ_{13} will also affect the decisions made regarding the next generation of neutrino projects, which will be built to measure δ (see Chapter 6).

T2K will also make a more precise measurement of the atmospheric mixing parameters θ_{23} and ΔM^2 than has been achieved by the MINOS and Super-K experiments. The Standard Model makes no predictions regarding the values of the mixing angles, but they are predicted by some GUT models [53], so precise measurement of θ_{13} and θ_{23} will put constraints on new physics beyond the Standard Model.

In addition to mixing measurements, the ND280 near detector will measure the cross-sections for various neutrino interaction processes on both plastic scintillator (carbon) and water. Measurements of the cross-sections for pion production processes are especially important, since these form significant backgrounds for the CCQE channels to be observed at Super-K in the oscillation searches. Specifi-

cally, CC-$\pi^{\pm}$ production (where either the pion or muon is undetected) and NC-$\pi^{\pm}$ production, can look like ν_{μ} CCQE. NC-π^{0} production where one γ is missed can imitate ν_{e} CCQE. Although the primary purpose of these cross-section measurements is to reduce systematic errors in the T2K oscillation searches, they will also have a broader value since the current data on neutrino cross-sections in the T2K energy range is extremely limited.

2.2 Primary beamline

The primary beamline for T2K is the proton accelerator complex at the J-PARC facility in Tokai, Japan. This facility is shown in Figure 2.1 - a pair of Linacs boost the protons to 600 MeV, and they are then fed into a 3 GeV synchrotron (RCS). This low-energy beam is used for several experiments, including feeding a spallation neutron source. Some of the 3 GeV protons are fed into the main synchrotron ring (MR), which has a design energy of 50 GeV, and is currently running at 30 GeV. The design beam power is about 0.75 MW.

The MR beam is shared between several experiments, including T2K. The nominal T2K run plan is for 120 days of beam per year for 5 years, with a total exposure of 5×10^{21} 50 GeV protons on the T2K target (p.o.t.), or equivalent. The beam intensity is gradually being ramped up — The projected beam output, as a function of time, is shown in Figure 2.2. In order to reach the design intensity, the magnet system will have to be improved in order to allow an increase in the spill repetition rate, and the beam power delivered by the RCS must also increase. The improvements in power and stability from gaining experience with running the

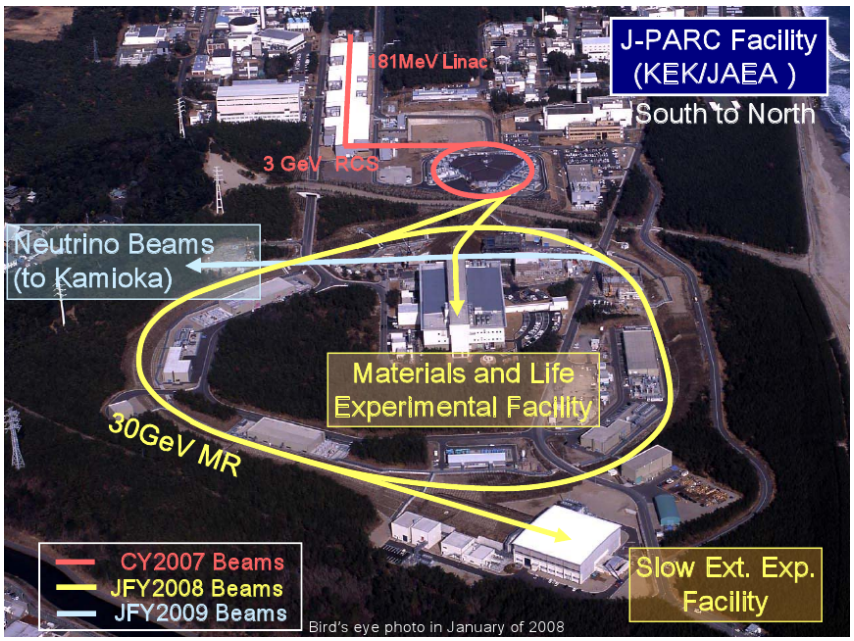

Figure 2.1: A bird's-eye view of the J-PARC accelerator complex, with the main accelerator components identified.

accelerator system will also play an important role in reaching this goal.

After the proton beam has been extracted from the main synchrotron, it passes through another section of beamline, which contains monitors to measure the beam's phase space profile and intensity using electromagnetic and optical effects. The beam then impacts on a fixed 900 mm graphite target. The bunch structure of the proton beam at the target is shown in Figure 2.3. Each beam spill is split into 8 bunches, with a repetition rate of around 0.3–0.5 Hz which will be increased as the beamline ramps up. In the first beam run (to summer 2010), only the first 6 bunches were present, since the kicker magnets required for fast extraction of the beam from the MR did not have a sufficiently fast rise time to extract from a fully occupied ring. These have now been upgraded.

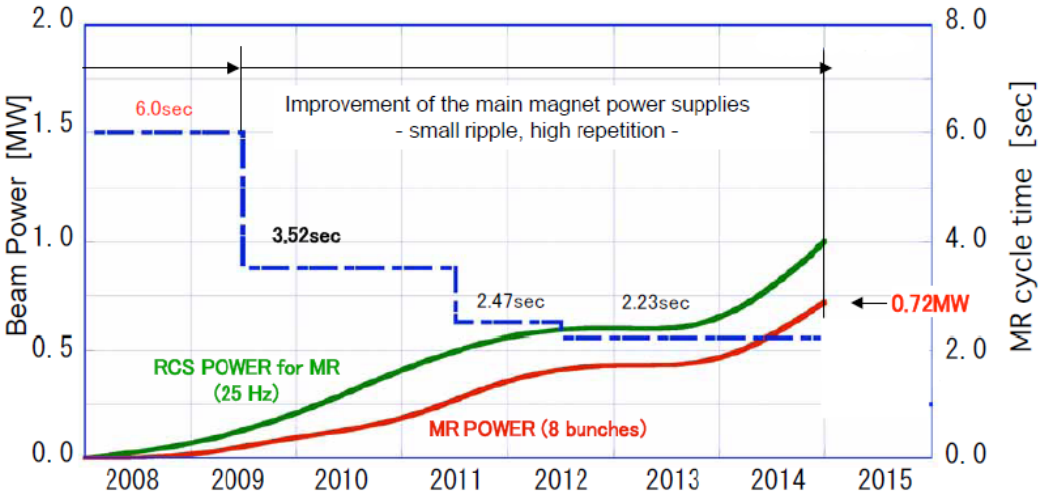

Figure 2.2: Planned ramp-up of the power supplied by the MR for the T2K beam. This depends on the input power from the RCS and the MR cycle time (time between beam spills), which are also shown [54].

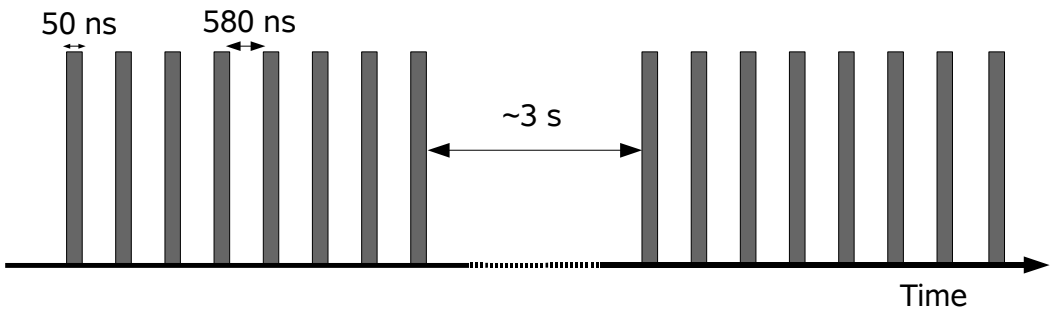

Figure 2.3: Illustration of the bunch structure for the T2K beam. Each beam spill consists of 8 bunches.

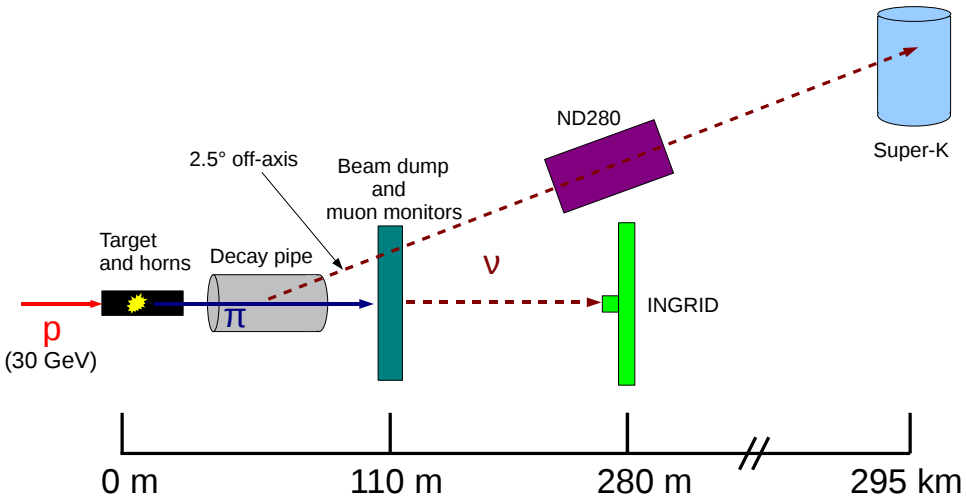

Figure 2.4: A simplified schematic of the T2K experiment.

2.3 Secondary beamline and detector configuration

A simplified schematic of the T2K configuration, showing the locations of the secondary beamline components and detectors, is shown in Figure 2.4. A more detailed schematic, showing the target and horn assembly, is shown in Figure 2.5. The proton beam interacts with the target, after being collimated by a baffle upstream to remove the beam halo and prevent damage to the target station. The proton interactions in the target produce large numbers of pions, and also some kaons. These are focused by a series of three magnetic horns, and then pass into a decay pipe. Neutrinos are produced in the decay pipe mainly by the following

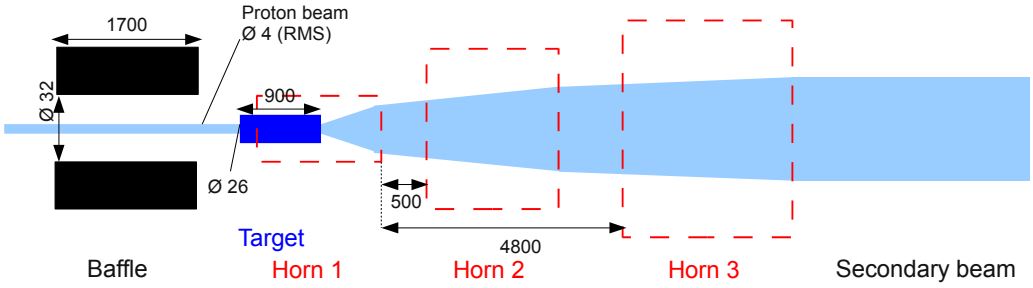

Figure 2.5: Diagram showing the proton target and horns for the T2K beam. The diagram is not to scale but the dimensions shown are accurate. All dimensions are in mm.

processes:

$$\pi^+ \rightarrow \mu^+ \nu_\mu \tag{2.1}$$

$$K^+ \rightarrow \mu^+ \nu_\mu \tag{2.2}$$

$$K^+ \rightarrow e^+ \nu_e \pi^0 \tag{2.3}$$

$$K^0_L \rightarrow e^+ \nu_e \pi^- \tag{2.4}$$

$$\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu \tag{2.5}$$

where the muons in (2.5) come from primary pion or kaon decays. The charge conjugate processes also occur, but in normal running the horn current focuses positive particles, so few negative hadrons are present in the beam. It is possible to reverse the horn polarity and produce a beam of mainly $\bar{\nu}_\mu$, but since more π^+ than π^- are produced in the target, the beam flux obtained will be lower. A beam dump downstream of the decay pipe absorbs any remaining hadrons and muons.

The beamline geometry is such that Super-K and ND280, the main near detector, are 2–2.5° from the beam axis. At this angle, the mean neutrino energy is lower than on-axis, putting the mean L/E of the beam at Super-K near the

first oscillation maximum. The off-axis beam flux is also much shorter-tailed than on-axis, reducing the rate of high-energy non-CCQE events which may be mis-reconstructed as CCQE events at lower energy. Since many of these high-energy events may be reconstructed as lying in the first oscillation peak, they form an important background to the ν_e appearance signal, and reducing the number of these events is the main motivation for using an off-axis beam. The Monte Carlo fluxes for 10^{21} p.o.t. are shown in Figure 2.6(a) for ND280, with the Super-K ν_μ flux shown for comparison.

In addition to the off-axis detectors, there is an on-axis neutrino detector (INGRID) in the same pit as the ND280 detector. A muon monitor in the beam dump also provides a cross-check of the beam flux by comparing the number of surviving muons to that expected for the observed number of neutrinos at INGRID. The expected on-axis neutrino flux is shown in Figure 2.6(b).

2.4 The ND280 detector

The ND280 detector [55] is a general-purpose particle detector which is placed 280 m downstream of the T2K proton target. Its main goals are to measure the neutrino flux at a similar off-axis angle to Super-K, and to determine the event rates for various background processes relative to CCQE. These measurements are necessary to control the systematic errors on the oscillation measurements at Super-K.

The detector, shown in Figure 2.7, consists of several modules. The central (basket) region contains a detector optimised for π^0 detection (P0D), and a

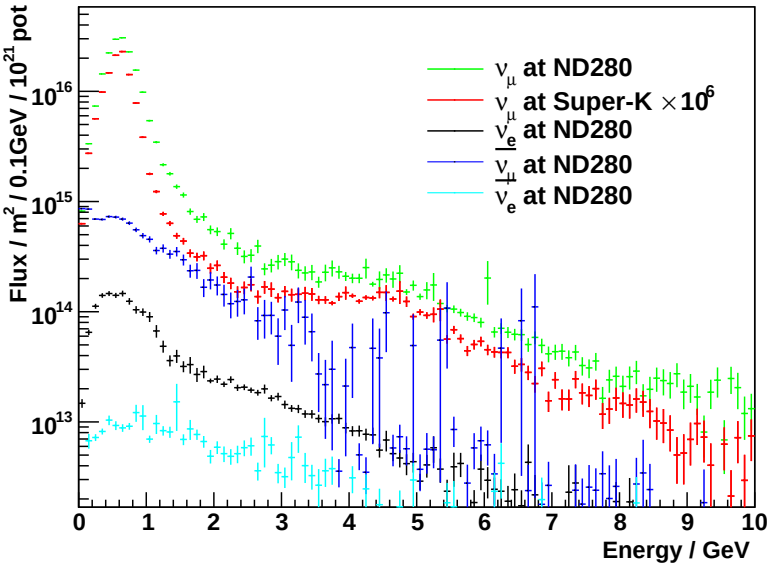

(a)

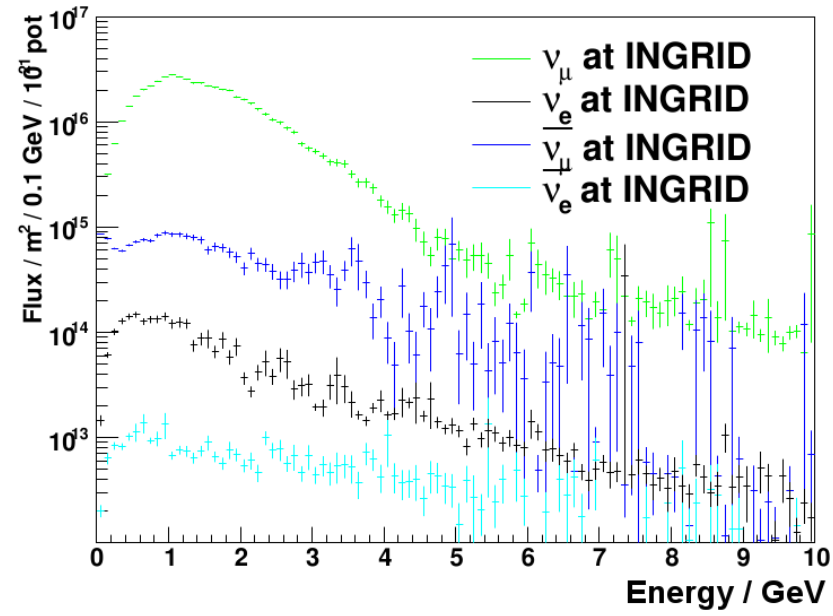

(b)

Figure 2.6: Monte Carlo neutrino fluxes at ND280 and Super-K (a), and at INGRID (b).

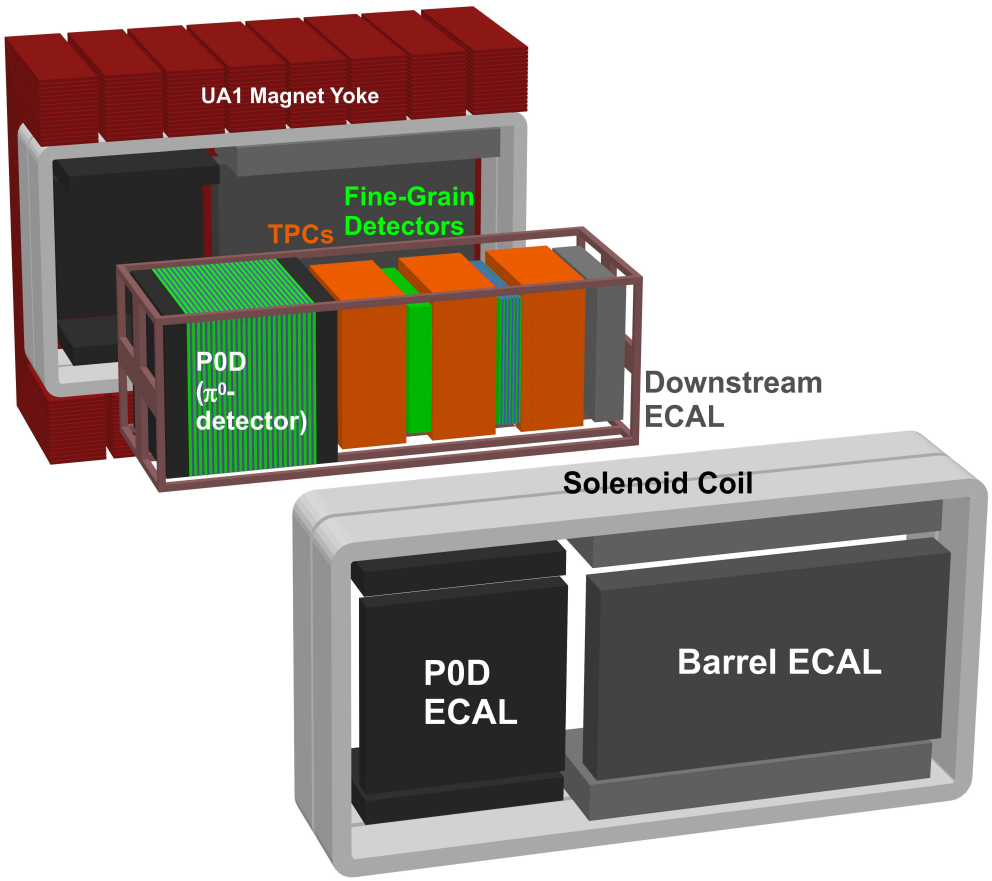

Figure 2.7: The ND280 detector.

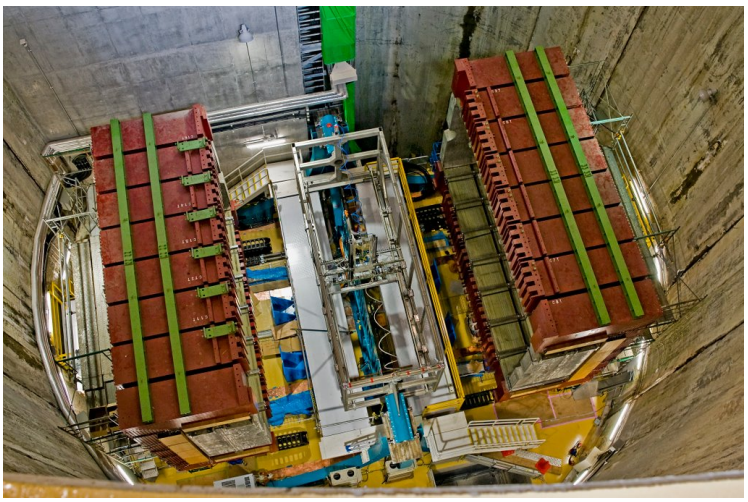

Figure 2.8: The ND280 pit containing the open UA1 magnet.

general-purpose tracker consisting of three time projection chambers (TPCs) separated by two fine-grained scintillator detectors (FGDs). Electromagnetic calorimeters (ECals) surround the basket on all sides except the upstream edge, and the P0D also contains submodules optimised for calorimetry. The detector is enclosed in a magnet, recycled from the UA1 experiment at CERN and providing a field of 0.2 T with a current of 2.9 kA. Side muon range detectors (SMRDs) are built into the magnet yokes in order to measure the energy of outgoing muons and trigger on incoming cosmics and magnet interactions. A view of the ND280 pit, with the magnet open before the installation of detectors, is shown in Figure 2.8.

All of the modules apart from the TPCs are scintillator-based. The active detector volumes are made up of bars of plastic scintillator extruded at Fermilab, coated with reflective TiO_2. The predominantly blue scintillator light is collected using Y11 wavelength-shifting fibres[1], which re-emit the light in the green part

[1]The specific fibre type used is Y11(200) S-35, 1 mm diameter and double-clad, with a WLS dye concentration of 200 ppm. S-35 fibres are optimised somewhat for flexibility at the expense

of the spectrum. The light is coupled into silicon Geiger-mode APDs (MPPCs manufactured by Hamamatsu — see Chapters 3 and 4) and these are read out using custom front-end boards. In all subdetectors apart from the SMRD, the scintillator bars are co-extruded with central holes to hold the wavelength-shifting fibre. All of the Y11 fibres, except for some ECal and all SMRD components, are read out at one end only, and mirrored at the other end.

2.4.1 The P0D

The ND280 Pi-0 detector sits upstream of the other detector components. It is designed to detect the γs from π^0s produced by ν_μ-NC interactions, and hence measure the cross-section for NC-π^0 production. It consists of three sections or "Super P0Dules" — a main fiducial volume in the centre, with modules optimised for calorimetry upstream and downstream (Figure 2.9). All modules are instrumented with triangular scintillator bars (33 mm base, 17 mm height, 2.14 m long), with two perpendicular layers glued together to give both X and Y information. The central module has cavities between every second X-Y block, which can be filled with water, and also thin lead sheets between the X and Y layers to convert photons, as in Figure 2.10. The calorimeter modules contain no water cavities, and have thicker lead to ensure EM shower containment. The mass of the fiducial region of the detector is 5.1 tonnes, of which 2 tonnes is water. A photograph of the P0D is shown in Figure 2.11.

π^0s in the P0D are identified by finding a pair of showers which extrapolate

of light yield. The fibres are manufactured by Kuraray Co., Ote Center Building,1-1-3, Otemachi, Chiyoda-ku, Tokyo 100-8115, Japan.

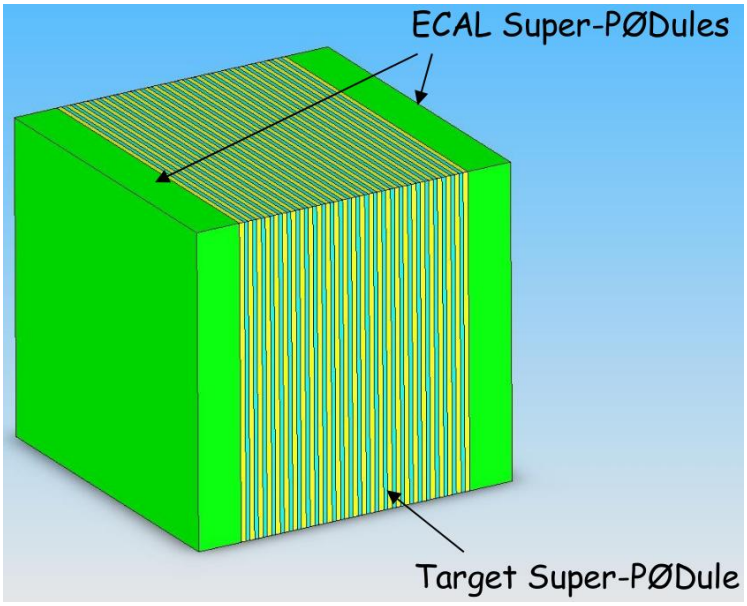

Figure 2.9: Schematic showing the structure of the P0D detector. The interleaving of active detector layers with water volumes in the central module is shown.

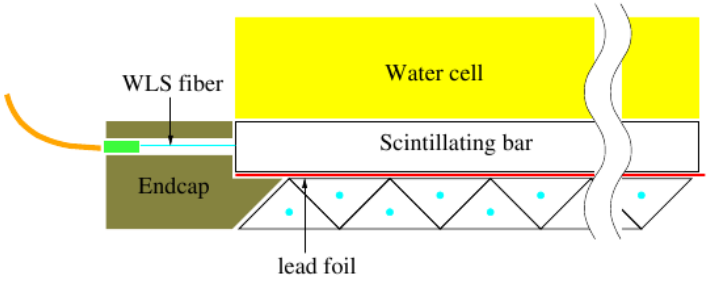

Figure 2.10: Schematic of a single P0D X-Y layer, showing the lead plane and water region.

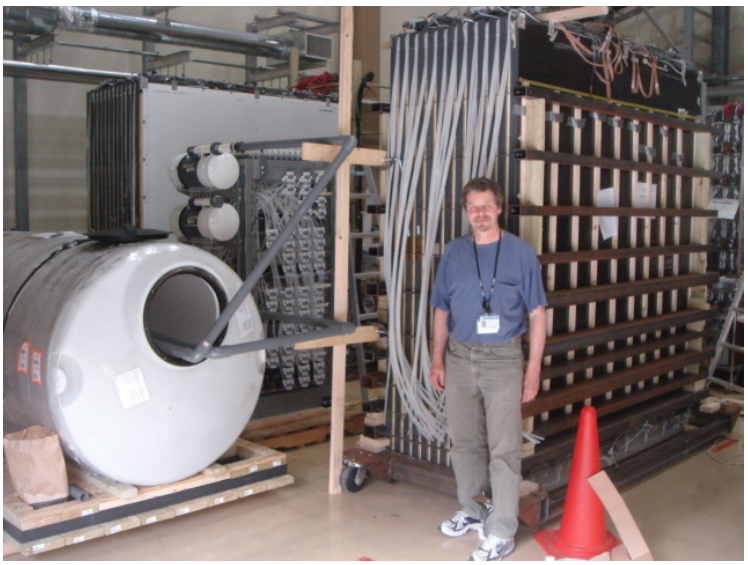

Figure 2.11: The P0D detector, open at the side so that the water system is visible.

back to a common vertex, and have an invariant mass close to the pion mass. This analysis relies on good tracking and energy resolution. The cross-section on water, which is important to reduce systematics on the ν_e appearance measurement at Super-K, is extracted by comparing the rates with the water cavities full with the rate when they are empty.

2.4.2 The Tracker

The main tracker region of ND280, shown in schematic in Figure 2.12, is downstream of the P0D, and is responsible for measuring the ν_μ and ν_e fluxes and energy spectra via the respective CCQE event rates, as well as the cross-sections for charged pion production. In order to distinguish these channels efficiently, and reconstruct neutrino momenta precisely, the detector must have good tracking and

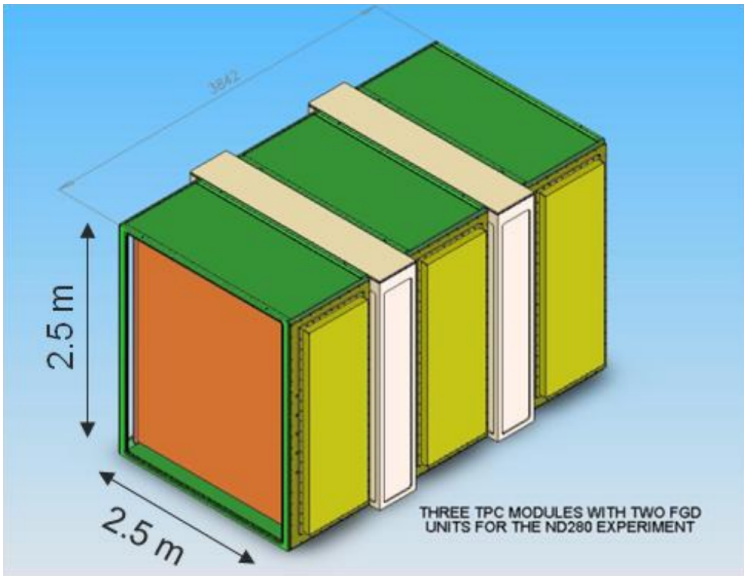

Figure 2.12: Diagram showing the structure of the ND280 main tracking region.

momentum measurement capabilities, while having a large enough mass to collect sufficient statistics.

The target mass is provided by the FGDs, which are made up of crossed X-Y layers of scintillator bars with a cross-section of $10 \times 10\,\mathrm{mm}^2$, and a length of 184 cm. The fine granularity of the FGDs allows the separation of the short-range recoiling hadronic system from the lepton in CC interactions. One of the FGDs is made up entirely of scintillator and has 15 X-Y layers; the other has 7 X-Y layers, interspersed with 6 water targets. The difference between the results obtained in the two FGDs allows to extract cross-sections on water. Each FGD has a mass of around 1 tonne; the FGD mass is somewhat lower than the P0D and tensions the need to collect sufficient statistics with the requirement that the leptons from FGD interactions escape into a TPC before depositing much energy, to allow momentum measurement. A photograph of one of the FGDs is shown in

Figure 2.13: Photograph of an FGD in its shipping jig.

Figure 2.13.

Most of the tracker volume is occupied by three Time Projection Chambers, which provide fully three-dimensional particle tracking, as well as momentum measurement from track curvature. The TPCs are shown in schematic in Figure 2.14, and an in-assembly photograph is shown in Figure 2.15. They are mainly filled with gaseous argon, and work by using a highly uniform electric field, produced by applying a voltage across the gas volume, to drift ionisation electrons and positive ions, produced by charged particles, towards Micromegas readout planes. These planes consist of a micromesh separated from a segmented anode by $\sim 100\,\mu$m, with a large electric field in the gap which allows avalanche amplification of the ionisation charge. The size of the charge deposit allows electrons and muons to be distinguished via their different specific energy deposition. The position of the

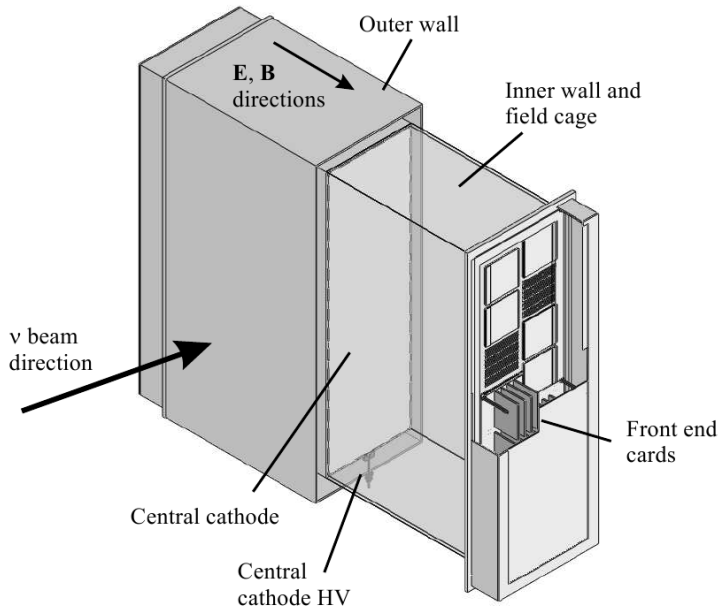

Figure 2.14: Schematic diagram of the TPC, showing the alignment of the readout planes (side with electronics boards) with respect to the beam and magnet field, and the central cathode dividing the TPC into two chambers.

anode segment hit provides two coordinates of the particle location. The position along the drift direction is calculated from the electron drift time, which requires matching TPC tracks to timestamped hits in the FGDs or ECals in order to get the electron production time.

2.4.3 The SMRD

The Side Muon Range Detector consists of many elements distributed in the air gaps of the magnet yoke, providing complete coverage of the sides of the detector. It is able to measure the energy of muons leaving the detector via the number of SMRD layers crossed, since the SMRD layers are interspersed with sections of iron magnet, where significant ionisation loss occurs. This is important, since muons

Figure 2.15: Photograph of a TPC during assembly.

produced at large angles to the beam are more likely to escape and not have their momenta measured by the TPCs — if their momenta coincide closely with the magnetic field direction their momentum cannot be measured by track curvature in any case. The SMRD is also able to tag and veto incoming particles from neutrino interactions in the magnet, and cosmic ray muons.

The SMRD elements are scintillator panels $0.7 \times 17.5 \times 87\,\mathrm{cm}^3$ in dimension. A wavelength-shifting fibre sits in an S-shape groove in the panel (Figure 2.16), since the panel is too wide for a straight fibre to collect sufficient light from tracks at all points. The fibres are $255\,$cm long and are read out at both ends, both to ensure a sufficient signal size and to give some position information. The regions of the magnet yoke to the left and right of the inner detectors are instrumented with six SMRD layers in the regions where the most muons from fiducial events will penetrate to the yoke. Fewer layers are used at the top and bottom where the

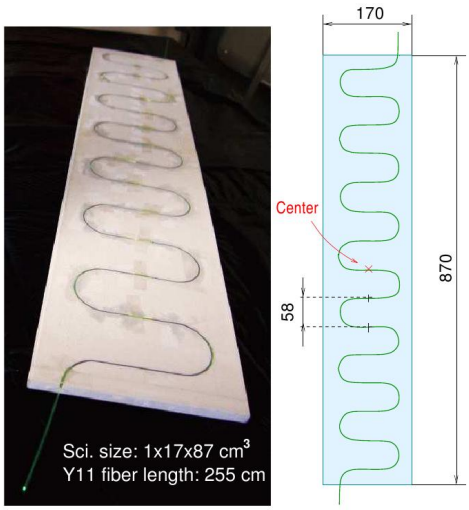

Figure 2.16: Photograph and diagram of an SMRD bar, showing the S-shaped groove carrying the wavelength-shifting fibre.

event rate is reduced by the magnet coils, and in the upstream part of the detector where fewer fiducial muons will hit.

2.4.4 The ECal

The Electromagnetic Calorimeter exists to enable the complete reconstruction of fiducial events which are not fully contained in the P0D and tracker. It also functions as a veto for events occurring outside the detector's fiducial volume. As the name suggests, the detector is optimised for calorimetry and is designed to fully contain all secondary particles from fiducial events except muons. The escaping particles will be mainly photons and electrons, so the sampling scale of the detectors is set by the electromagnetic shower scale.

