

Determining the Basic Properties of a New Particle Detector Concept for the T2K Experiment and Implementing the T2K Calculation for a Three-Neutrino Global Fit with GAMBIT

Wilf Shorrock

High Energy Physics,
Imperial College London

A Dissertation Submitted to Imperial College London
for the Degree of Doctor of Philosophy

Determining the Basic Properties of a New Particle Detector Concept for the T2K Experiment and Implementing the T2K Calculation for a Three-Neutrino Global Fit with GAMBIT

Abstract

Neutrino oscillations are of significant interest to the physics community as studying them could reveal much about the origins of the Universe. Tokai-to-Kamioka (T2K) is a long baseline neutrino oscillation experiment dedicated to measuring the neutrino oscillation parameters. A proposal to upgrade the near-detector of T2K to improve its systematic uncertainties has been made and research is ongoing to develop and build these detectors for installation in 2022. The work presented herein gives the methods and results of the analysis regarding the new detector design for the near-detector upgrade, the SuperFGD, and its multiple beam tests. Using the data from these beam tests, this thesis shows how some of the basic properties of the detector were measured. For example, the time resolution of the detector, measured as 1.14 ± 0.06 ns, which is in agreement with the predicted result from combining the time resolution of the readout electronics and the scintillator material. The results in this thesis regarding the SuperFGD Prototype have been an invaluable resource for the development and simulation of the SuperFGD detector to be installed in T2K. Also detailed in this work is the progress of a three-neutrino global fit implementation within the global fitting software GAMBIT. GAMBIT aims to improve upon previous neutrino global fits by sourcing information from more reliable resources than theses and presentations. The results shown here demonstrate how using certain internal information from T2K can help improve global fits, with the hope that T2K can publicly release this information in the future.

Declaration

This dissertation is the result of my own work with some contribution and collaboration from others as part of a larger group. All figures and their contents are produced by me unless a citation to the original work is included in the figure caption. Chapters 1 to 4 contain information I have gathered to give context to my work, but there are no analyses in those chapters directly performed by me. All analyses in Chapter 5 are my own except for the afterpulse and light attenuation analyses, which were mainly performed by other members of the SuperFGD Prototype analysis group and are included here to give a more complete picture of the prototype. The contents of Chapter 6 is all my own work, but was developed in the GAMBIT framework which was created by GAMBIT collaboration.

This dissertation has not been submitted for another qualification to this or any other university and does not exceed the word limit for the respective Degree Committee.

Wilf Shorrock

©The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial-No Derivatives 4.0 International Licence (CC BY-NC-ND). Under this licence, you may copy and redistribute the material in any medium or format on the condition that; you credit the author, do not use it for commercial purposes and do not distribute modified versions of the work. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Acknowledgements

There are so many people to thank from the course of my PhD that I'm worried I will forget a few. I'll start by thanking everyone on T2K. It was a pleasure to work on such an amazing experiment, not just because of the fascinating science and the cutting-edge detectors, but the community as a whole.

I would also like to thank everyone at Imperial College London for their input, ideas and company over the past few years. I loved working with everyone and it was great to hang out in Blackett 530 with you all.

Of course, massive appreciation to my supervisor, Yoshi, who managed to find time for me despite so many responsibilities. I always tried to make sure my plots were to his rigorous standards and I hope the plots in my thesis reflect that.

Some of the people I worked with most closely were my colleagues working on the SuperFGD. Several weeks and months were spent together at CERN and LANL and it was all made easier working with people like them. A particular thanks to my co-beamrun coordinators Guang Yang and David Last from the LANL beam test. Also, the writing team for the CERN beam test paper, including Etam Noah Messomo, Dana Douqa, Cesar Jesús-Valls and Kamil Skwarczynski. Etam especially for leading us so well and for his tireless efforts getting the SuperFGD Prototype to work.

And I would be hard put to forget to thank the wonderful people I have worked with on GAMBIT. Although I may never even have seen most of your faces, you've all helped me so much and always given amazingly useful advice and assistance. Anders Kvellstadt was a knight in shining armour with his ability to always be able to help with problems and even run scans for me despite his many responsibilities with the GAMBIT community. Speaking of running scans, thank you to the Joliot-Curie supercomputer cluster in France for allowing Anders to run my scans there. Without them, I would be running scans to this day trying to get acceptable plots.

All my love and affection goes out to my family as well. They have lent me so much support throughout the years, despite maybe not quite understanding what my work is about. I hope they know how much I appreciate what they have done for me in helping me get to where I am now. Without them, I would most certainly not have the confidence to finish a PhD!

Finally, the biggest thanks and appreciation goes to my partner, Emily, forever being the voice of reason and the greatest person in the world to live with.

Contents

Abstract	2
Declaration	3
Acknowledgements	4
1 Introduction	15
2 Neutrino Physics	18
2.1 The Neutrino: A History	18
2.2 Neutrino Interactions	20
2.2.1 CC0 π	20
2.2.2 CC1 π	21
2.3 Neutrino Oscillations	22
2.4 Neutrino Oscillations in Matter	24
2.5 Neutrino Oscillation Experiment Overview	25
2.6 Remaining Questions	27
3 The T2K Experiment	28
3.1 Neutrino Beam and Flux Constraints	29
3.2 Super-Kamiokande	31
3.3 ND280	33
3.3.1 π^0 Detector (PØD)	34
3.3.2 Fine Grain Detectors (FGDs)	36
3.3.3 Time Projection Chambers (TPCs)	36
3.3.4 Electronic Calorimeter (ECal)	38
3.4 ND280 Performance	38

4 T2K-II	42
4.1 Beamline Upgrade	42
4.2 ND280 Upgrade	43
4.2.1 High Angle TPCs	44
4.2.2 Time-of-Flight Detectors	45
4.2.3 The SuperFGD	46
4.3 Predicted ND280 Upgrade Performance	47
4.4 Simulated Performance of the SuperFGD	50
5 The SuperFGD Prototype	54
5.1 Prototype Scintillator Cubes	54
5.2 WLS Fibres and Photosensors	55
5.3 Readout Electronics	57
5.4 Assembly of the SuperFGD Prototype	59
5.5 Calibration	65
5.6 The CERN Particle Beam Test	69
5.7 CERN Beam Test Data Analysis: Basic Detector Properties	71
5.7.1 Particle Events	72
5.7.2 Hit Time Structure	74
5.7.3 Hit Amplitude Thresholds	77
5.7.4 Channel Response	79
5.7.5 Optical Crosstalk Between Adjacent Cubes	84
5.7.6 Light Attenuation in WLS Fibre	90
5.7.7 Cube Response	92
5.7.8 Signal Travel-Time and Time-Walk Corrections	93
5.7.9 Time Resolution	96
5.8 The LANL Particle Beam Test	100
5.9 Temperature Calibration at LANL	101
5.10 LANL Beam Test Data Analysis: Neutron Response	103

5.10.1	Observed Beam Structure and Energy Calculation	103
5.10.2	Particle Events and Crosstalk Identification	105
5.10.3	Event Reconstruction and Neutron Selection	109
5.10.4	Neutron Energy and Cross-section Measurements	111
5.11	Further Analysis	115
6	Neutrino Oscillation T2K Fit with GAMBIT	118
6.1	The GAMBIT Software	119
6.2	NuFit’s Global Fit of Neutrino Oscillation Data	120
6.3	GAMBIT’s Global Fit of Neutrino Oscillation Data	120
6.4	Likelihood Calculation for T2K	121
6.5	Including Systematic Uncertainties	131
6.6	The T2K Likelihood Scanning Algorithm: Differential Evolution	132
6.7	Constraints on the Oscillation Parameter Space	134
6.7.1	No Variation of Systematic Nuisance Parameters	135
6.7.2	Flux Parameters Variation	137
6.7.3	Reconstruction Function Parameters Variation	139
6.7.4	Flux and Reconstruction Function Parameters Variation	142
6.8	Further Work	143
6.9	Next Generation Experiments	146
6.9.1	Impact of T2K-II	147
7	Conclusion and Remarks	150
A	SuperFGD Prototype FEB and Channel Mapping	152
B	SuperFGD Prototype DAQ Settings	155
C	GAMBIT T2K Likelihood Reconstruction Functions	157
C.1	T2K Resources Used in GAMBIT	159
	Bibliography	159

List of Tables

2.1	The sensitivities of the different neutrino oscillation experiments	26
2.2	The current values of the different oscillation parameters and their uncertainties	27
3.1	Systematic uncertainties at SK and how they are affected by ND280 constraints	34
4.1	The number of selected events and sample purities for the simulated analysis of the current and upgraded ND280 detector	49
4.2	Predicted uncertainty contributions for flux and cross-section constraints using the upgraded ND280	50
5.1	Main parameters for the three MPPC types installed on the SuperFGD Prototype	56
5.2	The trigger combinations used to trigger on specific particles and the purity of the resulting samples	71
5.3	The properties of afterpulse hits for the different MPPC types	76
5.4	Mean and standard deviation of the light yields obtained from the different MPPC types and fibre lengths	80
6.1	The Asimov A neutrino oscillation parameters	125
6.2	The neutrino oscillation parameters used when comparing to T2K's MC	125
6.3	The scanning ranges and prior constraints on the neutrino oscillation parameters used in the GAMBIT global fit	134
B.1	The DAQ settings used for the CERN and LANL particle beam tests .	155
C.1	CCQE nominal reconstruction function parameters	158
C.2	CCQE nominal reconstruction function parameters for the ν_e FHC CCQE data sample	158
C.3	$2p2h$ nominal reconstruction function parameters	158
C.4	CC1Pi nominal reconstruction function parameters	159
C.5	An exhaustive list of T2K resources used in the GAMBIT T2K likelihood calculation	160

List of Figures

2.1	A diagram of the AGS neutrino experiment at Brookhaven	19
2.2	Feynman diagrams of CC0 π processes.	21
2.3	Feynman diagrams of CC1 π processes.	22
3.1	The path of the neutrino beam produced in the T2K experiment	28
3.2	An overview of the T2K beamline and a side view of the secondary beamline	29
3.3	The latest publicly released flux predictions at T2K’s far detector for the FHC and RHC modes	30
3.4	A schematic of Super-Kamiokande	31
3.5	Two event displays from Super-Kamiokande showing a muon event and an electron event	32
3.6	An expanded view of the ND280 detector	35
3.7	An example ND280 event display	35
3.8	A sketch of the FGD structure	37
3.9	The reconstructed momentum versus the scattering angle for detected muons at ND280 and SK	39
3.10	The proton reconstruction efficiency in ND280	40
3.11	Time-of-flight cuts and reconstruction efficiencies for different angled tracks	40
3.12	The distribution of reconstructed energies of ν_e particles in the ND280 tracker	41
4.1	A diagram of the proposed upgrade of ND280	43
4.2	A schematic view of the HATPC design	44
4.3	A schematic view of the TOF detector situated around the HATPCs and scintillator target	45
4.4	A rendering of the SuperFGD to be used in the upgraded ND280	47
4.5	A comparison of the number of detected muons from simulations for the current ND280 and the upgraded ND280	48

4.6	The simulated detection efficiencies with respect to lepton scattering angle and momentum for both the current and upgrade ND280 sub-detectors	49
4.7	An event display of a simulated ν_μ CCQE event within the SuperFGD	51
5.1	A diagram of the optical coupling interface	57
5.2	Sketch of the three different signal outputs of the SuperFGD Prototype	59
5.3	The three main stages of the cube assembly for the SuperFGD Prototype	60
5.4	Photograph of the SuperFGD Prototype during construction	61
5.5	WLS fibres with the fibre ferrule installed at one end	62
5.6	The SuperFGD Prototype with WLS fibres installed and mounted in an aluminium frame	63
5.7	Distribution of the three types of MPPCs around the six faces of the SuperFGD Prototype	64
5.8	The SuperFGD Prototype with all readout electronics attached	65
5.9	Calibration plots used to extract ADC/p.e. ratios for the HG and LG readout channels	66
5.10	ToT calibration plots for the SuperFGD prototype	67
5.11	Gain distributions for different MPPC types	67
5.12	The pedestal finding method used when calibrating an MPPC	68
5.13	The algorithm used to decide which output channel was used to represent the signal amplitude	69
5.14	The set-up of the CERN particle beam test in the T9 beam-line	72
5.15	2D displays of hits detected in a muon event at the SuperFGD Prototype (no charge cut)	73
5.16	2D displays of hits detected in a muon event at the SuperFGD Prototype (with 20 p.e. charge cut)	73
5.17	2D displays of hits detected in a stopping proton event at the SuperFGD Prototype (no charge cut)	74
5.18	Hit distribution in the SuperFGD Prototype at the CERN particle beam test	75
5.19	SuperFGD Prototype hit amplitude thresholds as measured at the CERN particle beam test	78
5.20	Type I MPPC and 24 cm WLS fibre channel responses to muon particles	81
5.21	Type I MPPC and 48 cm WLS fibre channel responses to muon particles	81

5.22 Type I MPPC and 8 cm WLS fibre channel responses to muon particles	82
5.23 Type II MPPC and 8 cm WLS fibre channel responses to muon particles	82
5.24 Type III MPPC and 8 cm WLS fibre channel responses to muon particles	83
5.25 A diagram showing the distribution of light about the cube from which the light originated	86
5.26 Light yield measurements from main hits and crosstalk hits	87
5.27 Light yields measured using the z -fibres of the SuperFGD Prototype	89
5.28 The percentage of light leakage from a single cube as a result from crosstalk in the SuperFGD detector	90
5.29 Measurements of light attenuation along WLS fibres in the SuperFGD Prototype	91
5.30 Cube light yields in the SuperFGD Prototype with attenuation corrections.	93
5.31 A diagram showing how signal amplitude can affect the rising edge timing (time-walk)	94
5.32 A fit of the time-walk model to a single channel	95
5.33 Mapping of the measured channel time resolutions in the SuperFGD Prototype	98
5.34 Histograms of the measured channel time resolutions	99
5.35 Temperature readings with the SuperFGD Prototype	103
5.36 The beam structure of the LANL neutron beam as observed by the SuperFGD Prototype	104
5.37 Neutron event where all charged particles were contained within one cube and no crosstalk cuts are applied	106
5.38 A neutron event where a charged particle is emitted from the interaction cube and no crosstalk cuts are applied	106
5.39 MC and real data before and after the crosstalk algorithm is applied	107
5.40 The algorithm to determine if a light signal originates from the main charge deposit of a charged particle or from the optical crosstalk of said light deposit into neighbouring cubes	109
5.41 Neutron event where all charged particles were contained within one cube and crosstalk cuts are applied using the crosstalk identification algorithm	110
5.42 A neutron event where a charged particle is emitted from the interaction cube and crosstalk cuts are applied using the crosstalk identification algorithm	110
5.43 Neutron energy spectra from real data and MC data	112

5.44	The neutron beam attenuation as it travels through the SuperFGD Prototype	114
5.45	The energy dependent cross-section for all neutron-scintillator interactions as measured using the LANL beam test data	114
5.46	The transverse spread of the simulated beam along the SuperFGD prototype with different invisible scattering weights	116
6.1	Simulated data showing the distribution of reconstructed energies for neutrino events of a single true energy	124
6.2	SK energy spectra of neutrino events as predicted by GAMBIT for parameter space point Asimov A	126
6.3	SK energy spectra of neutrino events as predicted by GAMBIT for parameter set 1	127
6.4	SK energy spectra of neutrino events as predicted by GAMBIT for parameter set 2	128
6.5	SK energy spectra of neutrino events as predicted by GAMBIT for parameter set 3	129
6.6	SK energy spectra of neutrino events as predicted by GAMBIT for parameter set 4	130
6.7	The run1-10 T2K oscillation analysis contours compared to the GAMBIT contours when nuisance parameters are kept fixed using the NO mass hierarchy	135
6.8	The run1-10 T2K oscillation analysis $\delta_{\text{CP}}\text{-}\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when nuisance parameters are kept fixed using the NO mass hierarchy	136
6.9	The run1-10 T2K oscillation analysis $\Delta m_{32}^2\text{-}\sin^2(\theta_{23})$ contours compared to the GAMBIT contours when nuisance parameters are kept fixed using the IO mass hierarchy	136
6.10	The run1-10 T2K oscillation analysis $\delta_{\text{CP}}\text{-}\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when nuisance parameters are kept fixed using the IO mass hierarchy	137
6.11	The run1-10 T2K oscillation analysis contours compared to the GAMBIT contours when flux nuisance parameters are profiled over using the NO mass hierarchy	138
6.12	The run1-10 T2K oscillation analysis $\delta_{\text{CP}}\text{-}\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when flux nuisance parameters are profiled over using the NO mass hierarchy	138
6.13	The run1-10 T2K oscillation analysis $\Delta m_{32}^2\text{-}\sin^2(\theta_{23})$ contours compared to the GAMBIT contours when flux nuisance parameters are profiled over using the IO mass hierarchy	139

6.14	The run1-10 T2K oscillation analysis $\delta_{\text{CP}}\text{-}\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when flux nuisance parameters are profiled over using the IO mass hierarchy	139
6.15	The run1-10 T2K oscillation analysis contours compared to the GAMBIT contours when reconstruction function nuisance parameters are profiled over using the NO mass hierarchy	140
6.16	The run1-10 T2K oscillation analysis $\delta_{\text{CP}}\text{-}\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when reconstruction function nuisance parameters are profiled over using the NO mass hierarchy	140
6.17	The run1-10 T2K oscillation analysis $\Delta m_{32}^2\text{-}\sin^2(\theta_{23})$ contours compared to the GAMBIT contours when reconstruction function nuisance parameters are profiled over using the IO mass hierarchy	141
6.18	The run1-10 T2K oscillation analysis $\delta_{\text{CP}}\text{-}\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when reconstruction function nuisance parameters are profiled over using the IO mass hierarchy	141
6.19	The run1-10 T2K oscillation analysis contours compared to the GAMBIT contours when flux and reconstruction function nuisance parameters are profiled over using the NO mass hierarchy	142
6.20	The run1-10 T2K oscillation analysis $\delta_{\text{CP}}\text{-}\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when flux and reconstruction function nuisance parameters are profiled over using the NO mass hierarchy	143
6.21	The run1-10 T2K oscillation analysis $\Delta m_{32}^2\text{-}\sin^2(\theta_{23})$ contours compared to the GAMBIT contours when flux and reconstruction function nuisance parameters are profiled over using the IO mass hierarchy	143
6.22	The run1-10 T2K oscillation analysis $\delta_{\text{CP}}\text{-}\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when flux and reconstruction function nuisance parameters are profiled over using the IO mass hierarchy	144
6.23	A comparison of the GAMBIT calculated contours when systematic uncertainties are left unvaried and when they are varied	144
6.24	GAMBIT's predictions for T2K's parameter space confidence limits with current and upgraded POT values and systematic uncertainties	148
6.25	GAMBIT's confidence limits for the δ_{CP} oscillation parameter using the current and upgraded POT values and systematic uncertainties	149
A.1	The mapping of FEB number to mounted MPPCs	152
A.2	Mapping of channel numbers and cable bundles	153

Chapter 1

Introduction

Progressing beyond the Standard Model (SM) of particle physics is a necessary step towards understanding our universe. The SM has been the foundation of particle physics for decades, but there are still several fundamental questions left unanswered. We cannot explain the presence of dark matter or the imbalance between matter and anti-matter that we currently see in our Universe. To find answers to these questions, there are several experiments searching for Beyond the Standard Model (BSM) physics. There are also many experiments developing our understanding of known BSM physics phenomena. An example of such phenomena are neutrino oscillations, where a neutrino with a particular flavour is observed to have changed to a different flavour after some propagation. These oscillations have long been accepted as real processes by the scientific community, but the ethereal nature of the neutrino with respect to all other forms of matter makes it difficult to examine these flavour changes.