The scintillator bars used for the calorimeter have a cross-section of $4 \times 1\,\mathrm{cm}^2$. The bars range in length from 152 cm to 384 cm depending on module and align-

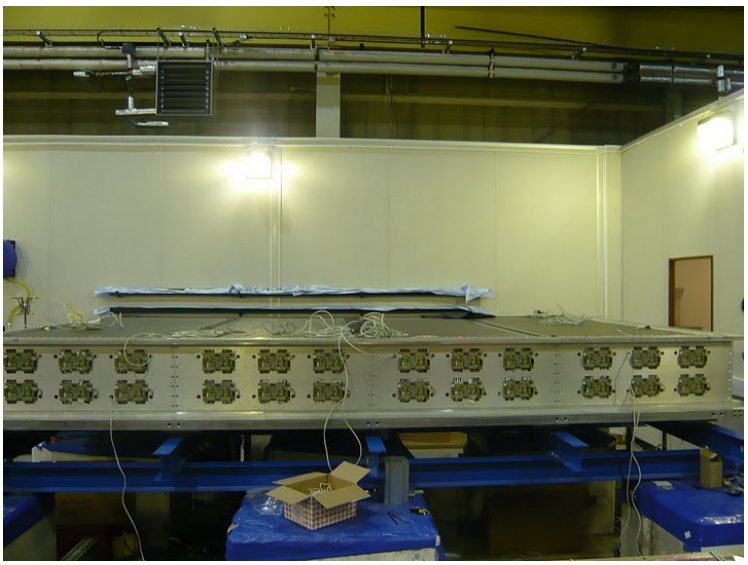

Figure 2.17: Photograph of a tracker ECal module during construction. The front-end electronics boards can be seen mounted on the module side.

ment; the longer bars in the tracker ECal, and all bars in the downstream ECal, are read out at both ends, while all other bars are read out at one end and mirrored on the other. The downstream and tracker ECals have 34 and 32 layers respectively, with alternate layers crossed, and interleaved with 1.75 mm lead sheets — these configurations are equivalent to 11 and 10.5 radiation lengths respectively. The P0D ECal has just 6 layers, all pointed along the beam direction, and 4 mm of lead between each layer, equivalent to 3.6 radiation lengths. A photograph of a tracker ECal module is shown in Figure 2.17.

The tracker and downstream ECals aid in the reconstruction of both ν_μ and ν_e CCQE events. High-angle muons not seen in the TPCs can be positively identified using the ECal — this complements the SMRD, which gives a good energy measurement, by allowing more precise tracking and unambiguous particle identification. The TPC measurement of electron energies is quite poor at energies

above about 1 GeV, so the ECal is needed to reconstruct high-energy ν_e CCQE events. These ECals will also identify gammas from π^0 events, which are likely to be undetected by the FGDs and TPCs, allowing a measurement of the NC-π^0 cross-section independent of the P0D, and also reducing the background on the ND280 CCQE sample.

The P0D ECal, with fewer sampling layers and 2D-only readout, will provide a somewhat cruder measurement than the other modules. It will be able to tag muons and electrons escaping the P0D, reducing the background rate for the main NC-π^0 search. It will also tag gammas which are not converted in the P0D — the ECal measurement will not be sufficiently clean to use these events in the NC-π^0 sample, but the escaping gamma rate will provide a useful check on the efficiency of the P0D itself. For gammas which convert in the P0D but are not fully contained, the P0D ECal will be able to provide calorimetry for the shower tail.

In addition to these roles, neutrino interactions in the ECals themselves will provide an independent measurement of the off-axis neutrino flux profile, since the large separation of the left/right and top/bottom ECal modules provides a lever arm to measure the flux as a function of direction. Since the ECals make up a large part of the active mass of ND280, the statistics of this measurement will be large.

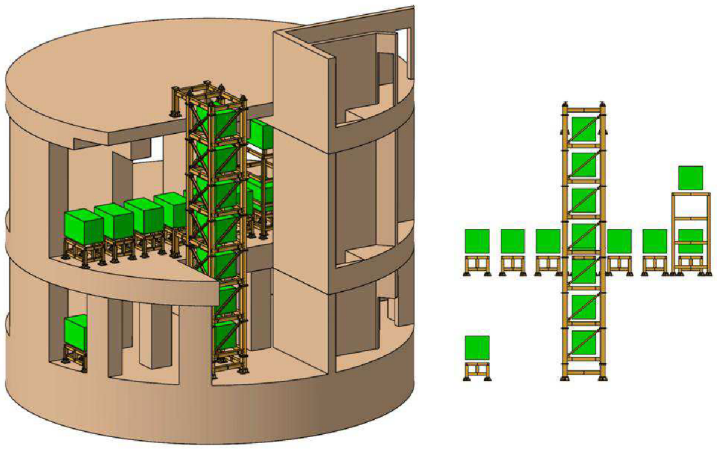

Figure 2.18: The INGRID detector.

2.5 The INGRID

The "Interactive Neutrino GRID" (INGRID) detector is designed to monitor the on-axis neutrino flux, and the precise direction and spread of the beam. The detector, shown in Figure 2.18, consists of 16 identical modules, with 7+7 arranged in a horizontal and vertical cross formation centred on the nominal beam axis, and two more in diagonal positions. Each axis of the cross spans 10 m in length. A photograph showing the support structure containing the vertical modules is shown in Figure 2.19.

Each INGRID module (Figure 2.20) consists of eleven crossed X-Y layers of scintillator bars, with 10 cm thick iron plates in between X-Y blocks. The module area perpendicular to the beam direction is 1.2 m^2. Each X or Y plane consists of 24 scintillator bars with a cross-section of $5 \times 1\,cm^2$. Like the ND280 scintillator detectors, the INGRID bars are read out using Y11 fibre and MPPCs.

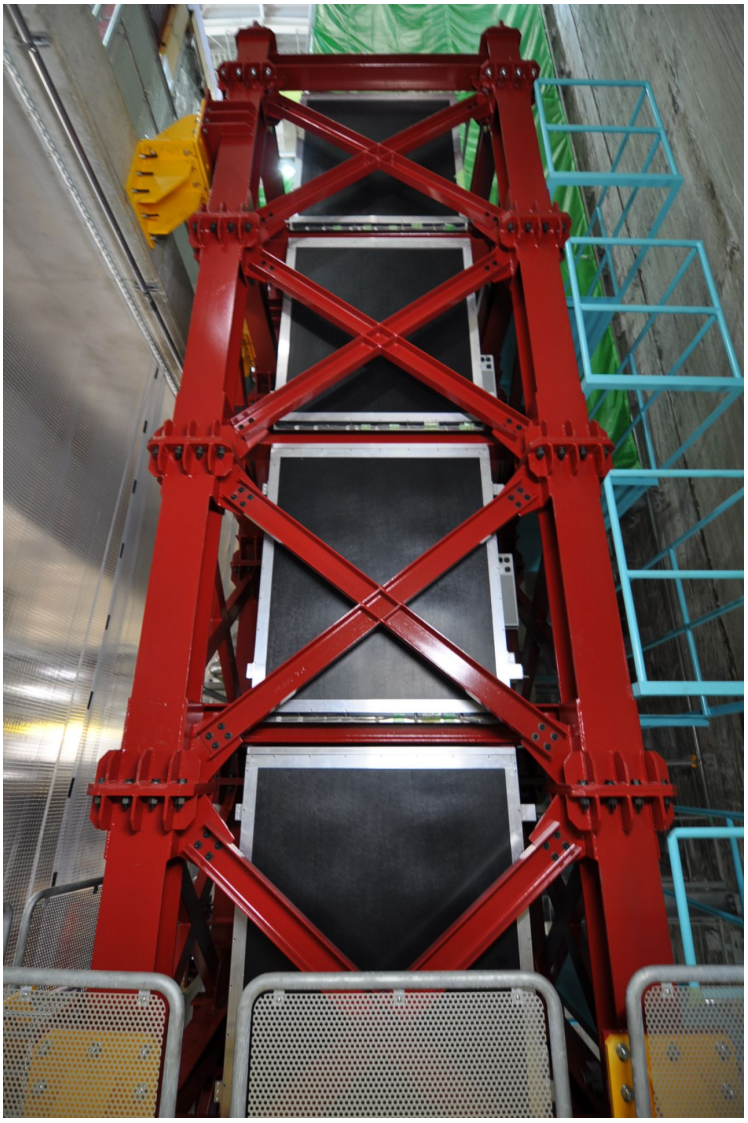

Figure 2.19: Photograph showing the vertical INGRID module stack.

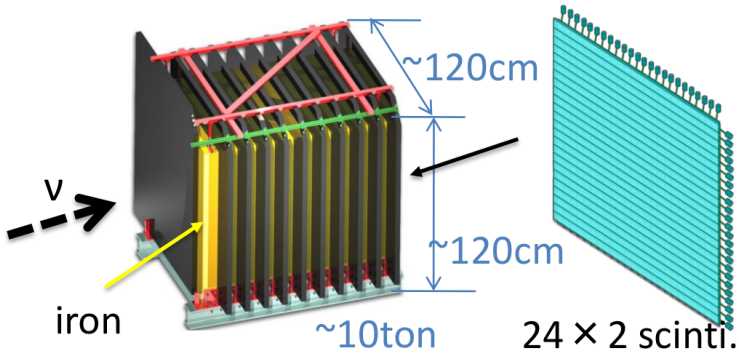

Figure 2.20: Diagram showing a single INGRID module, and a single X-Y plane with the MPPC photosensors attached.

Veto scintillator planes surround the main "tracker" volume of the module, to remove backgrounds due to interactions outside the module.

The simple design and fairly sparse sampling of the INGRID modules means that they will be used essentially as counters, but the large iron target mass will allow high statistics to be collected, and the configuration of the modules provides a large lever-arm to measure the beam direction and spread. Sufficient statistics will be collected each day to constrain the beam position with an error less than 1 mrad. These measurements are essential to drive down systematics related to the beam flux.

2.6 Super-K

The Super-Kamiokande detector has been described in Section 1.3.3; we will restrict ourselves here to a description of the new features of Super-K relevant to the T2K project. The front-end electronics have been completely replaced for the newest data-taking period, both for T2K data taking and to improve the quality of

the atmospheric data still being taken. The new electronics have a larger dynamic range, and, crucially, are free of dead-time, allowing T2K to maximise its statistics and catch stopping muon events with decay electrons in full, reducing the ν_μ-CC background in the ν_e appearance search. [56]. GPS synchronisation, using timestamps from J-PARC sent to Super-K, will be used to associate Super-K events with the T2K beam.

Because of the importance of background reduction in the appearance channel, a new algorithm has been developed to aid rejection of π^0 events. This involves fitting a two-ring hypothesis to events which have been reconstructed as single-ring, and then calculating an invariant mass, based on the momenta of the rings assuming they are photons. Events with a two-ring invariant mass near the pion mass are rejected. This cut has been seen to significantly improve π^0 rejection in Monte Carlo studies.

2.7 Analysis and physics reach

The T2K θ_{13} and θ_{23} measurements are based on the observation of CCQE ν_e and ν_μ events at Super-K respectively. As discussed above, the near detectors help to constrain the beam parameters (and hence the neutrino flux), and measure cross-sections for background processes, reducing systematics in the Super-K measurement. The presence of water in the ND280 detector provides good cancellation of cross-section uncertainties when predicting the far detector spectrum.

2.7.1 ν_e appearance

The main ν_e appearance analysis will use events reconstructed as ν_e-CCQE (single-ring electron-like events) at Super-K to estimate the values of $\sin^2 2\theta_{13}$ and ΔM^2, though in practice the latter will be better determined by the ν_μ analysis. The estimated ν_μ flux at Super-K without oscillation, multiplied by the oscillation probability (1.26), will be fitted to the ν_e flux estimated from the observed ν_e-CCQE event spectrum (with estimated backgrounds from ν_μ and beam ν_e subtracted), to give the best-fit values for the oscillation parameters. Systematic errors are associated with the neutrino fluxes, and with the measured event rates (due to detector efficiencies, cross-sections etc.) — these will be included in the fit.

The expected signal and background spectra passing Super-K analysis cuts for the ν_e appearance search, for $\sin^2 2\theta_{13}$ =0.1 and 0.01, are shown in Figure 2.21. The values of the other parameters used are ($\sin^2 2\theta_{12} = 0.8703$, $\sin^2 2\theta_{23} = 1.0$, $\Delta m^2 = 7.6 \times 10^{-5}\,\mathrm{eV}^2$, $\Delta M^2 = 2.4 \times 10^{-3}\,\mathrm{eV}^2$, $\delta = 0$). The normal mass hierarchy is assumed.

The sensitivity of T2K to θ_{13} depends on the integrated beam flux delivered by J-PARC, and on the total systematic error level of the ν_e appearance search. It also depends on the true values of ΔM^2, $\sin^2\theta_{23}$ and δ. Plots showing the 90% CL exclusion sensitivity to $\sin^2 2\theta_{13}$ are given in Figure 2.22, as a function of θ_{23} and ΔM^2, for 5 years running at 0.75 MW. All parameter values are the same as those used above except where stated. The T2K sensitivity as a function of exposure, compared with other projects, was shown in Figure 1.8.

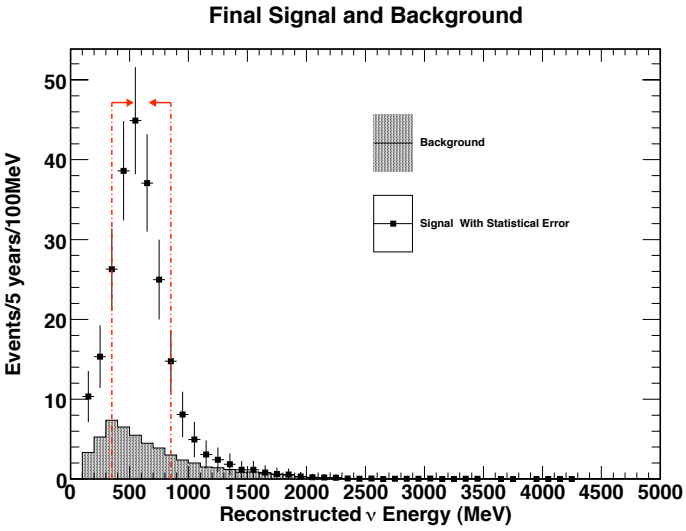

(a) $\sin^2 2\theta_{13} = 0.1$

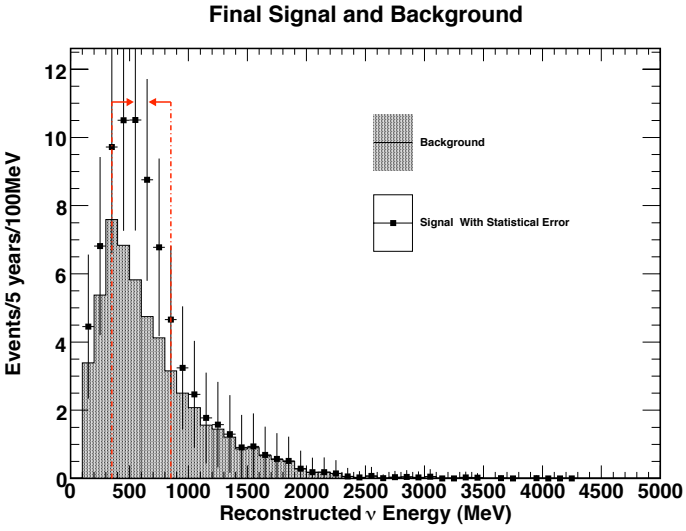

(b) $\sin^2 2\theta_{13} = 0.01$

Figure 2.21: Expected signal and background events for ν_e appearance at Super-K, for $\sin^2 2\theta_{13} = 0.1,\ 0.01$. The values of the other parameters used are given in the text. The arrows indicate the region used for analysis. Data is for five years nominal running, and errors shown are statistical only [57].

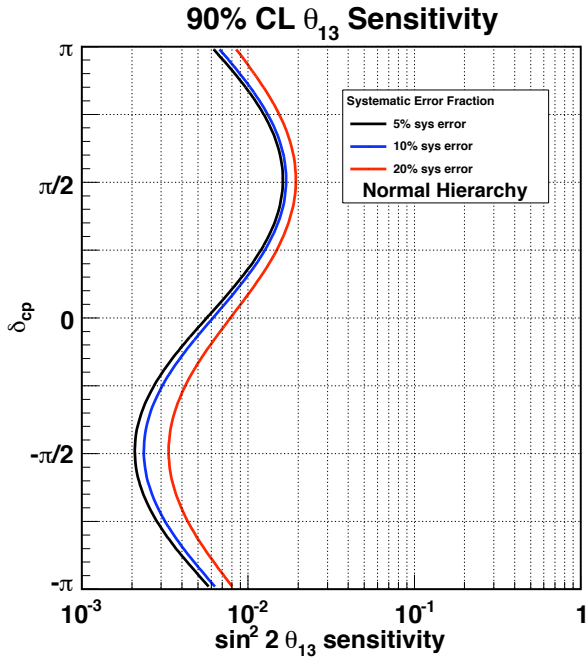

(a)

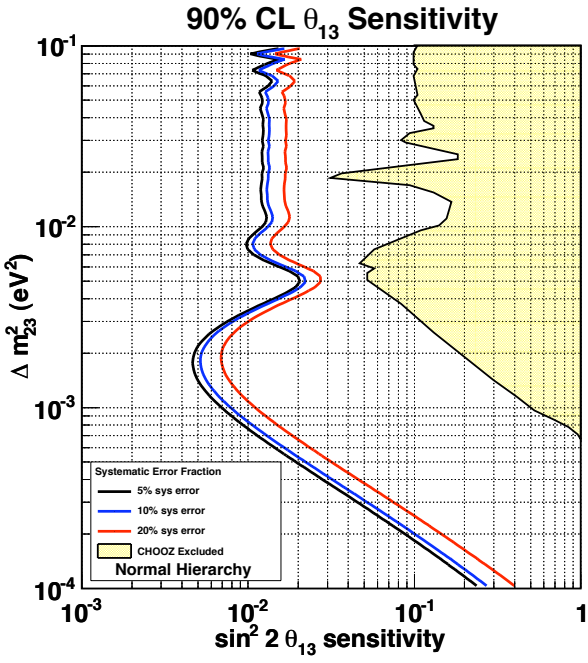

(b)

Figure 2.22: Plots showing T2K's sensitivity to $\sin^2 2\theta_{13}$, as a function of δ (a), and ΔM^2 (b). From [57].

2.7.2 ν_μ disappearance

The ν_μ disappearance search will compare the observed ν_μ-CCQE event energy spectrum at Super-K with that expected from the neutrino flux predicted by the near detector measurements, and fit the ν_μ survival probability (1.25) to extract the values of ΔM^2 and $\sin^2 2\theta_{23}$. Although the ν_μ-CC signal is large, in contrast to the ν_e search where a small signal must be extracted by aggressively removing backgrounds, T2K aims to make a very precise measurement of the mixing parameters, and it is therefore necessary to understand the backgrounds extremely well to reduce systematic errors. The expected reconstructed energy spectrum is shown in Figure 2.23, split into signal and the various background contributions. CC1π production dominates the background, and the ND280 detector will therefore need to measure the charged pion cross-section with high precision in order to reduce the uncertainty on this background.

The expected sensitivity of the T2K ν_μ disappearance search is shown in Figure 2.24 for 5 years running, with statistical errors only. At this stage of the experiment, the analysis should be systematics-dominated, so the errors will be larger than this and determined by the systematic error level achieved.

2.8 Summary

The T2K experiment will make more precise measurements of the atmospheric mixing parameters ($\sin^2 2\theta_{23}$, ΔM^2) than achieved by the previous generation of long-baseline and atmospheric neutrino experiments. It will also search for θ_{13} with

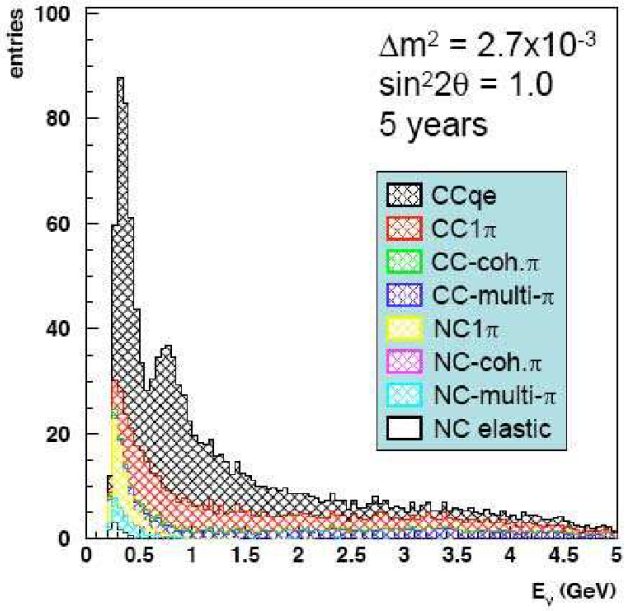

Figure 2.23: Expected signal and backgrounds for the Super-K ν_μ-CC event sample, as a function of reconstructed neutrino energy. The oscillation dip is not seen in the background channels, as these events are reconstructed with the wrong neutrino energy. Events with one or more final state pions (1π and multi-π respectively) are mostly resonant interactions. Coherent processes are shown separately. From [55].

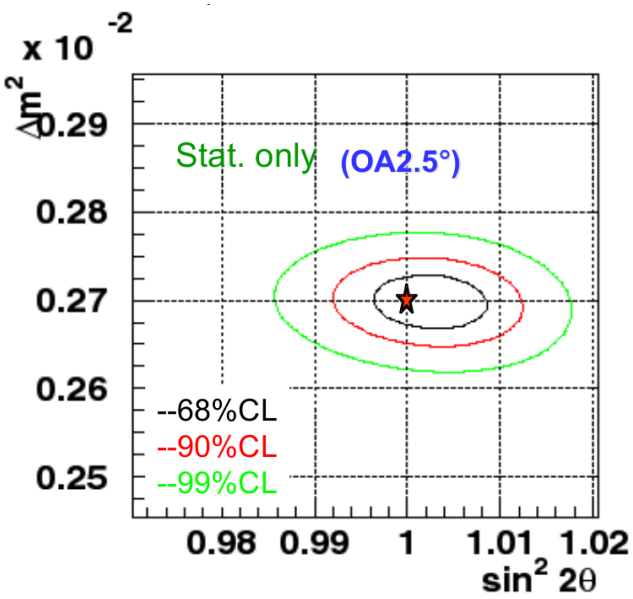

Figure 2.24: $(\sin^2 2\theta_{23},\ \Delta M^2)$ sensitivity for the full T2K proposal, including statistical errors only [58].

a sensitivity around an order of magnitude better than the CHOOZ experiment. T2K is competing with the NOνA, Daya Bay and Double CHOOZ projects to make the first unambiguous measurement of θ_{13}, but the experiments are also highly complementary because they sample different parts, and different projections, of the oscillation parameter space, as discussed in Section 1.4.3. It should also be noted that T2K has already started taking data, and is therefore well-placed to make a discovery before the other projects, which are running to later timescales (see Figure 1.8).

The measurements made by T2K will help to define the future direction of neutrino physics, and depending on the value of θ_{13}, T2K may itself by upgraded to allow a measurement of the CP phase δ. This is discussed further in Chapter 6.

Chapter 3

Characterisation and Monte Carlo simulation of MPPCs

The MPPC (Multi-Pixel Photon Counter) [59] is a novel solid-state photodetector, manufactured by Hamamatsu Photonics[1]. The ECal, FGD, P0D and SMRD subdetectors of ND280, as well as the INGRID on-axis detector, are read out using MPPCs. A unified effort has been made across the ND280 collaboration to characterise detailed MPPC properties, by making precision measurements on a small number of devices. This characterisation work has enabled the development of a detailed Monte Carlo simulation of the MPPC.

This chapter will begin with a qualitative description of the structure and function of the MPPC, and proceed to discuss the framework of the Monte Carlo model. The characterisation of specific MPPC properties, and their implemen-

[1]Hamamatsu Photonics K.K., 325-6, Sunayama-cho, Naka-ku, Hamamatsu City, Shizuoka Pref., 430-8587, Japan.

tation in the model, will then be covered. Finally, results from the Monte Carlo model will be presented, and compared with data where appropriate.

3.1 Description of the MPPC

3.1.1 Device structure and basic properties

The MPPC (Figure 3.1(a)) is an electrically parallel array of avalanche photodiodes (APDs) on a single silicon wafer, with the APD "pixels" arranged in a square grid. The model used in ND280[2] has active dimensions $1.3 \times 1.3\,\mathrm{mm}^2$, with pixel pitch $50\,\mu\mathrm{m}$; there are 9 missing pixels in one corner for an electrical connection, giving 667 pixels in total. The semiconductor is placed into a larger ceramic casing under a transparent epoxy window, with two pins for electrical connection.

Each pixel consists of a p-n diode junction, in series with a resistive layer (Figure 3.1(b)). The device is operated in Geiger mode just above the reverse breakdown voltage V_{bd}. A charge carrier liberated by an incoming photon is accelerated quickly enough by the electric field in the depletion region of the junction to gain enough energy to free further carriers upon collision with the lattice, producing avalanche charge multiplication.

Because the device is operated in Geiger mode, the avalanche is self-sustaining. Multiplication is quenched by the resistive layer in series with the diode junction, which limits the pixel current causing a buildup of charge on the

[2]The ND280 devices are a custom model, S10362-13-050C. They are similar to production model S10362-11-050C, but have a larger active area, and thus have a larger dynamic range and are more tolerant to a slight offset in the fibre-MPPC connection.

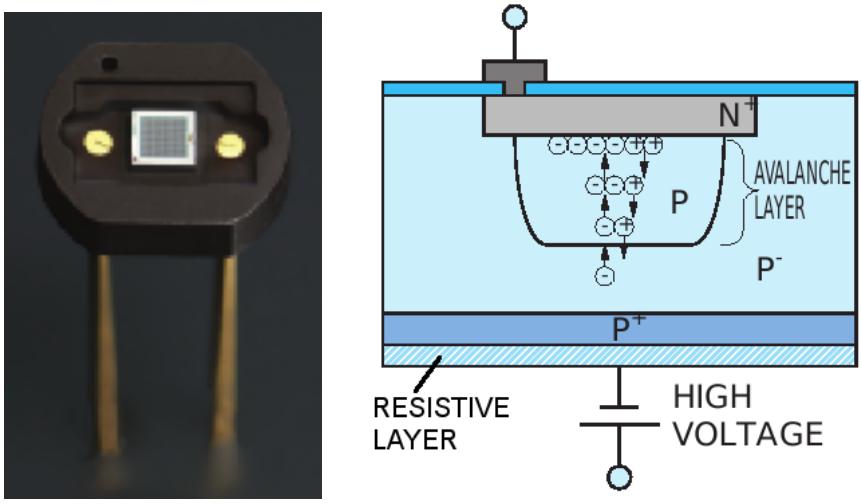

(a) (b)

Figure 3.1: (a) Photograph of an MPPC sensor, as used in ND280. (b) Schematic of a single MPPC pixel.

pixel which pulls the junction voltage down below $V_{\rm bd}$. The gain of the device (charge output per avalanche in units of e), is given by

$$G = \frac{C}{e}(V_{\rm bias} - V_{\rm bd}), \tag{3.1}$$

where C is the capacitance of the pixel diode, around 90 fF for the MPPC. Gains are typically in the region $10^5 - 10^6$, dependent on operating voltage. The actual voltage set will be a compromise between gain and noise level which is appropriate to the application. Since the pixel overvoltage ($V_{\rm o} = V_{\rm bias} - V_{\rm bd}$) is completely depleted by a single avalanche event, in the short-time limit a pixel operates as a binary counter. This is in contrast to a conventional APD, operated in proportional mode just below $V_{\rm bd}$, where avalanches are not self-sustaining, and the total current is proportional to the incident light level with a gain of around 100.

Since the MPPC pixels are connected in parallel, the output charge is a

sum of the charges for each fired pixel. The output at low light level is therefore almost proportional to the number of incident photons, since there is unlikely to be multiple photons incident on a single pixel. The gain is sufficiently stable for individual photons to be counted, and we often quote the MPPC output signal in pixel equivalents (p.e.), the number of single pixel charges that the signal corresponds to. MPPC response saturates under more intense light because of the finite number of pixels, as discussed in Section 3.1.3. The dynamic range of an MPPC is determined by the number of pixels.

The photon detection efficiency (PDE) of the MPPC, i.e. the probability for an incoming photon to produce an avalanche, is given by

$$PDE(V_{\mathrm{bias}}) = \epsilon_{\mathrm{geom}} \times \epsilon_{\mathrm{quantum}} \times \epsilon_{\mathrm{Geiger}}. \tag{3.2}$$

This is a product of ϵ_{geom} (proportion of nominal pixel area which is active), the quantum efficiency $\epsilon_{\mathrm{quantum}}$ (probability for an incident photon to produce an electron-hole pair), and $\epsilon_{\mathrm{Geiger}}$, the probability for this pair to initiate a Geiger discharge. The PDE is a strong function of the overvoltage V_{o} via $\epsilon_{\mathrm{Geiger}}$. The fill-factor ϵ_{geom} for ND280 devices is known to be $\sim$0.6.

Because the PDE is less than unity, the MPPC response N_{fired} for N_{phot} incoming photons, neglecting noise and saturation, is a random variable on a binomial distribution, with N_{phot} trials and a success probability equal to the PDE. This distribution has a mean and sigma of

$$\bar{N}_{\mathrm{fired}} = N_{\mathrm{phot}} \times PDE \tag{3.3}$$

$$\sigma = \sqrt{\bar{N}_{\mathrm{fired}}(1 - PDE)} \tag{3.4}$$

so while the avalanche gain is very stable, the response distribution is rather wide at low light levels. The response is further complicated by the addition of correlated noise, discussed below. The number of incoming photons from a given signal source will also have a significant statistical variation at low light levels, normally described by a Poisson distribution.

3.1.2 Spontaneous and correlated noise

The semiconductor used in the MPPC is silicon. At room temperature, the rate at which charge carriers are thermally excited in the semiconductor is significant, and these carriers cause avalanches (dark noise) in the same way as those freed by incoming light. The rate of thermal avalanches depends on the temperature — the number of free carriers is related to temperature by the approximate relation

$$N_{\text{carriers}} \propto e^{-\epsilon_{\text{gap}}/2k_{\text{B}}T} = e^{-A/T}, \tag{3.5}$$

where ϵ_{gap} is the band gap of silicon (around 1.11 eV) [60], giving a value for the temperature scale of A=6430 K.

The thermal noise rate also depends on the overvoltage $V_{\text{bias}} - V_{\text{bd}}$ since this affects the Geiger efficiency of the device. The observed whole-device dark count rate (DCR) (as quoted by the manufacturer) varies considerably between devices, but is on the order of 0.1–1 MHz [59] for normal operating conditions.

Correlated noise effects — additional pixel triggers stimulated by a primary avalanche — are also observed in the MPPC. Two distinct types of noise have been identified. *Crosstalk* (CT) refers to the process of avalanches causing secondary avalanches in other pixels. This is effectively instantaneous, giving a double-size

peak in the observed output charge. It is thought to be due to the emission of photons into other pixels during the avalanche process — this is supported by the observation that crosstalk is lower from avalanches near the edge of the MPPC (see Section 3.3.3).

Afterpulsing (AP) refers to the triggering of secondary avalanches, some time after the first, in the same pixel as the primary. The mechanism for this effect is charge-trapping by impurities in the semiconductor, with the decay of the trapped state releasing a carrier [61]. Very fast afterpulses have somewhat lower gain than a normal pulse, since the pixel fires again before its bias voltage has been fully recharged.

Given the observed DCR, thermal noise is not expected to be a big problem in ND280. The charge integration period for the detector is 580 ns, corresponding to only $\sim$0.5 dark pulses per gate (although some pile-up is caused by correlated noise). This is to be compared to the mean MIP signal in the ECal bars, expected to be around 20–30 pixels.

Correlated noise effects are potentially more problematic, since these are important regardless of the signal size. A priori, it is not clear whether the effect of correlated noise on the MPPC response will worsen the device resolution — this question is addressed in MC in Section 3.5.2. Regardless of the effect on resolution, the presence of correlated noise, especially afterpulsing, is still undesirable, since it adds to the complexity of quantitatively understanding and simulating MPPC behaviour. Afterpulsing introduces some non-trivial time dependence to the device behaviour — for example in a gated charge measurement, the expected signal for

a given light pulse varies with the pulse's arrival time since the proportion of the afterpulses caught decreases as the pulse position approaches the end of the gate.

3.1.3 Saturation and recovery

Because an avalanche causes the affected pixel to discharge to below V_{bd}, a fired pixel cannot fire again immediately. The recovery of pixel bias voltage depends on the details of the external circuit, but if only a small proportion of the pixels fire, their capacitances can be recharged by leeching charge from other pixels, on an exponential timescale set by the pixel RC constant $\tau_{\mathrm{I}} = R_{\mathrm{pix}} C_{\mathrm{pix}}$. For our devices, $R_{\mathrm{pix}} = 150\,\mathrm{k}\Omega$ and $C_{\mathrm{pix}} = 90\,\mathrm{fF}$, giving $\tau_{\mathrm{I}} = 13.5\,\mathrm{ns}$.

If a large proportion of the pixels fire, then the voltage across the device as a whole will drop significantly. Recovery of this voltage may happen on a longer timescale, and may not follow a simple exponential, depending on the details of the external circuit. The specific case of the front-end boards used for the ND280 ECal is discussed in Section 3.4.

A consequence of the finite recovery time of the MPPC is saturation of the device response under increasing light input, since each pixel can only fire once in a short time period. It is illustrative to consider an analytic approximation for the device saturation. For a given pixel in a device with N_{pix} pixels, the probability for a single photon incident on the device to fire that pixel is $\frac{PDE}{N_{\mathrm{pix}}}$. The probability for a given pixel to be fired by N_{phot} photons is then $1 - (1 - \frac{PDE}{N_{\mathrm{pix}}})^{N_{\mathrm{phot}}} \simeq 1 - e^{-PDE \times N_{\mathrm{phot}}/N_{\mathrm{pix}}}$. The mean number of pixels fired is

then given by

$$\bar{N}_{\mathrm{fired}} = N_{\mathrm{pix}}(1 - e^{-PDE \times N_{\mathrm{phot}}/N_{\mathrm{pix}}}). \tag{3.6}$$

This approximation is only valid for an instantaneous light pulse, where each pixel can only fire once. For the same number of incident photons distributed over an extended period of time, the saturation will be less harsh because the pixels will recharge during the light pulse and may be fired again. For light pulses much longer than the device recharging timescale, no saturation will be seen. The form of the saturation will also change if the spatial distribution of incoming light is not uniform — for example, if only a part of the device is illuminated then the number of pixels is effectively reduced.

The approximation (3.6) also neglects correlated noise effects. Crosstalk can be included to a very good approximation, by substituting the PDE in (3.6) with an "effective" value which is higher than the intrinsic number. Unfortunately, it is not possible to include afterpulsing in the same way, since this involves fired pixels refiring at a later time, and so will lead to a less harsh saturation, in a similar way to an extended light pulse.

The crosstalk and afterpulse probabilities, as well as the PDE and device gain, depend on pixel voltage, making the saturation behaviour of the device too complex to describe analytically. It must therefore be measured on a test stand, or extrapolated from low-light measurements using a Monte Carlo simulation of the device — Monte Carlo predictions for saturation behaviour are presented in Section 3.4.

3.2 Monte Carlo framework

The purpose of the MPPC Monte Carlo is to provide an empirically accurate simulation of the MPPC response to signals at all light levels. Such a simulation is a crucial part of the overall detector MC chain. It also allows us to check that our understanding of the MPPC behaviour is complete — the various measurements made to characterise the devices attempt to isolate certain device features and measure them individually, and the model enables us to compare the overall predicted response with real lab measurements. The simulation has also been used to extrapolate low-light measurements to predict the saturation behaviour of the device, though this prediction will ideally be replaced with, or tuned by, lab measurements once they are available.

The Monte Carlo [1] models the behaviour of a finite two-dimensional array of APD pixels at a given applied bias voltage and temperature, storing the actual voltage across each pixel, and simulating the effects of incoming photons, dark noise, afterpulsing and crosstalk. The recovery of pixel bias voltage over time is also simulated, and so saturation effects are also modelled. The MPPC simulation has been implemented as part of the overall ND280 MC code, though a standalone C++ version also exists.

The structure of the simulation is shown diagrammatically in Figure 3.2. The basis of the simulation is a list of potential pixel hits — these are anything which can cause an avalanche, i.e. incident photons, thermal carriers, crosstalk photons or trapped charges (for afterpulsing). For each event, a list of incident photons is provided by the code calling the simulation, allowing arbitrary time

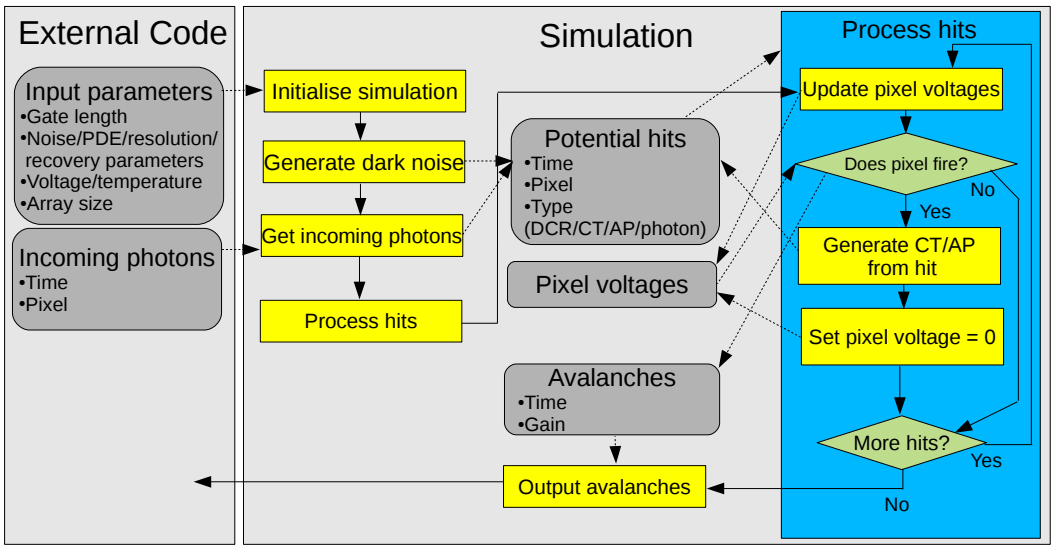

Figure 3.2: Flow chart showing the structure of the MPPC simulation. Processes are shown in yellow, data in grey and decisions in green. Solid arrows denote the process flow, and dashed ones the flow of data.

and position distributions for incident photons to be simulated as required. The simulation itself generates dark noise for the time period to be simulated, and both the incident photons and dark noise are used to populate the list of potential hits at the start of event processing.

Once the initial list is populated, the simulation steps through the potential hits in time order, using the physics models described in Section 3.3 to decide whether to trigger an avalanche for each one in turn. If an avalanche is triggered, the overvoltage of the firing pixel is set to zero, the avalanche is added to the output list, and a check is made to see whether crosstalk or afterpulses should be generated. If so, these triggers are added to the list of potential hits and are in turn processed in time order. This framework allows cascades of correlated noise to be dealt with simply and naturally, since any secondary noise effects will be

calculated when the primary crosstalk or afterpulse trigger is processed.

The measured output charge, and any error associated with it, are dependent on the properties of the output circuit, so some processing of the generated avalanches is required to convert them into an electronics response. However, the simulation does contain the capability to apply an adjustable Gaussian smearing to the charge of each avalanche, proportional to the avalanche size, to give a more realistic output. The effect of variations in the electronics pedestal (zero-signal response) must be modelled by the code calling the simulation.

Before processing each hit, the simulation updates the pixel voltages from those at the time of the previous hit, based on these previous pixel voltages and the time elapsed. Since pixel recovery behaviour depends on the external readout circuit, the software uses C++ class inheritance to allow a recovery model specific to the relevant electronics to be implemented.

3.3 Characterisation measurements and MC modelling for avalanche processes

3.3.1 Dark noise

As discussed in Section 3.1.2, the dark noise rate is expected to vary exponentially with the MPPC temperature (via the free carrier density) and also to vary with the overvoltage (via the Geiger probability). The dark noise rate has been measured at the Warsaw University of Technology, by a team led by M. Ziembicki and M. Dziewiecki, as a function of both voltage and temperature. The measurement

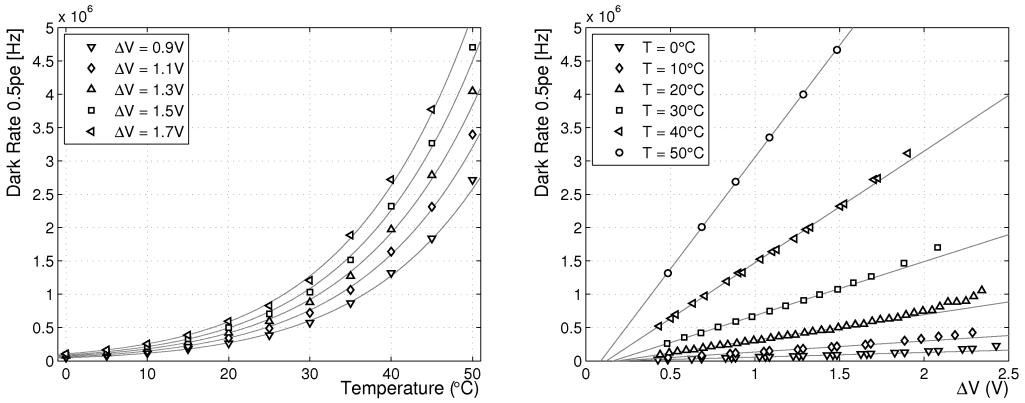

(a) (b)

Figure 3.3: Dark rate for an MPPC as a function of temperature and overvoltage. The data has been fitted with a function $(A+B\Delta V)e^{-C/T}$, with A, B and C free, and T in Kelvin. Measurements made at the Warsaw University of Technology.

was made by taking integrated output charge measurements with a 200 ns gate, and calculating the proportion of events with charges corresponding to more than 0.5 p.e.. As shown in Figure 3.3, the rate is seen to change exponentially with temperature and linearly with overvoltage to a high precision. These relations are therefore used to parameterise the dark noise rate in the simulation. The dark noise data predict a temperature scale of $A = 6900$ K. The voltage dependence is seen to deviate from linear at high overvoltages — this deviation can be attributed to afterpulses inside the gate caused by thermal triggers before the gate period, since the afterpulse probability is large at these voltages.

It should be noted that, from the manufacturer's figures, dark rates for nominally identical devices vary by a factor of $\sim$2. For the results presented in this chapter, the dark rate used was that measured for the sensor which is compared with MC in Section 3.5.1. The measured rate (from the same results that are

compared with MC) was

$$DCR(\mathrm{kHz}) = [577.3 \times (V_{\mathrm{bias}} - V_{\mathrm{bd}}) - 72.1] \times e^{-6900 \times (1/T - 1/295)}, \qquad (3.7)$$

where the reference temperature is 295 K or 22 °C. Since dark rates vary so much, it is desirable to know the individual dark rates for each sensor used in applications where dark noise is expected to be important.

When generating dark noise pulses at the start of event processing in the simulation, the expected number of dark pulses in the gate period is calculated for the applied bias voltage by multiplying (3.7) by the simulation period. The actual number of pulses generated for each event is decided using a Poisson distribution. The fired pixel and the time for each pulse are chosen randomly on uniform distributions. Because pixels may be at a lower voltage than V_{bias}, we need to account for the lower probability of these producing thermal avalanches, so when a dark pulse is processed, the pixel only fires with probability $DCR(V_{\mathrm{pix}})/DCR(V_{\mathrm{bias}})$. This method works correctly as long as the recovery model used does not allow pixel voltages to rise above the applied bias voltage.

3.3.2 Afterpulsing

In order to accurately simulate afterpulsing in the MPPC, we must understand the probability for an afterpulse to be caused by a primary avalanche, and also the time distribution of the afterpulses with respect to the primary pulse. Afterpulsing parameters have been measured with a waveform analysis method at TRIUMF, by a team led by F. Retière. The measurement was made using a fast digitiser to capture the waveform output of the MPPC, with the trigger set to take data

on an initial avalanche signal. Events were selected offline to keep only events with single-avalanche trigger pulses, in order to remove any events containing crosstalk. The delay between the initial trigger pulse and first subsequent pulse was then recorded for each event. This subsequent pulse may be due to either afterpulsing or thermal noise uncorrelated with the trigger pulse, and the distribution of delay times therefore depends on both the probability and time dependence of afterpulses, and on the dark noise rate. Since the measured quantity is the distribution of delays before the *first* subsequent pulse, we must do some work to obtain the actual afterpulse time distribution from the data — the method used was to assume a form for the afterpulse behaviour, derive the corresponding delay distribution, and fit to the data.

For a single kind of trapped state, with up to one afterpulse per avalanche, we expect the afterpulse time distribution to be exponential. In general, however, we may have multiple trapped states with different decay times, and more than one carrier may be trapped during a single avalanche. For i carriers trapped in a state with decay time τ, the probability that no afterpulse has yet occurred at time t is given by $P_{\text{no AP}}(t) = e^{-t\times\frac{i}{\tau}}$. If we assume that the number of trapped carriers is on a Poisson distribution with mean λ, then we instead have

$$P_{\text{no AP}}(t,\lambda,\tau) = \sum_{j=1}^{\infty} \text{Pois}(j,\lambda) e^{-t\times\frac{j}{\tau}}. \tag{3.8}$$

Similarly, we can obtain the probability that the first afterpulse is at time t. For i carriers, this is equal to the probability of not yet having an afterpulse, multiplied by the decay rate for i carriers, $R = \frac{i}{\tau}$. Again taking a Poisson-weighted

sum over the number of trapped carriers, we get

$$\frac{\mathrm{d}P_{\text{1st AP}}(t,\lambda,\tau)}{\mathrm{d}t} = \sum_{j=1}^{\infty} \mathrm{Pois}(j,\lambda) e^{-t\times j/\tau} \times \frac{j}{\tau}. \tag{3.9}$$

The corresponding probabilities for dark noise are

$$P_{\text{no DN}}(t, DCR) = e^{-t\times DCR} \tag{3.10}$$

$$\frac{\mathrm{d}P_{\text{1st DN}}(t, DCR)}{\mathrm{d}t} = DCR \times e^{-t\times DCR}. \tag{3.11}$$

When both dark noise and afterpulsing are included, then the probability of getting the first pulse at time t is given by

$$\frac{\mathrm{d}P_{\text{1st pulse}}(t)}{\mathrm{d}t} = \frac{\mathrm{d}P_{\text{1st DN}}(t)}{\mathrm{d}t} P_{\text{no AP}}(t) + \frac{\mathrm{d}P_{\text{1st AP}}(t)}{\mathrm{d}t} P_{\text{no DN}}(t). \tag{3.12}$$

The observed distribution is not well-fitted by (3.12). The agreement can be improved by using a model with two trapped states with different time constants and Poisson means (λ_S, τ_S) and (λ_L, τ_L), corresponding physically to two different metastable states created by impurities in the silicon. If the two afterpulsing processes (labelled APS, APL) are independent, the distribution of delay times should be given by

$$\begin{aligned}\frac{\mathrm{d}P_{\text{1st pulse}}(t)}{\mathrm{d}t} &= \frac{\mathrm{d}P_{\text{1st DN}}(t)}{\mathrm{d}t} P_{\text{no APS}}(t) P_{\text{no APL}}(t) \\ &+ P_{\text{no DN}}(t) \left[\frac{\mathrm{d}P_{\text{1st APS}}(t)}{\mathrm{d}t} P_{\text{no APL}}(t) + \frac{\mathrm{d}P_{\text{1st APL}}(t)}{\mathrm{d}t} P_{\text{no APS}}(t) \right].\end{aligned} \tag{3.13}$$

The afterpulse data for a single voltage are shown in Figure 3.4, with fits to models (3.12) and (3.13). It can be seen that the latter model provides a better fit to the data. The result from the MC tuned on the double-AP fit is also shown — the simulation reproduces the data very precisely. The fitted time constants and Poisson means for the afterpulse analysis, as a function of voltage, are

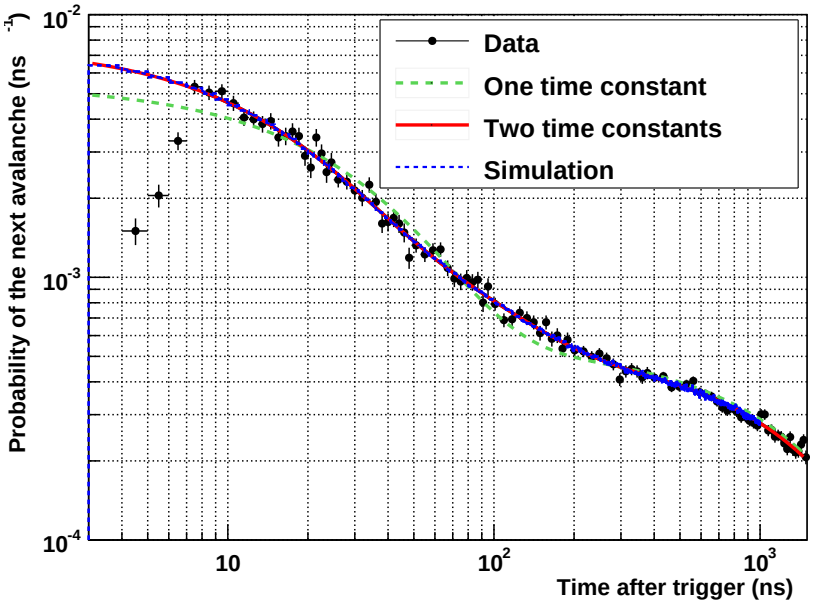

Figure 3.4: Plot showing the secondary pulse delay distribution for the TRIUMF afterpulsing analysis. Fits for one- and two-time constant models are shown, along with the result from the MPPC simulation. The discrepancy before 7 ns is due to pulse finder inefficiency for very small delays, and these points are not included in the fit.

shown in Figure 3.5. The best fits for these parameters are $\tau_{S,L}$=(17.7 ns, 70.9 ns), $\lambda_{S,L}$=(0.043 $V_{\rm o}^2$, 0.043 $V_{\rm o}^2$), for $V_{\rm o}$ in volts.

The afterpulsing model used in the simulation is based closely on the results of the TRIUMF analysis. Potential triggers for afterpulsing are generated whenever a pixel fires, and the numbers of short and long triggers to generate are decided using Poisson distributions. Each trigger is generated with a delay taken from an exponential distribution on the appropriate timescale.

The analysis above has assumed that every trapped carrier will generate an avalanche. In reality, the Geiger efficiency for a released charge will be less than unity, but the Poisson means $\lambda_{S,L}$ are calculated based on the observed

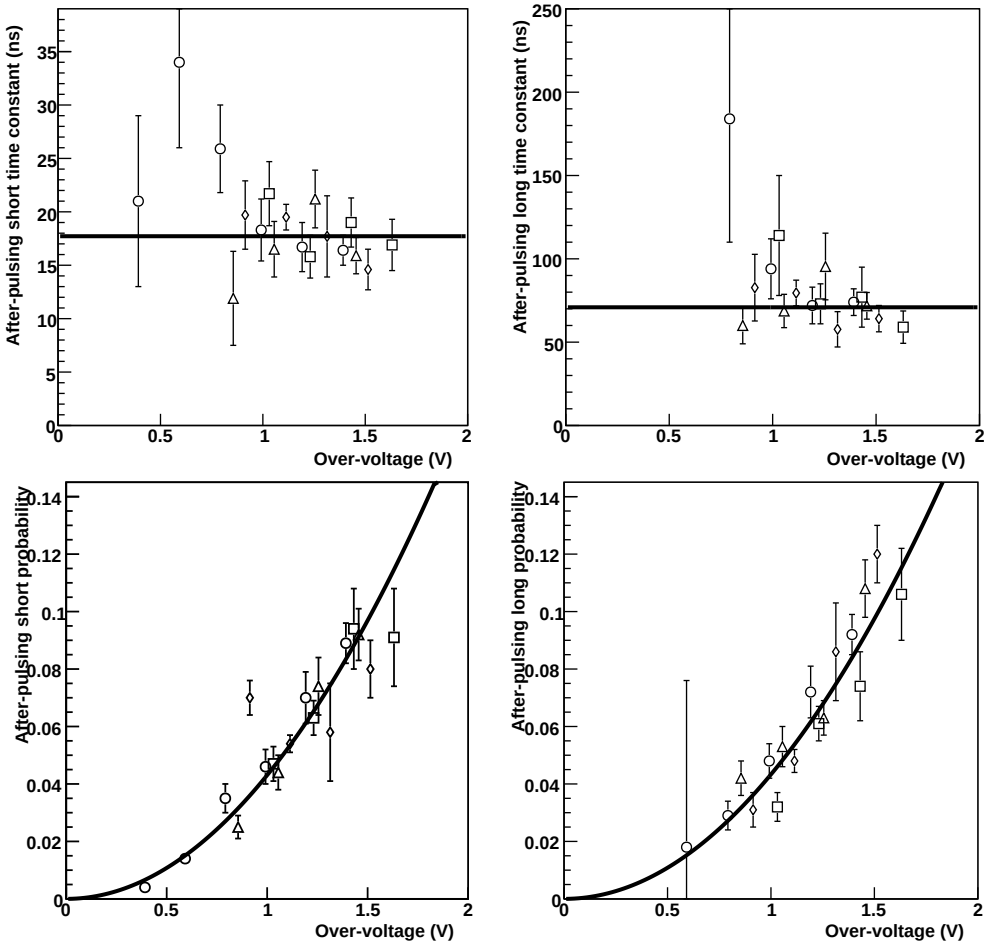

Figure 3.5: Final results for the TRIUMF afterpulsing analysis, for a sample of four MPPCs, over a range of voltages at a temperature of 25 °C. The different point shapes correspond to different devices. Errors shown are from fitting to the observed time distributions.

afterpulsing distributions, which already include this effect. Because the analysis already provides a complete description of afterpulsing behaviour, we do not need to account separately for efficiency effects when processing the afterpulse triggers in the simulation. An afterpulse trigger therefore produces an avalanche with a probability of unity. It should be noted that in principle, the exponential distribution for afterpulse decay should be modified to take into account pixel recovery — this would reduce the probability of the pixel firing at shorter times. A model taking this effect into account has been tested against the data but is disfavoured; the reason for this is unknown. If such a model were favoured, it could be accounted for in the simulation by factorising the afterpulsing voltage dependence into effects on both the mean number of trapped charges, and the probability of a trapped charge triggering an avalanche.

3.3.3 Crosstalk

The crosstalk probability for the MPPC has been measured by A. Vacheret at Imperial College London by using a pulsed laser which can be focused onto a single MPPC pixel. Since only one pixel is illuminated by the laser, any peaks larger than a single avalanche can be attributed to crosstalk with neighbouring pixels. These events can be differentiated from afterpulses, where several peaks are visible rather than a single large peak.

The probabilities for an avalanche to induce 1, 2 and 3 crosstalks, as well as the inclusive CT probability, have been measured. This was done for a pixel in the centre of the device ("normal" pixel), one on the edge, and one in the

corner. Because the pixels have different numbers of neighbours, the data put fairly detailed constraints on any crosstalk model for the device.

Implementation of a crosstalk model in the simulation is a compromise between achieving good agreement with the data, consistency with the underlying physics believed to be taking place, and computational efficiency. In principle, crosstalk photons may be rescattered between pixels before causing an avalanche, and the probability of getting crosstalk will depend on the position of the initial avalanche within the primary pixel (this is supported by data taken at Imperial by scanning the laser across the pixel faces). Including a detailed microscopic model would be computationally prohibitive, however, and is not required to achieve good agreement with the data. We instead considered a model where the probability of generating a crosstalk photon from a primary avalanche is a quadratic function of the device overvoltage. The pixel fired by this photon is determined by assuming a finite range R for the crosstalk photon, and choosing the secondary pixel using a probability distribution $e^{-r/R}$, where r is the distance from the centre of the primary pixel to the centre of the candidate pixel. This model accounts for the lower crosstalk probability for pixels near the device edge, by allowing the photon to be lost by scattering into "pixels" off the edge of the device.

The quadratic parameters for the crosstalk probability in the model were derived by fitting the inclusive crosstalk probability for a normal pixel, giving

$$P_{\mathrm{CT}}(V_{\mathrm{o}}) = 0.0356V_{\mathrm{o}} + 0.0598V_{\mathrm{o}}^2, \tag{3.14}$$

for V_{o} in volts. A toy MC was then used to calculate the CT probabilities for the different pixels, for several different values of the range parameter R. A range

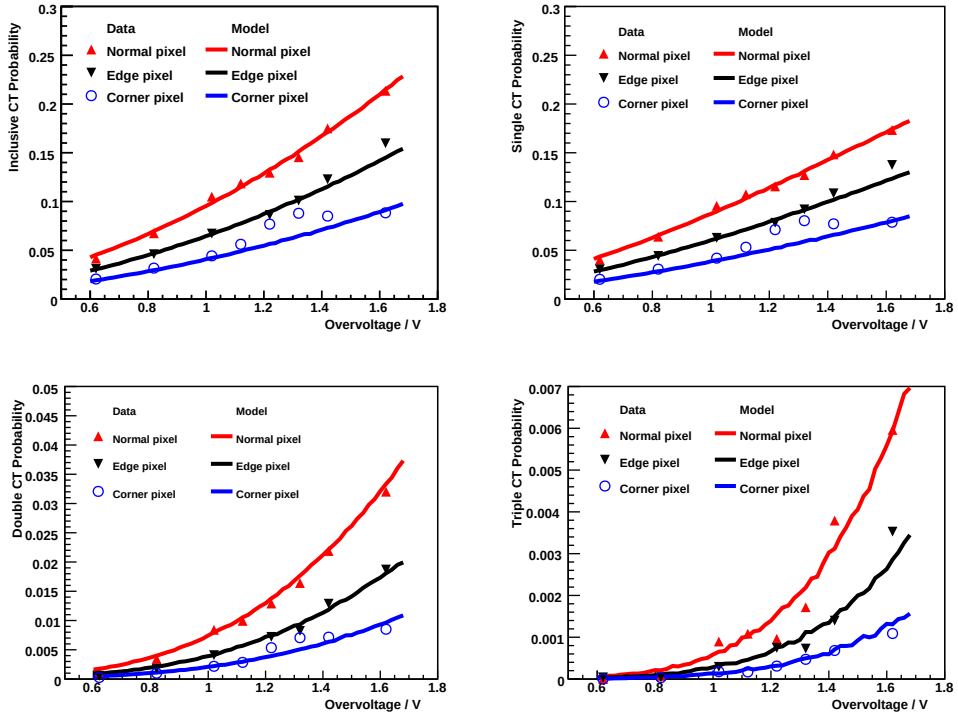

Figure 3.6: Inclusive, single, double and triple crosstalk probabilities (differentiated on plot axes), as a function of voltage for normal, edge and corner pixels. Results are from the Imperial test stand. The simulation results are also shown.

of 0.4 times the pixel pitch was found to best match the data — this value was chosen by eye since no error information was available to allow fitting the data. The results for this range value are shown in Figure 3.6. There is a plateau in the crosstalk probabilities for the corner pixel which is not matched by the simulation — this is not completely understood, but for empirical purposes it is not very important, since these pixels form a very small subset of all pixels in the device.

Another model has also been considered, where the number of crosstalk photons produced is decided using a Poisson distribution, rather than generating only one photon with a given probability. When this model is tuned to the in-

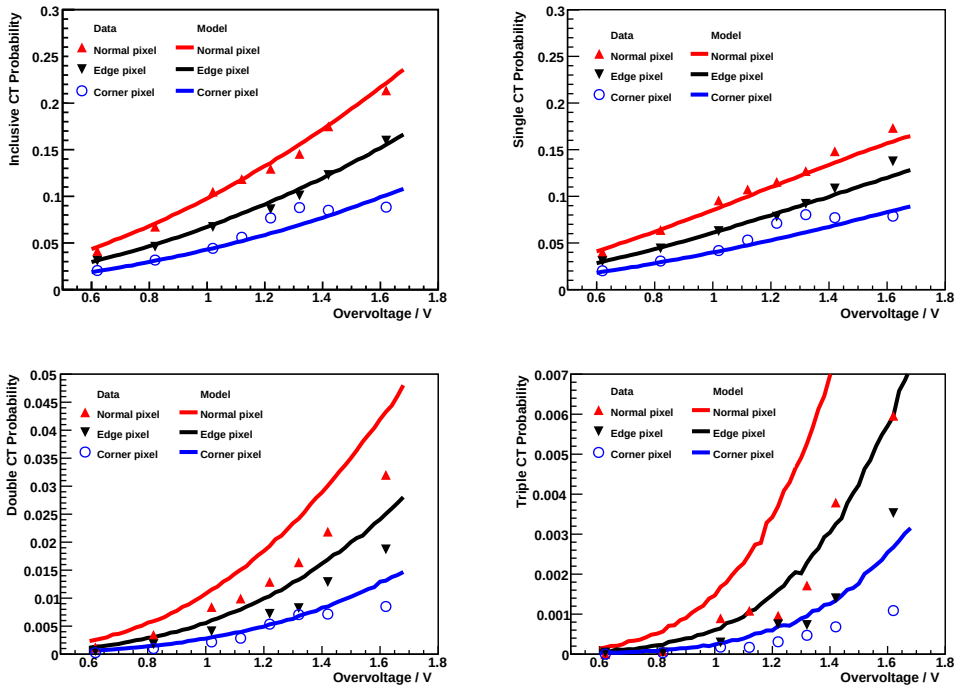

Figure 3.7: Crosstalk probability plots as in Figure 3.6, with the preferred model replaced by a model with the crosstalk photons generated on a Poisson distribution. Agreement with data for higher crosstalk multiplicities is poor.

clusive probability as before, however, the probabilities for high CT multiplicities are greatly overestimated, as shown in Figure 3.7. This is true for all values of the range parameter. The disagreement of this model with the data may be due to not taking into account the position of the primary avalanche within its pixel — it is possible that crosstalk photons are much more likely to scatter into the neighbour closest to the avalanche, and so even if several photons are produced, the probability for them to scatter into different pixels is low.

The data presented here are from a preliminary version of the laser analysis, which is currently in its final stages. They give a rather larger value for crosstalk

in a "normal" pixel than is compatible with total correlated noise measurements from, e.g. the UK QA tests in Chapter 4, once afterpulsing is subtracted. This is thought to be because the incident light may be rescattered by the epoxy window on the MPPC and hit adjacent pixels, causing additional crosstalk. This is supported by results for crosstalk from dark noise (also from Imperial), which give an inclusive crosstalk probability of 0.09 ± 0.01 at an overvoltage of 1.33 V. Because of the discrepancy between the light-induced and dark noise crosstalk measurements, and the large prediction of the former with respect to total correlated noise measurements, the crosstalk probability in the simulation is scaled so that it agrees with this dark noise measurement, but other parameters are as tuned on the laser data. This method is not quite correct as the probabilities for higher crosstalk multiplicities scale non-trivially with the inclusive probability, but this represents the current state-of-the-art for this measurement.

3.3.4 Photon detection efficiency

The PDE of the MPPC as a function of voltage has been measured at the Institute for Nuclear Research in Moscow, by a team led by Y. Kudenko. A green LED with a wavelength profile similar to the Y11 fibre was used to illuminate the MPPC, through a 0.5 mm collimator. The number of incident photons was measured using a calibrated photomultiplier. An analysis similar to that described for the QA in Section 4.2.3, measuring the probability for no pixels to be fired, was used to calculate the PDE as a function of temperature and overvoltage. The PDE was found to increase with overvoltage, but saturated as the voltage was increased, as

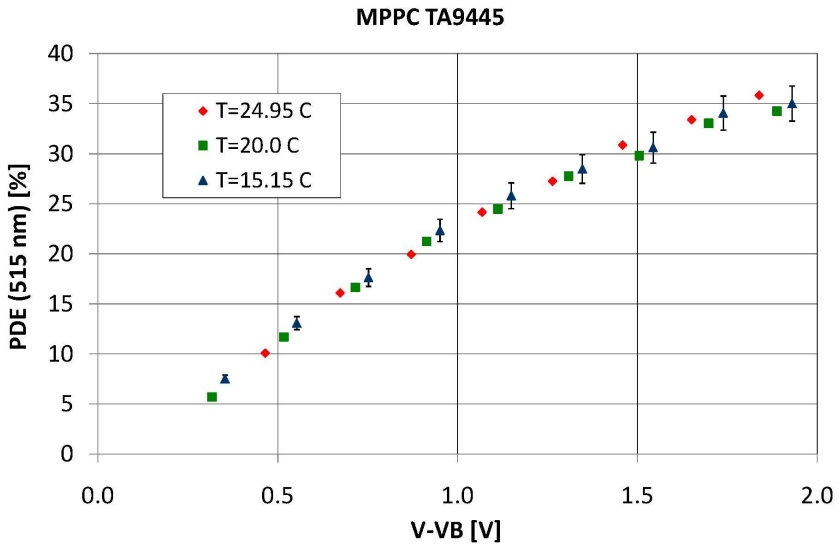

Figure 3.8: PDE measurements made at INR as a function of overvoltage, for three temperatures.

shown in Figure 3.8; no temperature dependence was observed. The PDE curve is well-fitted by the function

$$PDE(V_{\rm o}) = -0.0366 + 0.325V_{\rm o} - 0.0644V_{\rm o}^2 \tag{3.15}$$

for $V_{\rm o}$ in volts, and this function is used in the simulation. The results of a similar measurement, made at Sheffield to calibrate sensors for the UK photosensor QA, are shown in Section 4.1.3.

Procedurally, modelling of the PDE in the simulation is very simple. All incident photons are stored in the list of potential triggers, but only fire the pixel with probability $PDE(V_{\rm pix})$.

3.4 Recovery and saturation modelling

Unfortunately, no reliable data has yet been taken for the saturation behaviour of the MPPC. As discussed in Section 3.1.3, once correlated noise and finite-length light pulses are considered, the saturation behaviour of the device is too complicated to describe analytically, but it is possible to simulate it if a model for pixel voltage recovery is available.

The important features of the TFB-MPPC interface (described in detail in [62]), for a single channel, are shown in Figure 3.9. Each MPPC is read out by two amplifier channels, in order to expand the dynamic range of the readout, and charge is split between them in proportion to their input capacitors C_{HI} and C_{LO}. A single HV rail supplies the bias voltage to all channels, but a small trim voltage may be set independently for each MPPC, to set the channel gain. A large resistance $R_{\mathrm{B}} = 100\,\mathrm{k}\Omega$ between the bias voltage source and the MPPC limits the availability of charge from this source in the short-time limit; however a capacitor $C_{\mathrm{G}} = 330\,\mathrm{pF}$ connected to the MPPC by a very small resistance $R_{\mathrm{G}} = 50\,\Omega$ serves as a reservoir of charge over short times. This circuit has been modelled using a SPICE simulation program [63], the circuit schematic for which is shown in Figure 3.10. The MPPC is modelled as an electrically parallel array of N_{pix} RC series circuits, with $R_{\mathrm{pix}} = 150\,\mathrm{k}\Omega$, $C_{\mathrm{pix}} = 90\,\mathrm{fF}$.