The field of neutrino oscillation experiments has a rich history and a promising future with several next generation experiments, such as Hyper-Kamiokande [1] and DUNE [2], currently under development. For several decades, experiments have examined neutrino oscillations from one of four main sources: the Sun, Earth's atmosphere, nuclear reactors and particle accelerators. All these experiments share the same ultimate goal to build a greater understanding of neutrinos and their behaviour. This understanding hinges on several smaller goals: determining if neutrinos and anti-neutrinos oscillate differently, finding the correct mass ordering of the neutrino masses, accurately measuring all parameters in the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix that characterise neutrino oscillations and determining if the PMNS matrix is sufficient to describe observed properties [3–5].

One of these goals may be realised in the near future, with the T2K neutrino experiment about to enter a new phase of data-taking named T2K-II [6, 7]. This new phase aims to

bring clarity to the matter-antimatter asymmetry problem [8]. Such an asymmetry could be explained by charge conjugation parity (CP) violation within the leptonic sector. The phenomenon of CP violation, where an interaction is not observed to be symmetrical to inverting both charge and spatial coordinates, has so far only been confirmed in hadronic interactions. These include kaon interactions [9–11], B meson interactions [12–14] and charmed D^0 decays [15], the latter being a relatively recent discovery. CP violation in quarks alone is not enough to explain the matter-antimatter imbalance problem [16, 17]. T2K recently published results showing CP conservation within neutrino oscillation excluded at 3σ and aim to improve this to 5σ after reducing systematic errors with increased statistics and upgraded detectors. The confirmation that neutrino oscillations differ for neutrinos and anti-neutrinos would be a significant step towards an accepted model of leptogenesis, which requires that neutrinos violate CP conservation [17, 18]. CP violations in the neutrino sector would still not give a full picture of baryon asymmetry, as another requirement is that neutrinos are Majorana particles [19]. There is still uncertainty as to whether neutrinos are Fermi particles or Majorana particles.

The need to reduce T2K and Hyper-Kamiokande’s detector systematics to a point where such a 5σ zero-valued δ_{CP} exclusion is possible has spurred the design of a completely new particle detector concept, the SuperFGD. The SuperFGD concept has the potential to be included in many other experiments in the future, so the work put into understanding the basic properties and limitations of the technology now can influence the progress of physics experiments for years to come. In this thesis, we will detail the beam tests of prototype SuperFGD technology, as well as a new global fit of three-neutrino oscillation data and the potential impacts of future neutrino experiments on global fits.

T2K has made some of the most precise measurements of the PMNS parameters and therefore plays an important role in global fits of neutrino oscillation data. In these global fits, results from several experiments are combined to get a broader view of parameter constraints. This thesis will present a method for incorporating T2K’s results into future global fits, using the GAMBIT global fitting framework as an example. The goal of this method is to improve upon previous global fitting techniques by a more rigorous approach to detector effects and near-detector fitting. This GAMBIT example will eventually be included in GAMBIT’s own global fit of neutrino oscillation experiments, pending on permission from T2K to use provided internal information.

We have structured this thesis to first present a detailed history of neutrino physics

before then turning to the physics of neutrino oscillations in Chapter 2. We then outline the T2K experiment—including its goals, beam production and detectors—in Chapter 3. In Chapter 4 we move on to the T2K’s new data taking phase, T2K-II, where the planned upgrades to T2K detectors and beam delivery are outlined with focus on the new SuperFGD design. We expand upon the SuperFGD in Chapter 5 by describing the construction and development of the SuperFGD Prototype, along with the analysis of particle beam data to determine the detector properties. Details of neutrino global fits are given in Chapter 6, with information on the GAMBIT global fitting software and the T2K likelihood calculation implemented in GAMBIT. A discussion of the impact of T2K-II on global fits is also given in Chapter 6. We then conclude the thesis in Chapter 7 with some closing remarks.

Chapter 2

Neutrino Physics

2.1 The Neutrino: A History

The construction of the theoretical framework predicting the neutrino particle and its subsequent discovery is a classic example of theory and experiment. As has been the case several times in particle physics, the neutrino particle was theorised many years before its discovery. In 1930, Wolfgang Pauli proposed the existence of a new unknown particle to explain the lack of energy and momentum conservation in observed β decays [20]. Then in 1932, James Chadwick discovered the neutron and realised it was the progenitor of β decay processes, thus setting out the theory that β decay was the decay of a neutron to a proton and electron, $n \rightarrow p + e^-$ [21]. This model from Chadwick was later combined with Pauli's new particle by Enrico Fermi in 1934 [22]. Fermi dubbed this particle the “neutrino” and he managed to develop a model incorporating the neutrino into β decay and accommodating for Dirac's positron particle [23] and Werner Heisenberg's neutron-proton model [24–26].

The actual discovery of the neutrino would not occur until 1956 when inverse beta decay, $\bar{\nu}_e + p \rightarrow n + e^+$, was observed by Reines and Cowan [27, 28]. The discovery was made by observing the production of positrons and free neutrons in a Cd-doped water solution near a nuclear reactor. Inverse beta decays produce a prompt signal in the detector through positron annihilation ($e^+ + e^- \rightarrow 2\gamma$). This prompt signal is followed by a neutron capture event ($n + {}^{108}\text{Cd} \rightarrow {}^{109\text{m}}\text{Cd} \rightarrow {}^{109}\text{Cd} + \gamma$) within 5 μs . The rate of detected inverse beta decay events dropped when the nuclear reactor was turned off. Said drop in event rate corresponded with the predicted event rate drop if the reactor was producing neutrinos. The delayed-coincidence method used in this initial discovery closely resembles present-day neutrino experiments utilising nuclear reactors.

Following the initial detection of neutrinos, a second type of neutrino was detected in

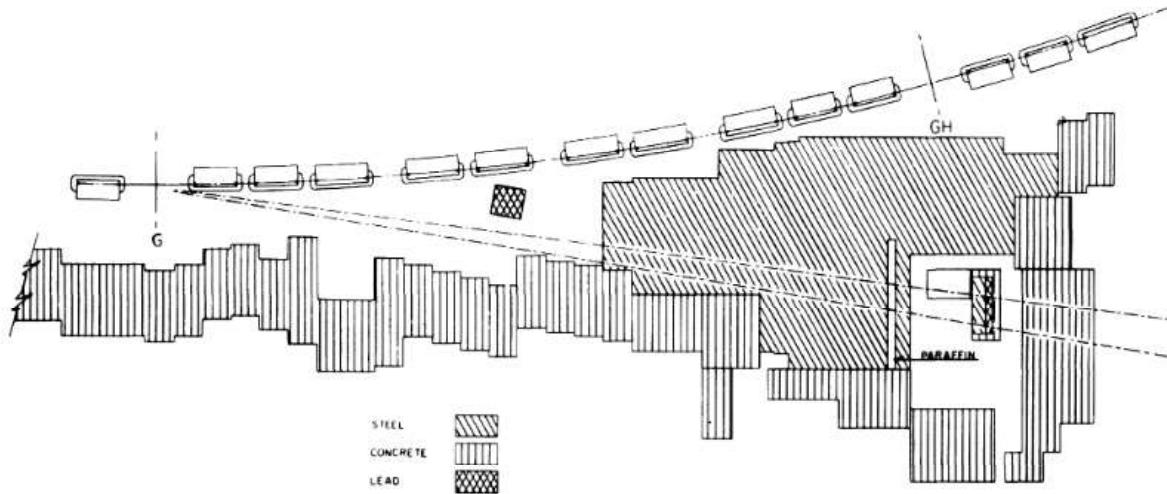


Figure 2.1: A diagram of the AGS neutrino experiment at Brookhaven [29].

1962 using the Alternating Gradient Synchrotron (AGS) at Brookhaven [29]. A plan view of the experiment is shown in Fig. 2.1. A beam of charged pions was produced by directing a beam of protons from the AGS to strike a beryllium target. These pions then decayed to produce muons and neutrinos ($\pi^\pm \rightarrow \mu^\pm + \nu(\bar{\nu})$). An aluminium spark chamber was placed 25 m from the target in the path of the beam of muons and neutrinos. The spark chamber was shielded by 13.5 m of steel, ensuring only the neutrinos could reach the spark chamber. The resulting neutrino interactions observed in the spark chamber only produced muons and not electrons, a behaviour distinctly different to the neutrinos produced by Reines and Cowan. This behaviour could be explained by assigning a muon lepton number of +1 (−1 for anti-neutrinos) to the neutrinos detected at AGS and an electron number of +1 (−1 for anti-neutrinos) to the neutrinos detected by Reines and Cowan. These lepton numbers are conserved in interactions, hence why the AGS neutrinos produced muons, also with a muon lepton number of +1. Given the clear difference between leptonic flavours in the interactions of the Reines and Cowan neutrinos and the AGS neutrinos, they have come to be identified as electron neutrinos (ν_e) and muon neutrinos (ν_μ) respectively.

Given that the electron and muon both had their own neutrino counterpart, it was only logical to hypothesise the existence of a tau neutrino after the discovery of the tau lepton in the 1970s at the Stanford Linear Accelerator Center (SLAC) [30]. Given the high mass of the tau particle and its short lifetime, it is much more challenging to detect tau neutrinos than electron and muon neutrinos. It wasn't until the year 2000 that a

direct observation of the tau neutrino was announced by the DONUT (Direct Observation of the Nu Tau) experiment at Fermilab [31].

2.2 Neutrino Interactions

Neutrinos can interact with matter in a number of ways, but the small cross-sections of the interactions lead to difficulty in cross-section measurements. As such, despite ongoing efforts to measure neutrino cross-sections, there is still significant uncertainty associated with the measurements. This also makes it challenging to decide on the correct neutrino interaction models to use in neutrino interaction generators. Because of these uncertainties, there is a significant contribution to experimental systematic errors in neutrino experiments from the neutrino cross-section values.

The numerous different interaction channels a neutrino can take can also lead to uncertainties in neutrino energy reconstruction within neutrino experiments. In this section, we will show some of the more common neutrino interaction channels and comment on how they can affect neutrino experiments.

2.2.1 CC0 π

Events labelled as CC0 π are charged-coupled (CC) interactions between a neutrino and a hadron resulting in zero pions in the final states. These interactions can be further broken down into CCQE (charge-coupled quasi-elastic) interactions and 2p2h (two protons two holes) interactions. The Feynman diagrams for these interactions are shown in Fig. 2.2, although only one possible 2p2h interaction is shown of many possibilities.

CCQE interactions (Fig. 2.2b) are one of the simplest and most common neutrino interactions. Given their relatively large cross-section, neutrino experiments often aim to build up samples of these interactions. In Fig. 2.2c, we show a 2p2h interaction where the nucleon interacting directly with the neutrino is influenced by another nucleon in the nucleus. 2p2h and CCQE interactions have exactly the same final state particles, so 2p2h interactions usually pass data cuts used to isolate CCQE interactions. Because these two interaction channels are often mixed together, we see a measurable difference

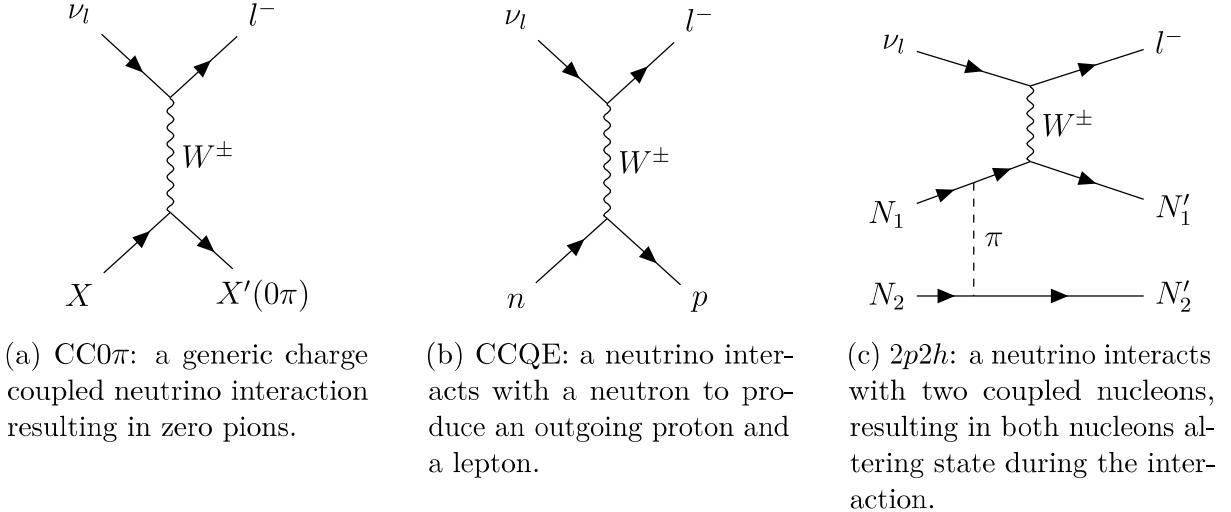


Figure 2.2: Feynman diagrams of CC0 π processes.

in cross-sections between neutrino interactions with individual nucleons and neutrino interactions with nuclei. This was seen experimentally in the first double differential cross-section measurements of neutrino CCQE interactions with carbon from the MiniBooNE experiment [32]. MiniBooNE showed that the neutrino CCQE carbon cross-section per nucleon is about 20% higher than that of neutrino-hydrogen scattering.

Not only do 2p2h interactions mean that the neutrino cross-section differs depending on the nucleus, they also affect neutrino energy reconstruction in neutrino experiments. In neutrino energy reconstruction, the norm is to create a sample of events with CCQE interaction properties and reconstruct neutrino energies assuming CCQE physics. Yet, because 2p2h interactions are so similar to CCQE and produce the same final state particles, 2p2h interactions are often included in the CCQE sample and lead to a bias in the energy reconstruction. Identifying 2p2h interactions or accounting for the introduced error is desirable to reduce systematic errors in experiments.

2.2.2 CC1 π

Shown in Fig. 2.3 are several charged-current neutrino-nucleon interactions resulting in the production of a single pion meson. The interactions shown are the dominant ones at the neutrino energies used by the T2K experiment.