Because the flow of charge from the bias source is limited, we consider that in the short-time limit, this can be excluded from the circuit to a high level of accuracy. We can also remove the Trip-t model, on the grounds that the high-gain channel will quickly saturate and not allow much current to pass, and C_{LO} is much

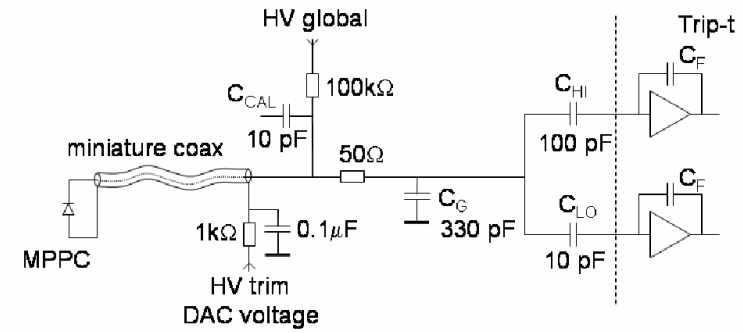

Figure 3.9: Simplified schematic of the MPPC-TFB interface.

smaller than C_{G}. We then have a circuit including only the MPPC itself and the capacitor C_{G}, as shown in Figure 3.11. If the voltage across each pixel V_i, and the voltage across the capacitor V_{G}, are known at $t = 0$, the future evolution of the circuit can be calculated analytically. The result for the voltage across pixel i at time t is given by:

$$\begin{aligned} V_i(t) =& V_i(0)e^{-t/\tau_{\mathrm{I}}} + (V_{\mathrm{G}}(0) - \frac{A\tau'}{C_{\mathrm{G}}})(1 - e^{-t/\tau_{\mathrm{I}}}) \\ &+ \frac{A(\tau'/C_{\mathrm{G}} - R_{\mathrm{G}})}{1 - \tau_{\mathrm{I}}/\tau'}(e^{-t/\tau'} - e^{-t/\tau_{\mathrm{I}}}) \qquad (3.16) \\ \text{where } \tau_{\mathrm{I}} = R_{\mathrm{pix}}C_{\mathrm{pix}}, &\qquad \tau' = \frac{R_{\mathrm{pix}}/N_{\mathrm{pix}} + R_{\mathrm{G}}}{1/C_{\mathrm{G}} + 1/C_{\mathrm{pix}}N_{\mathrm{pix}}}, \\ A = & \frac{V_{\mathrm{g}}(0) - \sum_j V_j(0)/N_{\mathrm{pix}}}{R_{\mathrm{pix}}/N_{\mathrm{pix}} + R_{\mathrm{G}}}. \end{aligned}$$

A similar expression exists for the voltage $V_{\mathrm{G}}(t)$. For our component values, the characteristic time for charge exchange with the external capacitor, τ', is very close to the intrinsic pixel recovery time τ_{I}, and so the last term in (3.16) is small. To a good approximation, therefore, the pixel voltages converge exponentially with time constant τ_{I}, but the voltage converged at is a capacitance-weighted mean of

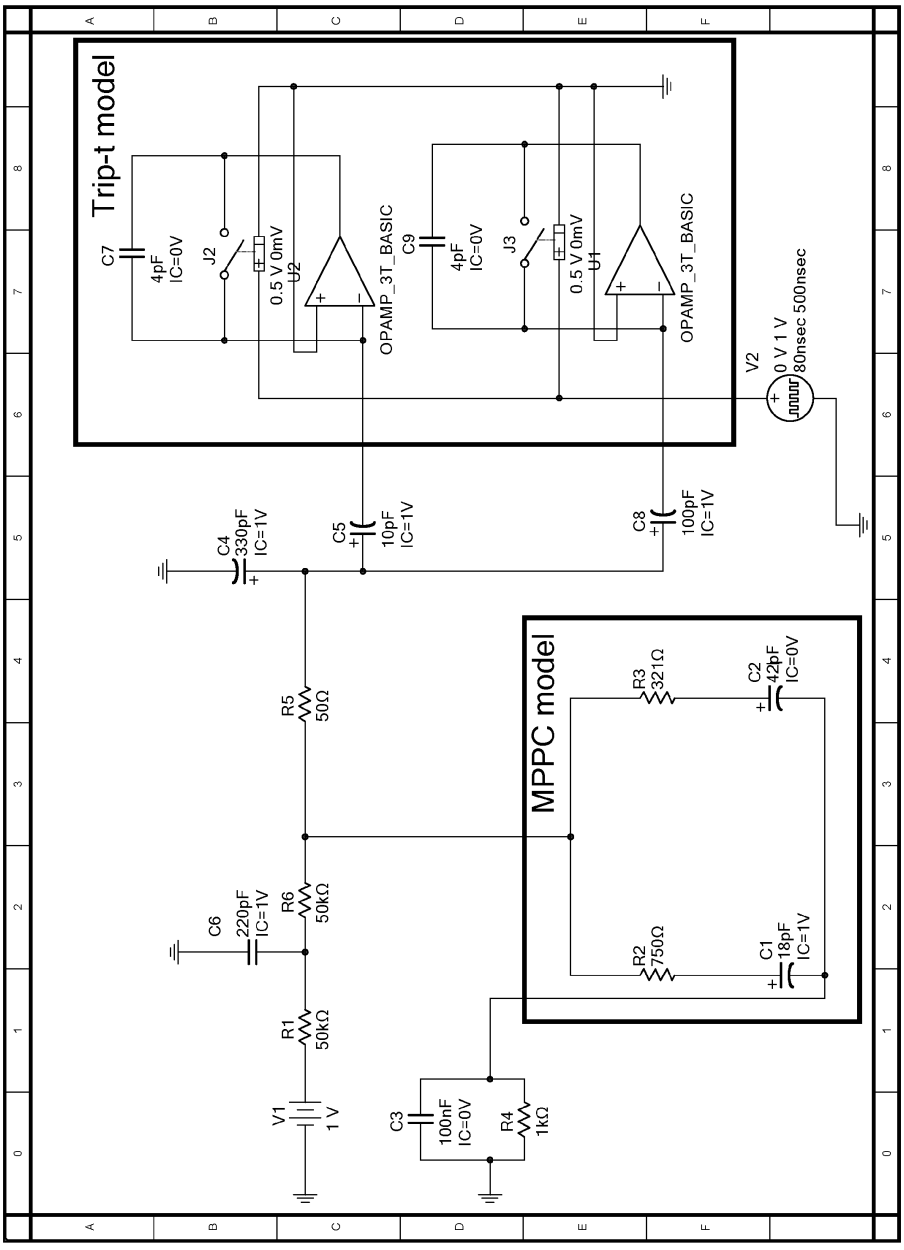

Figure 3.10: The SPICE simulation model for the MPPC-TFB interface, to simulate voltage recovery.

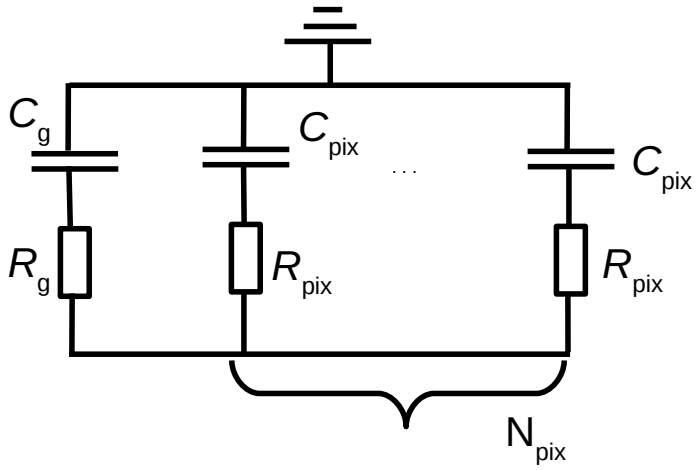

Figure 3.11: Analytic model used to simulate the MPPC recovery in the short-time limit.

the initial pixel voltages and $V_{\rm G}$:

$$V_{\rm final} = V_{\rm G}(0) - \frac{A\tau'}{C_{\rm G}} = \frac{V_{\rm G}(0)C_{\rm G} + \sum_j V_j(0)C_{\rm pix}}{C_{\rm G} + C_{\rm pix}N_{\rm pix}}. \tag{3.17}$$

The predictions of this model have been compared to the SPICE simulation. The results are shown in Figure 3.12. The agreement at short times is very good — the small discrepancy observed is due to current flow into the amplifier inputs. A larger transient effect is seen in the SPICE simulation as the amplifiers are reset between readout cycles; this is not included in the analytic model. In addition, the analytic model will not work in the long-time limit as it does not allow for recharging from the bias source. However, since the ND280 simulation is concerned with signals within a single timeslice, the simple model is sufficient for our requirements, and has been implemented in the MPPC MC code. An analytic model is needed for this purpose as a numerical recovery simulation would not be

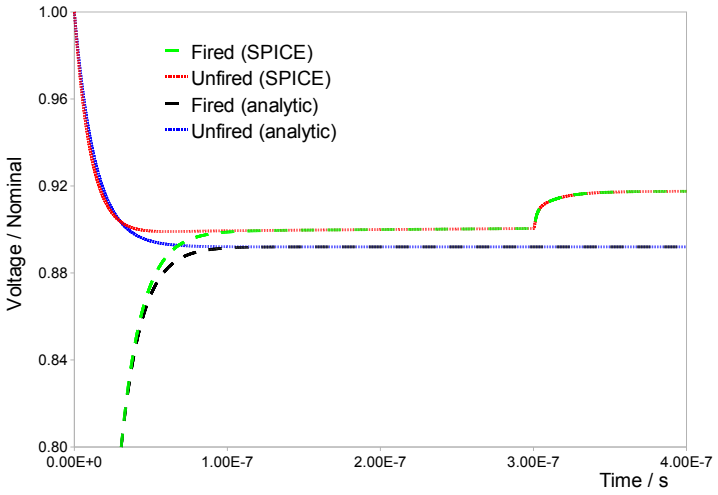

Figure 3.12: Plot showing the MPPC recovery predicted by SPICE, and by the analytic model. 70% of the pixels have been fired at $t = 0$.

computationally feasible.

3.5 MC results and comparison to data

3.5.1 Low-light data

The MPPC Monte-Carlo code, as tuned using the characterisation measurements described above, has been compared to data taken at Imperial using an LED tuned to various intensities. The setup used was similar to the ECal QA configuration, described in detail in Chapter 4; in particular, a TFB board with an integration time of 540 ns was used to take data. As in the QA, the output measurement was a histogram of the integrated charge from the MPPC over the gate length.

Because some parameters needed by the simulation to produce a charge spectrum are functions of the system configuration external to the MPPC, some

additional tuning was done for this comparison. The width of the pedestal peak at nominal overvoltage ($V_o = 1.32\,\mathrm{V}$) was used to smear every MC event by a Gaussian — this Gaussian was assumed to correspond to a constant value in charge units across the dataset, and so decreases as a proportion of the peak separation as the gain increases. The intrinsic width of the 1 p.e. peak was also measured for the nominal overvoltage, and this was used to smear each avalanche in an event in the MC. The overall light level was not calibrated for the dataset, and so the measured light level for nominal overvoltage was used to calculate the number of photons for each LED intensity. This means that the overall scale of the MPPC PDE is not tested in this comparison, but an error in the form of the PDE as a function of voltage will affect the results.

The data and MC results for a range of overvoltages, and two light levels, are shown in Figure 3.13. Good agreement is seen at all light levels and voltages — the peak shapes differ slightly between data and MC in a few cases, but the overall distribution of avalanches can be seen to agree very well in Figure 3.14, where the same data is plotted with one bin per p.e.. The largest discrepancies are seen in the bottom right plot, for the highest light level and overvoltage — this is due to the peak spacing increasing at high signal levels in the data, which may be due to a calibration problem, since the pedestal subtraction and linearity corrections (calibration steps 1 and 2 as described in Section 4.6) were performed using a now-obsolete calibration routine. No attempt was made to fit the simulation parameters to the response data, since the purpose of this comparison was to check the MC agreement with the overall response data, using the simulation as tuned on the

individual parameter measurements.

Taken as a whole, these results show that the MC is able to reproduce very precisely the MPPC response at low light levels, using non-sensor-specific parameters derived from characterisation measurements designed to measure a single parameter. This indicates that both our understanding of the MPPC avalanche processes, and our modelling of them, is quantitatively accurate.

3.5.2 Energy and timing resolution

Since our low-light level results indicate that the MPPC behaviour is well-understood, we do not expect to have large systematic uncertainties in the modelling or calibration of the devices. However, the MPPC response will show some statistical variation, and this will have an effect on the energy and timing resolution of the device. We have studied these resolution effects using the MPPC Monte Carlo.

The theoretical best energy resolution for a device is given by the spread in incident photons for a given energy deposit in the scintillator bar. Since we expect the photon distribution to be Poisson, then this spread should be the square root of the mean photon number, $\sigma_{\mathrm{phot}} = \sqrt{\bar{N}_{\mathrm{phot}}}$. A sensor with no spread in gain, no noise, and a PDE of unity, will give this energy resolution. For a real sensor, some additional error will be introduced. We can parameterise this as an excess noise factor F_{MPPC}:

$$F_{\mathrm{MPPC}} = \frac{\sigma_{\mathrm{meas}}}{\sigma_{\mathrm{phot}}}, \tag{3.18}$$

i.e. the ratio of the spread in reconstructed photon number, to the Poisson error on the photon number. We used the MPPC MC to calculate F_{MPPC} as a function

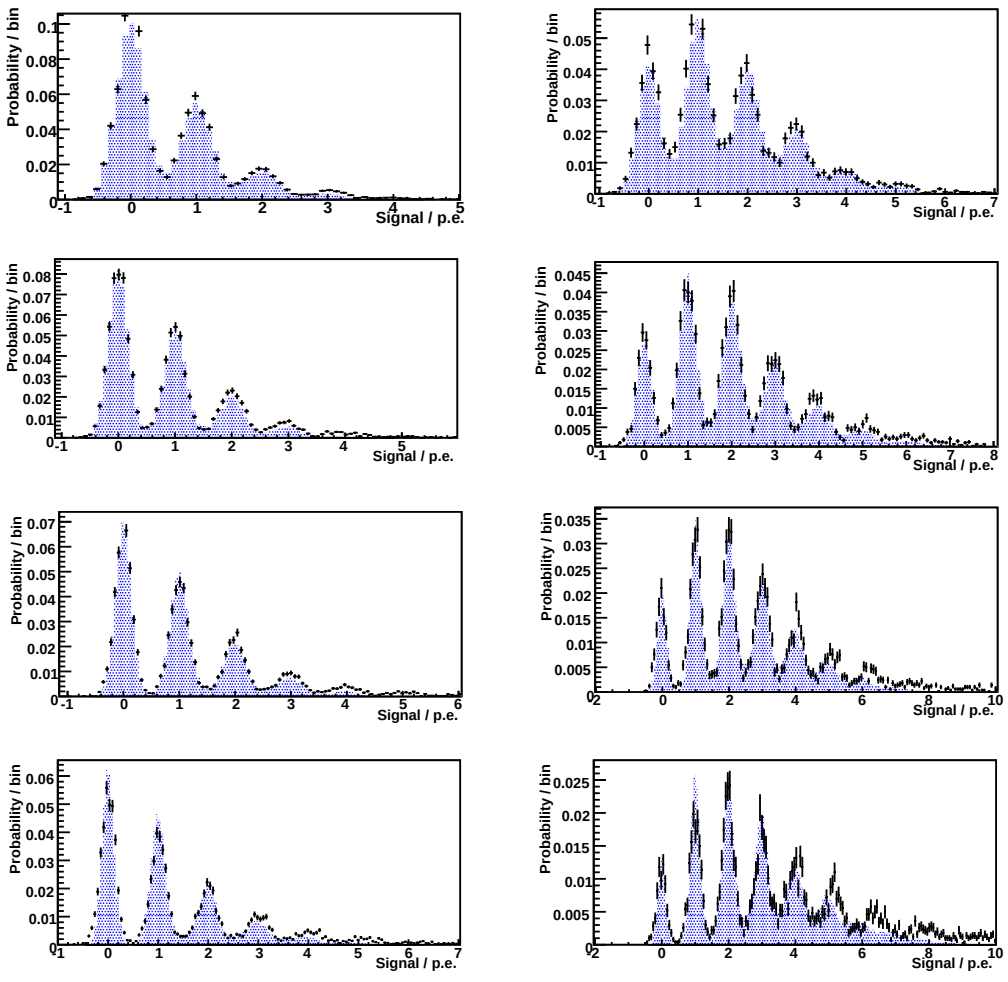

Figure 3.13: Data-MC comparison for low-light data taken with a TFB. Filled areas are MC, points are data.

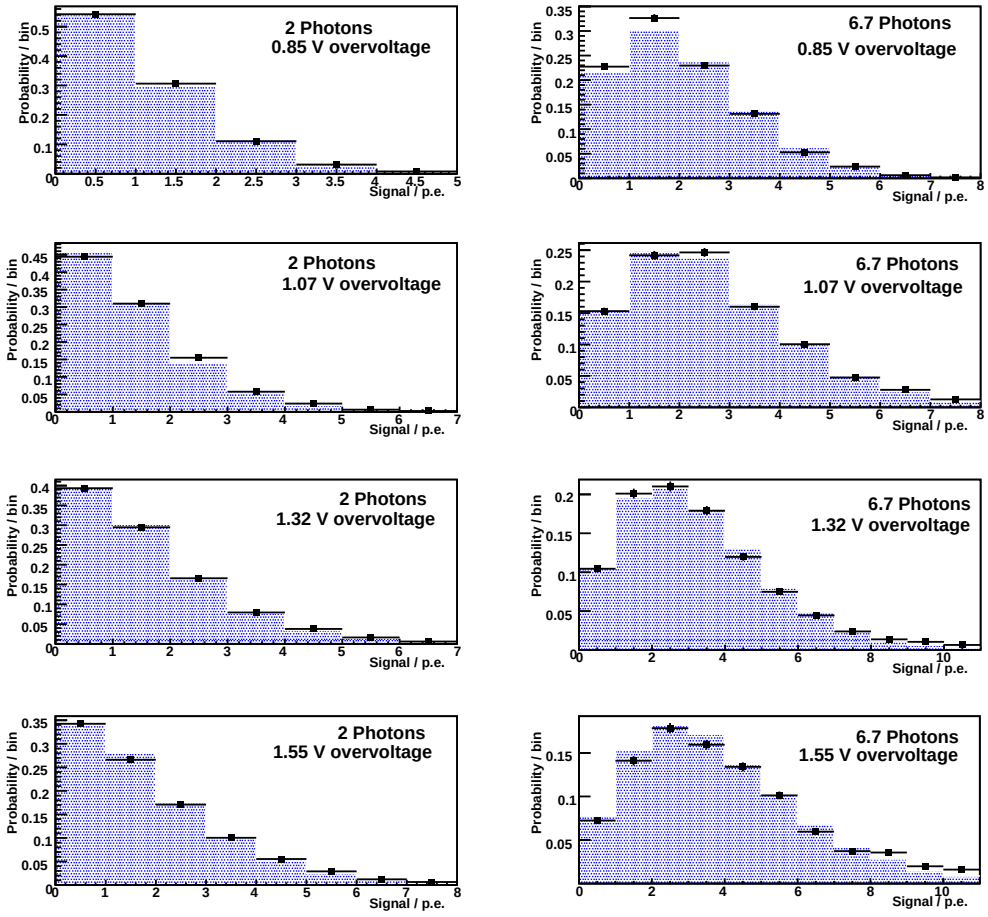

Figure 3.14: Data-MC comparison plots showing the same data as Figure 3.13, with one bin per p.e.. Filled areas are MC, points are data.

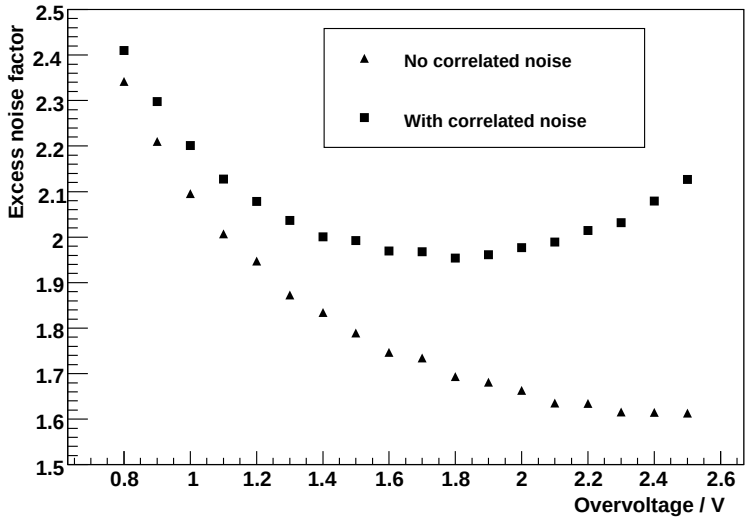

Figure 3.15: The MPPC excess noise factor F_{MPPC} for a MIP signal, as a function of overvoltage. Results are shown with and without correlated noise.

of overvoltage, for a MIP-level signal of $\bar{N}_{\mathrm{phot}} = 100$. The reconstructed photon number for each event was obtained by using a fit of mean response against photon number. The results of this simulation are shown in Figure 3.15. F_{MPPC} can be seen to drop with increasing overvoltage, as the PDE increases and so the effective statistics become larger. At higher overvoltages, F_{MPPC} starts to increase due to high levels of correlated noise. The increase in F_{MPPC} due to correlated noise at a gain of 7.5×10^5 is around 10%.

The timing resolution of the MPPC has also been estimated using the MC. The TFB boards used in ND280 provide a timestamp for each event by recording when the integrated output charge reaches a certain threshold. We estimated the timing resolution by generating data for several different light levels, with a light pulse 250 ns into the integration gate. The timescale of the light pulse was set

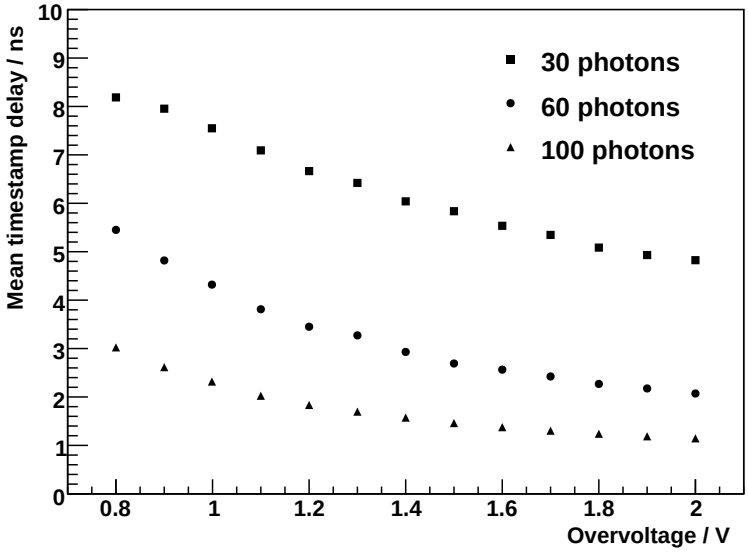

Figure 3.16: Timing resolution of the MPPC as a function of overvoltage, plotted for several light levels.

to 7 ns, to match the Y11 fibre decay time, and the timestamp was set to the time that the integrated output went above a 2.5 p.e. signal. For the MPPC pulse shape, we ignored the risetime, which is negligible compared with the pulse length, and used an exponential decay function with a timescale equal to the MPPC recovery time. The results are shown in Figure 3.16 — for 100 incident photons, the timing resolution is around 2 ns, better than the resolution of 2.5 ns set by the timestamping in the electronics. We therefore conclude that the MPPC does not significantly affect the timing resolution of the detector. However, with a threshold set to 2.5 p.e. and a gain of 7.5×10^5, around 1% of events were observed to have an incorrectly set timestamp before the light pulse due to dark noise. This is a problem which must be considered in reconstruction of detector events.

It should be noted that in our case, the timing resolution is limited by the use of a single discriminator to produce a timestamp based on the integrated output charge. In time-critical applications, where the output waveform itself could be read out, the timing resolution would be limited by the jitter on the delay between the light pulse and the MPPC peak current. This has been measured by the manufacturer to be less than 300 ps [59].

3.5.3 Saturation prediction

Because the maximum signals expected in the ND280 detectors correspond to only around 1000 photons, or 200–300 fired pixels, saturation effects are not expected to be large. However, in order to achieve the maximum precision possible, it is still important to include these effects in the simulation and calibration. Characterisation data for the MPPC response to large signals is not yet available, and so the MPPC simulation, as tuned on low-light characterisation measurements, has been used to predict the response in this region.

Figure 3.17 shows the predicted MPPC response for large signals, for both the TFB model used in the simulation, and a model where the MPPC recovers directly from a voltage source (so that the pixels recover independently with their intrinsic recovery time $\tau_{\mathrm{I}} = 13.5$ ns). It can be seen, as expected, that saturation is harsher for the TFB, since the amount of charge available to recharge the pixels is limited by the external capacitance available in the circuit. The saturation also becomes harsher for faster light pulses, since for longer pulses the pixels may recharge and be fired again by a later photon. In the region important for ND280

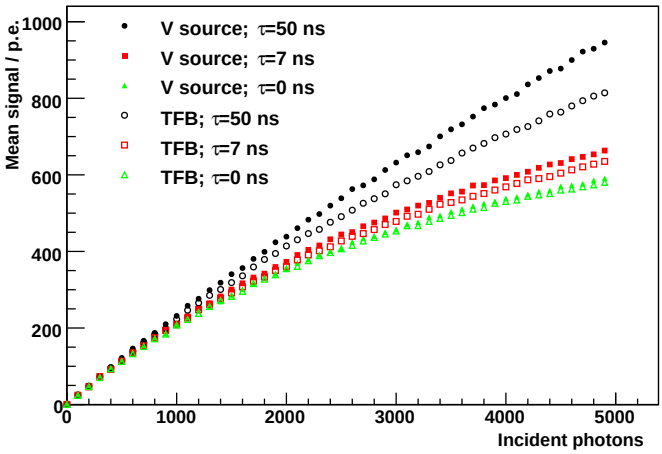

Figure 3.17: Plot showing the MPPC saturation predicted by the MC, for both the TFB recovery model and voltage source model.

physics, below around 1,000 photons, the response per photon is stable to about 10%. This level of variation is sufficiently large that a correction for saturation will be needed in the calibration to account for it (see Section 4.6).

3.6 Summary

The characterisation measurements made by the T2K-ND280 collaboration have led to a quantitative understanding of all MPPC properties which are empirically important at low light levels. The individual characterisation tests have been validated by incorporation into a complete MPPC Monte Carlo, which has been used to simulate the overall response of an MPPC exposed to low light levels and has been shown to agree very well with lab measurements. The MPPC simulation has also been used to give estimates of the timing and energy resolution of the MPPC, and to predict its response in the saturation regime.

Chapter 4

Testing and calibration of MPPC photosensors for the ND280 ECal

The downstream component of the ND280 ECal (DsECal) for the ND280 uses 3,400 MPPC sensors for readout. Since the T2K experiment is pioneering the large-scale use of these devices, and the DsECal was among the first set of subdetectors to be assembled, the UK T2K group has performed QA tests on all of these devices, prior to their installation in the calorimeter. In the present chapter, the method and results of these QA tests will be discussed. A similar analysis to that used in the QA is also used for in situ calibration of the MPPC gains, and initial results from this calibration, for data taken while testing the DsECal at CERN, will be presented.

4.1 The QA test stand and procedure

3717 MPPCs for the downstream component of the ND280 ECal were tested at Warwick (1820) and Imperial College London (1897), measuring the breakdown voltage, pixel capacitance, dark noise rate, correlated noise signal contribution, and relative PDE. The specified requirements for the QA, based on the requirement of having good energy resolution, and low enough noise to avoid large numbers of spurious detector hits, are given in Table 4.1. The primary purpose of these tests was to remove badly-performing devices, subject to these requirements, but they also provided useful data about the consistency of parameters between devices. This is important since significant variation in parameters between devices would complicate calibration of physics data.

Parameter	Requirement
Gain	7.5×10^5
PDE	>15%
Dark count rate	<1.2 MHz
Afterpulsing probability	<20%
Crosstalk probability	<5%

Table 4.1: Requirements for ECal photosensors.

4.1.1 Mechanical configuration

The test stand configuration is shown in Figure 4.1. The MPPCs were tested using the same plastic housings as in the final detector configuration (Figure 4.2). These housings include a small printed circuit board (PCB) to connect the MPPC pins to a miniature coaxial cable, and so by testing the whole assembly, the integrity

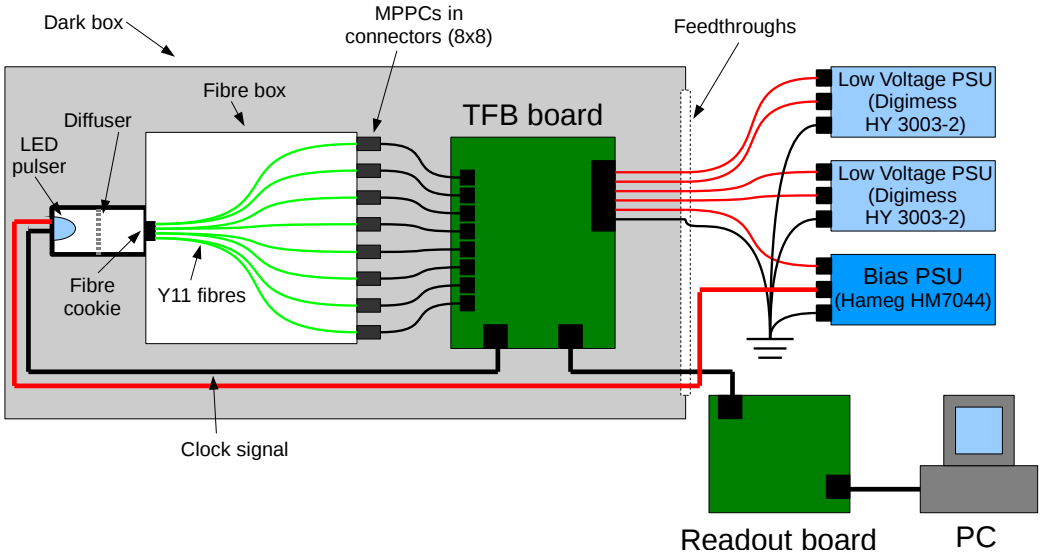

Figure 4.1: A schematic of the test stand used for MPPC QA in the UK.

of the electrical connections was also ensured.

Light was delivered to the MPPCs by the same Y11 fibres used in the detector. A pulsed LED array, triggered by a timing signal from the TFB, was directed through a diffuser onto bare fibre ends. These ends were glued into a plate (the "cookie") in a compact circular configuration, in order to minimise any non-uniformity in LED illumination. The fibres were strung through a junction box and terminated with ferrules, identical to those used in the final ECal design, glued into the front plate of the box. While gluing the fibres, the ferrules were clipped on to a gluing jig which ensured that the fibre protruded 200 μm from the ferrule to avoid light loss.

The ferrules were clipped onto the MPPC housings, which included a foam spring behind the MPPC to hold the MPPC's epoxy window in contact with the fibre, ensuring a good optical coupling. There were 64 fibres with ferrules in an 8x8

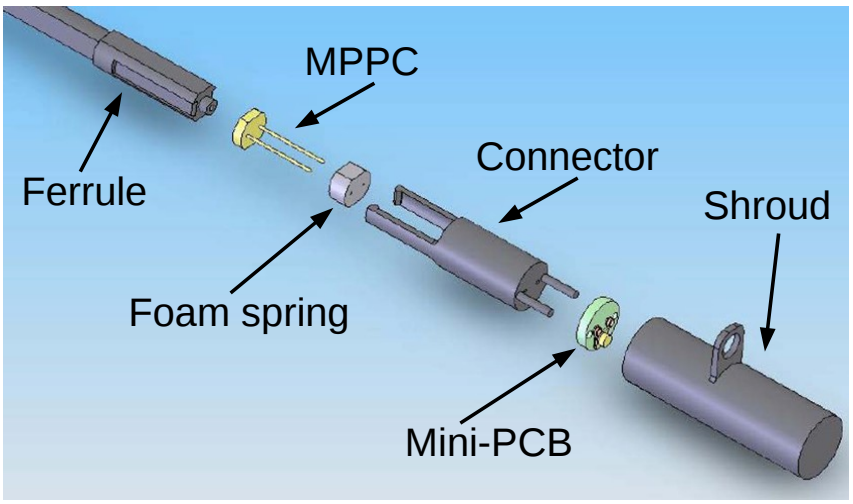

Figure 4.2: An exploded view of the MPPC and connector assembly used for the ECal and the QA setup.

configuration, so 64 sensors were tested at once. This is the maximum number that can be read out using a single TFB.

4.1.2 Power and DAQ systems

The MPPC bias voltage was provided by a Hameg HM7044 power supply. Scanning through a range of bias voltages was enabled by a downward voltage trim of up to 5 V, provided by the TFB and set using a byte value, giving a step of about 0.02 V corresponding to a gain step of around 1×10^4. The TFB is configured to integrate the charge on each channel 23 times in quick succession, to enable data to be taken for every J-PARC beam bunch, and for some time later to catch Michel electrons. For the QA tests the integration time for each bunch window, or timeslice, was set to 540 ns and the reset time between timeslices to 50 ns, matching quite closely the J-PARC beam timings (see Section 2.2). For each

timeslice, the TFB fed the charge on each channel into an amplifier and ADC, and the ADC value was stored.

The TFB channels were calibrated using a built-in charge injection mechanism, so that the ADC readings could be converted to a charge in Coulombs. Software provided by the ND280 DAQ group for the QA fitted and subtracted the pedestal (mean ADC value for zero pixels fired) and applied this charge calibration to the data. For readings taken with no LED, all timeslices were merged, and for LED data, only the timeslice containing the LED flash was used. The pedestal subtraction is important when combining timeslices, since the pedestals for different timeslices vary slightly, smearing the combined data sample if they are combined without pedestal subtraction. The final output for analysis consisted of histograms of calibrated charge measurements for each MPPC, for a range of bias voltages and with the LED on and off.

4.1.3 Golden MPPCs

A small number of MPPCs (labelled "golden" MPPCs) had their absolute PDE value measured as a function of bias voltage by the T2K group at Sheffield [64]. A calibrated optical power meter was used to measure the light intensity from a Y11 fibre in an ECal ferrule illuminated by a blue LED ($\lambda_{\mathrm{mean}} = 473\,\mathrm{nm}$). The optical power meter was then replaced with an MPPC clipped onto the ferrule in an ECal housing, and the output charge spectrum measured. A custom jig was used to locate the MPPC and optical power meter, and a simulation was used to calculate the acceptance ratio between the two detectors, so that the number of

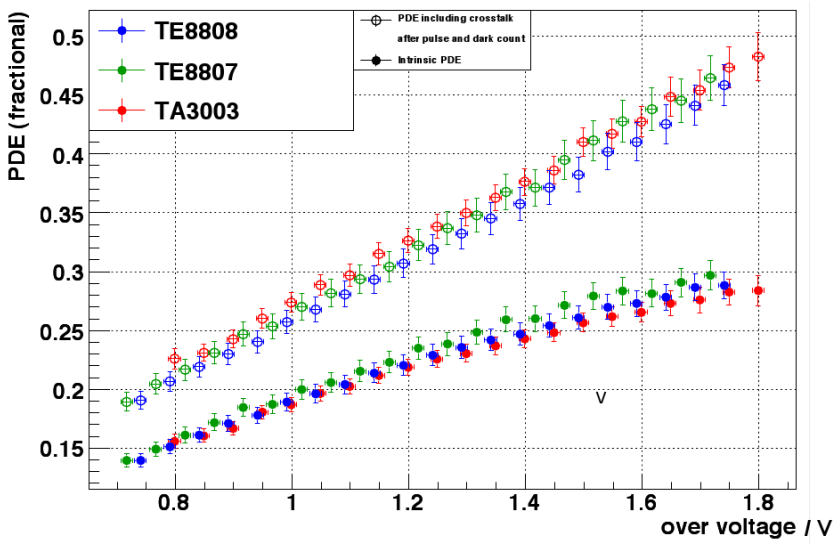

Figure 4.3: Measured PDEs for golden MPPCs TE8808, TE8807 and TA3003, calibrated at Sheffield. The lower curves, which were used for the ECal QA, have correlated noise effects removed.

photons incident on the MPPC could be calculated precisely.

The results of the Sheffield measurements are shown in Figure 4.3. The data was analysed in two ways — the first is the same as that used in the QA (see Section 4.2.3 below), and the second involved fitting the populations of peaks in the charge spectra to a Poisson distribution and taking the Poisson mean. The latter is sensitive to afterpulsing and crosstalk, so the former method, which gives the intrinsic device PDE, was used to obtain the PDE values for the purposes of the QA.