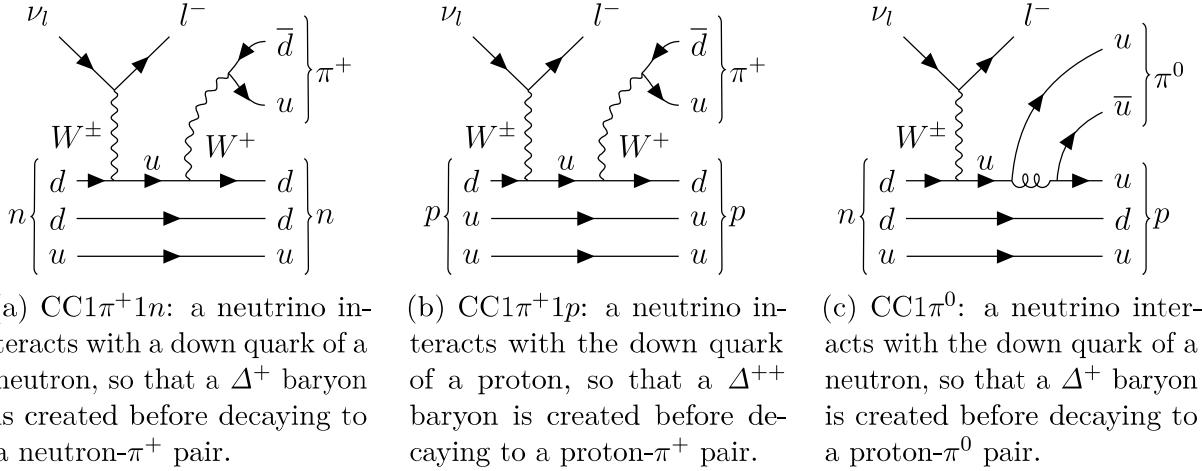


Figure 2.3: Feynman diagrams of CC1 π processes.

CC1 π interactions are often misinterpreted as CCQE interactions in experiments due to undetected pions. For example, in T2K about 20% of events in CCQE samples are CC1 π events. Similar to the inclusion of $2p2h$ interactions in CCQE samples, this impurity of CC1 π events in the CC0 π sample can lead to problems with the neutrino energy reconstruction where the CCQE interaction is assumed.

2.3 Neutrino Oscillations

For neutrino oscillations to exist, the neutrino flavour eigenstates, $(\nu_e, \nu_\mu, \nu_\tau)$, must be different to the neutrino mass eigenstates, (ν_1, ν_2, ν_3) . In fact, it must be the case that the flavour eigenstates are coherent superpositions of the mass eigenstates. Hence, they are linked by the equation

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{bmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{bmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}, \quad (2.1)$$

where the matrix enclosed by square brackets is the PMNS matrix, \mathbf{U} [3, 4]. The matrix elements $U_{\alpha i}$ represent the mixing amplitudes of mass state i contained within the neutrino flavour state ν_α .

\mathbf{U} is a unitary matrix, so it can be parametrised using three real angles, $(\theta_{12}, \theta_{13}, \theta_{23})$ and a real phase, δ_{CP} , to give

$$\mathbf{U} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & C_{23} & S_{23} \\ 0 & -S_{23} & C_{23} \end{bmatrix} \begin{bmatrix} C_{13} & 0 & S_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -S_{13}e^{i\delta_{CP}} & 0 & C_{13} \end{bmatrix} \begin{bmatrix} C_{12} & S_{12} & 0 \\ -S_{12} & C_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (2.2)$$

where $C_{ij} = \cos \theta_{ij}$ and $S_{ij} = \sin \theta_{ij}$. These mixing angles and phases are what oscillation experiments attempt to measure. A non-zero value of the δ_{CP} phase, given that $S_{13} \neq 0$, would mean that neutrino oscillations are CP violating. There may be two extra phase parameters required if neutrinos are Majorana particles (Majorana particles are their own anti-particles) but these have no effect on oscillation measurements [33].

Eq. (2.1) encapsulates much of the physics behind neutrino oscillations. It implies that, following the weak decay $W \rightarrow l_\alpha \nu_\alpha$, where W is a charged weak boson, ν_α is a neutrino flavour eigenstate and l_α is the corresponding charged lepton state with flavour α , the neutrino will propagate as a mixture of the mass eigenstates that interfere quantum mechanically. That is to say, the waveforms of each mass state are superimposed upon each other and these waveforms are time dependent. Due to the time dependent aspect, the probability of the neutrino's flavour, when measured, is dependent on the distance the neutrino has travelled [33].

For a detailed derivation of the neutrino oscillation probabilities, see Kayser [34]. Here, only the end result is quoted, which is the probability of a (anti-)neutrino with initial flavour α and energy E_ν propagating a distance L and being detected with flavour β :

$$P\left(\overleftrightarrow{\nu}_\alpha \rightarrow \overleftrightarrow{\nu}_\beta\right) = \delta_{\alpha\beta} - 4 \sum_{i < j}^3 \Re[U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}] \sin^2\left(\Delta m_{ij}^2 \frac{L}{4E_\nu}\right) \\ \mp 2 \sum_{i < j}^3 \Im[U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}] \sin\left(\Delta m_{ij}^2 \frac{L}{2E_\nu}\right), \quad (2.3)$$

where $\Delta m_{ij}^2 = m_i^2 - m_j^2$ is the difference between the squared masses of mass eigenstates i and j , often called “mass splittings”. Δm_{ij}^2 are another set of parameters that oscillation experiments measure, in lieu of the absolute masses of each neutrino flavour, which are as yet unmeasured.

2.4 Neutrino Oscillations in Matter

Equation 2.3 can only be applied to neutrinos travelling through a vacuum. When neutrinos propagate through matter, as they do in long baseline accelerator experiments, one needs to account for the coherent forward scattering the neutrinos undergo with the matter particles they encounter along the way. Electron neutrinos and anti-neutrinos are more sensitive to these particles than other flavours, as they can interact via W boson exchange with atomic electrons. Neutrinos other than electron neutrinos can only interact with the atomic electrons via Z boson exchange. This difference in interactions between neutrino flavours can lead to resonance effects, making electron neutrino oscillations to other flavours more likely than if the neutrinos were in a vacuum. This was initially thought to be the reason why so few electron neutrinos were observed to originate from the Sun, but it has since been discovered that this is more attributable to the large value of the mixing angle θ_{12} [33].

The matter effects do not change the form of the neutrino oscillation probability in Eq. (2.3) and only affect the mixing angles and mass splittings. The changes for each mass splitting and mixing angle can be expressed through the equations

$$\Delta m_M^2 \equiv \Delta m^2 \sqrt{\sin^2 2\theta + (\cos 2\theta - x)^2} \quad (2.4)$$

and

$$\sin^2 2\theta_M \equiv \frac{\sin^2 2\theta}{\sin^2 2\theta + (\cos 2\theta - x)^2}, \quad (2.5)$$

where Δm_M^2 and θ_M are the effective mass splitting and mixing angle in matter and Δm^2 and θ are the mass splitting and mixing angle in a vacuum. x is a parameter that measures the importance of matter effects with respect to the neutrino mass-squared splitting. This x parameter is proportional to $N_e E / \Delta m^2$, where N_e is the density of electrons in the matter the neutrino is passing through and E is the neutrino energy. Also, x has the opposite sign between neutrinos and anti-neutrinos. This fact can be exploited to ascertain the sign of the squared mass-splitting. The opposite sign will cause neutrinos and anti-neutrinos to have different oscillation parameters in matter. By comparing the different values, and if the sign of $\cos 2\theta$ is known, then one can ascertain the sign of x for neutrinos and anti-neutrinos, which indicates the sign of the squared mass-splitting, seeing as x is inversely proportional to it. This has been done for the solar mass splitting Δm_{23}^2 ,

showing it has a positive value. It is hoped that long-baseline accelerator experiments, such as T2K, will allow the measurement of the sign of Δm_{12}^2 , although the relative low density of matter the neutrinos travel through compared to that of solar neutrinos makes it a much more difficult task that requires precise results [33].

2.5 Neutrino Oscillation Experiment Overview

The field of neutrino oscillation experiments has a rich history and a bright future. Kajita of the Super-Kamiokande Observatory and McDonald of the Sudbury Neutrino Observatories were awarded the 2015 Nobel Prize in Physics for their discovery of neutrino oscillations, a discovery which dated back to the 1990s for Super-Kamiokande and the 2000s for Sudbury Neutrino Observatories. Super-Kamiokande, based in Japan, used a massive water tank to detect two samples of atmospheric muon neutrinos, one where the neutrinos were produced on the other side of the planet and travelled upwards through the Earth to reach the detector and one where the neutrinos were produced in the atmosphere above the detector and did not pass through the Earth [35]. Theoretical predictions with no neutrino oscillations calculated that the muon neutrino flux from both samples should be equal. However, a significant deficit of muon neutrinos for the upward going sample was detected, suggesting over the long baseline of the Earth's diameter the muon neutrinos were changing flavour. Super-Kamiokande is still running today and is used as the far detector in the long baseline accelerator neutrino experiment T2K, which we will detail later on.

The Sudbury Neutrino Observatories, or SNO, located in Canada, used heavy water to detect solar neutrinos [36]. Heavy water can be utilised to separately measure the solar flux of electron neutrinos and all neutrinos, as solar electron neutrinos can undergo charged current or elastic scattering interactions with heavy water but other neutrino flavours cannot, or at least have a significantly smaller cross-section. All neutrino flavours are equally likely to interact with heavy water via neutral current interactions. By identifying the different interaction types, SNO was able to show that there was a significant non-electron flavour component to the solar neutrino flux, implying the electron neutrinos produced by the Sun were changing flavour on their way to the Earth.

Since the discovery of neutrino oscillations, several experiments have reported measure-

Experiment (Examples)	Dominant	Important
Solar (SNO, SK)	θ_{12}	$\Delta m_{21}^2, \theta_{13}$
Reactor Long Baseline (KamLAND)	Δm_{21}^2	θ_{12}, θ_{13}
Reactor Medium Baseline (Daya Bay, RENO, Double Chooz)	$ \Delta m_{32,31}^2 , \theta_{13}$	
Atmospheric (SK, IC-DC)		$ \Delta m_{32,31}^2 , \theta_{23}, \theta_{13}, \delta_{\text{CP}}$
Accelerator LBL $\nu_\mu/\bar{\nu}_\mu$ Disappearance (K2K, MINOS, T2K, NO ν A)	$ \Delta m_{32,31}^2 , \theta_{23}$	
Accelerator LBL $\nu_e/\bar{\nu}_e$ Appearance (MINOS, T2K, NO ν A)	δ_{CP}	θ_{13}, θ_{23}

Table 2.1: The sensitivities of the different neutrino oscillation experiments [44].

ments of the neutrino oscillation parameters. What parameters an experiment is sensitive to depends on the distance the neutrinos travel before detection and their energy, due to the L/E_ν dependence in Eq. (2.3). Nuclear reactor experiments—of which some notable names are RENO [37], Double Chooz [38], KamLAND [39] and Daya Bay [40]—often utilise short baselines and are most sensitive to the θ_{13} parameter. Solar neutrino experiments like SNO are sensitive to the Δm_{21}^2 parameter, also known as the “solar mass splitting”, and θ_{12} , the “solar mixing angle”. Atmospheric neutrino experiments like Super-Kamiokande are most sensitive to the “atmospheric mass splitting”, $|\Delta m_{31}^2| \approx |\Delta m_{32}^2|$, and the “atmospheric mixing angle”, θ_{23} . Accelerator neutrino experiments such as K2K [41], T2K, MINOS [42] and NO ν A [43], which have a large range of neutrino energies and baselines, have given significant contributions to the measurements of θ_{12} , θ_{23} and $|\Delta m_{32,31}^2|$. T2K has also made some of the first measurements of the δ_{CP} parameter, but there are still significant uncertainties. Tab. 2.1 summarises the dominant and important neutrino oscillation parameter sensitivities for different oscillation experiments. Tab. 2.2 shows the currently accepted values of the neutrino oscillation parameters as reported by the Particle Data Group [44].

Oscillation Parameter	Observed Value
$\sin^2(\theta_{13})$	$(2.18 \pm 0.07) \times 10^{-2}$
$\sin^2(\theta_{12})$	0.307 ± 0.013
$\sin^2(\theta_{23})$ NO	0.545 ± 0.021
$\sin^2(\theta_{23})$ IO	0.547 ± 0.021
Δm_{21}^2	$(7.53 \pm 0.18) \times 10^{-5} \text{ eV}^2$
Δm_{32}^2 NO	$(2.453 \pm 0.034) \times 10^{-3} \text{ eV}^2$
Δm_{32}^2 IO	$(-2.546 \pm 0.034) \times 10^{-3} \text{ eV}^2$
δ_{CP}	$1.36 \pm 0.17 \pi \text{ rad}$

Table 2.2: The current values of the different oscillation parameters and their uncertainties obtained by combining results from different experiments [44]. NO and IO stand for the Normal and Inverted neutrino mass Ordering respectively.

2.6 Remaining Questions

Despite decades of experimental research into neutrinos, there are still several unanswered questions. For one, the fundamental masses of the neutrinos are not known. The masses have been constrained using astronomical and nuclear observations, but there has been no direct measurement as of yet [45]. Furthermore, the ordering of the neutrino masses is not known. We have been able to measure the mass splittings for the neutrinos and ascertained that the sign of Δm_{21}^2 is positive, indicating that $m_2 > m_1$, but the signs of Δm_{32}^2 and Δm_{31}^2 are still unknown. As such, neutrino experiments often present results corresponding to the Normal Ordering (NO) regime where $m_3 > m_2$, as well as the Inverted Ordering (IO) regime where $m_2 > m_3$. Accelerator experiments have shown a slight preference for the NO regime in their results, but no definitive conclusions have yet been drawn [46].

One of the last remaining neutrino oscillation parameters to be accurately measured is the δ_{CP} parameter. T2K recently published an analysis with CP conserving values of δ_{CP} excluded at 3σ [47]. Improving this to a 5σ exclusion would be a vital piece of the puzzle towards solving the matter-antimatter asymmetry problem. Such a measurement is the target of the next phase of the T2K experiment, but first, let us describe the T2K experiment in more detail.

Chapter 3

The T2K Experiment

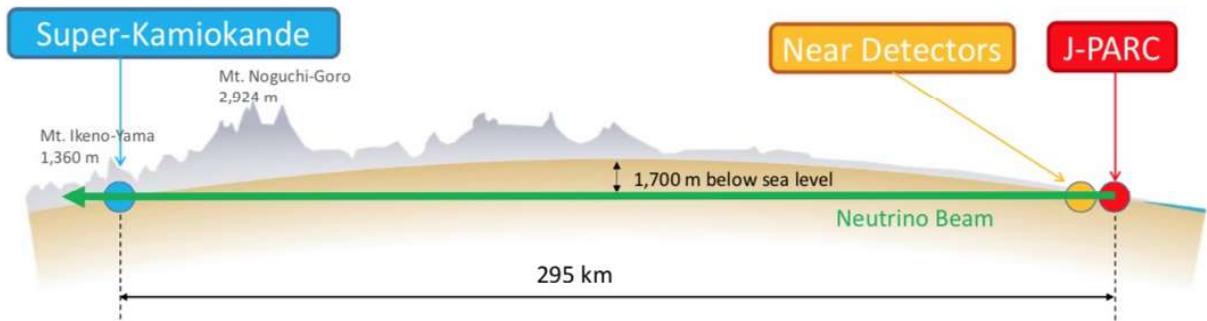


Figure 3.1: The path of the neutrino beam produced in the T2K experiment [48]. Neutrinos are produced at the J-PARC facility and travel through the Earth towards the far detector, SK. There are also two near-detectors, ND280 and INGRID.

The T2K (Tokai-to-Kamioka) experiment spans 295 km from the village of Tokai on the east coast of Japan to Kamioka on the west coast. The size of the experiment and its capability of generating a beam of neutrinos or anti-neutrinos make it one of the most sensitive current experiments to the δ_{CP} parameter. It has also published some of the most precise measurements of the Δm_{23}^2 and θ_{23} parameters [49].

In Tokai, T2K has two near-detectors—INGRID and ND280—that are 280 m from the origin of the neutrino beam. At the other end of the experiment, in Kamioka, the far detector—Super-Kamiokande (SK)—observes interactions from neutrinos in the beam. An illustration of the experiment is shown in Fig. 3.1.

The INGRID detector lies on the axis of the neutrino beam, but ND280 and SK are 2.5° off-axis. This off-axis angle sets the peak of the neutrino flux for ND280 and SK at 600 MeV, the energy around which muon neutrino disappearance is most likely after travelling 295 km [48]. This idea of an off-axis beam to tune particle energies was first used at the Brookhaven National Laboratory [50]. The focus on 600 MeV neutrinos also decreases the number of interactions in SK resulting in meson production, which occur

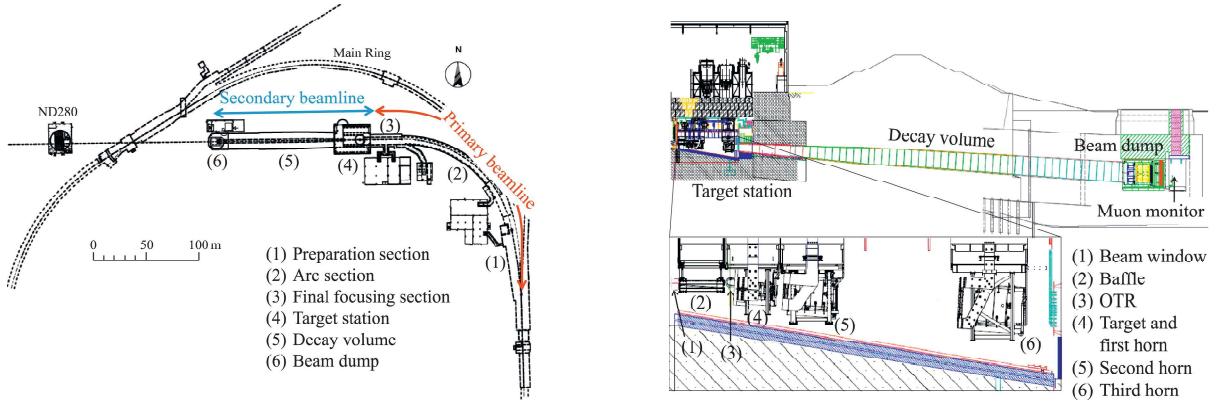


Figure 3.2: An overview of the T2K beamline (left) and a side view of the secondary beamline (right). In the primary beamline, protons are extracted from the Main Ring and directed towards the target. The secondary beamline consists of the graphite target and the secondary beam of pions [48].

at higher energies. These interactions are treated as background in T2K's oscillation analysis.

3.1 Neutrino Beam and Flux Constraints

The neutrino beam of T2K is spawned at the Japan Proton Accelerator Research Complex (J-PARC) in Tokai. Protons accelerated up to 30 GeV are fired from J-PARC's main ring synchrotron and collide with a graphite target, producing charged pions that are focused into a beam using magnetic horns. This secondary beam of pions and the layout of the beamline after the graphite target are shown in Fig. 3.2. Before the protons reach the target they are collimated by a baffle that prevents the protons damaging the magnetic horns. The Optical Transition Radiation (OTR) monitor measures the two-dimensional profile of the proton beam just prior to the target. The pions produced by proton collision rapidly decay as they pass through the decay volume, a 96 m long steel tunnel, to give muons and muon neutrinos. The muons and mesons that have not yet decayed are stopped by a concrete beam dump, which just leaves the muon neutrinos [48]. These neutrinos travel through the Earth towards Kamioka.

The content of the neutrino beam used in T2K can be switched between muon neutrinos (ν_μ) and muon anti-neutrinos ($\bar{\nu}_\mu$). Reversing the current passed through the magnetic

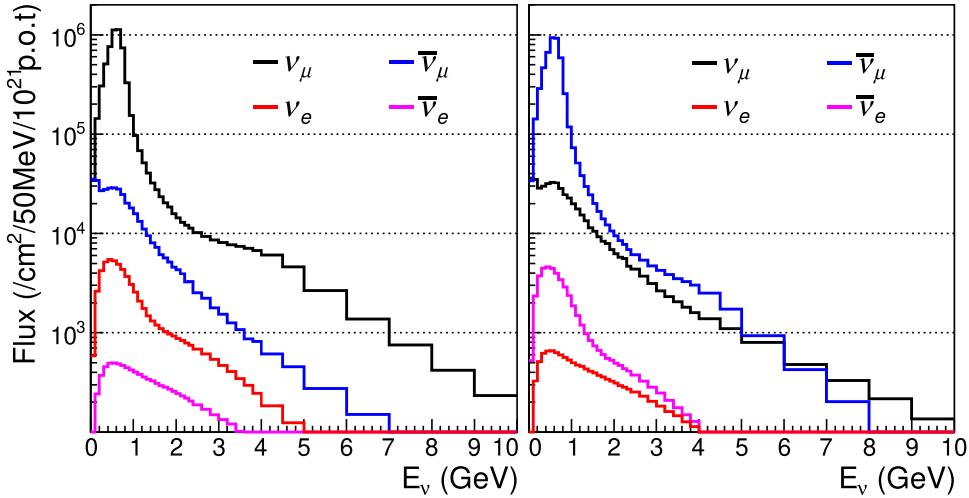


Figure 3.3: The latest publicly released flux predictions at T2K’s far detector for the Forward Horn Current (left) and Reverse Horn Current30 (right) modes [51].

horns focusing charged pions causes pions of opposite charge to be focused and pions with the original charge to be deflected. Hence, one can focus positive pions (π^+) or negative pions (π^-) depending on the current’s direction. The mode in which π^+ s are focused to give a ν_μ beam is the forward horn current (FHC) mode. To focus π^- s and create a $\bar{\nu}_\mu$ beam the reverse horn current (RHC) mode is needed [48].

The neutrino beam flux is one of the largest contributors to uncertainty in oscillation parameters measurements at T2K. The flux predictions for each of the detectors at T2K are made using simulations of the J-PARC beam-line and its interactions with the beam-line instruments. The simulations take as inputs the primary proton beam profile measurements, measurements of the horns’ magnetic fields and hadron production data, including external measurements from the NA61/SHINE experiment [51, 52]. Fig. 3.3 shows the nominal flux estimations for both the FHC and RHC beams for T2K data run1–10. Run1 was taken back in 2010 and run10 is the latest publicly released data, which was collected in 2019 [53].

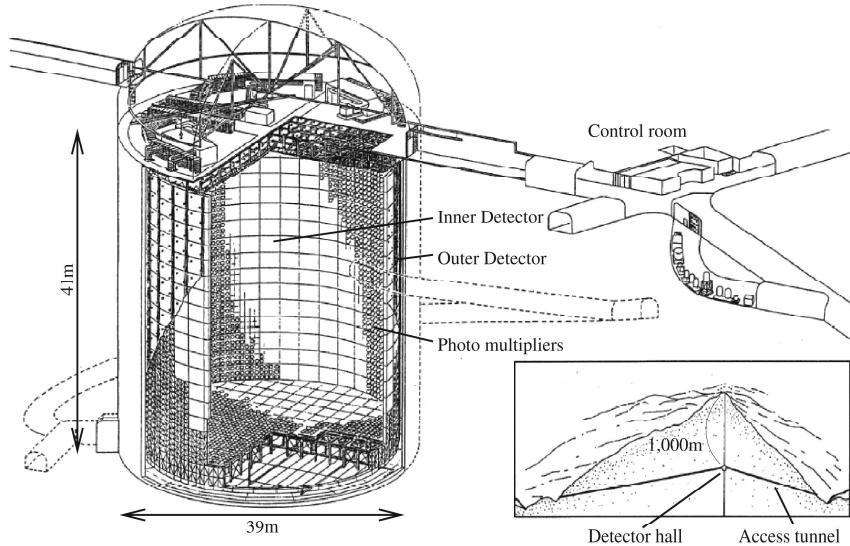


Figure 3.4: A schematic of Super-Kamiokande and its location within Mozumi Mine [54].

3.2 Super-Kamiokande

The fiducial volume of the cylindrical water Cherenkov detector Super-Kamiokande measures 29.8 m in diameter and 32.2 m in height, with the whole tank measuring 39.3 m in diameter and 41.4 m in height. The fiducial volume is often referred to as the Inner Detector (ID) and the rest of the detector's volume is known as the Outer Detector (OD). A sketch of the detector is shown in Fig. 3.4. The detector itself is a stainless steel tank with a supported stainless steel structure inside the tank separating the ID and the OD. The whole tank is filled with 50,000 tons of ultra-pure water. This sealed tank of water resides 1000 m underground in Mozumi Mine to shield the detector from external sources of radiation. The inner walls of the ID are lined with around 11,146 20-inch Photo-Multiplier Tubes (PMTs) used to detect single photons with roughly 40% coverage of the ID surface area. The OD is lined with 1885 outward facing 8-inch PMTs providing veto and shielding for the ID. The massive volume of the detector increases the interaction probability of neutrinos within the detector and increases the probability of fully containing high-energy neutrino events. The detector is capable of measuring muons and electrons produced in CC interactions of ν_μ and ν_e particles respectively. If these leptons have enough energy—above 0.26 MeV for electrons and 54.1 MeV for muons—they travel at speeds faster than the speed of light within water. This results in a cone of Cherenkov light being produced along the trajectory of the lepton. The Cherenkov light

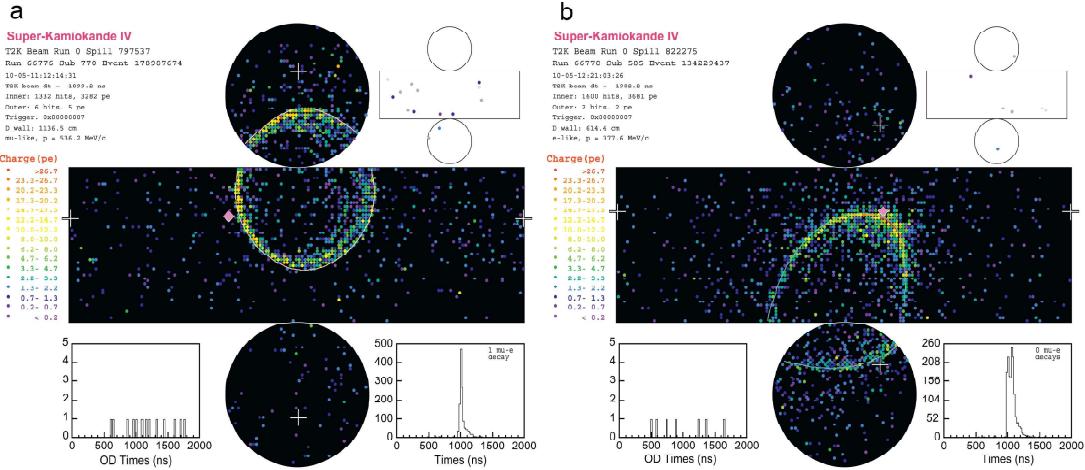


Figure 3.5: Two event displays from SK. The left display shows the well defined light ring produced by a muon inside the detector. The right display shows the fuzzier features seen in the detected Cherenkov light ring produced by an electron inside the detector [48].

is in turn detected by the PMTs on the inside of the detector. The PMT array of SK is sensitive enough and has a wide enough dynamic range to detect and differentiate events from 4.5 MeV to 1 TeV [55–57].

Signals produced by muons can be differentiated from electron signals by the different properties of the particles’ detected Cherenkov signals. The higher mass of muons lets them propagate through the detector with little scattering, resulting in relatively smooth and well-defined rings of light. Electrons, on the other hand, experience significantly more scattering as well as interactions with the water molecules resulting in electron showers. This less direct path of the electron results in a much “fuzzier” light ring. An example of the differences between the signals of muons and electrons is shown in Fig. 3.5.

An estimate of the parent neutrino’s energy for a neutrino interaction at SK can be made by measuring the final state lepton’s momentum, P_l , and angle with respect to the neutrino beam, θ_{beam} . The formula used to calculate this reconstructed energy depends on the assumed interaction channel of the neutrino interaction. T2K collect samples of both CCQE and CC1 π^+ interactions. The CCQE formula is

$$E_{\nu, \text{CCQE}}^{\text{rec}} = \frac{(M_n - V_{\text{nuc}})E_l - m_l^2/2 + M_n V_{\text{nuc}} - V_{\text{nuc}}^2/2 + (M_p^2 - M_n^2)/2}{M_n - V_{\text{nuc}} - E_l + P_l \cos(\theta_{\text{beam}})}, \quad (3.1)$$

and the CC1 π^+ formula is

$$E_{\nu, \text{CC1}\Delta}^{\text{rec}} = \frac{2M_n E_l + M_{\Delta^{++}}^2 - M_n^2 - M_l^2}{2(M_n - E_l + P_l \cos(\theta_{\text{beam}}))}. \quad (3.2)$$

In the above formulae, M_n (M_p) is the neutron (proton) mass, $M_{\Delta^{++}}$ is the Δ^{++} baryon mass, V_{nuc} is the binding energy of the neutron in oxygen (taken to be 27 MeV), m_l is the mass of the charged lepton and E_l is its total energy of the lepton. For the CC1 π^+ formula, since a two body process with Δ baryon recoil is expected, the nuclear binding potential is not considered in the energy reconstruction. Eq. (3.2) relies on the assumption that the target nucleon in the neutrino interaction is at rest, which mitigates the lack of available information on the π^+ particle produced in the interaction.

3.3 ND280

ND280 gathers data on the neutrino beam to estimate the beam's flux, energy spectrum and flavour content. This collected data significantly reduces uncertainties in the neutrino beam properties, leading to reduced systematic error in the experimental result [58]. Tab. 3.1 shows the systematic uncertainties for predicted events at SK with and without ND280 constraints.

As well as monitoring the beam characteristics, ND280 also takes data for neutrino cross-section measurements, providing the community of neutrino oscillation experiments with data to refine their interaction models.

The ND280 detector can be broken down into three sections: a π^0 detector (PØD), a Tracker consisting of two fine grain detectors (FGDs) and three time-projection chambers (TPCs), a system of electromagnetic calorimeters (ECals) and a side muon range detector (SMRD). The whole detector is subject to a 0.2 T magnetic field for charge identification and momentum determination. See Fig. 3.6 for a visualisation of the ND280 detector and its sub-detectors. There is also an example event display shown in Fig. 3.7 demonstrating how leptons are tracked through the ND280 detector.

Source of Uncertainty	ν_μ $\delta N/N$	ν_e CCQE-like $\delta N/N$	ν_e CC1 π^+ $\delta N/N$
Flux			
(w/ ND280 constraint)	3.6%	3.7%	3.6%
Cross-section			
(w/ ND280 constraint)	5.1%	4.0%	4.9%
Flux+Cross-section			
(w/o ND280 constraint)	10.8%	11.3%	16.4%
(w/ ND280 constraint)	2.9%	4.2%	5.0%
Flux+Cross-section+Detector			
(w/o ND280 constraint)	12.0%	12.7%	21.9%
(w/ ND280 constraint)	5.1%	5.5%	14.8%

Table 3.1: The change in event numbers at SK with 1σ variations of the associated systematic parameters for the indicated data samples [59]. Combining the uncertainties for flux and cross-section contributions with the ND280 constraints results in smaller overall uncertainties due to the anti-correlation between the two contributions.

3.3.1 π^0 Detector (PØD)

T2K’s far detector, SK, has difficulty discerning between electrons and gamma particles. Two gamma particles with small angular separation are produced in neutral pion (π^0) decay. This can lead to significant background for ν_e appearance samples where the neutral current interaction $\nu_\mu + N \rightarrow \nu_\mu + N + \pi^0 + X$ is misidentified as a CCQE or CC1 π interaction. All three of these interactions produce a single electron-like ring at SK. To account for this background signal, the PØD sub-detector of ND280 is used to measure the cross-sections of neutrino interactions, specifically with water, where a π^0 appears in the final state. Using this cross-section information, the number of π^0 s produced in SK can then be calculated and removed from the final samples.

The central section of the PØD consists of alternating layers of brass, plastic scintillator and fillable water bags. The brass acts as a radiator, causing electrons produced by pion decays to shower, whereas the water and plastic layers provide a target for neutrinos to interact with. Plastic scintillator has the useful property of emitting photons when a charged particle passes through it. This is exploited by using wavelength shifting (WLS) fibres to read off the light signal from a passing particle and using the data to reconstruct

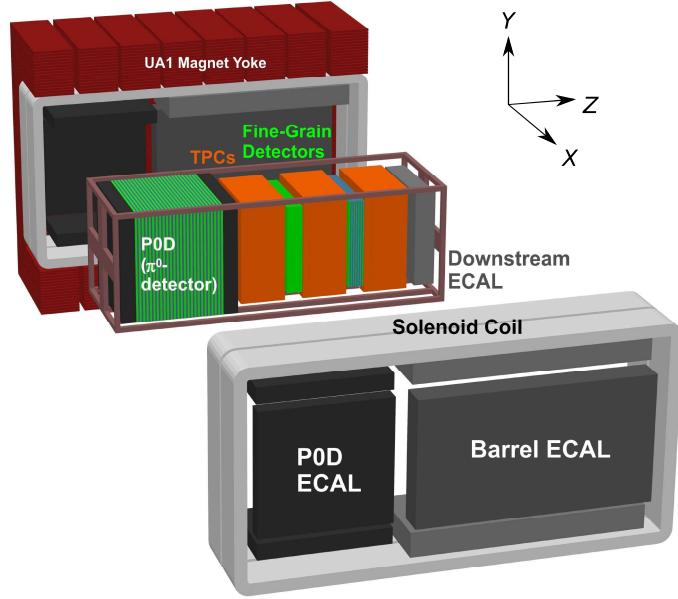


Figure 3.6: An expanded view of the ND280 detector showing the positions of each of the three main sections and sub-detectors. The entire detector is encased in the UA1 magnet with a 0.2 T magnetic field. The neutrino beam is parallel to the z -axis of the detector [60].

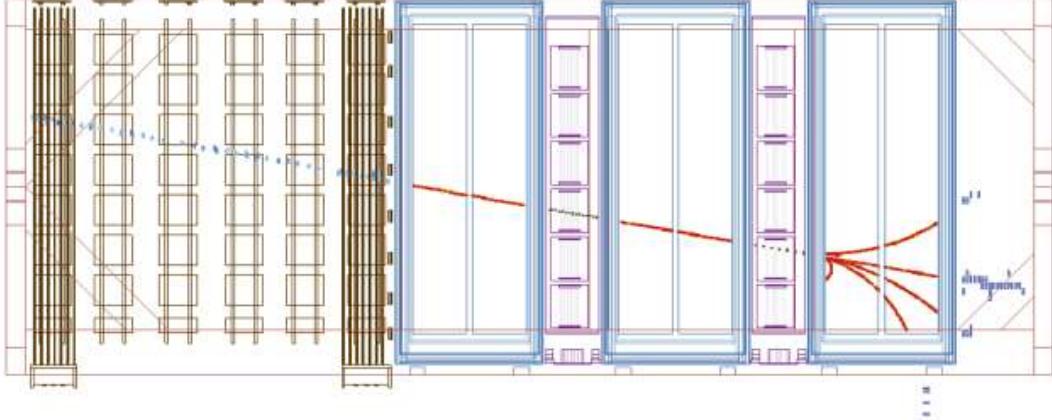


Figure 3.7: An event display of a neutrino interacting in the PØD and producing a muon that is tracked through all the sub-detectors.

the particle's trajectory.