One of these golden devices was used in every QA run at Warwick and Imperial to measure the incident light level and provide an absolute calibration of the PDEs of the other devices. The QA procedure itself was only able to measure the PDEs of the MPPCs relative to that of the calibrated device, rather than their

absolute values.

4.1.4 Temperature monitoring and correction

Some MPPC parameters are a function of temperature. In particular, the breakdown voltage increases by about 0.06 V/K around room temperature, and the dark rate for a given overvoltage depends exponentially on the temperature, as shown by the results in Section 3.3.1. It was therefore necessary to monitor the temperature at which readings were taken, so that known temperature variations could be removed from results and uncertainties evaluated.

A temperature measurement, using an LM92 sensor[1] attached to the QA box and read out using the TFB, was made for each data point, corresponding to a single bias voltage value for LED or dark noise. However, it cannot be assumed that over short times the MPPC temperatures tracked the LM92 measurements, since they were enclosed in plastic connectors with poor thermal conductivity. For this reason, only the average temperature was used in the analysis of a run. Ambient air conditioning was used at Warwick to control the temperature at around 24 °C, but variations on a scale of up to $\sim 1^{\circ}\mathrm{C}$ over a run were still observed. However most runs had variations of less than half this, and the Imperial measurements were carried out in a temperature-controlled environment so were not subject to this problem.

Data was taken over a bias voltage range of around 2.5 V, but for each device, only about a third of data was used in the analysis, corresponding to a

[1] Rated accuracy, around 30 °C, of ± 0.33 °C [65]. Manufactured by National Semiconductor, 2900 Semiconductor Dr., P.O. Box 58090, Santa Clara, California, USA 95052-8090.

range of 0.8 V around the operating voltage of that device. The worst-case change in temperature over the analysis period was therefore about 0.33 °C, corresponding to a relative error in capacitance measurements of about $0.33\,^{\circ}\mathrm{C} \times \frac{\mathrm{d}V_{\mathrm{bd}}}{\mathrm{d}T}/0.8\,\mathrm{V} = 2.5\%$. Assuming that the temperature of the device in the analysis period was up to 0.5 °C away from the average measured run temperature, then once the error in the LM92 measurement is included, the error in the measured operating voltage is 0.036 V.

The measured dark count rate was corrected for temperature differences using an exponential formula based on the characterisation results. The change in dark rate due to a 0.5 °C change in temperature (at a base temperature of 25 °C) is around 3% — this is the worst-case error in dark rate associated with temperature instability. Because the point at which the device gain is nominal can be identified accurately regardless of temperature variations, temperature instability does not affect the measurement of other device properties.

4.2 Analysis procedure

The goal of the QA analysis was to extract the MPPC parameters (gain, dark noise rate, correlated noise level, and PDE) for each device at each voltage point, and use these to find the capacitance and breakdown voltage for each device, and its properties at the nominal operating point.

The analysis package was written in C++, and makes extensive use of the ROOT object libraries [66] for data analysis and I/O. The calibrated charge histograms for MPPCs show clear peaks corresponding to different numbers of

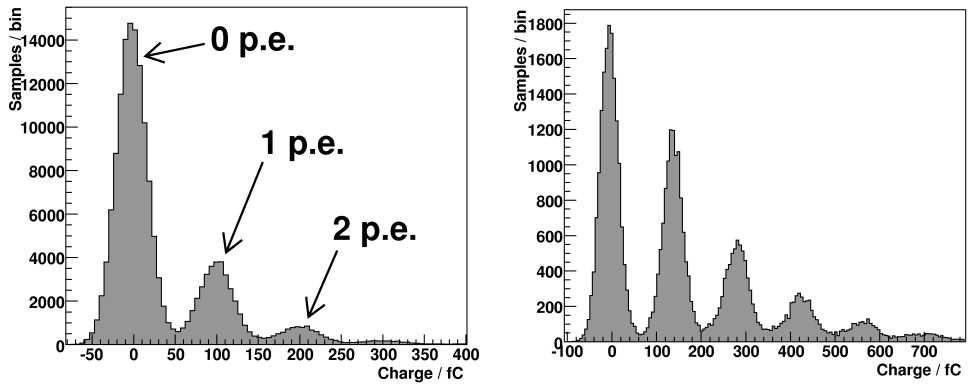

(a) Dark noise only (b) With LED pulse

Figure 4.4: MPPC charge spectra from the QA test stand, with and without the LED pulser on. Dark noise runs have higher statistics as all timeslices can be used.

avalanches (referred to as pixel equivalents or p.e.) in the MPPC, as in Figure 4.4, and analysis of these peaks is the basis of the parameter extraction methods. Data from the dark runs is dominated by the 0th peak since any signal comes only from dark noise. Higher peaks are visible in the LED data — the LED brightness was tuned to give $\sim$1 p.e. per pulse.

4.2.1 Gain

The MPPC gain is given by the separation between adjacent peaks in an MPPC charge histogram, if the charge values are calibrated. For the QA, the peaks in the histograms were found using the TSpectrum ROOT class, which is a C++ implementation of Morhac et al.'s algorithm [67]. This algorithm uses smoothed second derivatives to separate peaks from background, assuming that the local form of the background can be described by a linear function. It finds peaks by looking in the histogram for regions where the second derivative profile matches

that expected for a Gaussian. It estimates the peak locations, but does not fit them. The algorithm is rather robust, but occasionally associates two or more peaks with a single real peak in the histogram. Since a single mis-interpreted histogram can lead to large errors in extracted device properties, it was necessary to identify these mistakes and either correct for them or reject the spectrum. This was accomplished by looking for pairs of peaks much closer together than the rest, and replacing the pair with an average value. Only the first three peaks were used in gain calculations, and a spectrum was rejected if the second peak to be analysed was not within the central 25% of the region between the other two peaks.

If the first three peaks were located successfully, a localised Gaussian fit was performed on each peak. The fitted peak parameters were then used as initial parameter values for a global fit to a sum of three Gaussians. This global fit constrained the separation between each pair of adjacent peaks to be the same. The fitted widths of the peaks, σ_i, were also constrained to follow the pattern

$$\sigma_i = \sqrt{\sigma_0^2 + i \times \sigma_{\mathrm{av}}^2}, \tag{4.1}$$

where i is the number of fired pixels, σ_{av} is the additional width (quadratically) contributed by each avalanche, and the width of the 0 p.e. peak, σ_0, is finite due to electronics noise. σ_0 and σ_{av} were free parameters in the fit. The fitted peak separation in fC, divided by the electron charge e, gives the device gain. This fitting method was chosen on the assumption that the level of electronics noise is independent of the MPPC output, and that the spread in MPPC output charge for each avalanche does not change with the number of avalanches produced. Figure 4.5 shows an MPPC charge spectrum with the peak finder and fitter results

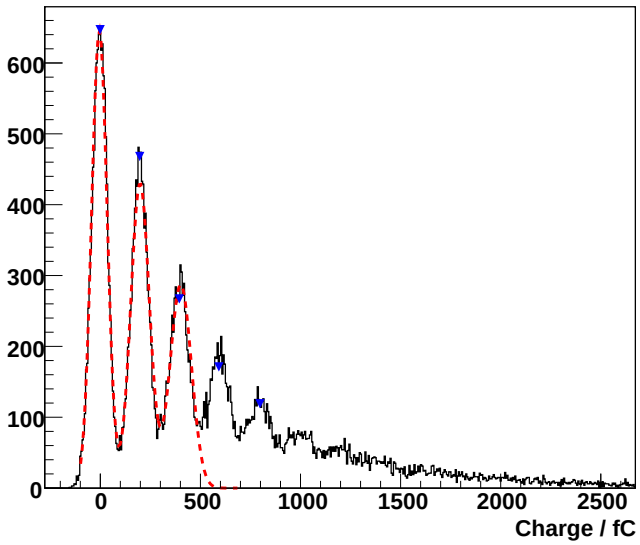

Figure 4.5: An MPPC charge spectrum with peak finder and fitter results marked, in blue and red respectively. Only the first three peaks are included in the fit.

marked.

4.2.2 Dark noise and correlated noise

In the absence of correlated noise (afterpulsing and crosstalk), the populations of the charge spectrum peaks, with LED off, should follow a Poisson distribution with mean $\lambda =$ dark count rate $\times$ gate length. Correlated noise distorts this spectrum by adding extra avalanches to events, but since there can be no afterpulsing or crosstalk without an initial trigger, the population of the zeroth peak $P(0)$ (i.e. the probability of no pixels firing) is unaffected. The intrinsic dark rate (DCR) was therefore calculated from the population of this peak only:

$$DCR = \frac{-\ln(P^{\text{dark}}_{\text{data}}(0))}{\text{gate length}}. \tag{4.2}$$

Conversely, the contribution to the overall noise rate from afterpulsing and crosstalk, $F_{\mathrm{AP+CT}}$, can be calculated from the difference between the total noise signal in the data, and that expected from a Poisson distribution:

$$F_{\mathrm{AP+CT}} = \frac{\sum_{i=1}^{\infty} i \times [P_{\mathsf{data}}(i) - P_{\mathsf{pois}}(i)]}{\sum_{i=1}^{\infty} i \times P_{\mathsf{data}}(i)}. \tag{4.3}$$

Here $P_{\mathrm{data}}(i)$ is the measured population of the ith peak, and the summation $\sum_{i=1}^{\infty} i \times P_{\mathsf{data}}(i)$ is the sum of the peak populations, weighted by the number of p.e. which each peak corresponds to — essentially the total number of avalanches over the entire data sample. $P_{\mathrm{pois}}(i)$ is the population of the ith peak, as predicted from the Poisson distribution defined by $P_{\mathrm{pois}}(0) = P_{\mathrm{data}}(0)$. The corresponding sum, $\sum_{i=1}^{\infty} i \times P_{\mathsf{pois}}(i)$, is therefore the total number of avalanches caused by intrinsic dark noise. The numerator in (4.3) is the difference between these two sums, i.e. the number of avalanches due to afterpulsing and crosstalk. The denominator normalises $F_{\mathrm{AP+CT}}$ to the total number of avalanches observed.

In practice, calculating the populations of the higher peaks was difficult, since the presence of afterpulses with lower gain (due to incomplete voltage recovery), and the clipping of some avalanches which fall near the edge of the gate period, caused the peak shapes to deviate from Gaussians. Also, peaks due to more than a few p.e. were smeared into a continuum, and could not be individually fitted. The raw mean value of the histogrammed data, subtracting the zeroth peak position and dividing by the peak separation, was therefore used to give the average number of pixels fired, and compared with the mean of the Poisson distribution. This gave the proportion of the average signal (continuous) which is due to crosstalk and afterpulsing, rather than the proportion of avalanches (discrete);

the former is slightly smaller due to the lower gain of some afterpulses. The QA configuration did not permit separate measurements of crosstalk and afterpulsing — in principle this could be done using integrated charge measurements by moving the LED pulse around within the gate period, but in practice the time required to take the data would have been prohibitive, and the precision achieved by the analysis was not sufficient to make this measurement useful in any case.

4.2.3 PDE

The device PDE, for low light levels, is given by the average number of light-induced avalanches for a gate period, divided by the average number of incident photons for the same period. As discussed in Section 4.1.3, a calibrated device was included in each test run, which was be used to measure the number of incident photons and allow the calculation of the other devices' PDEs.

The average number of pixels fired by an LED pulse, $\bar{N}_{\mathrm{light}}$, can be found by using the same Poisson zero probability method as for the dark rate. We must subtract off the average number of dark noise pulses, so a dark noise spectrum is also required for this calculation. The zero probability for an LED spectrum, $P^{\mathrm{LED}}_{\mathsf{data}}(0)$, is the probability that we have no light induced avalanches, and also no dark noise:

$$P^{\mathrm{LED}}_{\mathsf{data}}(0) = P^{\mathrm{light}}(0) P^{\mathrm{dark}}_{\mathrm{data}}(0). \tag{4.4}$$

$P^{\mathrm{LED}}_{\mathsf{data}}(0)$ and $P^{\mathrm{dark}}_{\mathsf{data}}(0)$ are measured quantities, so we have enough information to

calculate $\bar{N}_{\mathrm{light}}$:

$$\bar{N}_{\mathsf{light}} = -\ln P^{\mathrm{light}}(0) = -\ln\left(\frac{P^{\mathrm{LED}}_{\mathsf{data}}(0)}{P^{\mathrm{dark}}_{\mathsf{data}}(0)}\right)$$
$$= -\ln P^{\mathrm{LED}}_{\mathrm{data}}(0) - \mathsf{DCR} \times \mathsf{gate\ length}. \qquad (4.5)$$

$\bar{N}_{\mathrm{light}}$ was calculated using (4.5), and used to calculate the PDEs of the MPPCs by dividing by the calibrated light level from the golden MPPC. Some variation in light level was observed between channels, which can be attributed to variation in LED intensity across the fibre array, in fibre length and bending radii, and in fibre/ferrule gluing and ferrule/MPPC connection. These variations were accounted for by a self-calibration procedure — the relative light level in each channel was adjusted so that the mean PDE over all QA runs (with different sensors) was the same. Variations in the ferrule/MPPC coupling between runs for the same channel were not corrected for by this procedure, so these variations contribute to the spread in PDE values obtained. The self-calibration procedure was performed using the first batch of sensors only, so the second batch is statistically independent from that used to calibrate the test stand. The number of runs used for the self-calibration was 14 (Warwick) and 17 (Imperial), so the error in the channel-channel calibration was around $\sqrt{15} \simeq 4$ times smaller than the spread in the PDEs of individual devices, and therefore did not significantly affect the spread in observed PDE values.

Unfortunately, the golden MPPCs were left in the same QA box channels for the whole of the first batch at Warwick, and for both batches at Imperial. This means that the light level of the golden MPPC channel relative to the other channels was not measured for these batches — this is further discussed in Section

4.5.4.

4.3 Validation of the analysis code

The analysis routines used for the QA testing have been applied to data from the MPPC Monte Carlo simulation (see Chapter 3), in order to ensure that they can reproduce the simulation input parameters. This was a useful validation of both the simulation and analysis packages.

The simulation was run at 25 °C, with a gate length of 540 ns. The number of incident photons was generated on a Poisson distribution, and the photons were distributed exponentially with a time constant of 10 ns, near the start of the gate. The simulation models device behaviour quite closely, but it does not model the finite time-domain distribution of charge from an avalanche, so "clipping" of avalanches which occur near the gate edges does not occur.

The results of the analysis are shown in Figure 4.6. The analysis achieves excellent agreement with the input parameters for gain, dark count and PDE. The dark count plot shows simulation data with and without a "pre-gate" period simulated — adding this period gives more realistic performance because afterpulses due to dark counts before the gate start will be included. These orphaned afterpulses look like dark noise pulses, and so lead to an overestimation of the dark noise rate. This effect will be present in QA data.

It is more difficult to validate the performance of the analysis in measuring crosstalk and afterpulsing, since the number calculated by the analysis ($F_{\mathrm{AP+CT}}$, the proportion of total signal due to CT+AP) is not a simple function of the input

parameters to the simulation. Because of this, it is best to interpret $F_{\mathrm{AP+CT}}$ as a figure of merit to measure the general stability of correlated noise levels between devices, rather than a direct measurement of a specific device property. It is clear, from the deviation of the points from a smooth curve, that there is a significant error in the calculation of $F_{\mathrm{AP+CT}}$. It is not surprising that the error in this quantity is larger than for the other parameters, since, as described in Section 4.2.2, it is dependent on the difference between the raw histogram mean and Poisson mean used to calculate the dark rate, so for $F_{\mathrm{AP+CT}} = 0.2$, an error in DCR of 1% leads to an error in $F_{\mathrm{AP+CT}}$ of about 4%. Since the histogram mean is also pedestal subtracted and divided by the peak separation, it also depends sensitively on the gain fit.

4.4 Test results

4.4.1 Capacitance and breakdown voltage

The capacitance and breakdown voltage of each MPPC were calculated by fitting the measured gains to Eq. (3.1). The required operating voltage for a gain 7.5×10^5 was calculated in the same way. The latter was also measured by the manufacturer for each device before dispatch. The QA operating voltage measurements, corrected for temperature difference, are compared to the manufacturer's numbers for both Warwick batches in Figure 4.7(a). There is a clear discrepancy which, since it is different for the two batches, is likely to be due to a problem with our test stand rather than the manufacturer's measurements. It is plausible

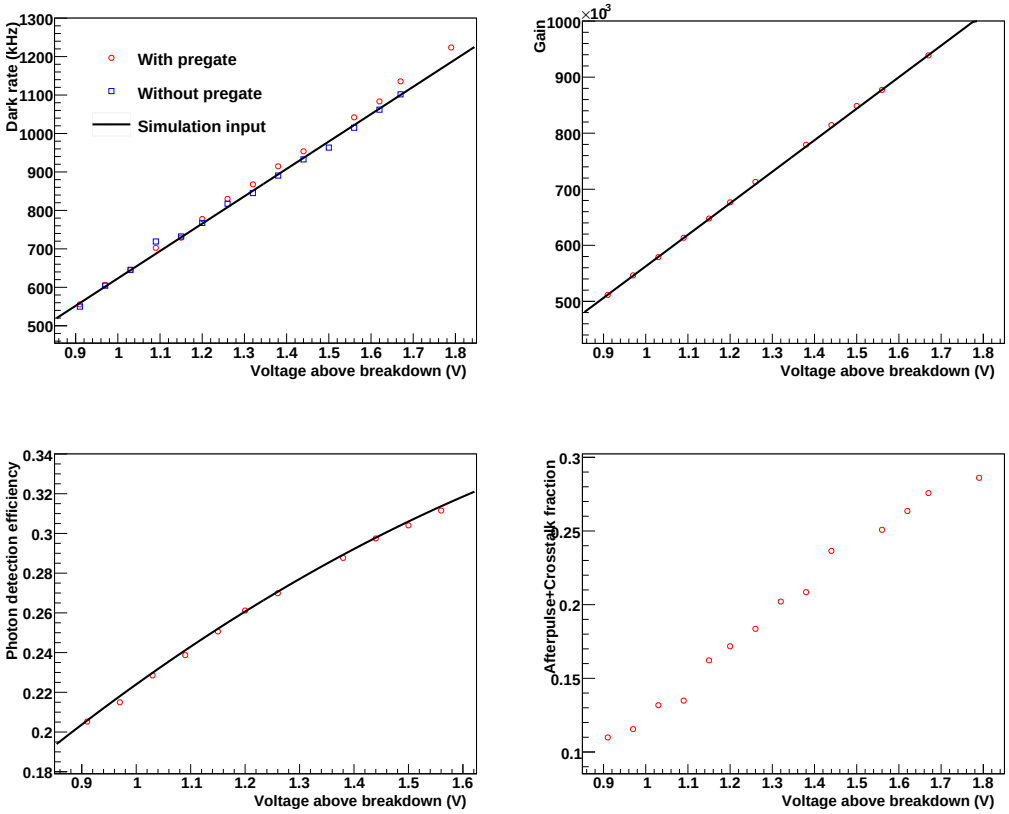

Figure 4.6: Results from the QA analysis of Monte Carlo data. Simulation input parameters are included for comparison where appropriate. Statistical errors are negligible (smaller than data points).

that the two batches are different — they were tested using different amplifier gain settings, approximately halving the number of ADC bins/p.e. for the second batch, since the gain to be used for ECal operation was decided by the Electronics Group between the testing of the two batches. Unfortunately time did not allow to test samples from either batch with the settings used for the other batch.

The measured capacitances are also different from the manufacturer's quoted value of 90 fF. The capacitances are a function of channel number, as shown in Figure 4.7(b) — this suggests that there is a channel-dependent scaling error associated with the charge-injection calibration of the TFB amplifier and ADC. This is unsurprising since a different capacitor, with a tolerance of 5%, is used to inject charge into each channel. This tolerance is compatible with the observed channel-to-channel variation.

The scaling between batches and channels has been corrected for in the data by applying a per-channel scaling to all gain measurements (one number per channel per batch) such that the mean measured capacitance for that channel matches the manufacturer's quoted value. It was then possible to estimate the stability of capacitances between devices, but an absolute measurement of the device capacitance has not been made. The corrected gain values were used to calculate device operating voltages since this allows data from all channels and batches to be compared.

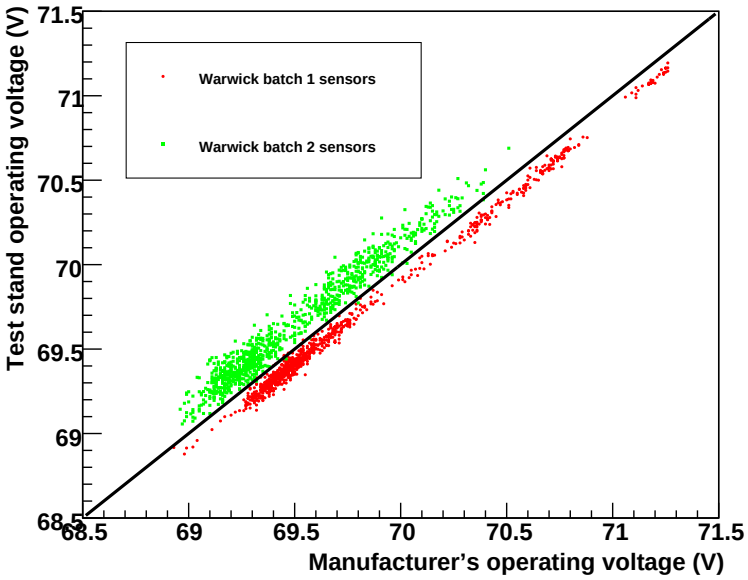

(a)

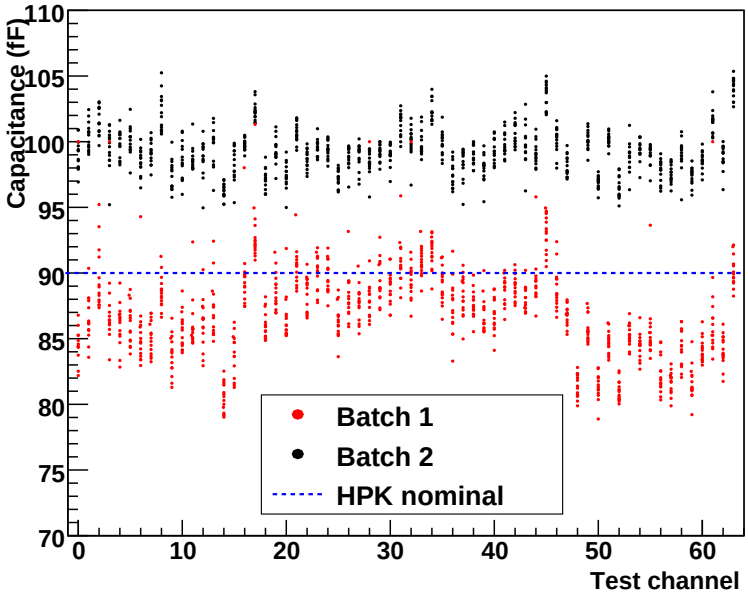

(b)

Figure 4.7: (a) Comparison of the operating voltages quoted by the manufacturer, and those measured on the test stand. (b) Capacitance measurements for both MPPC batches at Warwick, for each test channel.

4.4.2 Parameter values and stability

The gain-corrected results of the QA, for both Warwick and Imperial, are shown in Figure 4.8. The PDE, DCR and AP+CT fraction (estimated by $F_{\mathrm{AP+CT}}$) are quoted at the nominal gain. It can be seen that there are differences in the values of, and spreads in, several parameters between batches and institutions. As mentioned above, the TFB settings changed between batches, which affected the number of ADC bins per fC, and the ancillary hardware (specifically power supplies) was different at Warwick and Imperial. Also, fewer voltage points (one sample per 0.08 V, rather than one every 0.04 V) were taken for the second batch at Imperial due to time constraints — the discrepancies between batches may therefore be caused by these effects rather than a real difference in device properties. The reduced statistics for the second Imperial batch mainly affect the PDE spread — this is because with fewer voltage points, failures in fitting the LED spectra have a serious effect on the number of data points used to fit the PDE.

It is useful to consider the results for the "golden" MPPCs. Since the parameters for the golden MPPC at each institution should not change, the standard deviation in the parameter measurements for this device should give an estimate of the error associated with the measurement of each parameter, and hence whether the spread in values obtained is due to a physical difference between devices, or to measurement error or bias. The results for these devices are shown in Figure 4.9.

Table 4.2 shows the mean parameter values and their standard deviations, for each QA batch, for all devices and for the golden MPPCs only. The results, and the conclusions which can be drawn from them, are discussed below.

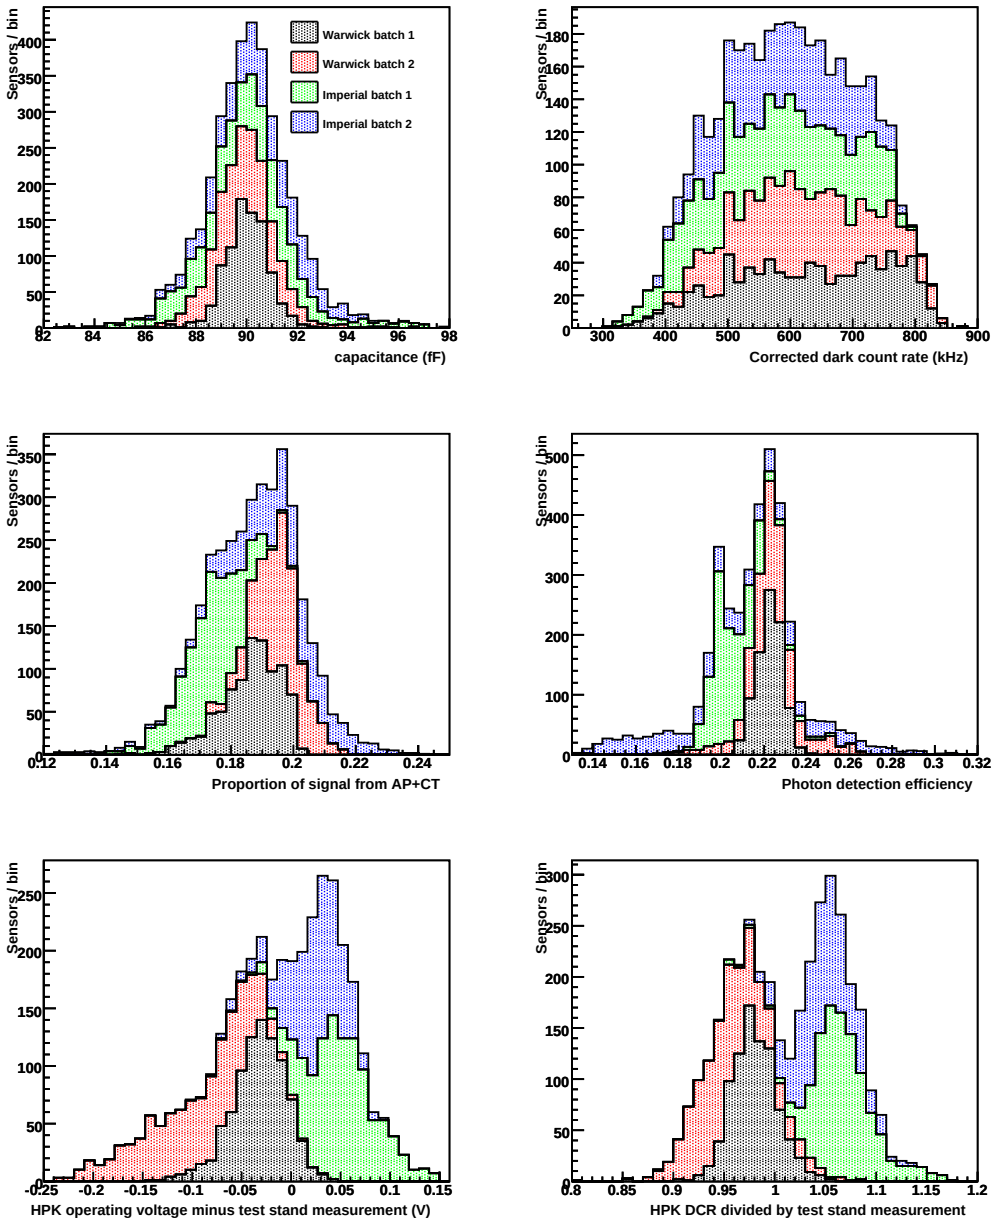

Figure 4.8: Distribution of device parameters as measured on the QA test stands. Data is for nominal (7.5×10^5) gain and extrapolated to $T = 22\,^{\circ}$C (except for comparisons with HPK which are at $T = 25\,^{\circ}C$). The histograms are stacked, i.e. the total bar height in each bin is equal to the sum of the bin values for all batches, and the bar is divided into coloured sections representing the contribution from each batch. The plot comparing DCR measurements with the manufacturer uses the uncorrected DCR, which includes correlated noise effects.

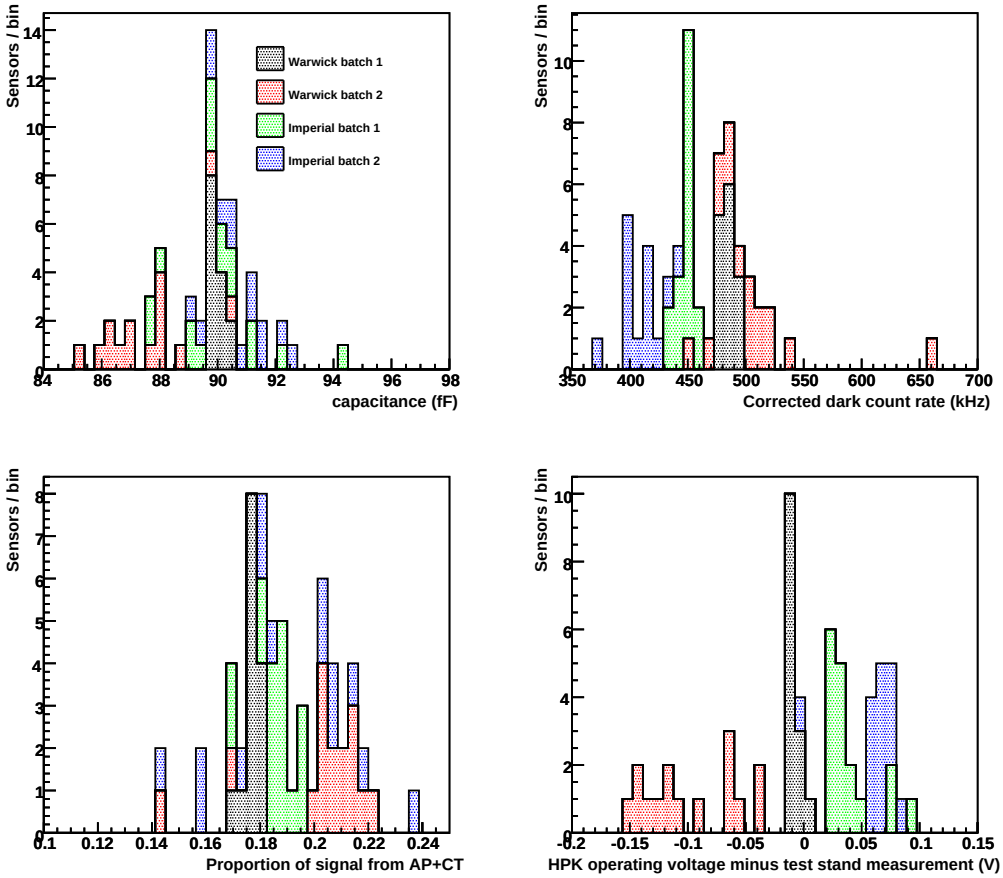

Figure 4.9: Distribution of device parameters for the golden MPPCs only, as measured on the QA test stands. Data is for nominal (7.5×10^5) gain and extrapolated to $T = 22\,^{\circ}C$ (except $V_{\rm op}$ comparison, which is at $T = 25\,^{\circ}$C). The histograms are stacked. Since a single golden MPPC was used in each run, the histograms have one entry per run.

	Warwick Batch 1 (W1)	Warwick Batch 2 (W2)	Imperial Batch 1 (I1)	Imperial Batch 2 (I2)
Capacitance (fF)	90.0±0.85 (90.00±0.25)	89.9±1.41 (87.49±1.39)	90.0±2.35 (90.01±1.58)	90.0±1.80 (90.00±1.98)
$V_{\rm op}^{\rm HPK} - V_{\rm op}^{\rm meas}$ (V)	−0.032±0.028 (−0.010±0.005)	−0.098±0.050 (−0.097±0.038)	0.050±0.034 (0.055±0.021)	0.022±0.029 (0.064±0.021)
Dark Rate (kHz)	618.4±120.4 (483.7±5.6)	610.6±100.5 (506.6±44.1)	573.4±112.9 (448.2±7.7)	587.6±93.0 (409.7±15.9)
$DCR^{\rm HPK}/DCR^{\rm meas}$	0.968±0.023 (0.989±0.013)	0.941±0.032 (0.934±0.058)	1.040±0.031 (0.989±0.021)	1.023±0.027 (1.108±0.013)
AP+CT Fraction	0.186±0.010 (0.177±0.003)	0.196±0.008 (0.203±0.019)	0.173±0.010 (0.186±0.008)	0.190±0.019 (0.180±0.033)
PDE	0.222±0.006	0.224±0.016	0.205±0.011	0.203±0.040

Table 4.2: Mean measured MPPC parameters for each test batch. Numbers shown after ± are standard deviations, not errors. The results for the golden MPPCs are shown in brackets. The nominal PDEs measured at Sheffield for the golden MPPCs were 0.222 and 0.205, for Warwick and Imperial respectively.

4.5 Discussion of QA results

4.5.1 Capacitance and operating voltage

The mean capacitances for each batch are set to 90 fF by construction (batch W2 is very slightly different as the golden MPPC was moved between channels on each run). For all batches except W1, the measured spread in capacitances is comparable to the spread observed for the golden MPPC; this indicates that the observed spread represents only an upper limit, since it may be entirely due to measurement error. The spread observed in the first batch tested at Warwick is considerably larger than that seen for the golden MPPC, so the true spread in device capacitance, $\sigma_{\rm gain}^{\rm MPPC}$, can be estimated, from the standard deviation for the batch and the golden MPPC, as

$$\sigma_{\rm gain}^{\rm MPPC} = \sqrt{\left(\sigma_{\rm gain}^{\rm W1}\right)^2 - \left(\sigma_{\rm gain}^{\rm golden}\right)^2} = 0.81\,{\rm fF}. \tag{4.6}$$

Although the spreads are expected to be different due to variations in running conditions, there is no particular reason why we would expect batch W1 in particular to have the smallest errors.