The central target is preceded and followed by electromagnetic calorimeters comprised of alternating sheets of lead and scintillator. These ECals are used for rejecting particles

entering the PØD that originated from outside the detector and also to contain and measure electromagnetic showers.

Cross-section values with respect to the water target can be deduced by taking measurements when the target bags in the central section are filled with water and comparing them with when they are filled with air [48, 61].

3.3.2 Fine Grain Detectors (FGDs)

The FGDs provide 2.2 tonnes of mass with which incoming neutrinos can interact and precision tracking to map out the interaction vertices. A sketch of the FGD structure is shown in Fig. 3.8. The FGDs consist of bars of plastic scintillator oriented perpendicular to the neutrino beam and stacked side-to-side to form sheets. Two sheets are placed one after the other to form a module, with the bars of one layer being arranged vertically (along the y -axis) and the other layer horizontally (along the x -axis). The bars are read-off by WLS fibres, so the position of a charged particle incident on a module can be determined in two dimensions. Placing several modules in sequence then gives the trajectory of a particle through the detector.

There are two FGDs in ND280. The downstream FGD also contains layers of water in between the scintillator layers. Comparing interaction rates between the two FGDs allows calculation of the ratio of the cross-sections of neutrinos on carbon and water [48, 63].

3.3.3 Time Projection Chambers (TPCs)

There are three TPCs in ND280, sandwiching the two FGDs. The bulk of each TPC consists of 3000 litres of an argon-based gas. To track charged particles passing through the gas, the TPCs utilise Micro-Mesh Gaseous Structure (MicroMegas) technology. This technology involves applying an electric field across the argon gas using an anode and cathode. The anode is comprised of a fine metallic mesh. When a charged particle passes through the gas, it will ionise the gas molecules. The ionisation electrons drift towards the anode, pulled by electrostatic forces. Upon reaching the anode, the ionised electrons pass through the mesh and enter an even more intense electric field created by the readout electrode of the TPC next to the metallic mesh. This intense electric field causes the

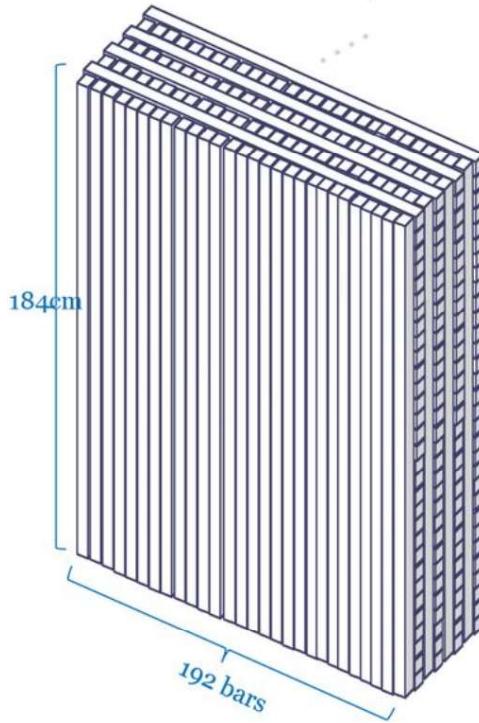


Figure 3.8: A sketch of the FGD structure. Scintillator bars are arranged parallel to each other in sheets. These sheets are then layered along the beam direction with adjacent sheets having bars perpendicular to each other and the beam direction [62].

electrons to accelerate, ionising the gas as they travel, amplifying the charge roughly 2000 fold. These charges are detected by the readout plane of the TPC. The location and arrival time of the electrons on this plane can be used to construct a 3D image of the particle's trajectory.

TPCs are ideal for identifying the number and trajectory of charged particles in the detector, although their relatively small mass makes it unlikely for a neutrino to interact with them, hence the need for FGDs. The information provided by TPCs allows selection of low-contamination interaction samples, as well as momentum measurements from curved tracks caused by the magnetic field present in the detector. Combining a particle's momentum along with the ionisation of the particle's track, one can identify the type of particle that made the track. Hence, when combined with the FGD data, the TPCs can be used to help identify neutrino interactions in ND280 and to determine the electron neutrino contamination in the muon neutrino beam [48, 64].

3.3.4 Electronic Calorimeter (ECal)

As well as the ECals either end of the PØD, there are also ECals surrounding the entire detector. A “barrel” ECal encompasses the tracker around the z -axis, another ECal is situated downstream of the tracker and a third encircles the PØD about the z -axis.

The main uses of the ECals are to complement the inner detectors by measuring the energies of photons missed by the inner detectors, reject particles from interactions outside the detector and help with particle identification.

Alternating layers of scintillator and lead absorber make up each ECal. The scintillator layers consist of several bars of plastic scintillator. These bars are oriented perpendicular to the scintillator bars of the adjacent layers, similar to the FGDs [48].

3.4 ND280 Performance

ND280’s ability to significantly reduce systematic errors in the T2K experiment, as previously shown in Tab. 3.1, is testament to the detector’s excellent PID and momentum measurement capabilities. The combination of FGDs and TPCs in the ND280 Tracker is imperative when identifying the momenta and charge of particles inside the detector, as particles of opposite charge will have oppositely curved trajectories. The tightness of the trajectory curve also indicates the particle momenta. This is particularly useful in identifying the neutrino beam contamination, for example, as neutrinos and anti-neutrinos can be identified by the leptons and anti-leptons produced in interactions from each particle respectively. The detector’s reliability in differentiating muons and electrons also allows measurement of the ν_e contamination in the ν_μ beam produced at J-PARC.

Despite its excellent performance in T2K thus far, ND280 has a few limitations that will be addressed for T2K’s next phase of data-taking in order to further reduce systematics. One of the main limitations is the near-detector’s inferior angular efficiency compared to SK. ND280 relies on a particle passing through at least one of its TPCs in order to identify the particle and reconstruct its properties. Due to the orientation of the TPCs and FGDs, if a particle is produced in an FGD at an angle roughly 50° or more with respect to the neutrino beam, it is likely to miss the adjacent TPC or only pass through

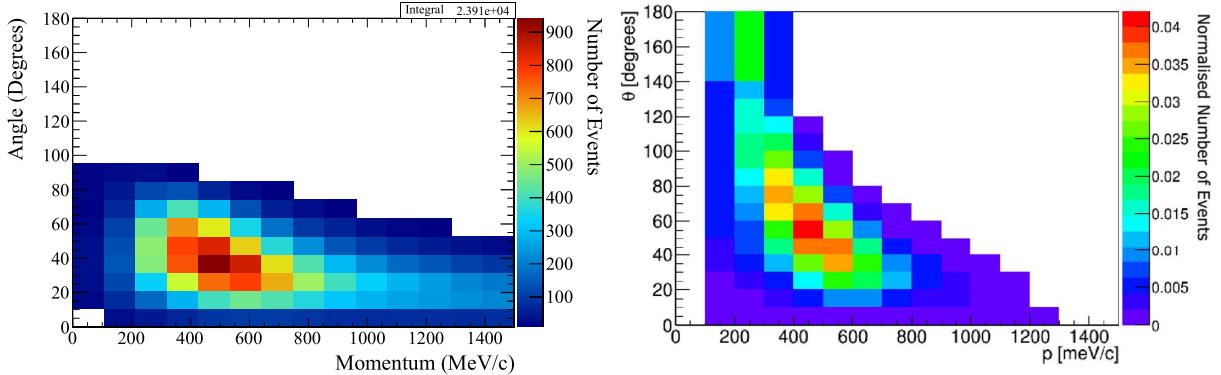


Figure 3.9: The reconstructed momentum versus the scattering angle for detected electrons at ND280 (left) and SK (right) [59, 65].

a small section of it. In such a situation, it is impossible to reconstruct the momentum of said particle. Conversely, at SK, particles at any angle with respect to the neutrino beam can be reconstructed accurately. Fig. 3.9 clearly shows the difference in angular acceptance between the two detectors, with ND280 showing a preference for low scattering angles. This difference between the two detectors introduces a reliance on theoretical cross-section models to describe the dependency between lepton momentum and angle.

The reliance on TPCs for accurate momentum reconstruction also applies to protons scattered in neutrino interactions. While protons can be tracked in the FGDs, only 2D track information is available, leading to limited tracking efficiency and a high momentum threshold. Fig. 3.10 demonstrates how proton reconstruction has nearly zero efficiency below 400 MeV. Low momenta protons and pions are essential to the identification of nuclear effects in neutrino interactions, which will be discussed more in Sec. 4.2.

We've mentioned the low reconstruction efficiency in ND280 for high scattering angles of leptons. Reduced efficiencies are also seen for backwards travelling leptons, ones that are travelling in the opposite direction to the neutrino beam. The efficiencies for these backwards-going tracks are slightly higher than high-angled tracks, as if the backwards-going lepton passes through two scintillator detectors then its time-of-flight can be calculated, indicating that the lepton is travelling upstream. However, the reliance on the lepton passing through two high-density detectors results in a low reconstruction efficiency. The left plot of Fig. 3.11 shows how backwards going leptons are identified by their time-of-flight, but it is clear from the relative size of the two peaks there is a smaller efficiency for backwards-going leptons as opposed to forwards-going leptons. This

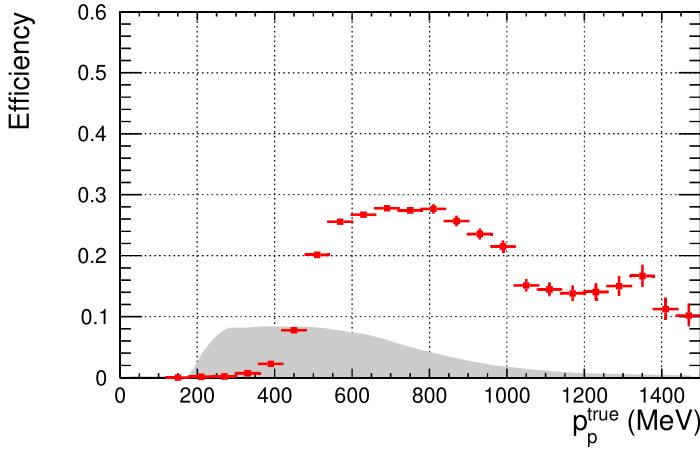


Figure 3.10: The proton reconstruction efficiency in ND280. The red data points correspond to real data. The grey histogram is the predicted spectra of protons according to MC [59]. A dip in efficiency is seen around 1000 MeV due to the lack of trajectory curvature for high energy particles making it difficult to identify the particle.

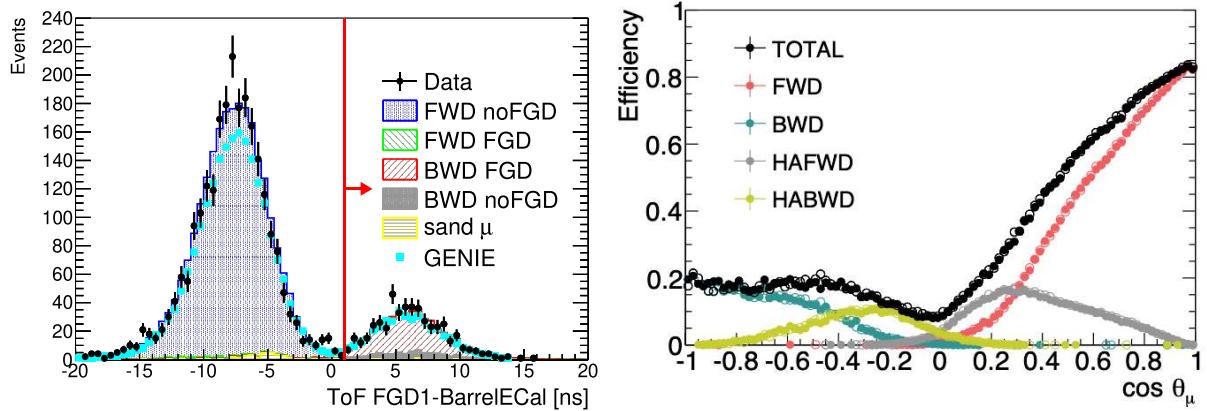


Figure 3.11: The left plot shows how the time-of-flight (ToF) between FGD1 and the barrel ECal in ND280 can be used to identify backwards-going tracks. The right plot shows reconstruction efficiencies for forward-going (FWD), backward-going (BWD), high-angled forward-going (HAFWD) and high-angle backward-going (HABWD) tracks in ND280 [59].

is demonstrated further in the left plot of Fig. 3.11, where the reconstruction efficiencies of forwards-going and backwards-going leptons are shown. The efficiencies for high-angled forwards- and backwards-going tracks are also shown. There is an undeniable preference for forward-going tracks.

Another limitation of the current ND280 is its low selection efficiency for sub-1 GeV electron neutrinos. This low efficiency is partly related to the poor reconstruction efficiency

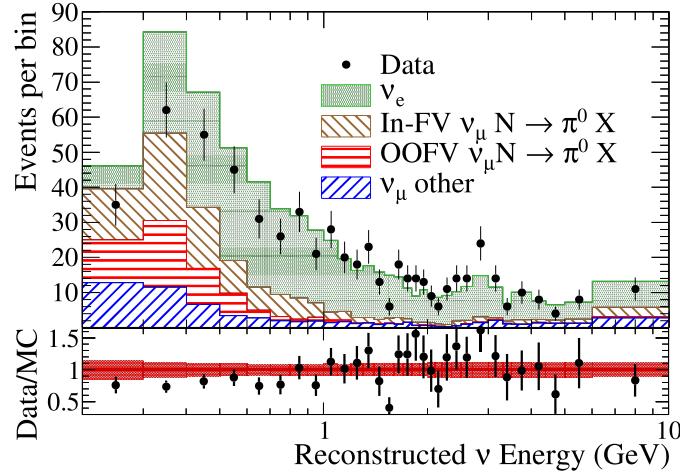


Figure 3.12: The distribution of reconstructed energies of ν_e particles in the ND280 tracker. There is a clear deficiency in the data compared to MC for events below 1 GeV [59].

of high-angled tracks, as well as the large contamination due to converted gamma particles. Fig. 3.12 shows this low efficiency by comparing ND280 tracker data to MC predictions. There is a clear lack of data events compared to MC events for reconstructed energies below 1 GeV.

These limitations of the ND280 detector, amongst other sources of systematic uncertainties, are to be addressed in the next phase of T2K, T2K-II.

Chapter 4

T2K-II

As of 2021, T2K is entering a new phase of data taking with aims to further improve the measured precision of oscillation parameters and to exclude a zero value of the δ_{CP} parameter at 3σ . To achieve this, significant modifications are being made to reduce systematic and statistical errors, including upgrades to the beamline and detectors. The proposed upgrades aim to reduce systematic errors from 6% to around 4% by introducing more fiducial mass to the near-detector, improving the reconstruction efficiency for high-angled and backwards-going particles and increasing the accumulated protons-on-target by almost six-fold to 20×10^{21} protons [59]. The resulting increase in statistics, coupled with a new granular scintillator detector that can distinguish between electrons and gammas, will allow a new dataset of low energy electron neutrino interactions. Such a dataset will lead to improvements in constraining nuclear effects in cross-section models, leaving the dominant contribution to systematic uncertainty due to differences between muon neutrinos and electron neutrinos, which is roughly 3%.

4.1 Beamline Upgrade

In February of 2020, the J-PARC Main Ring proton beam reached a power of 515 kW. After a shutdown of the beam at the start of 2021, the beam restarted with 750 kW of power, which will steadily grow to reach 1.3 MW in 2029. The higher beam power will increase the neutrino beam flux, effectively reducing statistical errors within the T2K experiment. This increase in beam power is achieved by both shortening the time between proton cycles and also increasing the number of protons per pulse. The 2021 shutdown was performed to install new hardware in the beamline to engender the reduction in proton cycle time; mainly new power supplies for the J-PARC main ring. Upgrades to the cooling systems for the beamline were also made to accommodate for the higher beam

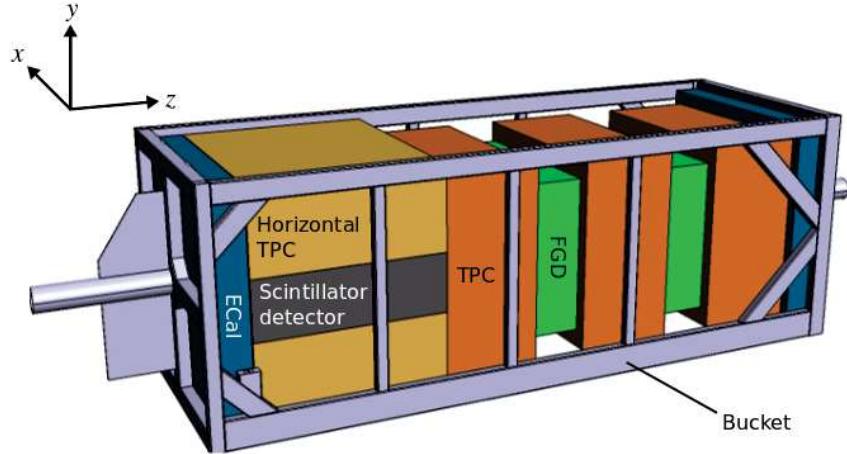


Figure 4.1: The proposed upgrade of ND280. The new additions are the horizontal TPCs and the scintillator detector, which replace the PØD. The “bucket” is a framework that supports the detectors. Not shown are the barrel ECals that will surround the detectors, similar to the original design, and the six new time-of-flight detectors surrounding the new structure inside ND280. The neutrino beam propagates in the positive z -direction [66].

power.

Furthermore, the magnetic horns in the beamline focussing charged pions will receive stronger currents, leading to roughly 10% more neutrinos of the correct sign being focussed and 10% fewer wrong-sign neutrinos (e.g. in FHC mode there will be more neutrinos and fewer anti-neutrinos). The installation of an additional power supply for the magnetic horns will make this increase in current possible.

4.2 ND280 Upgrade

The ND280 detector as it currently operates has suited the T2K experiment well. With the proposed upgrades to the beam and the increased statistics that come with them, it gives an opportunity to address some of the few shortcomings of ND280. As mentioned in Sec. 3.4, the current layout of ND280 is well-designed to detect forward-going leptons, but has low reconstruction efficiency regarding leptons with high scattering angles with respect to the neutrino beam direction. It also has limitations regarding its energy threshold, meaning particles other than leptons produced in neutrino interactions are often left undetected within ND280. With this loss of additional particles, some neutrino

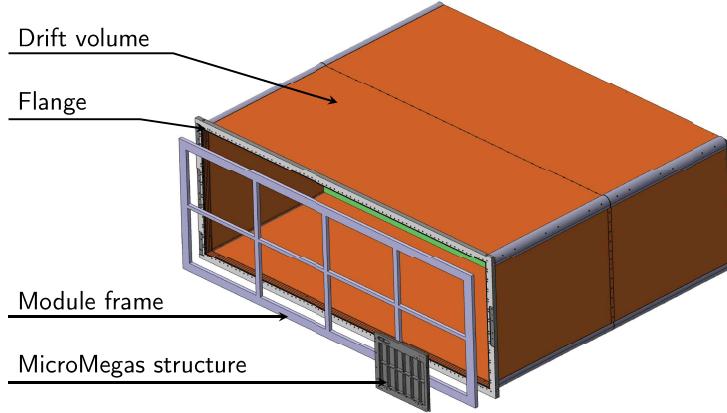


Figure 4.2: A schematic view of the HATPC design [59].

interactions can be misidentified as CCQE, leading to a bias in the neutrino energy reconstruction.

To address ND280’s current limitations, new subdetectors will be installed in ND280 as part of the T2K-II phase. The upgrades aim to lower the energy thresholds of the detector, leading to increased efficiency in the reconstruction of final state particles; increase the detection efficiency of high-angled tracks and backwards-going tracks; and increase the active volume of the detector, allowing for more neutrino interactions to occur.

The proposed upgraded ND280 is shown in Fig. 4.1. The downstream end of the detector will remain unchanged, with three TPCs separated by two FGDs followed by an ECal. At the upstream end, however, the central PØD module (not the ECal endcaps) will be replaced by a scintillator detector—the SuperFGD—sandwiched vertically between two horizontal High Angle TPCs (HATPCs). The reason for the PØD’s removal is the large value of θ_{13} , making its measurements to estimate the double gamma backgrounds at SK less necessary than previously thought, as the background is significantly smaller than the observed signal. Not shown in Fig. 4.1 are the six time-of-flight detectors that will surround the new ND280 detectors.

4.2.1 High Angle TPCs

The new scintillator detector will act as a target material, with two High Angle TPCs sandwiching the scintillator detector along the y -axis. This orientation of HATPCs

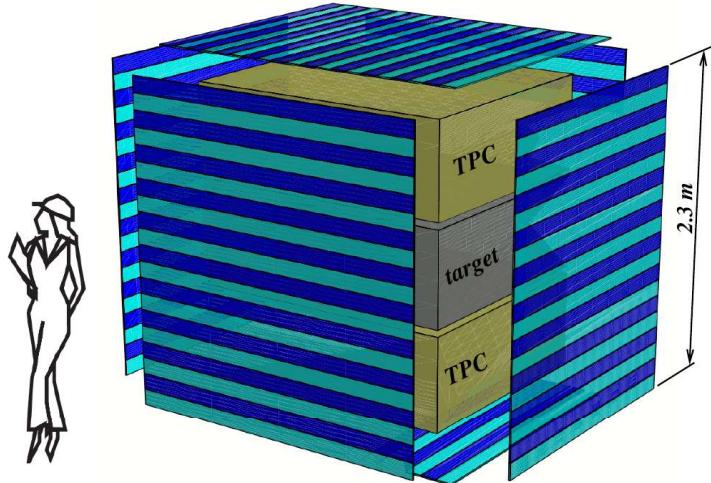


Figure 4.3: A schematic view of the TOF detector situated around the HATPCs and the SuperFGD (target) [59].

significantly increases the number of high-angled tracks seen within ND280. When combined with interactions within the old TPCs and FGDs, ND280 has nearly 4π solid angle coverage. This will address the biased tracking efficiency towards particles with low scattering angles.

The HATPCs themselves use very similar technology to the pre-existing TPCs in ND280; both detectors using MicroMegas modules. The HATPCs go a step further in utilising resistive MicroMegas technology that allow a more even spread in charge across the anode plane. In principal, this should result in fewer required readout channels for the same spatial resolution. A schematic of one of the TPCs can be seen in Fig. 4.2.

As the HATPCs are not the main subject of this thesis, we will not go into detail on the design and performance of the proposed HATPC detectors. For more information, see Abe et al [59].

4.2.2 Time-of-Flight Detectors

The Time-of-Flight detectors to be installed for the ND280 Upgrade aim to address the low efficiencies for backwards travelling particles, as well as to provide a more reliable veto for background events that originated from outside the detector. The TOF setup will be six planes of scintillator bars, with one plane covering each face of the new detector

module, as shown in Fig. 4.3. With a time resolution of less than 500 ps, this setup will ensure a reliable and unambiguous TOF for any particle that is not fully contained within the detector by combining timing measurements with the SuperFGD. PID will also be improved with time resolutions below 100–200 ps. The proposed design for the TOF detectors is predicted to give a time resolution of 150 ps. As with the HATPC, the TOF detectors are not the main subject of this thesis, so for more information consult Abe et al [59].

4.2.3 The SuperFGD

The proposed scintillator detector for the ND280 upgrade will have a first-of-its-kind “SuperFGD” design that incorporates scintillator cubes rather than bars, which can be read out from three orthogonal directions by WLS fibres. To accommodate the fibres, each cube will have three holes drilled through them, with each hole orthogonal to the other two holes and perpendicular to two of the cube’s surfaces. These holes will be positioned so they do not intersect with the other holes and will be drilled in the same position on each cube, so that the holes are aligned when the cubes are lined up. Hence, when cubes are arranged in a cuboid shape, fibres can be threaded through each line of cubes in the x , y and z directions, resulting in each cube having three WLS fibres threaded through it. These fibres will have photosensors attached to one end so that scintillation light from a particle passing through the detector can be detected and used to reconstruct the particle trajectory using information from three dimensions.

The overall size of the SuperFGD detector will be $192 \times 192 \times 56$ cm. The cubes’ side length will be 1 cm, allowing for significant granularity but requiring a large number of cubes and readout channels, 2,064,324 and 58,368 respectively. The size of the cubes is comparable to the height and width of the scintillator bars used in the current design, but the bars only read out in two directions. The extra information provided by the cubes helps to reconstruct the paths of multiple particles, making it easier to discern electrons produced in neutrino interactions from electron-positron pairs produced by gamma particles. It will also significantly increase the reconstruction efficiency for high angled tracks within the scintillator detector. Fig. 4.4 shows an illustration of the SuperFGD concept. Note that the illustration does not have the same dimensions as the SuperFGD proposed for T2K, which will have significantly more cubes and add almost two tons of active material to

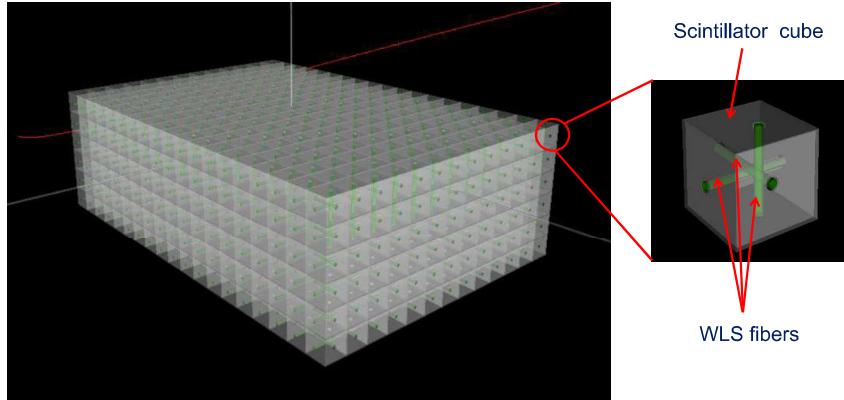


Figure 4.4: A rendering of the SuperFGD to be used in the upgraded ND280. The detector comprises many small plastic scintillator cubes that are read off in three orthogonal directions by wavelength shifting fibres. There are fewer cubes depicted here than there will be in the actual detector to make the diagram clearer [66].

ND280.

4.3 Predicted ND280 Upgrade Performance

The performance of the ND280 detector after the proposed upgrades are installed has been simulated extensively using a Monte Carlo simulation. The neutrino interactions within the detector are modelled using the particle simulation programme GENIE [67], whereas the detector geometry and particle trajectories are simulated using Geant4 [68]. Both the current ND280 and the upgraded ND280 were simulated for performance comparisons. The official T2K flux simulation was used as an input to the GENIE neutrino generation. For more information on the Monte Carlo simulation and validation of its accuracy, see [59]. Here, we will focus on the results of the performance comparison between the current ND280 and the upgraded ND280.

Runs of both the FHC and the RHC T2K neutrino beam were simulated using the Monte Carlo simulations of ND280. Using these simulations, selections of neutrino events were made. The event selections were for ν_μ -CC inclusive events from the T2K FHC beam and $\bar{\nu}_\mu$ -CC inclusive events from the T2K RHC beam. The selection itself assumes the most energetic particle detected with the correct charge is the expected muon from the interaction. If the muon candidate enters a TPC and deposits a track longer than 20 cm

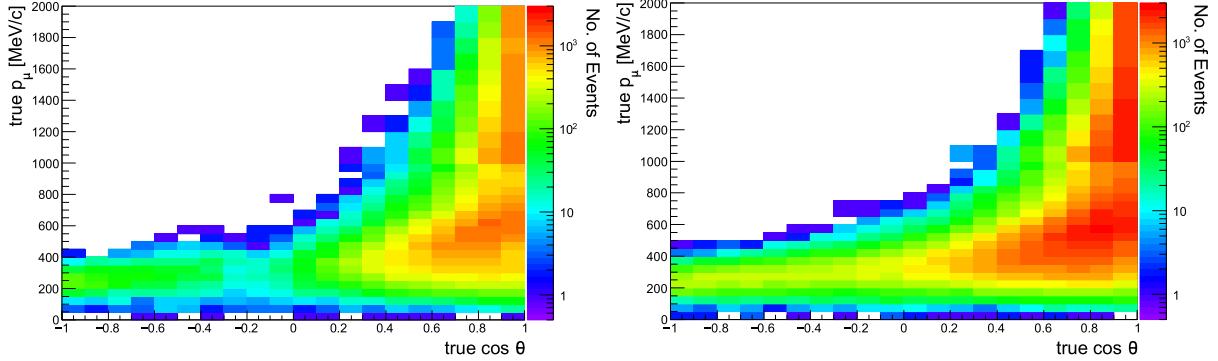


Figure 4.5: A comparison of the number of detected muons from simulations for the current ND280 (left) and the upgraded ND280 (right). The colour bars represent the number of detected muons for the true momenta and scattering angles shown in each plot [59].

and the TPC PID algorithms identify the particle as a muon, the event is accepted. In the absence of a TPC track the event can be accepted if the muon candidate entered an ECal and was identified as a muon.

Fig. 4.5 is a comparison of the total number of events selected for the simulations of the current and upgraded ND280 using the FHC T2K beam. The number of selected events at different true scattering angles and momenta are shown. Fig. 4.6 shows plots of the selection efficiency for the same ν_μ -CC selection of the simulated FHC beam and directly compares the efficiencies of the current and upgraded scintillator detectors as a function of both scattering angle and momentum. Using evidence from both figures, it is clear that the upgraded ND280 has an improved angular acceptance for high-angled tracks due to the SuperFGD and TPC. Backwards-going tracks are also detected more efficiently thanks to the TOF detectors.

Due to the significant mass of the SuperFGD detector, the number of detected events within ND280 is expected to increase. Tab. 4.1 displays the number of selected events with the simulated ν_μ -CC selection as described previously for both the current and upgraded ND280. The total number of selected events is more than doubled for the upgraded ND280 detector mostly due to the presence of the SuperFGD, but the FGD detectors also see increased selection numbers because of the improved acceptance of backwards going tracks engendered by the new TOF detectors. The purity of certain sub-samples of the selection are also shown Tab. 4.1. Major improvements in sample purity are seen in the CC1 π and

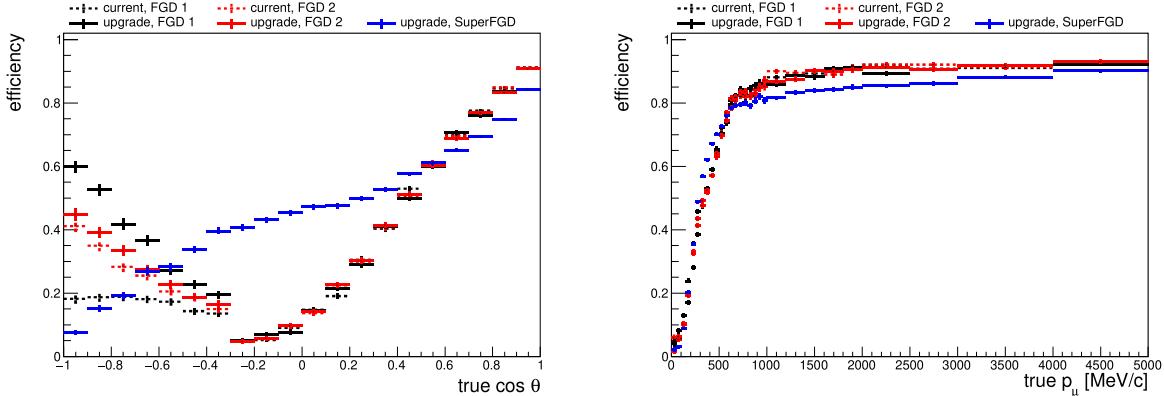


Figure 4.6: The simulated detection efficiencies with respect to lepton scattering angle (left) and momentum (right) for both the current and upgrade ND280 sub-detectors. The indicated scintillator detectors are the detectors in which the muon candidate originated. FGD 1 and FGD 2 are the FGD detectors already present in ND280, with FGD1 being the most upstream of the two [59].

Detector	No. of Events (/10 ²¹ POT)	Purity		
		CC0 π	CC1 π	CCother
Current	FGD1	50507	72.5%	64.0%
	FGD2	50125	71.5%	62.3%
Upgraded	FGD1	52655	72.9%	64.1%
	FGD2	51460	71.6%	62.9%
	SuperFGD	95490	72.5%	70.3%

Table 4.1: The number of selected events and sample purities for the simulated ν_μ -CC analysis of the current and upgraded ND280 detector using 1×10^{21} POT [59].

CCother samples, mainly due to the SuperFGD's low momentum thresholds.

The improvements brought to ND280 by the new sub-detectors aim to reduce the systematic uncertainties relating to the near-detector fit in the T2K oscillation analysis. As such, sensitivity studies have been carried out to estimate the effect of the ND280 upgrade upon the experiment's systematics. The near-detector fitter (named "BANFF" in other literature) that is used to constrain the flux and cross-section uncertainties in the T2K oscillation analyses is described in detail in Abe et al [69]. In short, it maximises a binned likelihood ratio as a function of the neutrino flux, cross-section and detector systematic parameters, with constraints reflecting prior knowledge of neutrino fluxes, cross-sections and detector systematic uncertainties. For the sensitivity studies, a tool

Source of Uncertainty	ν_μ	ν_e CCQE-like	ν_e CC1 π^+
	$\delta N/N$	$\delta N/N$	$\delta N/N$
Flux+Cross-section (constrained by upgraded ND280)	1.8%	1.9%	1.4%

Table 4.2: Predicted uncertainty contributions for flux and cross-section constraints using the upgraded ND280 with 8×10^{21} POT [59].

was developed that was effectively identical to the near-detector fitter but could fit the simulated samples of neutrinos and anti-neutrinos for the current and upgraded ND280.