As stated above, it has not been possible to make an absolute measurement of device capacitance; however the results from batch W2 suggest that the capacitance of the golden MPPC (which was taken from batch W1) may be a little smaller ($\sim$3%) than that of the devices in W2. Since this is only a single device, however, it is not sensible to draw strong conclusions about any drift in capacitances between batches — such a conclusion would require the testing of several devices from each batch alongside each other. It is also possible that the capacitance of the golden MPPC is simply quite low compared to other W1 devices; this is compatible with the observation that the operating voltage (relative to the HPK measurement) for the golden device in W1 is measured to be lower than the average, which could be due to an erroneous upscaling of the device capacitance to the assumed mean value.

The spread in measured operating voltages (relative to HPK) is fairly small for all batches, though a discrepancy does exist which changes sign between institutions. The absolute voltage at the MPPC pins is not known — the values used are calculated from the PSU readout and nominal TFB trim voltage — and given that different models of PSUs, and different TFBs, were used at the two institutions, a difference on this scale is not surprising. The LM92 calibration error (see Section 4.1.4) also introduces a temperature uncertainty equivalent to $\pm 0.2\,\mathrm{V}$ for each institution. The error in the manufacturer's measurement is not known,

and some spread in measurements may come from this source, since the spreads for the golden MPPCs were smaller than for the corresponding batches.

4.5.2 Dark rate

The dark noise rate for given gain is known to vary by a factor $\sim$2 between devices, so the large spread in values obtained for all batches is as expected and in agreement with the manufacturer's data. It is therefore more interesting to compare our measurements with the manufacturer, to see how closely they agree.

The manufacturer measures the dark rate using a discriminator and counter — this just measures the number of peaks above a certain threshold (set to 0.5 p.e.) in a given period of time. This means that crosstalk, which increases the height rather than the number of peaks, does not affect the measurement, but afterpulsing, which generates additional peaks, increases the measured rate. We would therefore expect that the rate measured by Hamamatsu would be slightly lower than the total dark rate measured on the test stand, with afterpulses and crosstalk included. This is true for the Warwick data, but not the Imperial data.

The mean temperatures for the Warwick and Imperial test stands were around 23.5 °C and 22.0 °C respectively, corresponding to a difference in dark rates of about 10%. This difference was corrected for in the analysis, but some residual error may be introduced from the temperature scaling which was used (from the Warsaw DCR results). The temperature uncertainty from the LM92 calibration ($\pm$0.33 °C) introduces a systematic error of around 2% to each institution's measurement. These components, taken together, can plausibly explain

the discrepancy between the institutions.

The spread in the results for each batch is similar to the spread for the golden MPPC, so it is likely that it is due to errors in the test stand measurement rather than in the manufacturer's tests.

4.5.3 Afterpulsing and crosstalk

The proportion of the total noise signal due to afterpulsing and crosstalk, $F_{\mathrm{AP+CT}}$, is fairly stable, but varies a little between batches and institutions. It is possible that the AP+CT level changed between batches — since afterpulsing depends on impurities in the silicon, a change in the density of impurities will affect the afterpulse level. However, the observed spreads within the batches are of a similar scale to the difference between batches, so batch-batch stability is not a big issue.

The batch spread in $F_{\mathrm{AP+CT}}$ is on the same scale as the spread in golden MPPC measurements except for W1. Hence, as with the capacitance above, we can estimate the intrinsic spread in $F_{\mathrm{AP+CT}}$ within this batch as 0.095 or 5.1%. Given that the spreads in the other batches appear to be dominated by measurement error, there is no reason to assume that these have a larger intrinsic spread, though we have not directly measured it.

4.5.4 PDE

Because the PDE of the golden MPPC was used to calibrate the incident light level, its PDE cannot be independently measured, and so we cannot use it to get an estimate of the measurement error as with the other parameters. In fact,

because the results from batches W1 and I1 were used to self-calibrate the channel-to-channel light level variations in the QA box, and the golden MPPC was kept in the same channel throughout, the mean PDE for these batches agrees with the nominal PDE of the golden MPPC by construction. This accounts for the difference between the Warwick results and batch I1, since the PDE of the golden MPPC used at Imperial (as measured at Sheffield) was lower. The absolute error in the PDE calibration of the golden device used at Warwick (at nominal overvoltage) was 3%, which is easily enough to explain the discrepancy.

The Warwick results are very stable between batches, and show a fairly small spread in PDEs. The number for batch W2 is more informative, since it is statistically independent of the data used to calibrate the channel light levels. This suggests a spread in PDEs of 7.1% (relative). This is an upper limit on the spread in MPPC intrinsic PDEs, since it includes the variation in optical coupling between the fibre end and MPPC, and also an unevaluated measurement error. A short study was conducted at Warwick to measure the reproducibility of the optical coupling, using an ECal connector and ferrule, Y11 fibre and a blue LED, and reading out the MPPC current with an ammeter. The current (corrected for dark noise) for a single MPPC, was found to be reproducible with a standard deviation of 2.3%. The optical coupling effect may therefore make some contribution to the spread observed in the QA.

The spread in PDEs for batch I2 is rather larger than for any of the other batches. As mentioned above, the number of data points in this batch was smaller than the others, and the binning for the second batches was coarser at both

institutions, leading to more fitting failures particularly in LED spectra. This loss of statistics may account for this spread. It is unlikely that it is due to an MPPC defect since the I2 batch is sampled from the same group of devices as W2, which does not exhibit such a spread. No overall bias in the PDE is seen for this batch relative to I1.

4.5.5 Conclusion

The MPPC QA tests for the downstream ECal were successfully completed at Warwick and Imperial and no devices were found which failed to function. The measured device parameters were compared to the requirements in Table 4.1. Some of the I2 sensors were measured to have a lower PDE than the 15% requirement, but these were not rejected, since there is good reason to believe this is due to measurement uncertainty. As discussed in Section 4.2.2, it was not possible to measure afterpulsing and crosstalk separately, but adding the allowed crosstalk and afterpulsing values together we obtain $F_{\mathrm{AP+CT}}$ <25%, which is satisfied by all devices. All devices were also found to have a DCR lower than 1.2 MHz at 22 °C, even when correlated noise is included. Since DCR is the only parameter which varied significantly between devices, and only 3,400 devices of the 3,717 tested were required for the DsECal, the devices with the lowest noise rate were used.

Because of issues with the calibration of the front-end board amplifiers, it was not possible to make an absolute measurement of device capacitance, but the spread in device capacitances has been measured (albeit using only batch W1) to be around 1%. Similarly, data from the same batch sets the spread in the AP+CT

fraction at around 5%, and data from W2 estimates a maximum spread in PDE of 7.1%. The spread in PDE is quite large, but this will be calibrated out as part of the "MIP-scale" calibration (see Section 4.6). The effect on saturation of delayed afterpulses will not be corrected for at a bar-by-bar level, but the difference between sensors, for a 5% difference in correlated noise, will be very small in the expected region for hits in the ND280 physics data ($N_{\mathrm{fired}} < 100$). The QA results, therefore, do not suggest that any calibration, beyond a per-channel gain correction and hardcoded saturation correction, are required for the MPPC.

While the QA, due to the constraints imposed by limited manpower and time when testing large numbers of devices, was not able to measure device properties with the same precision or detail as the dedicated characterisation measurements described in Chapter 3, it has provided useful information on the stability of parameters between devices and batches. Even where (as is the case with much of the QA data) the spread in results can be largely attributed to measurement error, an upper limit on the variability of parameters has been obtained, which gives a useful insight into the stability of the manufactured devices.

4.6 Calibration of ND280 data

In order to use the data from the ND280 scintillator detectors for analysis, it is first necessary to calibrate it. The purpose of the calibration is to convert the raw hit data from the front-end boards, in ADC counts, to an estimate of the energy deposited in the scintillator bars by through-going particles. Logically, this process consists of the following steps:

1. Subtract the electronics pedestal value (mean ADC for zero pixels fired) from the raw ADC. The pedestals vary between channels and timeslices and must be remeasured regularly.

2. Convert the pedestal-subtracted ADC value into a charge in Coulombs. This requires that an ADC-to-charge mapping be built by injecting charge via a dedicated capacitor on each channel.

3. Convert the charge in Coulombs into an equivalent number of pixels by dividing by the MPPC gain.

4. Convert the number of pixels to a number of photons. This mapping is not quite linear since there is some saturation at high light levels — the saturation correction requires an empirically-determined function, which will be developed using test stand measurements since it is not practical to measure it in situ for all devices. Until measurements are available, results from the MPPC Monte Carlo will be used. The linearity mapping is a function of gain via the PDE and noise probabilities.

5. Correct the photon number for the attenuation in the Y11 fibre. The attenuation as a function of distance has been measured at Warwick, but in order to perform the correction, it is necessary to know how far from the sensor the particle hit the bar. For detectors with crossed layers this can be found by looking at hits in the other view, but for the P0D ECal it is necessary to match the hits with tracks in another detector.

6. Convert the photon number to a charge deposit. Charge deposits are mea-

sured in "MIP Equivalent Units" — the amount of charge left by a minimum-ionising particle passing straight through the bar on a line from the front to the back of the detector. The number of photons produced by a MIP can be measured by looking at cosmic ray data, and this must be done for each bar separately. This measurement will be made using in situ cosmic data, and will be updated periodically.

This is rather close to what is done in practice, but there are a few points which should be remarked upon. Firstly, as discussed in Section 4.4.1, the linearity correction (2) carries a scaling error because of the tolerance on the charge injection capacitor. The Coulomb charge is therefore better thought of as a "linearised ADC" value which is proportional to the MPPC output charge. Because the MPPC gain used in (3) is calculated in terms of these linearised ADCs, the capacitance error cancels out between steps (2) and (3). Similarly, although the form of the linearity correction in (4) is similar for each MPPC, the PDE will vary between channels (due to fibre-MPPC coupling as well as device variability), but this unknown scaling will be corrected for in the MIP-scale correction in (6), since a device with lower PDE will appear to see fewer photons per MIP. Finally, although the correction in (5) is logically prior to (6), it is applied later on since it requires reconstruction-level data (positions of other hits) to perform, while the others use only stored calibration constants and single hit-level data. This does not affect the result since (5) and (6) are simple multiplicative corrections.

4.6.1 MPPC calibration

We will now focus on the steps in the calibration which relate directly to the MPPC photosensors. In order to predict the MPPC gain for a given run, the terms appearing in (3.1) — the device capacitance C and breakdown voltage V_{bd}, and the applied bias voltage V_{bias} — must be known. We also need to know the device temperature T, if the breakdown voltage was measured at a different temperature to that for the run being calibrated, so that V_{bd} can be corrected. The run temperature is monitored using LM92 sensors mounted on the detector bulkheads.

The device parameters C and V_{bd} are calculated in the same way as for the QA (see Sections 4.2.1 and 4.4.1), by taking dedicated runs which scan over a range of bias voltages, and measuring the gain at each voltage by observing the peak separation. In practice, it should not be necessary to perform these voltage scans very frequently, since the device gain can be measured periodically by taking some unsparsified data (where data for all channels is output without the threshold which is normally applied) in normal running, and using this to monitor the gains to check that they agree with the stored device properties. This is advantageous because scanning over device bias voltages is disruptive to normal running, and can only be carried out when no beam is available — this may mean doing voltage scans weekly, or even less frequently.

4.6.2 Testbeam results

The MPPC calibration has been tested using data from the Downstream ECal testbeam. A voltage scan run, taking data at seven bias voltages around nominal, was used to calculate the MPPC parameters, and these were used to calibrate some data taken several days later, after the module had been moved to a new position. The actual gain for the running period was monitored using unsparsified runs taken between physics data runs. The runs were taken over a few days, so the gains varied due to diurnal temperature variation (the applied bias voltages stayed constant for all runs).

Figure 4.10 shows the mean measured channel gain (for all MPPCs in the ECal) for every run against time in seconds — the peaks and troughs are due to temperature changes. By applying the MPPC calibration, including a correction for the run temperature, it is seen that the difference in gains can be compensated for. Figure 4.11(a) shows, for a single testbeam run, the difference between the measured gain for each MPPC, and that predicted by the MPPC calibration based on the device temperature. The measured spread in this difference corresponds to around 0.6% of the absolute gain value, equivalent to an error in V_{bd} of 8 mV, or an error in temperature of 0.13 °C. Figure 4.11(b) shows, with one entry for each MPPC, the spread in the MPPC gain over the sample of runs plotted in Figure 4.10, before and after T correction. The standard deviation after correction is on the same scale as the spread seen in Figure 4.11(a) for a single run, indicating that the temperature correction is effective on a channel-by-channel basis.

The second part of the MPPC calibration, converting the number of fired

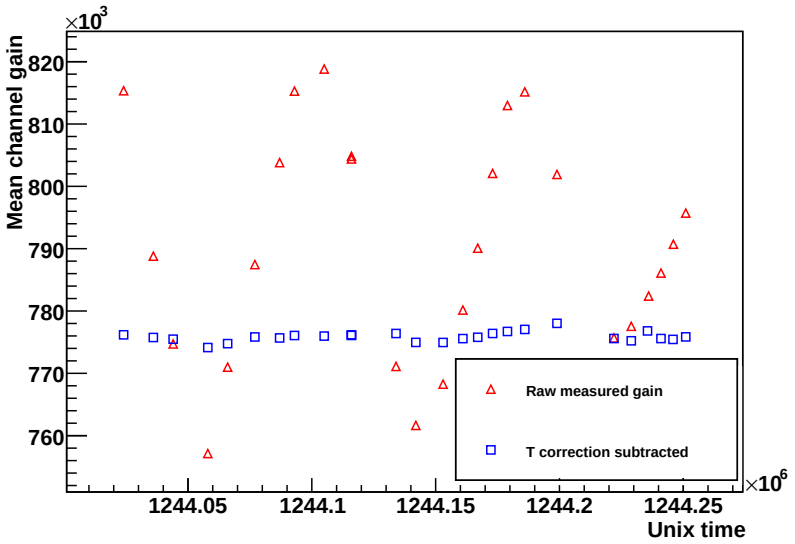

Figure 4.10: Plot showing the gains for a sample of CERN testbeam runs against run time. The blue points have had the gain shift due to temperature changes subtracted.

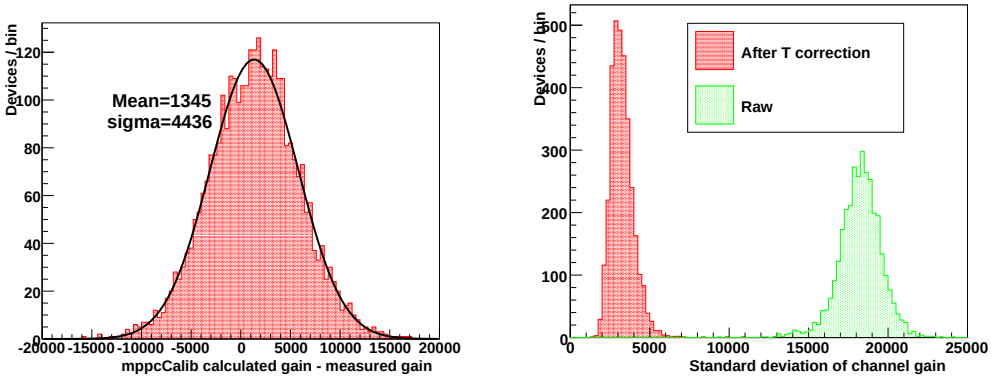

(a) (b)

Figure 4.11: (a) Difference in predicted and true gains for all MPPCs in a single testbeam run. (b) Spread in MPPC gains for each device over all runs considered, before and after calibration.

pixels to a number of photons, is still pending a lab measurement for the linearity curve. A preliminary calibration, using the linearity curve generated for the TFB electronics and 7 ns Y11 decay time, as shown in Chapter 3, has been used on the testbeam data. The voltage dependence of the form of the linearity curve was not included, but the PDE was scaled as a function of voltage, estimating the effective PDE as the intrinsic PDE multiplied by unity plus the sum of afterpulse and crosstalk probabilities. The form for all these parameters as a function of voltage was taken from the characterisation results in Chapter 3.

The results of applying this first-order calibration (which also included the gain calibration detailed above) to a sample of runs is shown in Figure 4.12, which shows the mean MIP signal for a run plotted against the run temperature. It can be seen that some residual temperature dependence exists which has not been fixed by the calibration, but the total response has been calibrated to within a few percent for a temperature shift of around 2 °C. The temperature at Tokai is expected to remain stable to around 1 °C, so it is anticipated that this calibration will allow correction of MPPC response to a precision of $\sim$1%. This should be further improved by an empirical measurement of the response curve as a function of voltage.

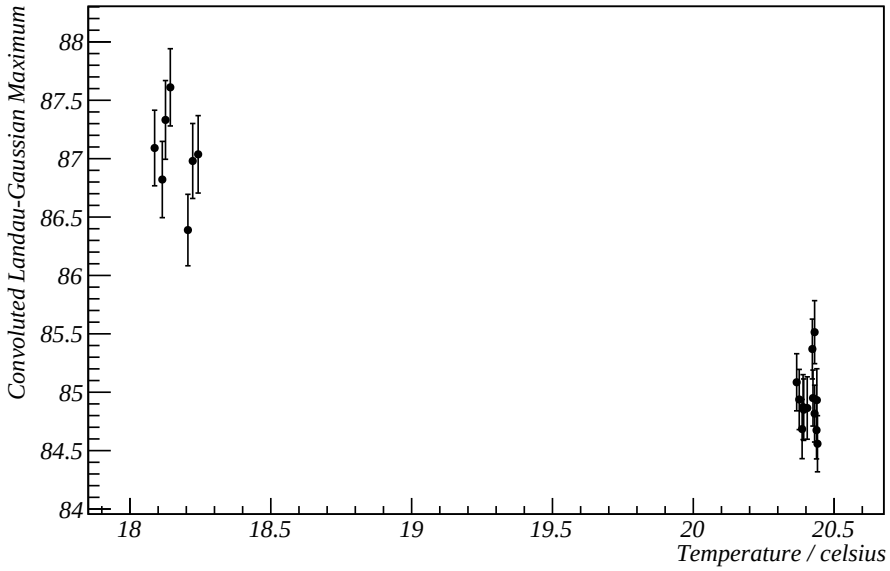

Figure 4.12: Mean predicted number of photons for a MIP signal, for a sample of testbeam runs, plotted against run temperature.

Chapter 5

Beam Extrapolation for T2K

One of the major sources of systematic error associated with long-baseline oscillation analyses is the uncertainty in predicting the neutrino flux at the far detector. For a configuration like T2K, systematic errors come from uncertainties in the geometries of the proton beamline and pion focusing system, and also from uncertainties in hadron production in the target. These effects will alter the shape and normalisation of the neutrino flux spectrum as well as the total flux, and one of the main reasons to employ near detectors is to enable these effects to be corrected for — a direct measurement of the unoscillated flux at a near detector means that the a priori uncertainties in the beam properties can be largely eliminated. However, the neutrino flux seen by a detector depends on its angular acceptance with respect to the beam direction, and it is therefore non-trivial to use the near detector data to predict the neutrino flux at the far detector.

Although the ND280 detector sits at a similar off-axis angle to Super-K, the solid angle which it samples is somewhat different — indeed for a detector

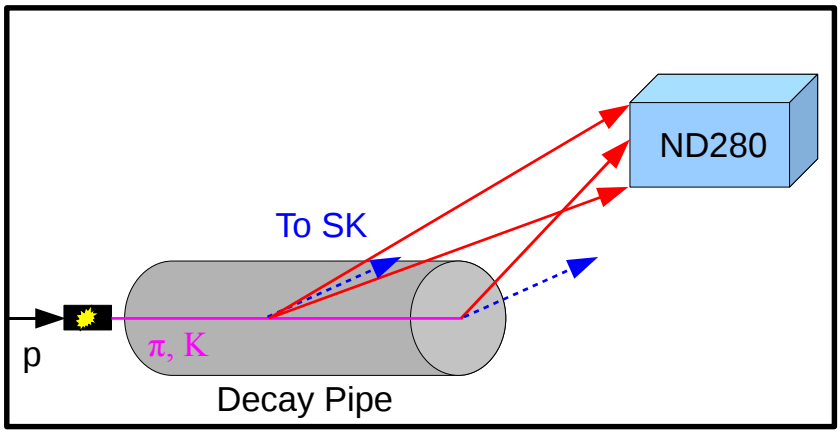

Figure 5.1: Cartoon illustrating the different angular acceptances of ND280 and Super-K, due to the former subtending significant solid angle and its distance from the target being comparable to the length of the decay pipe.

so close to the beam source, the concept of off-axis angle is itself ill-defined, as shown in Figure 5.1. The present chapter is concerned with the problem of flux extrapolation, and will describe the principle and implementation of a matrix method for doing this, as used by the MINOS experiment [36]. This will be compared with a simpler ratio method, as used by K2K [33].

It should be noted that, in addition to uncertainties in the flux-level extrapolation, systematic errors also arise from the process of unfolding the observed event spectra in the detectors to give a neutrino flux. These errors are largely due to uncertainties in the efficiencies for event channels, in energy resolution and shower energy scale, or in the neutrino cross sections. Evaluation of these errors is beyond the scope of the flux-level study presented here, but they are nevertheless central to a full oscillation analysis.

5.1 Ratio method

For each (binned) neutrino energy E_i, the simplest way in which the observed near detector flux N_i^{obs} can be used to predict the unoscillated far detector flux $F_i^{\mathrm{predicted}}$ is to define a ratio between the near and far fluxes in each bin R_i. Then

$$F_i^{\mathrm{predicted}} = R_i N_i^{\mathrm{obs}}, \tag{5.1}$$

where R_i is calculated from the Monte Carlo flux prediction

$$R_i = \frac{F_i^{\mathrm{MC}}}{N_i^{\mathrm{MC}}}. \tag{5.2}$$

The assumption implicit in this method is that any errors in the beam Monte Carlo will lead to an identical scaling of the near and far detector fluxes for a given energy bin. This will be the case if the parent phase space distributions (i.e. the distributions of neutrino parent decay positions and momenta) contributing to the neutrino flux at the near and far detectors, are identical for a given neutrino energy. This is unlikely, since the detectors do not sample identical solid angle segments. However, the ratio may also be a good approximation if the distributions contributing to each detector occupy regions in phase space which are close together. In this case, they are likely to be scaled in the same way by a realistic change in the beam properties, if this change leads to a scaling of the parent flux which does not vary rapidly over the phase space.

5.2 Beam matrix method

5.2.1 Definition of the matrix

A more sophisticated method of performing the extrapolation is to relate the near and far fluxes using a beam matrix B_{ij}, such that

$$F_i^{\text{predicted}} = \sum_j B_{ij} N_j^{\text{obs}}. \tag{5.3}$$

In contrast to the ratio method, this parameterisation allows to account for the fact that a neutrino parent which could produce a neutrino of energy E_i in the far detector, may produce a neutrino of different energy E_j if it decayed towards the near detector, depending on its position and momentum at the point of decay.

5.2.2 Calculation of the elements

From (5.3), it can be seen that any beam matrix such that $\sum_j B_{ij} N_j^{\text{MC}} = F_i^{\text{MC}}$ for all i, will give the correct far detector flux prediction in the case that the Monte Carlo matches the data exactly. However, we would like to have a matrix which is approximately correct even if the Monte Carlo does not exactly model the data. This can be done (following [36]) by considering, for a large sample of parents, the range of near detector neutrino energies E_n contributed to by parents which would give a neutrino of given energy E_f in the far detector. If the data does not agree with the MC, these correlations can be used to estimate how the far detector flux will change for a given change in the near detector flux.

The matrix elements are calculated on the above principle, using the ex-

pression

$$B_{ij} = \sum_{\text{parents}} \delta_{E_f,E_i}\delta_{E_n,E_j} \left[\frac{P_f}{P_n}\frac{P_n}{\sum_{\text{parents}} P_n \delta_{E_n,E_j}}\right] = \frac{\sum_{\text{parents}} \delta_{E_f,E_i}\delta_{E_n,E_j} P_f}{\sum_{\text{parents}} P_n \delta_{E_n,E_j}}. \quad (5.4)$$

Here E_n denotes the energy that a neutrino produced by a given parent would have if directed towards the near detector, and P_n the probability that the neutrino from the decay will be directed into unit area (defined with a normal along the neutrino direction vector) located at the near detector. $(E_f,\ P_f)$ denote the same quantities for the far detector. The interpretation of the expression is as follows:

1. For the matrix element B_{ij}, we consider only contributions from the parents giving binned far and near detector energies of $E_i,\ E_j$ respectively.

2. For each parent, we calculate the ratio of the probabilities to give a neutrino in the far detector and near detector. The result is the value which B_{ij} would take if this parent were the only one in the beam.

3. To calculate the matrix element for the whole ensemble of parents, we must combine the ratios from (2), weighting the ratio for each parent by the relative contribution of that parent to the near detector flux in energy bin E_j, i.e. the probability that a neutrino of energy E_j observed in the near detector came from this parent. This is done by taking the ratio of the neutrino flux from this parent and the total expected flux in this energy bin $\frac{P_n}{\sum_{\text{parents}} P_n \delta_{E_n,E_j}}$. In the limit of a large sample of parents then each weighting will be infinitesimal and the sum over parents is then equivalent to an integral over the region of parent phase space in which $\delta_{E_f,E_i}\delta_{E_n,E_j} = 1$.

5.2.3 Applicability

It can be seen analytically, by substituting (5.4) into (5.3), that the matrix is correct in the case that the data and MC are identical. If the MC does not accurately model the data, then the phase space distribution of the parents contributing to the neutrino flux in each near detector energy bin E_j may be different, and therefore the correlation between the contents of this energy bin and the energy bins at the far detector may change. If this is the case, the correct beam matrix describing the experiment will not be the same as that calculated from the MC. However, a change which rescales the absolute flux in each near detector energy bin, without changing the phase space distribution contributing to it, will be corrected for by the matrix. In general, the effect of a systematic error will be a combination of these two effects and will be compensated for at least partially by the matrix.

A change in the phase space distribution of the parent flux is also likely to be corrected for poorly by the ratio method, since this will affect the flux at a given energy differently for the near and far detectors. Because the beam matrix does not assume that the populations contributing to the same energy bin in the two detectors are the same (which is certainly not completely true), one would expect it to correct more effectively for systematic shifts than the ratio. We test this a priori prediction in Section 5.5, by evaluating the effects of realistic systematic errors.

5.3 Implementation of the beam extrapolation

5.3.1 Overview

A set of software tools to model the decays of neutrino parents in the T2K secondary beamline, written in C++, has been developed to study the effectiveness of the ratio and beam matrix methods in correcting for systematic errors in the beam properties. The input to the package is a sample of neutrino parents from the T2K beam MC (`jnubeam`)[1], giving the positions and momenta of the parents at the point of decay. The package uses the physics of each decay type, and the phase space coordinates of the decay, to calculate ($E_n,\ E_f,\ P_n,\ P_f$) as defined in Section 5.2.2, for each parent. E_n and P_n are calculated using a random target position in the near detector in order to obtain an average over the flux; the far detector is treated as point-like. The software then uses the energies and probabilities for each parent to generate a beam matrix or ratio vector, which can be used to extrapolate a near detector flux to the far detector.

The probabilities to decay towards the detectors are calculated analytically, rather than using a Monte Carlo method and choosing only neutrinos passing through the detector. It is therefore possible to use each parent only once, allowing the errors on the calculated beam matrix elements or ratios to be properly evaluated.

[1] Developed by the T2K Beam Group. The version used was 07a.

5.3.2 Two-body decays

We will first consider the case of pion decays, $\pi^+ \to \mu^+\nu_\mu$. Since the pion is spinless and the decay is two-body, the neutrino is produced isotropically and mono-energetically in the pion rest frame. From elementary Special Relativity, the probability for a neutrino produced by a pion decaying at position $\boldsymbol{x}_\pi$ with momentum $\boldsymbol{p}_\pi$, to pass through area ds centred on the point $\boldsymbol{x}_\nu$ is

$$\begin{aligned}\frac{\mathrm{d}P}{\mathrm{d}s} &= \frac{\mathrm{d}P}{\mathrm{d}\cos\theta_\mathrm{R}} \times \frac{\mathrm{d}\cos\theta_\mathrm{R}}{\mathrm{d}\cos\theta_\mathrm{L}} \times \frac{\mathrm{d}\cos\theta_\mathrm{L}}{\mathrm{d}s} \\ &= \frac{1}{2} \times \frac{1}{\gamma_\pi^2(1-\beta_\pi\cos\theta_\mathrm{L})^2} \times \frac{1}{2\pi(\boldsymbol{x}_\nu - \boldsymbol{x}_\pi)^2},\end{aligned} \tag{5.5}$$

where the lab decay angle $\cos\theta_\mathrm{L}$ is given by

$$\cos\theta_\mathrm{L} = \frac{(\boldsymbol{x}_\nu - \boldsymbol{x}_\pi)\cdot\boldsymbol{p}_\pi}{|\boldsymbol{x}_\nu - \boldsymbol{x}_\pi|\,|\boldsymbol{p}_\pi|}, \tag{5.6}$$

and θ_R is the rest decay angle. $(\beta_\pi,\ \gamma_\pi)$ are the usual relativistic boost parameters. Note that since this decay is isotropic, there is a trivial angular dependence in the centre-of-mass frame, $\frac{\mathrm{d}P}{\mathrm{d}\cos\theta_\mathrm{R}} = \frac{1}{2}$.

Similarly, the neutrino energy in the lab may be calculated by boosting from the pion frame:

$$E_\nu = \frac{m_\pi^2 - m_\mu^2}{2m_\pi}\frac{1}{\gamma_\pi(1-\beta_\pi\cos\theta_L)} \tag{5.7}$$

The two-body kaon decay $K^+ \to \mu^+\nu_\mu$ is exactly analogous to pion decay, and can be handled using the above method with m_π replaced by m_K.

5.3.3 $Ke3$ decays

The decays $K_L^0 \to e^+\nu_e\pi^-$ and $K^+ \to e^+\nu_e\pi^0$ ($Ke3$ decays) are isotropic, so the decay probability is described by the distribution in (5.5). However, the neutrino

is not mono-energetic in the kaon rest frame, since there are three decay products. The neutrino energy distribution cannot be calculated exactly, since the interaction has a hadronic component, so the neutrino energy is selected using a form-factor parameterisation of the Dalitz plot density for the decay. The parameterisation used is an approximation of that preferred by [7], giving the joint probability density for the muon and pion energies in the kaon rest frame as

$$\rho(E_\pi, E_\mu) \propto \left[f_+(0) \left(1 + \lambda_+ \frac{t}{m_\pi^2} \right) \right] \times \left[2E_e E_\nu - m_K \left(\frac{m_K^2 + m_\pi^2 - m_\mu^2}{2m_K} - E_\pi \right) \right], \tag{5.8}$$

where t is the squared momentum transfer to the leptonic system, and we use $f_+(0) = 0.98\,\text{GeV}^2$, $\lambda_+ = 0.0286$. These parameter values are chosen to match the model used by the `jnubeam` MC.

5.3.4 Muon decays

The case of muon decays is somewhat more complicated, because the muon has spin-$\frac{1}{2}$ and so the decay is not isotropic. The distribution of electron neutrinos from the decay $\mu^+ \to e^+ \nu_e \bar{\nu}_\mu$, in the μ centre-of-mass frame, is given by

$$\frac{\mathrm{d}^2 P}{\mathrm{d}x\, \mathrm{d}\cos\theta_\mathrm{R}} = 6x^2 (1 - x) (1 - p \cos\theta) \qquad \text{where } x = \frac{2E_\nu^{rest}}{m_\mu}, \tag{5.9}$$

where p is the muon spin polarisation, and θ_R is the angle at which the ν_e is emitted, relative to the muon polarisation axis [68]. We can easily derive the single derivative $\frac{\mathrm{d}P}{\mathrm{d}\cos\theta_\mathrm{R}}$, and the decay probability can then be calculated by substituting this into (5.5). The energy distribution in (5.9) is used to select a rest-frame

energy, using an acceptance-based MC — a value of x is chosen randomly, and accepted with a probability equal to the ratio of the probability density at this energy to that at the most probable value of x.

The same procedure can be used for the $\bar{\nu}_\mu$ from the decay; in this case the decay distribution is given by

$$\frac{\mathrm{d}^2 P}{\mathrm{d}x\,\mathrm{d}\cos\theta} = x^2\left[(3-2x) - p\,(1-2x)\cos\theta\right] \tag{5.10}$$

and the energy distribution depends on $\cos\theta$.

Note that since the three-body decays are not mono-energetic for a given lab decay angle, and also because the solid angle subtended by the near detector is significant, identical parents may contribute to different elements of the beam matrix. Over a large sample of parents, the method used will average out these possibilities to give the correct distribution.

5.4 Beam matrix and ratio for nominal Monte Carlo

The beam matrix and ratio corresponding to the nominal beam Monte Carlo, for ν_μ, are shown in Figure 5.2. It can be seen that the matrix is broadly centred on the diagonal. The spread around the diagonal is mainly due to pions, which can travel some distance down the beam pipe before decaying, and therefore see a range of solid angles for the near detector (see Figure 5.1). The large kaon mass means that kaon decays give high energy neutrinos, and have a smaller boost so that the neutrino energy is less angle-dependent. These decays are responsible

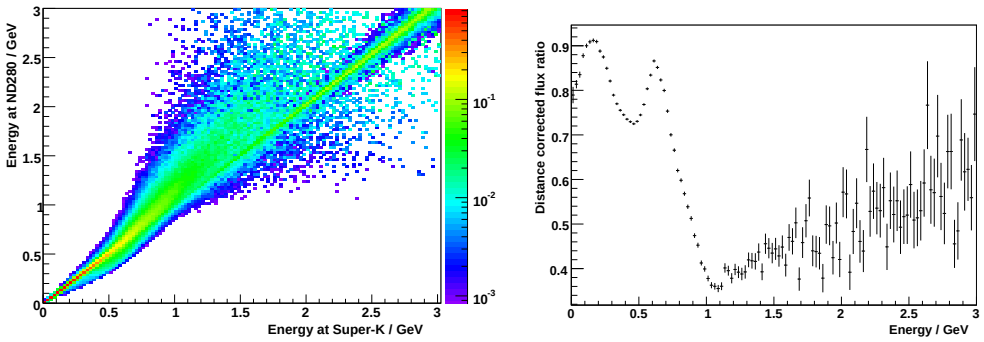

Figure 5.2: (a) Beam matrix for extrapolation of the ND280 ν_μ flux to Super-K. (b) MC ratio of Super-K and ND280 ν_μ fluxes. Numbers are corrected for the difference in distance between ND280 and Super-K.

for the continuing diagonal line at higher neutrino energies. The contribution of muon decays to the ν_μ flux is negligible. The corresponding ν_e matrix and ratio are shown in Figure 5.3. This matrix is much more diffuse than the ν_μ case; this is because ν_e are generated by three-body decays of kaons and muons, which are not mono-energetic, and a parent can therefore produce neutrinos of a wide range of energies for a given decay angle.