The uncertainties associated with the upgraded ND280 detector relating to the neutrino flux and cross-section constraints were evaluated by performing a fit of the Asimov data set, the most probable data set corresponding to the Monte Carlo expectation. The simulated FHC and RHC beam corresponded to 8×10^{21} POT each, which is roughly a third of the expected total data collected at the end of T2K-II. The end result of the study is shown in Tab. 4.2. This can be compared with Tab. 3.1 showing the systematic uncertainties associated with the current ND280 detector. The uncertainties are more than halved moving from the current ND280 to the upgraded version, with an average sample uncertainty of 4.0% for the current ND280 and 1.7% for the upgraded ND280. This is a clear demonstration of the impact the near-detector upgrade will have on the T2K experiment.

4.4 Simulated Performance of the SuperFGD

As a brand new and untested technology before the initial ND280 upgrade studies, the SuperFGD required accurate simulation studies to confirm its viability as a particle detector. Such studies were performed using the simulation described in Sec. 4.3. The SuperFGD-specific study focussed on particle tracks that were fully contained within the SuperFGD volume so there were no detected particle hits in the outermost cubes of the detector. For such a track to be selected in the study, it had to leave at least two detected particle hits in at least two views (xy , xz , yz) and the track's endpoint had to be separated by at least one cube width from any other concurrently observed tracks [59].

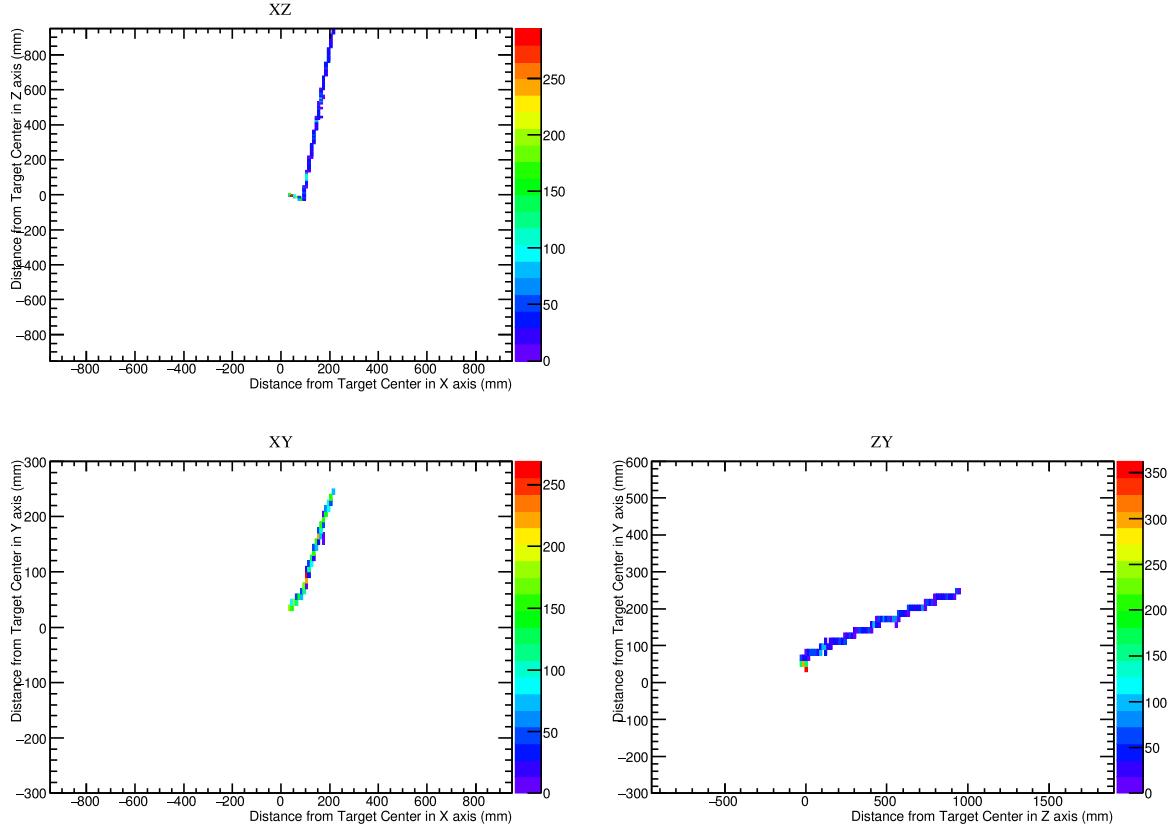


Figure 4.7: An event display of a simulated ν_μ CCQE event within the SuperFGD. Each plot shows the hits detected in a single plane of the detector. Colours represent the charge deposit detected in units of photoelectrons [59].

For context, a typical ν_μ CCQE event simulated within the SuperFGD is shown in Fig. 4.7. The event display shows hits from each view of the detector separately. The long track in the event corresponds to the muon produced in the interaction, whereas the short tracks at the interaction vertex and along the muon track belong to recoiling protons. The visibility of these protons attests to the reconstruction efficiency of low momentum protons within the SuperFGD.

Using the simulated selection in the SuperFGD, reconstruction efficiencies for different momenta and scattering angles were calculated. The results for reconstructed muons, protons and pions can be seen in Fig. 4.8. Also shown are the reconstruction efficiencies if hits from one plane of the detector are ignored. The plots in Fig. 4.8 demonstrate the significant improvement a three-dimensional readout gives to the track reconstruction.

In all cases, using all three readouts results in better reconstruction efficiencies. This improved efficiency over all momenta and scattering angles shows the benefit of using a SuperFGD design, rather than something like a simple XZ FGD design, where a normal FGD design is used with scintillator bars orientated along the x - and z -axes. Such a design was considered for the ND280 upgrade, but the study shown in Fig. 4.8 clearly shows the SuperFGD would be more desirable. The momentum threshold for protons is also improved with the SuperFGD design, giving a threshold of about $300 \text{ MeV}/c$ compared to about $500 \text{ MeV}/c$ when only two dimensions are used.

Another simulated study was performed to ascertain the potential momentum resolution of the SuperFGD. Using a particle gun firing 600 MeV muons, a momentum resolution of 3% was obtained [59].

The ability for the SuperFGD to identify protons, pions and muons was also evaluated. In the study, particle identification was performed using the ratio between the number of detected photoelectrons in a track and the track length. A stopping muon/pion is misidentified as a proton in 8% of the cases, while a proton is misidentified as a muon/pion in 17%/22% of the cases. These numbers are similar to the ones found for the FGD detectors in the same simulation.

The simulated studies of the SuperFGD have shown the many potential benefits of the SuperFGD design over a typical FGD design. However, being a brand-new design, real-life studies are also necessary to both corroborate the simulation findings and evaluate the difficulties in building and running the detector. In the next chapter, we will go into details of the extensive studies performed on a prototype of the SuperFGD.

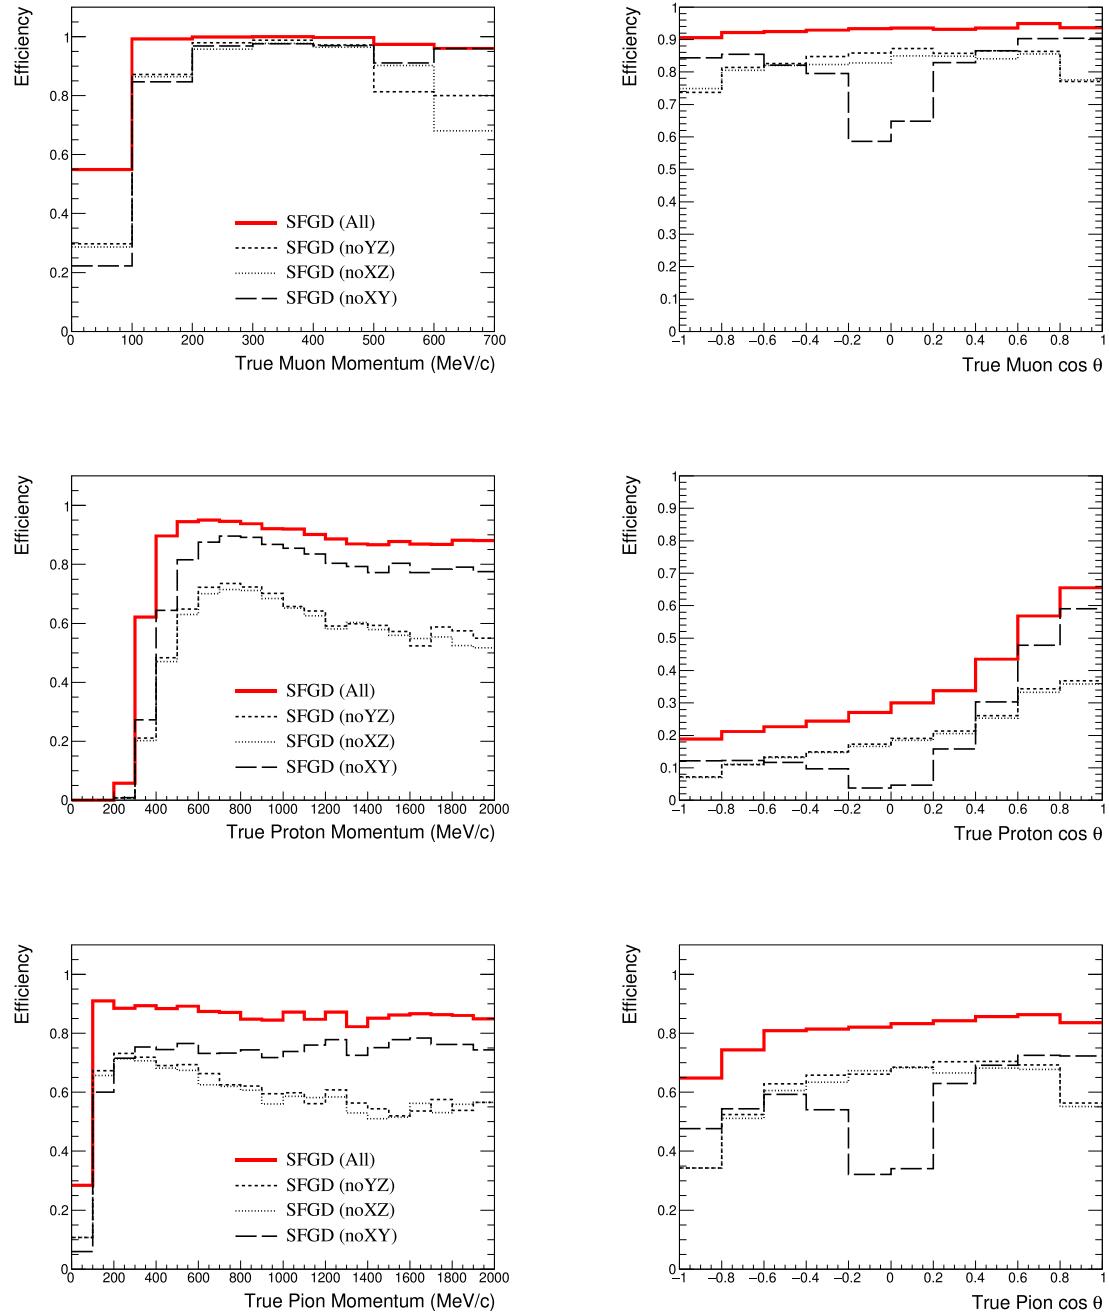


Figure 4.8: Reconstruction efficiencies for muons (top), protons (middle) and pions (bottom) as a function of true momentum (left) or true scattering angle (right) in the SuperFGD with three readout views (red) or with only two readout views (various black) [59].

Chapter 5

The SuperFGD Prototype

In 2018, a prototype of the SuperFGD was constructed for research and development of the SuperFGD concept. The prototype was tested using a charged particle beam at CERN in 2018 and a neutron beam at the Los Alamos National Laboratory (LANL) in 2019.

The SuperFGD prototype beam tests and the subsequent analysis of the beam test data have given a plethora of invaluable information and insight as to the operation, characteristics and limitations of the SuperFGD technology. In this chapter we will detail the design and construction of the SuperFGD prototype, as well as the beam test configurations and the results of the beam test data analyses.

5.1 Prototype Scintillator Cubes

The scintillator cubes of the SuperFGD prototype comprise polystyrene doped with 1.5% of paraterphenyl (PTP) and 0.01% of 1,4-bis benzene (POPOP). In 2018, more than 10,000 cubes were produced at UNIPLAST Co. (Vladimir, Russia) for the use in assembly testing and prototype construction. The cubes measured $10 \times 10 \times 10 \text{ mm}^3$ and were cut to size from long 10 mm-thick extruded square prisms.

This cube production was also used to collect information on the best cube construction method to use with the final SuperFGD detector. The variation of the cube sizes was tracked and shown to have a standard deviation, σ_w , of roughly 100 μm . This is relatively large, hence a new production method has been developed for the final SuperFGD design. The new method developed for the final SuperFGD design is based on injection moulding, which has a cube size standard deviation, σ_w , of less than 30 μm . This variation in the cube size is within the tolerance required to be able to fit the SuperFGD within the box

designed for the detector. However, note that cubes produced with the extrude and cut method, rather than the injection moulding method, were used to construct the SuperFGD Prototype detailed in this thesis.

Once the cubes for the SuperFGD prototype were constructed, a reflective layer was etched onto each of the cubes' surfaces with a chemical agent, leaving a 50–80 μm -thick white polystyrene micro-pore deposit [70]. This layer prevents light passing between cubes within the detector, although it is not completely effective, as seen in the data analyses.

For the WLS fibres to be threaded through the cubes, three orthogonal holes of 1.5 mm diameter were drilled through each cube using a Computer Numerical Control (CNC) milling machine. The average distance between the hole centre and the cube side was measured to be 3.11 ± 0.08 mm, larger than the specified value of 3.00 mm. Fortunately, deviations of the hole position less than 0.2 mm from the desired value should not significantly increase the cube position uncertainty. The gap between the 1 mm thick WLS fibre and the 1.5 mm thick hole gives us some tolerance to work with.

5.2 WLS Fibres and Photosensors

The SuperFGD Prototype uses the Y-11(200)MS model of WLS fibres, produced by Kuraray Co. [71]. This is the same fibre that is used in ND280's current FGDs and will be used in the final SuperFGD design. The fibre model is a 1 mm diameter, multi-clad, round shape fibre of S-type (increased flexibility). These fibres are used for their well-known reliability throughout the T2K experiment and other large experiments [72].

There are three different types of photosensor installed in the SuperFGD prototype. All photosensors are Multi-Pixel Photon Counters (MPPCs), which are a type of silicon photomultiplier that utilise multiple pixels so that more than one photon can be detected by a single photosensor at the same time. Three different MPPC models are used partly to save on construction costs of the prototype detector and partly for research purposes. Most of the MPPCs are of type S13360-1325CS, which are the same model to be used in the final SuperFGD design, except the prototype's MPPCs include a ceramic casing. The MPPCs for the final SuperFGD will be surface mounted for increased stability, eliminating the need for ceramic housing. The other two MPPC types used in the SuperFGD Prototype

Parameter	Type I	Type II	Type III
Manufacturer ref.	S13360-1325CS	S13081-050CS	S12571-025C
No. in Prototype	1152	384	192
Active area [mm ²]	1.3 × 1.3	1.3 × 1.3	1.0 × 1.0
Number of pixels	2668	667	1600
Pixel size [μm]	25	50	25
Operating voltage [V]	56–58	53–55	67–68
Photon detection eff. [%]	25	35	35
Dark count rate [kHz]	70	90	100
Gain	7×10^5	1.5×10^6	5.15×10^5
Crosstalk probability [%]	1	1	10

Table 5.1: Main parameters for the three MPPC types installed on the SuperFGD Prototype. Crosstalk probability is the probability of fired pixels creating secondary photons and firing neighbouring pixels and is unrelated to the optical crosstalk between scintillator cubes discussed in later sections. The parameters were retrieved from the manufacturer, Hamamatsu [75].

were spares from the Baby-MIND detector for the WAGASCI experiment [73, 74]. Due to several members of the SuperFGD Prototype development team also being involved with WAGASCI, there are several aspects of the SuperFGD Prototype derived from Baby-MIND, as will be made clear later on.

For simplicity, rather than use the model references to refer to the different MPPCs, we will name them Type I, Type II and Type III, as specified in Tab. 5.1, where the properties of the different MPPC models are also listed.

A custom optical coupler, also known as a ferrule, was used to mount the MPPCs to the ends of the WLS fibres. These ferrules were originally designed for the Baby-MIND detector. There are two main parts to a ferrule, the fibre ferrule and the MPPC housing, both of which are shown in the diagram of Fig. 5.1. An optical cement (EJ-500 produced by Eljen Technology) is used to glue the fibre ferrules onto the ends of the WLS fibres. The fibre is threaded through the hole at the end of the fibre ferrule before being glued. The protruding fibre is then cut so that the end is flush with the fibre ferrule. The cut end is subsequently polished using a drilling machine with a polishing wheel. Only one end of each WLS fibre has a fibre ferrule attached, the other end is left bare and unpolished.

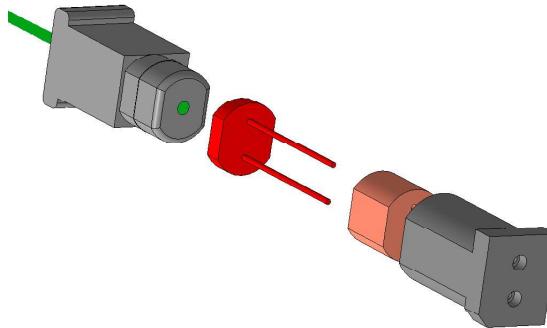


Figure 5.1: A diagram of the interface between the fibre ferrule (left) and the MPPC housing (right). The MPPC (red) is mounted into the MPPC housing with a foam spring (light red). The fibre ferrule is then plugged into the MPPC housing so that the polished end of the fibre makes contact with the MPPC [76].