At higher energies, it can be seen that the statistics for both the matrix and ratio become rather sparse, but at very low energies, where the flux is also small, no such problem exists. This is because the spectrum at high energy is dominated by a small sample of high-energy parents with a high probability of decaying towards the detectors. This presents a problem in producing enough MC statistics to constrain the spectrum at higher energies — the extrapolations shown are for 10^7 p.o.t., corresponding to about 700 hours running on a single CPU[2]. MC generation is readily parallelised, but these computing requirements are

[2]Tested on an Intel Pentium 4 running at around 3 GHz.

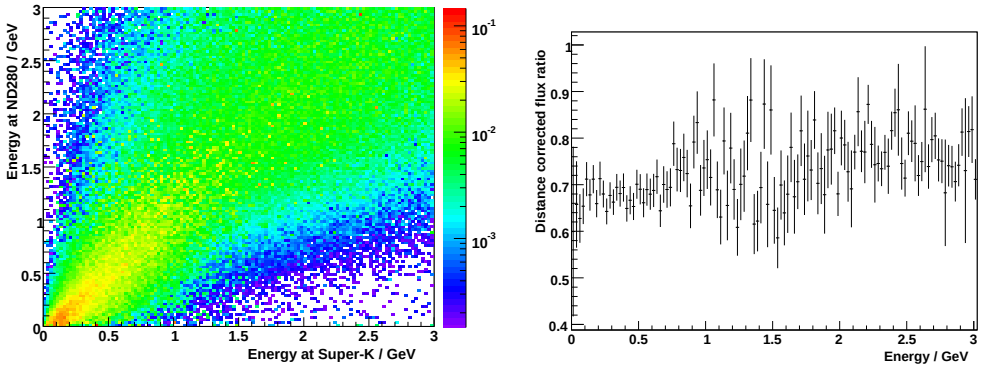

Figure 5.3: (a) Beam matrix for extrapolation of the ND280 ν_e flux to Super-K. (b) MC ratio of Super-K and ND280 ν_e fluxes. Numbers are corrected for the difference in distance between ND280 and Super-K.

nevertheless significant.

For contrast, the ν_μ beam matrix between the INGRID on-axis detector and Super-K is shown in Figure 5.4. The correlation between near and far detector energies for a given parent is seen to be negative except at low energies, since parents directed towards Super-K could produce a high-energy neutrino directed towards that detector, or a low-energy neutrino directed towards INGRID, and vice versa. The diagonal component at higher energies is due to kaons as for the ND280-Super-K matrix. The ratio is not shown for this case since it is not meaningful if the matrix does not indicate strong positive correlation between near and far detector energies.

5.4.1 Validation of the matrix code against jnubeam

Our matrix code uses its own implementation of the parent decay physics rather than the `jnubeam` code. However, if the physics is equivalent to that in `jnubeam`,

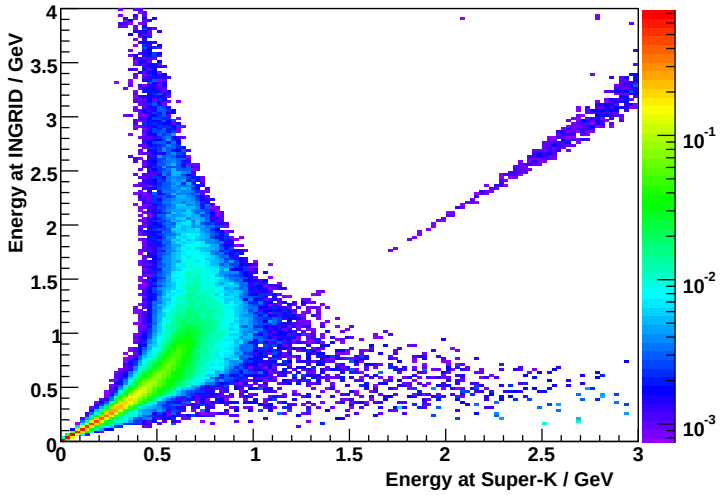

Figure 5.4: Beam matrix for extrapolation of the INGRID ν_μ flux to Super-K. Numbers are corrected for the difference in distance between INGRID and Super-K.

then the neutrino fluxes produced by our code and `jnubeam` should be the same, and the results from each can be cross-checked to ensure that they agree. We therefore compared the fluxes at ND280, for each kind of decay, predicted by `jnubeam` and by our code. The results of this cross-check are shown in Figures 5.5 and 5.6. All important decays in the beam (2-body kaon and pion decays, $Ke3$ decays and the ν_e from μ decays), are seen to agree very well between the beamline MC and our code.

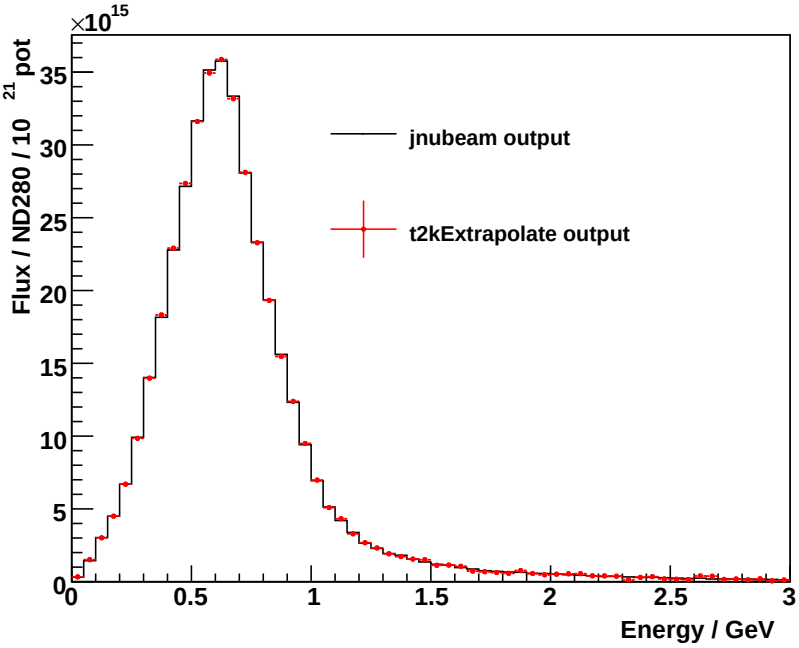

(a) ν_μ from π decay

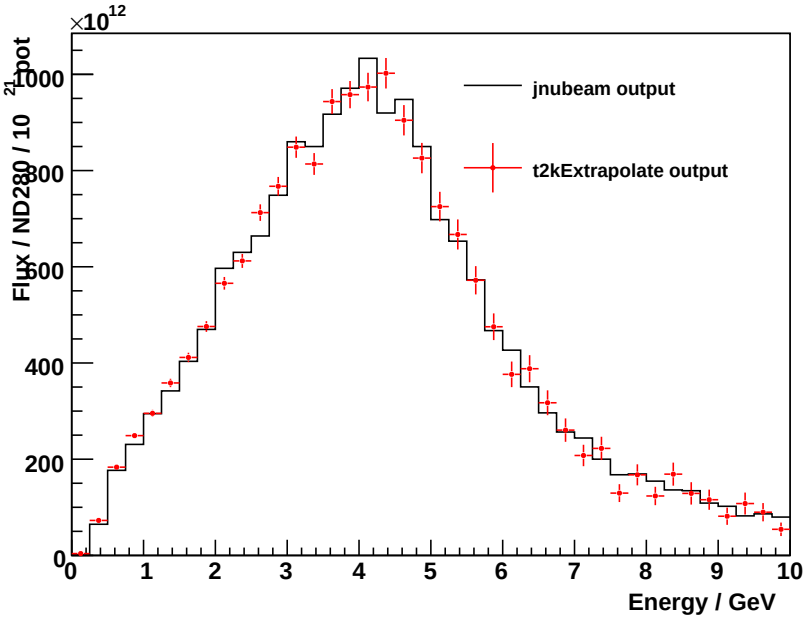

(b) ν_μ from K decay

Figure 5.5: Comparison of the ν_μ fluxes integrated over the ND280 fiducial region, predicted by the T2K beamline MC and by our extrapolation package.

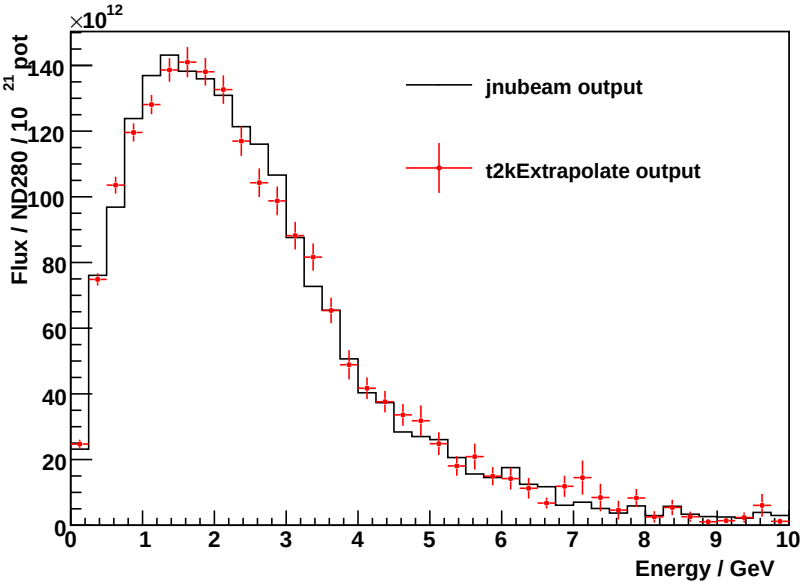

(a) ν_e from K decay

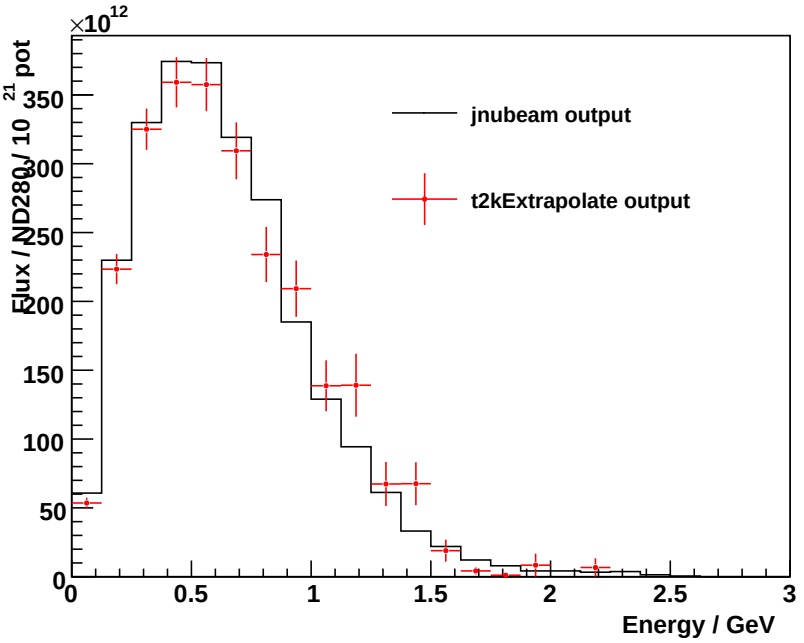

(b) ν_e from μ decay

Figure 5.6: Comparison of the ν_e fluxes integrated over the ND280 fiducial region, predicted by the T2K beamline MC and by our extrapolation package.

5.5 Study of systematic effects

5.5.1 Procedure

The purpose of a beam extrapolation is to enable a prediction of the far detector flux in the case that near detector observations do not match the Monte Carlo predictions. It is therefore important to test the robustness of our extrapolation methods, i.e. how precisely they predict the far detector flux when the beamline parameters are shifted by some systematic error. The sources of error considered here are

1. errors in the position and angle of the proton beam at the target;
2. errors in the hadron production model used in the target;
3. errors in the secondary horn alignments; and
4. surveying errors affecting the alignment of the beamline relative to the detectors.

We are particularly interested in studying geometrical changes in the beamline, since such changes are likely to change the shape of the neutrino spectrum rather than just scaling the overall flux. Overall changes in the flux scale will be completely corrected by either of the methods considered, and so factors like the proton beam intensity are not interesting for this study. We are interested in looking separately at vertical and horizontal shifts in the geometry — the effect of a change will be direction-dependent, since the displacement of ND280 and Super-K from the beam axis is mostly vertical. The alignment of the beam and detector centres,

with respect to the beam direction projected onto the horizontal plane at Tokai, are shown in Table 5.1.

	Downward angle (mrad)	Clockwise horizontal angle (mrad)
Beam	63.5	0
ND280	30.0	−10.9
Super-K	22.0	−13.7

Table 5.1: Alignment of the T2K beam centre, and the ND280 and Super-K detectors. Angle to ND280 is given at the target — at points downstream of this its angle relative to the beam is larger. The downward angle is relative to the horizontal plane at Tokai.

The `jnubeam` MC takes a specified phase space distribution for the proton beam at the target, and then uses the `GEANT3` particle transport simulation to simulate proton interactions in the target, and to transport the secondaries through the magnetic field of the horns. The `jnubeam` code was changed to generate data for systematically altered configurations, and the output was then used to produce fluxes for both the real T2K configuration, and an alternative configuration where the near and far detectors are positioned on the beam axis — this allows to see how the sensitivity to systematics changes for an off-axis beam. Different parent samples were used to generate the near and far detector fluxes, so that the two were statistically independent. The total data generated for each systematic shift was the same as that used for the nominal MC — half of this was used to generate the near detector flux and the other half for the far detector.

The ratio and matrix methods were used to predict the altered far detector spectrum based on the near detector flux, using the extrapolation parameters for nominal Monte Carlo (as shown in Figure 5.2). These predictions were then

compared to the true far detector flux, in order to evaluate the performance of the ratio and matrix. The agreement between the true flux and extrapolations, for the various systematics, is shown in the plots in Section 5.5.2 and evaluated quantitatively in 5.5.3.

5.5.2 Results

Hadron model

The GEANT3 package which is used to simulate hadron production contains interfaces to several hadron production models. The default model is CALOR [69], but FLUKA [70] and MICAP [71] can also be used. These models have elements in common — MICAP uses the FLUKA MC, but with new data-driven modelling for low-energy neutron transport. CALOR uses MICAP for low-energy neutrons and FLUKA for hadrons above 10 GeV — a different set of models are used for hadrons at intermediate energies.

In the final T2K analysis, it is likely that a parameterised model for hadron production on the target, based on recent hadron production experiments on a similar target [72], will be used rather than a Monte Carlo model. However, in the meantime it is possible to estimate the likely errors in hadron production by considering the difference between the fluxes predicted by the different models. To this end, datasets have been produced using both FLUKA and MICAP, and compared to the nominal MC, which uses CALOR.

The results for FLUKA and MICAP are shown in Figures 5.7 and 5.8 respectively. These figures (as well as the similar ones presented for other systematics)

were produced by extrapolating the near detector flux generated using the alternative model to the far detector, using the matrix and ratio. These extrapolated fluxes, as well as the true far detector flux generated with the altered model, were then divided by the nominal far detector flux — it is these ratios to the nominal flux which are plotted. The extrapolations are shown as fills representing the 1σ bound of the extrapolation prediction, and if an extrapolation corrects well for the systematic shifts, it should track the data points.

Both alternative models give higher fluxes than the default, with a similar scaling for on- and off-axis configurations. The flux scalings below 200 MeV are large, but this region is outside the area around the oscillation maximum which will be used to fit for the mixing parameters. Outside this region, the flux scaling does not change sharply as a function of energy.

Beam position on target

The position of the primary beam is monitored extremely precisely by monitors in the beamline itself. It is expected that the error on the beam impact point with respect to the target face will be dominated by the tolerance on the target position. The target itself is located to a precision of around 0.1 mm inside the primary horn, but the tolerance on the position of this horn is 1 mm, so an error of this size is introduced in the position of the beam on the target. This error is to be compared to the expected RMS beam size of around 4–6 mm. The target diameter is 26 mm.

Data has been generated for a shift in the beam position on the target

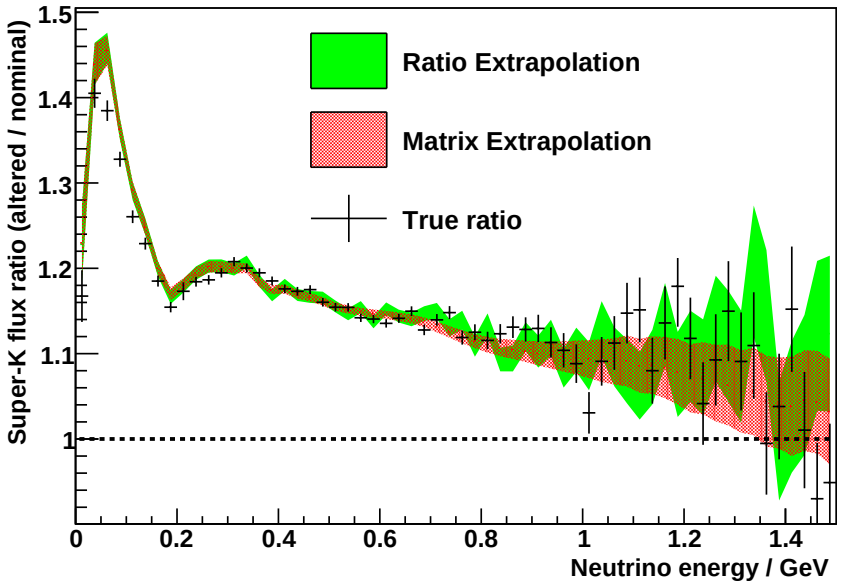

(a)

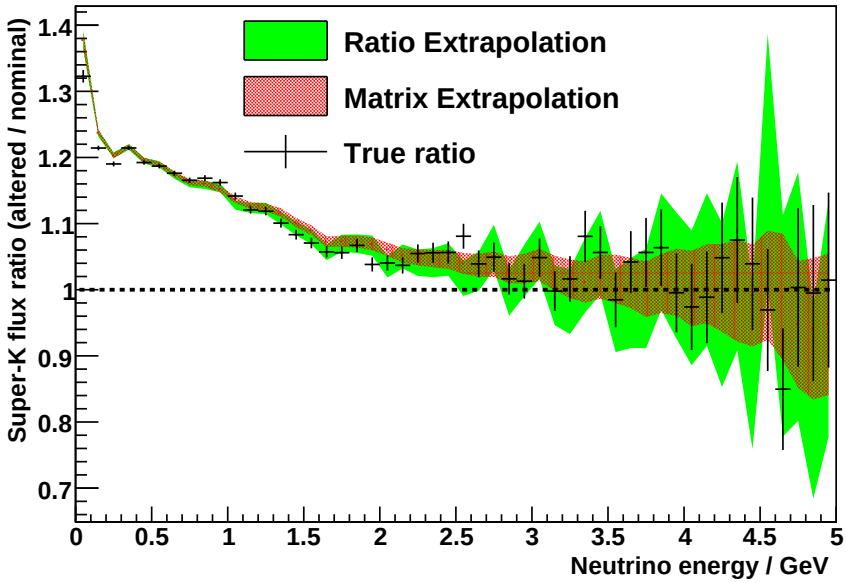

(b)

Figure 5.7: (a) Ratio of ν_μ flux at Super-K to nominal MC, for the FLUKA hadron model. Extrapolation predictions are also shown. All errors are statistical only. (b) As for (a), for an on-axis configuration.

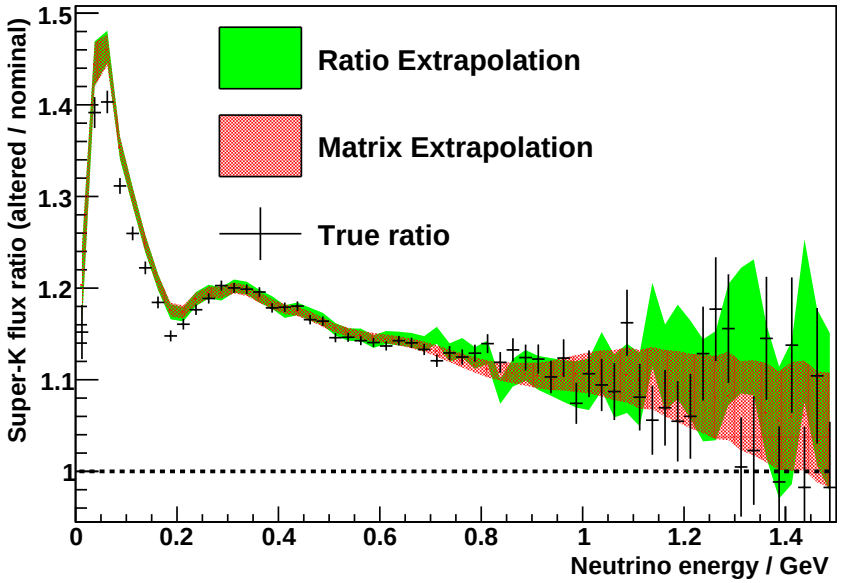

(a)

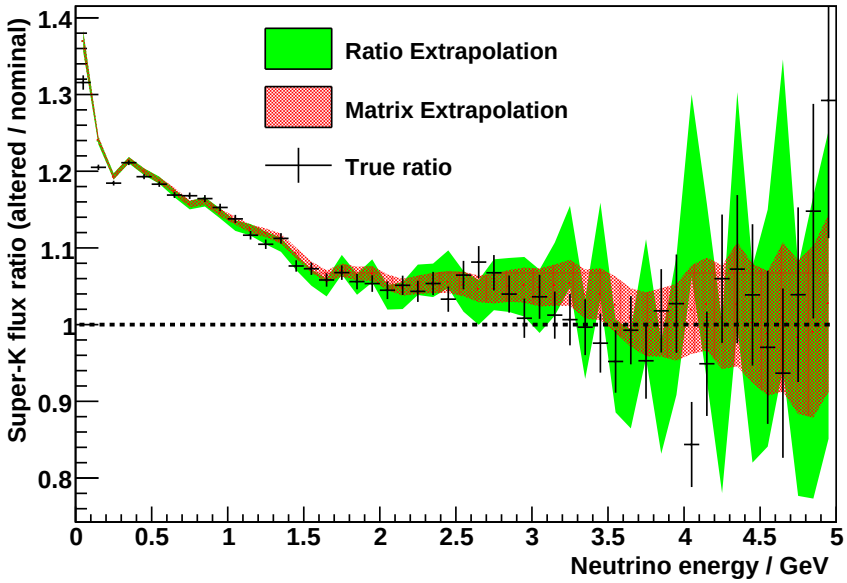

(b)

Figure 5.8: (a) Ratio of ν_μ flux at Super-K to nominal MC, for the MICAP hadron model. Extrapolation predictions are also shown. All errors are statistical only. (b) As for (a), for an on-axis configuration.

of 2 mm (2σ), in both the horizontal and vertical directions. These results are shown in Figures 5.9 and 5.10. It can be seen that for the on-axis configuration, a shift in the beam position produces a negligible effect on the beam flux. For the off-axis configuration, the effect of a horizontal shift is quite small, but a vertical shift produces a significant reduction in the flux at higher energies. This is not surprising, since the off-axis displacements of ND280 and Super-K are mainly in the vertical direction.

Beam angle on the target

Like the error on the beam position, the error in the beam angle at the target is dominated by the error on the target and horn alignment. If we take a tolerance of 1 mm, independently, on the positions of the front and back of the primary horn, the error in the transverse displacement of the horn front with respect to the back will be given by the sum in quadrature, i.e. $\sqrt{2}$ mm. The primary horn length is around 1.4 m, giving an error in the target angle of about 1 mrad.

We considered the effect of a rather larger angular shift of 5 mrad. Evaluating the effect of such a shift is useful, since maintaining the alignment of the target station at Tokai is made non-trivial by the instability of the ground at the site. It should be understood, however, that the shift considered represents a worst-case scenario. The effect on the fluxes from such a shift is shown in Figure 5.11. The effect on the on-axis flux is small, but the off-axis flux is reduced significantly at higher energies, with a rather sharp change in the flux ratio as a function of energy.

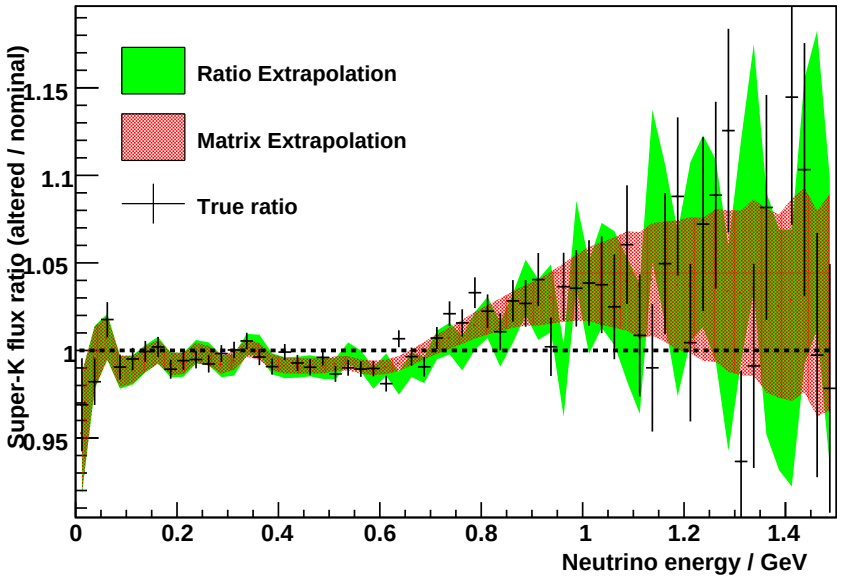

(a)

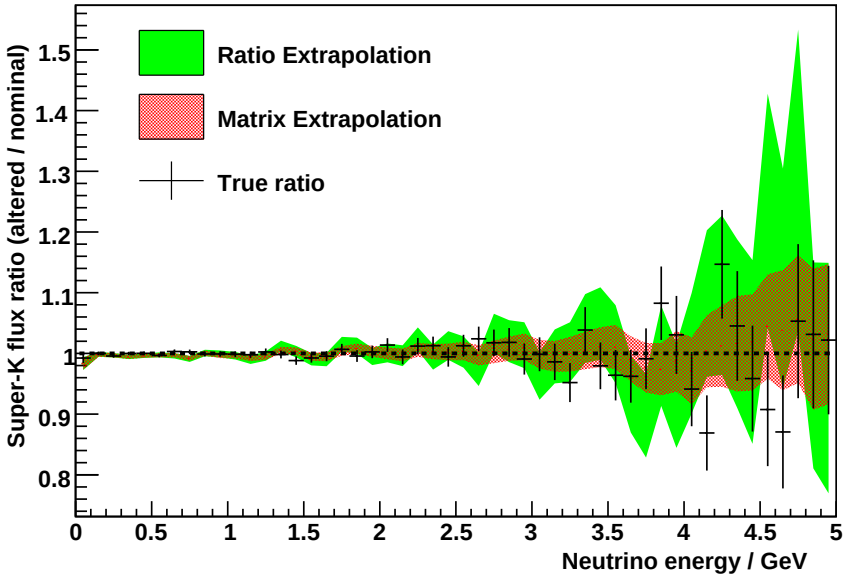

(b)

Figure 5.9: (a) Ratio of ν_μ flux at Super-K to nominal MC, for a horizontal shift of $+2$ mm of the proton beam on the target. Extrapolation predictions are also shown. All errors are statistical only. (b) As for (a), for an on-axis configuration.

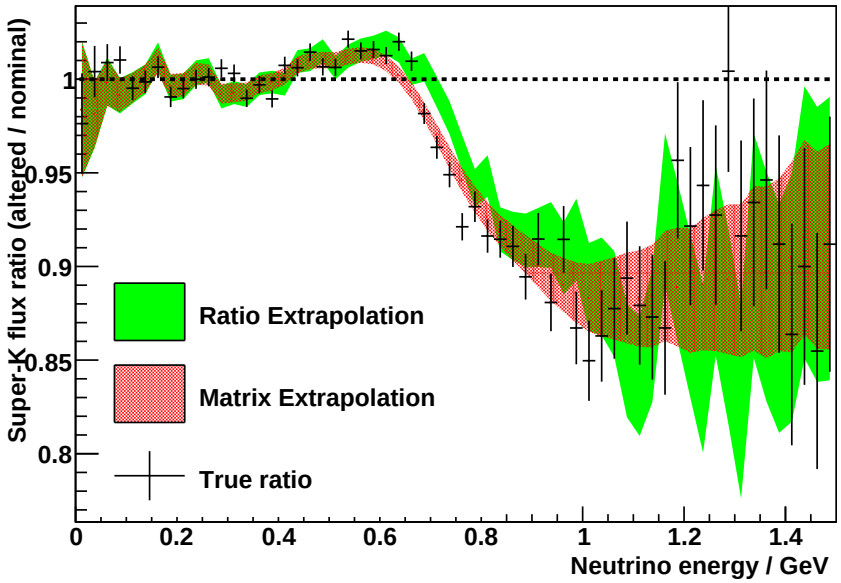

(a)

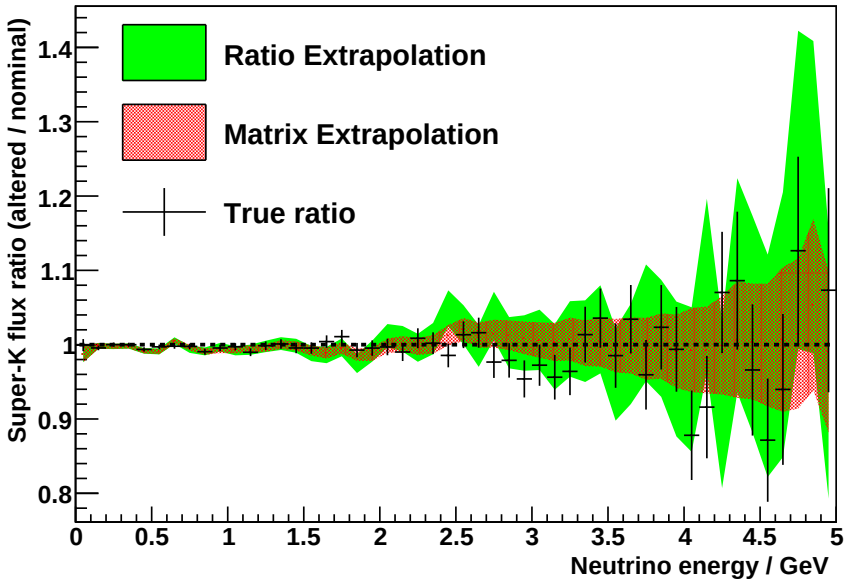

(b)

Figure 5.10: (a) Ratio of ν_μ flux at Super-K to nominal MC, for a vertical shift of $+2$ mm of the proton beam on the target. Extrapolation predictions are also shown. All errors are statistical only. (b) As for (a), for an on-axis configuration.

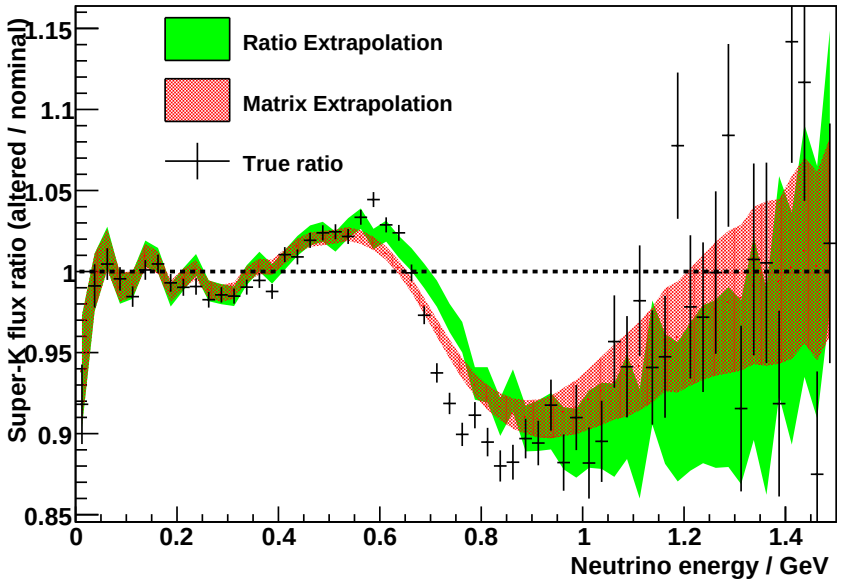

(a)

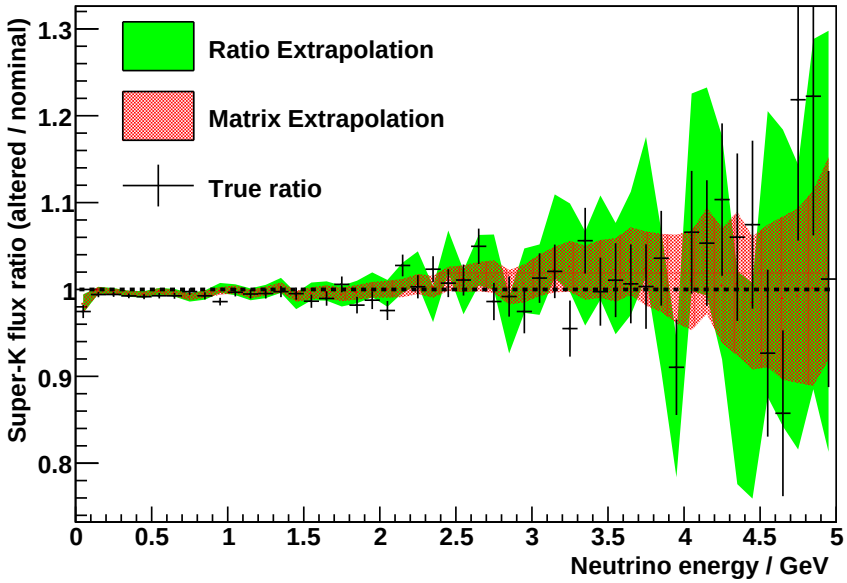

(b)

Figure 5.11: (a) Ratio of ν_μ flux at Super-K to nominal MC, for a proton beam with an upward angle of 5 mrad at the target. Extrapolation predictions are also shown. All errors are statistical only. (b) As for (a), for an on-axis configuration.

Beam angle relative to detectors

The angle of the neutrino beam centre relative to the detectors is measured by the INGRID on-axis near detector. The precision of this measurement is around 1 mrad for a single day's data at full beam intensity. We therefore consider a variation of 2 mrad in the beam direction, in both the horizontal and vertical directions — over a longer period of measurement INGRID will be able to constrain the direction much better than this, but while the beam intensity is ramping up then an error on this scale may be reasonable. This shift was accomplished in the simulation by moving the whole beamline geometry with respect to the detectors.

The results of the horizontal and vertical shifts are shown in Figures 5.12 and 5.13. The effect of the shift on the beam flux for the on-axis configuration is small, though it is non-zero at high energies. For the off-axis configuration, the effect on the ratio is is very large, especially for a vertical shift. It also changes rapidly with energy.

Displacement of secondary horns

The focusing of the pion beam by the secondary horns is important in determining the shape of the final neutrino beam. The tolerance on the transverse position of the second and third horns is 3 mm. We considered the case of both horns being shifted upwards by 6 mm; other permutations of horn displacements could also be considered, but a measure of the scale of errors due to the horn position tolerance can be drawn from our data.

The results from the horn movement are shown in Figure 5.14. At higher

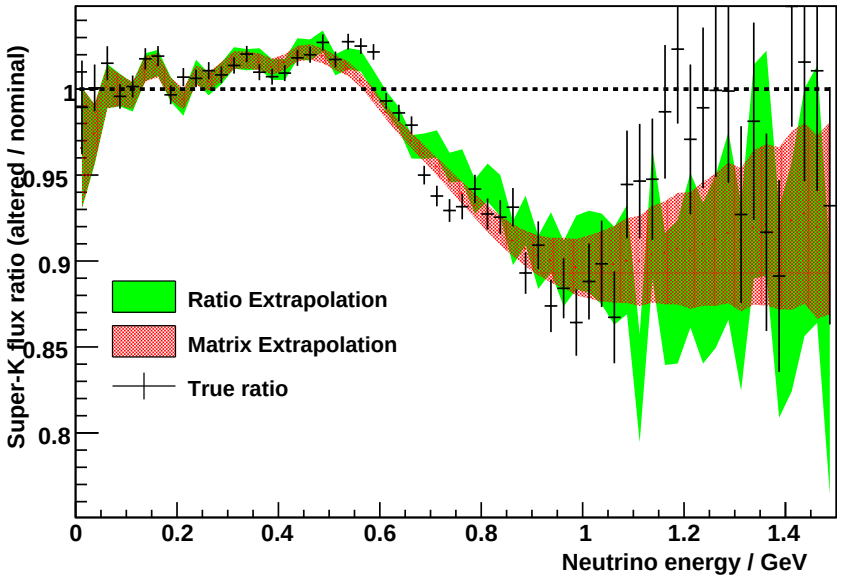

(a)

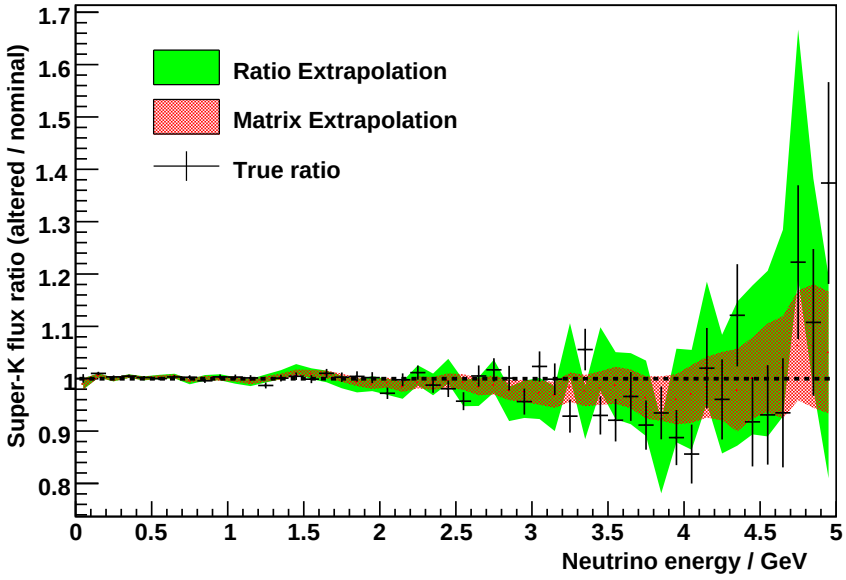

(b)

Figure 5.12: (a) Ratio of ν_μ flux at Super-K to nominal MC, for a $+2$ mrad horizontal shift in the neutrino beam direction. Extrapolation predictions are also shown. All errors are statistical only. (b) As for (a), for an on-axis configuration.

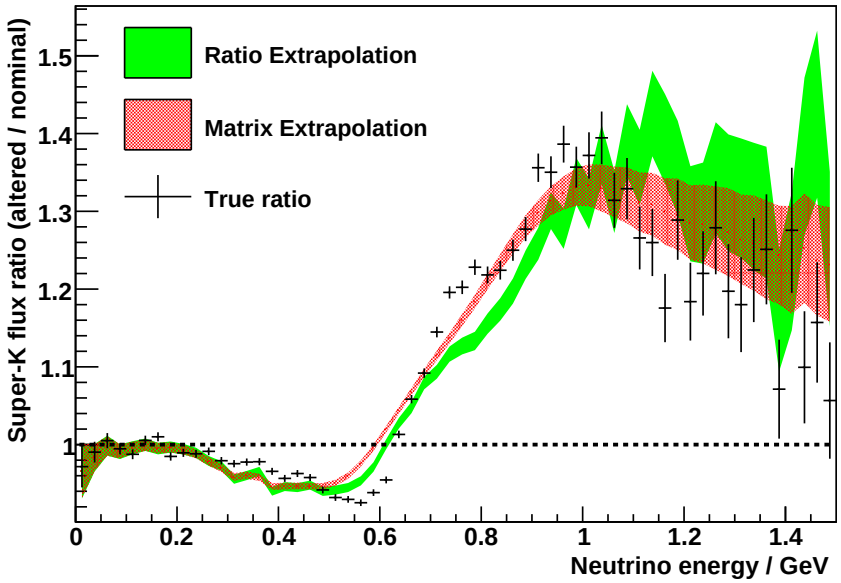

(a)

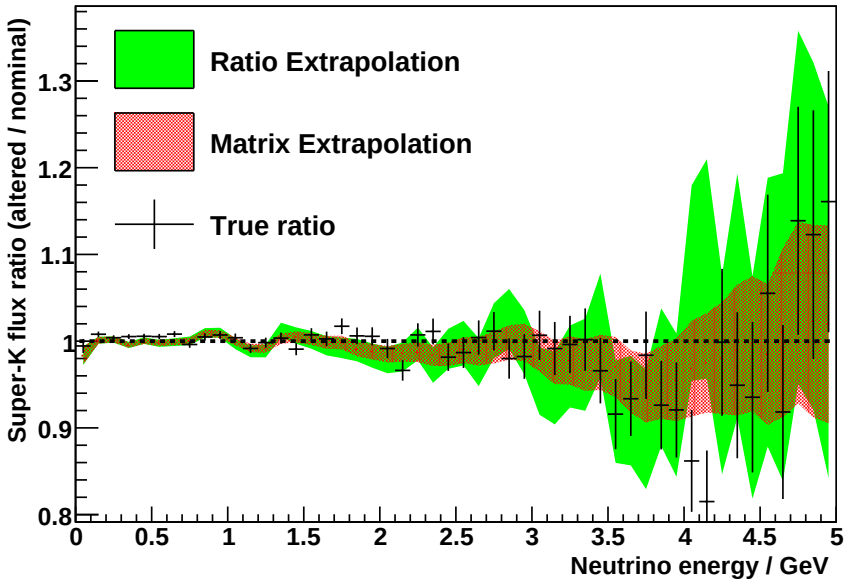

(b)

Figure 5.13: (a) Ratio of ν_μ flux at Super-K to nominal MC, for a $+2$ mrad vertical shift in the neutrino beam direction. Extrapolation predictions are also shown. All errors are statistical only. (b) As for (a), for an on-axis configuration.

energies, the effect of a horn shift on the off-axis flux is large.

5.5.3 Performance evaluation and discussion

Before making a quantitative evaluation of the performance of the extrapolations, a few general remarks can be made. The off-axis flux is seen to be much more sensitive to changes in the beamline geometry than the on-axis flux. This increased sensitivity is unsurprising, since the beam flux at all energies peaks on-axis, giving a locally flat flux as a function of angle. In contrast, the off-axis configuration sits at a point where the flux is changing rapidly as a function of angle, so small changes in beam configuration can have a large effect. Though for the most part the beam flux is well-corrected by both extrapolations, there are residual errors of up to a few percent where the flux ratio is changing rapidly, and this region is around 0.6–1.0 GeV for most of the systematics, in the part of the spectrum which will be used for analysis.

Geometrical changes in the beamline have been seen to affect the high-energy part of the neutrino flux more severely than the lower-energy part. This is expected, since the decay probability and neutrino energy depend more strongly on decay angle for more energetic neutrino parents. Small changes in angle can thus have a large effect on the flux from these parents. Changes in this part of the spectrum are also problematic for the T2K oscillation measurements — although this part of the spectrum is above the oscillation analysis region, it is higher energy neutrinos which are responsible for much of the non-CCQE background to the Super-K measurements. It is therefore important to constrain this region of

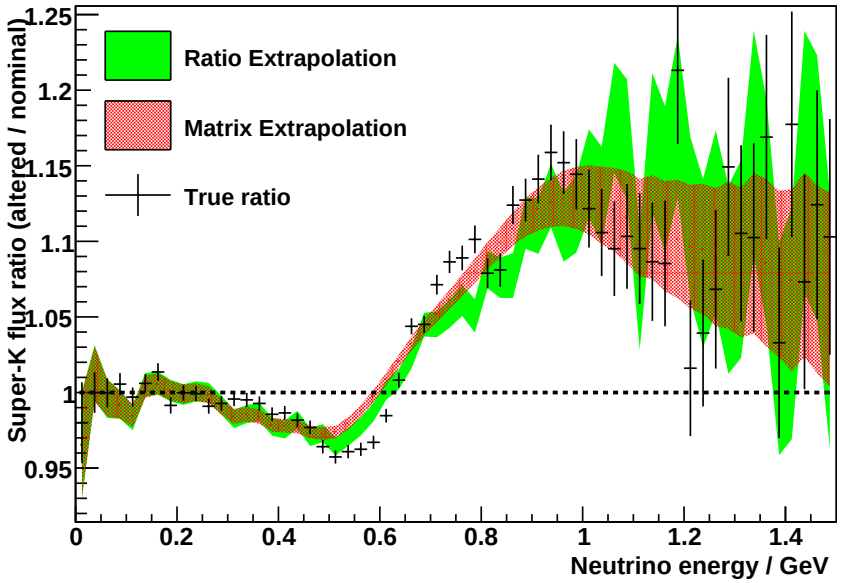

(a)

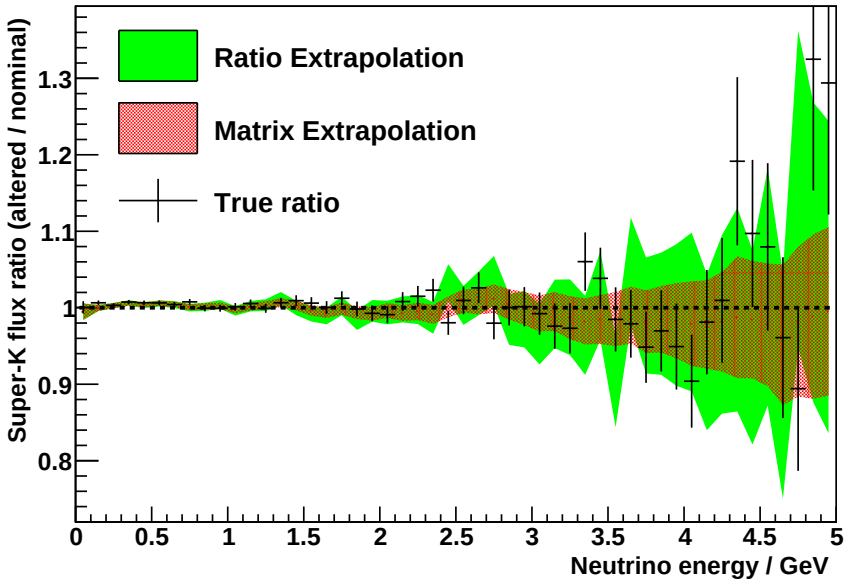

(b)

Figure 5.14: (a) Ratio of ν_μ flux at Super-K to nominal MC, for a $+6$ mm vertical shift in the positions of horns 2 and 3. Extrapolation predictions are also shown. All errors are statistical only. (b) As for (a), for an on-axis configuration.

the spectrum.

The errors in the matrix extrapolation are rather smaller than for the ratio, because the matrix method effectively has higher statistics as several near detector bins are used to estimate the flux for each bin at the far detector. These extra statistics are useful at high energies, where the MC is sparse. Because MC production is computationally intensive, generating larger samples would be difficult.

Because the purpose of the extrapolations is to facilitate precise oscillation measurements, the performances of the ratio and matrix have been evaluated using a simplified oscillation analysis, to see how systematic errors affect the measured oscillation parameters once they have been corrected for by the extrapolation. We multiplied the true neutrino fluxes for each systematic by the muon neutrino survival probability (1.25), using the parameters ($\sin^2(2\theta_{23}) = 0.95$, $\Delta M^2 = 2.5 \times 10^{-3}$ eV). We then divided this oscillated flux by the flux predictions given by the matrix and ratio, in order to give the survival probabilities $P_{\mu\to\mu}$ which would be predicted by these methods. These were then fitted using (1.25), to obtain the oscillation parameters which would be obtained using the matrix and ratio. The fit was performed over the range 0.35–0.85 GeV corresponding to the region around the first oscillation peak.

The results of this analysis for all systematics are shown in Table 5.2. It can be seen that the discrepancy between the fitted values and input parameters is less than 0.25% in all cases, though from the fit errors we see that this discrepancy is larger than expected from statistical fluctuations for many of the fits — this can be seen in the contour plot for the vertical beam angle shift in Figure 5.15. The

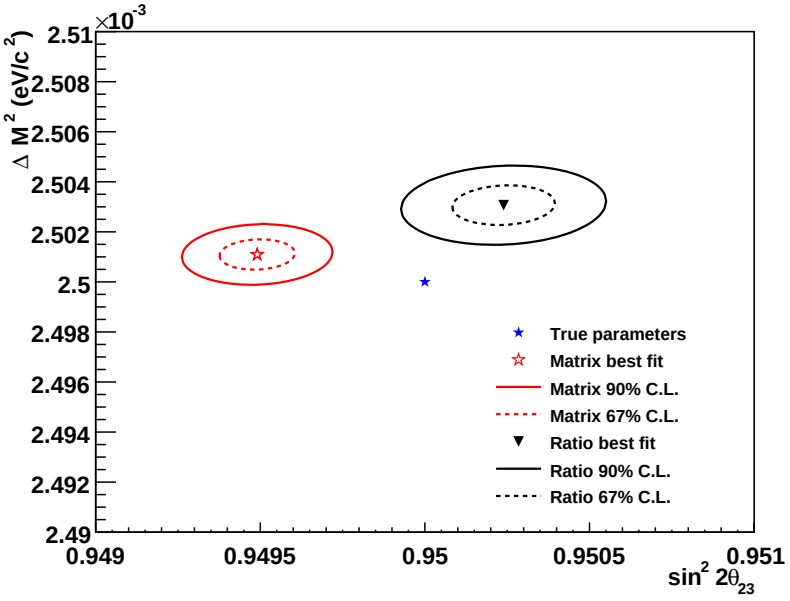

Figure 5.15: Contour plot for the matrix and ratio fit results, for changing the vertical beam angle by 5 mrad. It can be seen that the differences between the matrix, ratio and true parameters are statistically significant.

discrepancies are smaller than the statistical errors shown for the disappearance measurement in Figure 2.24 (where we expect systematic errors to be dominant), so given that the systematics considered represent 2σ or larger shifts, we conclude that the error introduced by the extrapolation is small, and that the performances of the ratio and matrix extrapolations are similar.

Although the results of this study indicate that the error introduced by the flux extrapolation is small, it is still important to obtain a more robust result by performing another study which includes an exhaustive and fully evaluated set of systematic errors. This study should take place within the framework of a full oscillation analysis, since the flux-level extrapolation cannot be cleanly separated from the unfolding of observed ND280 events into a neutrino flux spectrum (and

	$\delta(\sin^2(2\theta_{23}))_{\mathrm{Mat}}$ (%)	$\delta(\Delta M^2)_{\mathrm{Mat}}$ (%)	$\delta(\sin^2(2\theta_{23}))_{\mathrm{Rat}}$ (%)	$\delta(\Delta M^2)_{\mathrm{Rat}}$ (%)
FLUKA Model	-0.05 ± 0.01	$+0.04\pm0.02$	$+0.02\pm0.02$	$+0.12\pm0.03$
MICAP Model	$+0.02\pm0.01$	-0.02 ± 0.02	$+0.03\pm0.02$	$+0.00\pm0.03$
Beam position on target (vertical)	-0.03 ± 0.01	$+0.00\pm0.02$	$+0.03\pm0.02$	$+0.10\pm0.03$
Beam position on target (horizontal)	$+0.00\pm0.01$	-0.02 ± 0.02	-0.02 ± 0.02	-0.06 ± 0.03
Beam angle on target	-0.05 ± 0.01	$+0.04\pm0.02$	$+0.02\pm0.02$	$+0.12\pm0.03$
Beam angle (vertical)	$+0.22\pm0.01$	-0.08 ± 0.02	$+0.08\pm0.02$	-0.12 ± 0.03
Beam angle (horizontal)	-0.06 ± 0.01	$+0.06\pm0.02$	-0.01 ± 0.02	$+0.06\pm0.03$
Horn offset	$+0.09\pm0.01$	-0.10 ± 0.02	$+0.03\pm0.02$	-0.10 ± 0.03

Table 5.2: Discrepancies between the true parameters and values from the extrapolations, for our simplified oscillation analysis.

the reverse process at Super-K). In particular, we have not considered shower scaling or momentum resolution errors, which lead to energy smearing or rescaling in the unfolded neutrino flux, and therefore interact with the extrapolation non-trivially. We have also not included the effect of backgrounds from the high-energy tail of the flux. This further study will require a more complete understanding of the experimental configuration than is currently available, and also development of a full analysis framework. We conclude in the meantime that both the matrix and ratio correct for systematic shifts sufficiently well to not introduce a significant error in the oscillation analysis considered here.

5.6 Conclusion

A package has been developed which calculates the neutrino flux at a given detector location, from a set of neutrino parents generated by the T2K beamline MC. A largely analytic method is used for the flux calculations, maximising computational efficiency and producing an accurate measure of the statistical errors on the fluxes. The package is able to use the parent sample to produce a beam matrix encoding the correlations between the energy bins in the near and far detector fluxes, and also, trivially, to produce a ratio of far/near neutrino fluxes as a function of energy.

This code has been used to test the performance of the matrix and ratio extrapolation methods in correcting for systematic shifts in beam properties, by using the matrix and ratio to predict far detector spectra for systematically altered configurations. This study concluded that both the matrix and ratio correct well enough for large geometrical shifts in the beamline configuration to not form an important error in the ν_μ disappearance analysis, but this study should be supplemented by a thorough evaluation including other effects, e.g. uncertainties on interaction cross-sections, once a complete analysis framework is available.

Chapter 6

Conclusion and outlook

6.1 Conclusion

We have seen that the physics of neutrino oscillation is now entering an exciting new phase of high-precision measurement. We are now at the point where neutrino measurements can begin to put constraints on new physics, and perhaps even shed light on the origins of the Universe as we know it. The T2K experiment is at the forefront of the effort to advance our understanding of the neutrino.

We have discussed the configuration of the T2K experiment, which uses the Super-K water Čerenkov detector to measure neutrino oscillations. Two of the novel features of the project are the use of MPPC photosensors, and the off-axis positioning of the detectors to reduce high-energy beam contamination. This thesis has focused on these topics — we have shown that MPPCs are capable of comparable performance to PMTs, while being compact and insensitive to magnetic fields. We have also shown that the T2K collaboration has had sufficient

success in understanding these devices that it has been possible to develop a MC simulation of them which accurately models data. The QA effort made to test MPPCs bound for the ECal has also shown that the properties of the MPPCs are stable over large numbers of devices.

We have covered in some detail the important physics governing the production of neutrinos in the T2K beam, and how an accurate beamline MC can allow the flux of neutrinos at the near detector to be used to predict the flux at the far detector, using either a ratio or a more sophisticated method. T2K is the first long-baseline experiment to use an off-axis beam configuration, and we have found that the flux is rather more sensitive to systematic errors in this configuration than on axis. However, we have shown that to first order, the matrix and ratio methods of extrapolation can correct for these systematics, and that the residual errors will not limit the measurement potential of the experiment.

6.2 Beyond T2K

The future direction of the international neutrino program will depend on the results of the current generation of experiments, including T2K. If a positive measurement of θ_{13} is made, the next task will be to measure δ. If θ_{13} is reasonably close to the current limit, such a measurement could likely be made by an extended T2K experiment, either by building a detector at the second oscillation maximum [73] [74], or by building a larger detector at the first maximum and running with antineutrinos.

If θ_{13} cannot be measured by T2K or its contemporaries, it may be pos-

sible to make a measurement by extending T2K, with a more powerful beam, to collect more statistics. In order to reduce systematic errors associated with flux extrapolation, a second near detector would need to be built further away from the beam source than ND280, which would sample essentially the same solid angle as Super-K. Ultimately, however, the measurement which T2K can make will be limited by the errors associated with the far detector measurements, and residual uncertainties in the beam flux.

The errors associated with the far detector stem from the fairly crude reconstruction which is possible at Super-K. Water Čerenkov detectors are highly scalable and able to provide large fiducial masses at a manageable cost; however tracking, energy measurement and particle identification are limited. A next-generation neutrino detector, whether built as an extension to T2K or as part of a new project, is likely to be based on a different technology. The leading idea for such a detector is a Liquid Argon TPC [75] [76], which would use a combination of scintillation light and ionisation charge readout to reconstruct particle interactions. There are many technical challenges in developing economical and effective readout systems for LAr-TPCs, and in scaling them up to the required masses, but they would allow much more detailed reconstruction than water Čerekovs.

The systematics associated with flux uncertainties can be reduced by using a near detector; however another option is to produce a beam which is better understood to begin with. The next long-baseline project is likely to use a different kind of beamline, based on the decay of particles in a storage ring to produce neutrinos. The leading proposals for such a beam are a Neutrino Factory [77],

which would use the decays of muons, and a Beta Beam [78], using beta decays of radioactive ions.

These accelerator and detector R&D projects are currently very active, and, along with the results from the current set of neutrino projects, will determine where the field of experimental neutrino physics will go in the future.

Bibliography

[1] M. Haigh et al. Monte Carlo simulation of MPPC photosensors for the T2K experiment. In *Proceedings, International workshop on new photon detectors (PD09)*, 2009.

[2] J. Lesgourgues and S. Pastor. Massive neutrinos and cosmology. *Phys. Rept.*, 429:307–379, 2006.

[3] Ch. Weinheimer et al. High precision measurement of the tritium β spectrum near its endpoint and upper limit of the neutrino mass. *Phys. Lett. B*, 460:219–226, 1999.

[4] V. M. Lobashev. Study of the tritium beta-spectrum in experiment Troitsk ν-mass. *Prog. Part. Nucl. Phys.*, 48:123–131, 2002.

[5] C. I. Cowan et al. Detection of the free neutrino: A confirmation. *Science*, 124:103–104, 1956.

[6] G. Danby et al. Observation of high-energy neutrino reactions and the existance of two kinds of neutrinos. *Phys. Rev. Lett.*, 9:36–44, 1962.

[7] C. Amsler et al. (Particle Data Group). Review of Particle Physics. *Phys. Lett. B*, 667, 2008.

[8] K. Kodama et al. (DONUT Collaboration). Observation of tau neutrino interactions. *Phys. Lett. B*, 504:218–224, 2001.

[9] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *Sov. Phys. JETP*, 7:172–173, 1958.

[10] Z. Maki, M. Nakagawa, and S. Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870–880, 1962.

[11] R. Davis, Jr., D. S. Harmer, and K. C. Hoffman. Search for neutrinos from the Sun. *Phys. Rev. Lett.*, 20:1205–1209, 1968.

[12] J. N. Bahcall, N. A. Bahcall, and G. Shaviv. Present status of the theoretical predictions for the ^{36}Cl solar-neutrino experiment. *Phys. Rev. Lett.*, 20:1209–1212, 1968.

[13] J. N. Abdurashitov et al. (SAGE Collaboration). Results from SAGE (the Russian-American Gallium solar neutrino Experiment). *Phys. Lett. B*, 328:234–248, 1994.

[14] W. Hampel et al. (GALLEX Collaboration). GALLEX solar neutrino observations: results for GALLEX IV. *Phys. Lett. B*, 447:127–133, 1999.

[15] K. S. Hirata et al. (Kamiokande Collaboration). Real time, directional measurement of ^{8}B solar neutrinos in the Kamiokande II detector. *Phys. Rev. D*, 44:2241–2260, 1991.

[16] Q. R. Ahmad et al. (SNO Collaboration). Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301, 2002.

[17] P. Lipari and M. Lusignoli. On exotic solutions of the atmospheric neutrino problem. *Phys. Rev. D*, 60:013003, 1999.

[18] G. L. Fogli, E. Lisi, A. Marrone, and D. Montanino. Status of atmospheric $\nu_\mu \rightarrow \nu_\tau$ oscillations and decoherence after the first K2K spectral data. *Phys. Rev. D*, 67:093006, 2003.

[19] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 2007.

[20] Kai Zuber. *Neutrino Physics*. Taylor & Francis Group, 2004.

[21] B. Kayser. On the quantum mechanics of neutrino oscillation. *Phys. Rev. D*, 24:110–116, 1981.

[22] L. Wolfenstein. Neutrino oscillations in matter. *Phys. Rev. D*, 17:2369–2374, 1978.

[23] S. P. Mikheev and A. Yu. Smirnov. Resonant amplification of neutrino oscillations in matter and solar neutrino spectroscopy. *Nuovo Cim. C*, 9:17–26, 1986.

[24] C. Andreopoulos et al. The GENIE Neutrino Monte Carlo Generator. *Nucl. Instrum. Meth.*, A614:87–104, 2010.

[25] P. Lipari et al. The neutrino cross section and upward going muons. *Phys. Rev. Lett.*, 74:4384, 1995.

[26] T. Araki et al. (KamLAND Collaboration). Measurement of neutrino oscillation with KamLAND: Evidence of spectral distortion. *Phys. Rev. Lett.*, 94:081801, 2005.

[27] Y. Fukuda et al. (Super-Kamiokande Collaboration). Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.

[28] S. Fukuda et al. (Super-Kamiokande Collaboration). The Super-Kamiokande detector. *Nucl. Instr. Meth. A*, 501:418–462, 2003.

[29] Super–K website.
http://www-sk.icrr.u-tokyo.ac.jp/sk/index1.html.

[30] S. Fukuda et al. (Super-Kamiokande Collaboration). Tau neutrinos favored over sterile neutrinos in atmospheric muon neutrino oscillations. *Phys. Rev. Lett.*, 85:3999–4003, 2000.

[31] K. Abe et al. (Super-Kamiokande Collaboration). A measurement of atmospheric neutrino flux consistent with tau neutrino appearance. *Phys. Rev. Lett.*, 97:171801, 2006.

[32] J. Hosaka et al. (Super-Kamiokande Collaboration). Three flavor neutrino oscillation analysis of atmospheric neutrinos in Super-Kamiokande. *Phys. Rev. D*, 74:032002, 2006.

[33] M. H. Ahn et al. (K2K Collaboration). Measurement of neutrino oscillation by the K2K experiment. *Phys. Rev. D*, 74:072003, 2006.

[34] M. G. Catanesi et al. (HARP Collaboration). Measurement of the production cross-section of positive pions in p–Al collisions at 12.9 GeV/c. *Nucl. Phys. B*, 732:1–45, 2006.

[35] S. Yamamoto et al. (K2K Collaboration). Improved search for $\nu_\mu \rightarrow \nu_e$ oscillation in a long-baseline accelerator experiment. *Phys. Rev. Lett*, 96:181801, 2006.

[36] P. Adamson et al. (MINOS Collaboration). Study of muon neutrino disappearance using the Fermilab Main Injector neutrino beam. *Phys. Rev. D*, 77:072002, 2008.

[37] P. Adamson et al. (MINOS Collaboration). Measurement of neutrino oscillations with the MINOS detectors in the NuMI beam. *Phys. Rev. Lett.*, 101:131802, 2008.

[38] M. Apollonio et al. Initial results from the CHOOZ long baseline reactor neutrino oscillation experiment. *Phys. Lett. B*, 420:397–404, 1998.

[39] F. Boehm et al. Results from the Palo Verde neutrino oscillation experiment. *Phys. Rev. D*, 62:072002, 2000.

[40] F. Boehm et al. Final results from the Palo Verde neutrino oscillation experiment. *Phys. Rev. D*, 64:112001, 2001.

[41] M. Apollonio et al. Limits on neutrino oscillations from the CHOOZ experiment. *Phys. Lett. B*, 466:415–430, 1999.

[42] C. Athanassopoulos et al. (LSND Collaboration). Evidence for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillation from the LSND experiment at the Los Alamos Meson Physics Facility. *Phys. Rev. Lett.*, 77:3082–3085, 1996.

[43] A. A. Aguilar-Arevalo et al. (MiniBooNE Collaboration). A search for electron neutrino appearance at the $\delta m^2 \sim 1\ \mathrm{eV}^2$ scale. *Phys. Rev. Lett.*, 98:231801, 2007.

[44] A. A. Aguilar-Arevalo et al. (MiniBooNE Collaboration). A search for muon neutrino and antineutrino disappearance in MiniBooNE. *Phys. Rev. Lett.*, 103:061802, 2009.

[45] B. T. Fleming et al. (MicroBooNE Collaboration). *A Proposal for a New Experiment Using the Booster and NuMI Neutrino Beamlines: MicroBooNE.* *http://www-microboone.fnal.gov/Documents/MicroBooNE_10152007.pdf*, October 2007.

[46] Shi-Wen Li and Bo-Qiang Ma. Unitarity boomerangs of quark and lepton mixing matrices. *Phys. Lett.*, B691:37–45, 2010.

[47] The NOνA Collaboration. *Proposal to build a 30 kiloton off-axis detector to study $\nu_\mu \rightarrow \nu_e$ oscillations in the NuMI beamline.* *http://arxiv.org/abs/hep-ex/0503053*, 2005.

[48] X. Guo et al. (Daya Bay Collaboration). *A precision measurement of the neutrino mixing angle θ_{13} using reactor antineutrinos at Daya Bay.* $http://arxiv.org/abs/hep\text{-}ex/0701029$, 2007.

[49] F. Ardellier et al. (Double CHOOZ Collaboration). *Double CHOOZ: A search for the neutrino mixing angle θ_{13}.* $http://arxiv.org/abs/hep\text{-}ex/0606025$, 2006.

[50] Mauro Mezzetto. Next Challenge in Neutrino Physics: the θ_{13} Angle. *arXiv:0905.2842 [hep-ph]*, 2009.

[51] K. Nishikawa et al. (T2K Collaboration). *Letter of Intent: Neutrino Oscillation Experiment at JHF.* $http://neutrino.kek.jp/jhfnu/loi/loi.v2.030528.pdf$, 2003.

[52] M. Fukugita and T Yanagida. Baryogenesis without grand unification. *Phys. Lett. B*, 174:45–47, 1986.

[53] S. Antusch et al. Perturbative estimates of lepton mixing angles in unified models. *Nucl. Phys. B*, 820:32–46, 2009.

[54] T. Kobayashi. *J-PARC Neutrino Beam. Talk given at the Workshop on Next Generation Nucleon Decay and Neutrino Detectors (NNN09), Colorado, October 2009.*

[55] T2K-ND280 Collaboration. *Document for ND280 Review (Internal T2K Document)*, 2006.

[56] H. Nishino et al. The new front-end electronics for the Super–Kamiokande experiment. *Conference Record, IEEE Nuclear Science Symposium*, pages 127–132, 2007.

[57] Joshua Albert. *Plots of T2K Sensitivity to θ_{13} for the 30 GeV ν_μ beam (T2K Internal Note)*, 2009.

[58] K. Sakashita. *Status of the T2K Experiment. Talk given at the Europhysics Conference on High-Energy Physics (EPS HEP 2009), Krakow, June 2009.*

[59] Solid State Division Hamamatsu K.K. *MPPC Technical Information*, 2009.

[60] Charles Kittel. *Introduction to Solid State Physics*. John Wiley and Sons, eighth edition, 2004.

[61] S. Cova et al. Trapping phenomena in avalanche photodiodes on nanosecond scale. *IEEE Electron Device Letters*, 12:685–687, 1991.

[62] A. Vacheret et al. The front end readout system for the T2K-ND280 detectors. In *Proceedings, IEEE Nuclear Science Symposium*, 2007.

[63] National Instruments. *NI Multisim information.* $http://www.ni.com/multisim$.

[64] Mark Ward. *Multi-Pixel Photon Counter characterisation for the T2K near detector electromagnetic calorimeter*. PhD thesis, University of Sheffield, UK, 2009.

[65] National Semiconductor. *LM92 Datasheet*, 2005.

[66] R. Brun et al. *ROOT Users Guide 5.21.* *ftp://root.cern.ch/root/doc/Users_Guide_5_21.pdf*, Dec 2008.

[67] M. Morhac et al. Identification of peaks in multidimensional coincidence gamma-ray spectra. *Nucl. Instr. Meth. A*, 443:108–125, 2000.

[68] Alain Blondel. Muon polarisation in the neutrino factory. *Nucl. Instr. Meth. A*, 451(1):131 – 137, 2000.

[69] T. A. Gabriel et al. *ORNL technical note, unpublished*, 1977.

[70] A. Fass et al. The physics models of FLUKA: status and recent developments. In *Proceedings, CHEP03, La Jolla*, March 2003.

[71] C. Zeitnitz and T. A. Gabriel. *ORNL technical note, unpublished*, 1996.

[72] Boris A. Popov et al. (HARP and NA61 Collaborations). HARP and NA61 (SHINE) hadron production experiments. In *Proceedings, Sixth international workshop on neutrino-nucleus interactions in the few-GeV region (NUINT09)*, 2009.

[73] Masaki Ishitsuka, Takaaki Kajita, Hisakazu Minakata, and Hiroshi Nunokawa. Resolving the neutrino mass hierarchy and CP degeneracy by two identical detectors with different baselines. *Phys. Rev. D*, 72:033003, 2005.

[74] F. Dufour. T2KK, a Mton water Čerenkov detector in the line of the T2K beam. *Nuclear Physics B - Proceedings Supplements*, 188:201 – 203, 2009. Proceedings of the Neutrino Oscillation Workshop.

[75] C. Rubbia. The liquid-argon time projection chamber: A new concept for neutrino detector. *CERN-EP/77-08*, 1977.

[76] S. Amerio at al. (ICARUS collaboration). Design, construction and tests of the ICARUS T600 detector. *Nucl. Instrum. Meth. A*, 527:329–410, 2004.

[77] J.S. Berg et al. *International Scoping Study of a future neutrino factory and super-beam facility – Accelerator design concepts for future neutrino facilities*.

[78] P. Zucchelli. A novel concept for a $\bar{\nu}_e/\nu_e$ neutrino factory: the beta-beam. *Phys. Lett. B*, 532:166–172, 2002.