The MPPC housing is the female plug to the fibre ferrule's male plug. A single MPPC is inserted into the MPPC mounting by threading the two connector pins of the MPPC through the two holes of the MPPC housing. A small foam spring is also placed between the MPPC and the housing to allow the end of the fibre ferrule to press up against the MPPC when inserted but still allow the fibre ferrule to be fully inserted into the MPPC housing. The fibre ferrule and the MPPC housing latch together by means of a locking groove inside the MPPC housing and a matching ring protrusion on the fibre ferrule.

The MPPC connector pins protrude from the MPPC housing and attach to a mini-PCB to which an ultra-miniature, low-profile coaxial connector is soldered. A 1.5 m long micro-coaxial cable (U.FL-2LP-088) is attached to the coaxial connector and transmits the MPPC signal to the front-end electronics. A bias voltage is applied to the MPPCs via individual micro-controllers, which are connected to the MPPC cathodes via micro-coaxial cables [73].

5.3 Readout Electronics

For the SuperFGD prototype, and also the final SuperFGD detector, the readout electronics are heavily based on previous technology used in the Baby-MIND detector. Signals from MPPCs are processed by a CITIROC (Cherenkov Imaging Telescope Integrated Read Out Chip) front-end ASIC (Application-Specific Integrated Circuit). Each CITIROC

chip can process 32 MPPC signals. The CITIROCs are integrated into front-end boards (FEBs), which contain circuitry for manipulation and transmission of the signal processed by the CITIROCs. Each FEB contains three CITIROCs, so 96 MPPCs can be read by a single FEB. In total, there are 18 FEBs, distributed across four minicrates that also house cooling fans to prevent overheating.

In each CITIROC, there are three readout channels for each MPPC signal, as seen in Fig. 5.2. There is a High Gain (HG) channel, a Low Gain (LG) channel and a Time-over-Threshold (ToT) channel. The HG and LG channels are independent of each other and include a signal pre-amplifier and a slow signal-shaper. The ToT channel reads either the HG or LG pre-amplifier output (for the beam tests it was the HG pre-amplifier output) and utilises a fast signal-shaper. A discriminator threshold is imposed on the pulse shaped by the fast signal-shaper, such that if an MPPC pulse exceeds the threshold, then the acquisition window is triggered for all readout channels of all MPPCs connected to the CITIROC. For the HG and LG channels, the peak of the signal recorded within the acquisition window is used to calculate the signal amplitude. The ToT channel instead records the points where the signal crosses the discriminator threshold. These timestamps are then sampled by a 400 MHz clock in the FPGA, which records the rising edge time and the falling edge time. The rising edge time is the time at which the signal passed above the threshold and the falling edge time is when the signal passed back below the threshold.

The duration of the acquisition window was set to 10 μ s for the CERN and LANL beam tests. For the HG and LG channels, a peak detector identifies the peak of the signal for each MPPC during this acquisition window. The amplitude of this peak is then digitised using a 12-bit Analogue-to-Digital Converter (ADC). The signal amplitudes are then passed on to the FEB's FPGA, where they have to pass a threshold in order to suppress low amplitude hits.

As implied by their names, the HG and LG channels apply different levels of gain during the amplification stage of the signal readout. Thus, the HG and LG channels complement each other as the HG channel is ideal for reading low amplitude signals, whereas the LG channel can read high amplitude signals that would over-saturate the HG channel.

It should be noted that only one HG and one LG signal can be recorded per MPPC

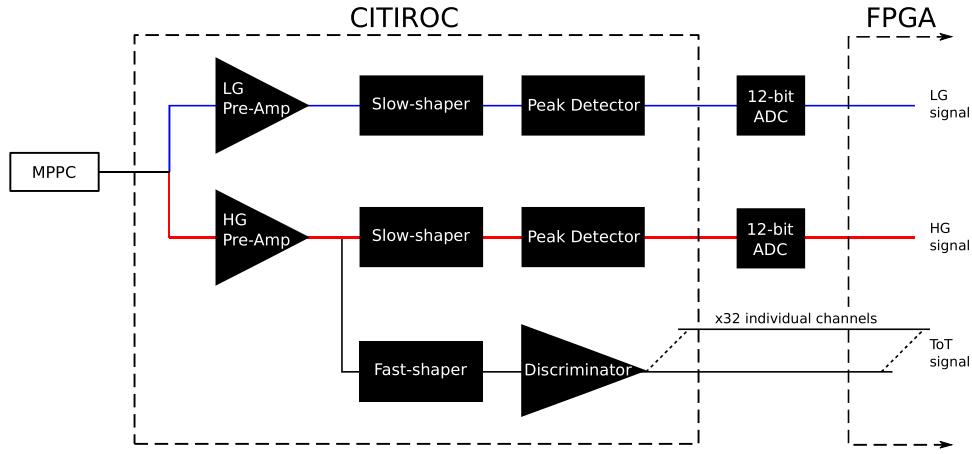


Figure 5.2: The three signal outputs of the SuperFGD readout and the different processing stages of each path. A High Gain (HG), Low Gain (LG) and Time-over-Threshold (ToT) signal is generated for each MPPC signal sent to the CITIROC. The HG and LG channels multiplex the signals from the 32 MPPCs connected to the CITIROC into one complex signal, whereas the ToT channel has 32 separate channels dedicated to each MPPC. There is some dead-time associated with the HG and LG channels. The ToT path has no dead-time and can be used to read signal amplitudes that otherwise wouldn't have been output by the HG and LG channels.

for each acquisition window. Also, the HG and LG channels multiplex the 32 MPPC signals into one signal, which is then digitised. These processes introduce a dead-time of $9.12\ \mu\text{s}$ after every acquisition window. As such, only one HG and one LG signal can be recorded per MPPC every $19.12\ \mu\text{s}$. This dead-time can be overcome by utilising the ToT channel readings, which are not multiplexed and instead have dedicated channels for each MPPC. The time period over which the signal is above the threshold, i.e. the difference in time between the falling and rising edges, is a function of the signal amplitude. As a result, for hits that are not recorded by the HG or LG channels, the time-over-threshold value can be used. The use of the ToT channel in this way effectively means there is no dead-time associated with the readout electronics.

5.4 Assembly of the SuperFGD Prototype

Given the nature of the SuperFGD design, comprising of many small cubes, the assembly of the detector is a challenging process. The assembly of the SuperFGD Prototype was an excellent run-through of the technique planned for use in assembling the final SuperFGD

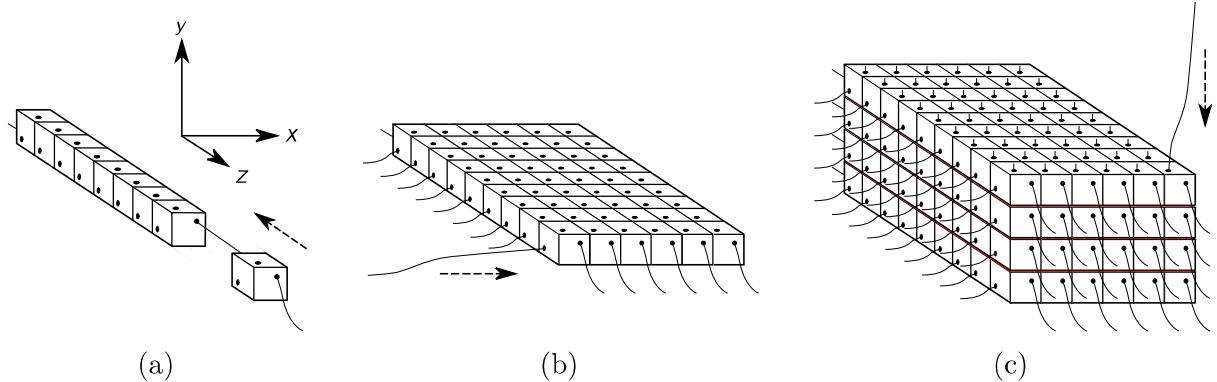


Figure 5.3: The three main stages for the cube assembly of the prototype. Images are not to scale and do not match the number of cubes in the prototype in order to make the pictures clearer. Stage (a) shows a single row of cubes being threaded along a fishing line. When enough single rows are made, in stage (b) they are threaded together using more fishing lines along the x direction to make several 2D planes. Stage (c) is where the assembled 2D planes are stacked vertically, separated by Tyvek sheets (shown in red), and more fishing line is threaded through the holes in the y direction..

detector for T2K.

The technique used to assemble the prototype is shown in Fig. 5.3. There were three main stages of assembly. The first stage was assembling single rows of 48 cubes along a fishing line fibre (see Fig. 5.3a). The cubes were not initially assembled on the WLS fibres due to the delicate nature of the WLS fibres, which could easily be damaged during the process.

Once 24 rows of cubes had been assembled, they were lined up next to each other in an array and fishing lines threaded through the aligned holes in the cubes in order to form a 2D plane (see Fig. 5.3b).

This process was repeated eight times to produce eight $24 \times 48 \text{ cm}^2$ planes. These eight planes were then stacked one on top of the other, with a sheet of Tyvek paper placed in between each layer. The Tyvek paper had holes punched into it to allow more fishing line to be threaded through the vertically aligned holes of the stacked planes (see Fig. 5.3c).

It should be noted that the Tyvek sheets were used in the detector design partly as research into the effectiveness of Tyvek as a prevention of light leakage between cubes. It was also useful in creating a flat surface between each layer of cubes. As mentioned in Sec. 5.1, the cubes used in the SuperFGD Prototype construction had a standard

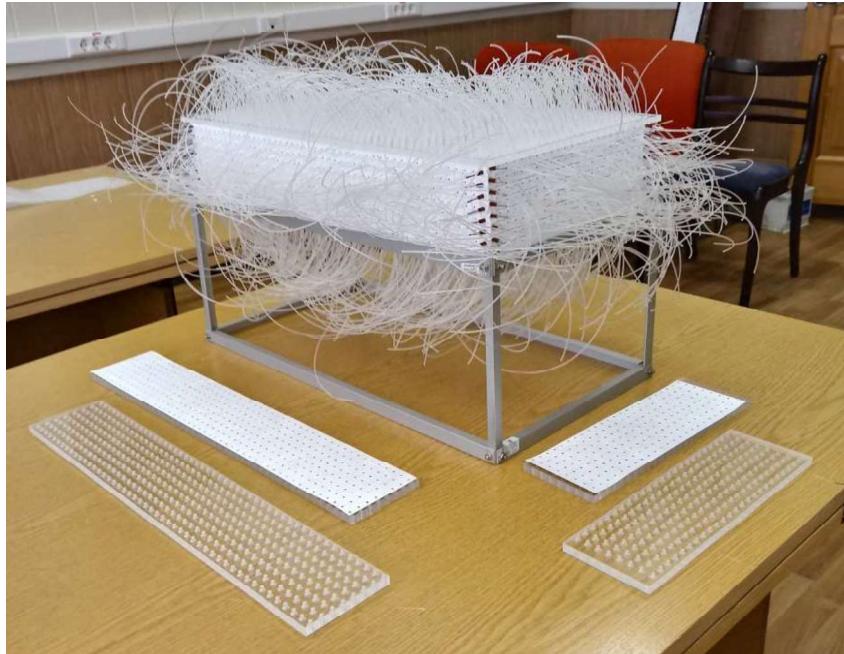


Figure 5.4: The SuperFGD Prototype volume assembled with fishing lines. Also shown are the plates of the acrylic box to be fixed to the detector to hold the cubes in place. The acrylic plates for the upper and lower surfaces have already been attached.

deviation of 100 μm in size. As such, without the Tyvek paper the cube layers would not have stacked together so neatly. However, in the final SuperFGD design, the standard deviation can in cube size should be no more than 30 μm , given the new cube construction technique. As a result, the use of Tyvek paper is not intended in the final SuperFGD.

A photograph of the assembled cubes for the prototype threaded with fishing lines is shown in Fig. 5.4. Before removing the fishing lines, 8 mm-thick acrylic plates, with drilled holes for the fishing lines to fit through, were fitted around the mass of cubes to form a box. The box fits tightly around the assembled cubes so that there is no air gap between the acrylic and the cubes. The plates were fixed together by screws at their edges. Once fixed in place, the cubes are held together by the pressure applied by the acrylic sheets, allowing the fishing lines to be removed and replaced with the WLS fibres.

The WLS fibres were inserted with the fibre ferrules attached to one end (see Fig. 5.5). To reduce the density of fibre ferrules along the sides of the detector, the WLS fibres were installed such that the fibre ferrules were not all on one side of the detector. Instead, fibres were inserted in the same direction for one column of cubes, then for the next column of



Figure 5.5: WLS fibres with the fibre ferrule attached to one end. The fibres were installed in the detector in this configuration.

cubes the fibres were inserted from the other side. This produced an even spread of fibre ferrules along the outside of the detector, making the job of attaching the MPPCs to the fibre ferrules much easier.

To keep the fibre ferrules fixed in place, 6 mm-thick strips of acrylic, with holes cut out to accommodate the fibre ferrules, were screwed in to the external box of the prototype. The prototype with the installed WLS fibres can be seen in Fig. 5.6, where there is a view of the acrylic strips holding the fibre ferrules in place.

In order to protect the mounted electronics on the detector, such as the photosensors and readout cables, the SuperFGD Prototype was mounted in an aluminium frame that would support readout cables and prevent contact between the electronics and any light isolation sheets to be wrapped around the detector when in use. This aluminium frame was fixed to a steel platform with a mechanism to allow the detector and its aluminium frame to be rotated about both the axis parallel to its long edge and the vertical axis. The ability to rotate the detector easily allowed testing of the prototype with particles travelling at different angles within the detector. The steel platform was also attached to aluminium bars that extended by about a metre past both ends of the detector. The minicrates housing the sensitive readout electronics were later mounted to the ends of these aluminium bars so that they were kept out of the magnetic field the detector was

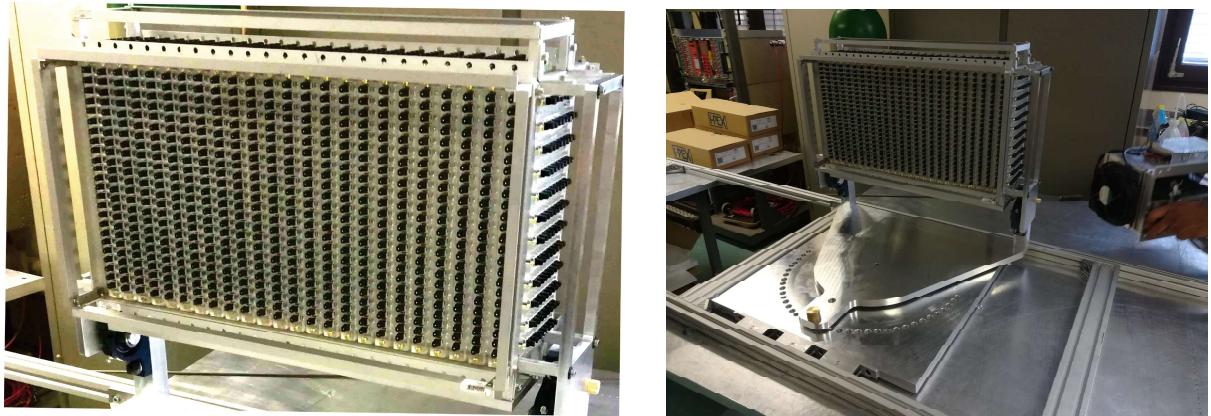


Figure 5.6: The SuperFGD Prototype with WLS fibres installed and mounted in an aluminium frame. Left: a close-up of the face of the detector. Right: An overview of the detector and the aluminium frame. The frame is mounted onto a steel platform that allows the detector to be rotated about the axis parallel to its long edge and the vertical axis. The steel platform is also attached to aluminium poles that allow the minicrates holding the electronics to be mounted a decent distance from the detector itself. This meant the minicrates did not sit in the magnetic field the detector was placed in during the beam tests. The detector is shown here rotated 90° about the z -axis compared to its normal configuration.

placed in during the CERN beam test. The detector in its aluminium frame and the steel platform can be seen in Fig. 5.6.

Once the SuperFGD detector was mounted in its aluminium frame and attached to the steel platform, the MPPCs were connected to the fibre ferrules via the mechanism shown in Fig. 5.1. As mentioned in Sec. 5.3, there were three different types of MPPCs installed on the SuperFGD Prototype. The distribution of the MPPC types around the faces of the prototype is shown in Fig. 5.7. The co-ordinate axes used to reference cube position and beam direction for the beam test data is also shown in the figure, with the z -axis pointing along the default direction of the particle beam during the beam tests.

There was significant strategy to the placement of MPPCs around the detector. Type I MPPCs had the highest dynamic range, so we placed them on the front and back faces of the detector where they would read signals from fibres collecting light from several cubes, assuming the particle beam was directed along the z -axis. The left and right sides of the detector were also fully instrumented with Type I MPPCs to guarantee we captured full track profiles with just one MPPC type. The upper and lower faces of the detector were the only faces to have more than one MPPC type. The remaining Type I MPPCs were

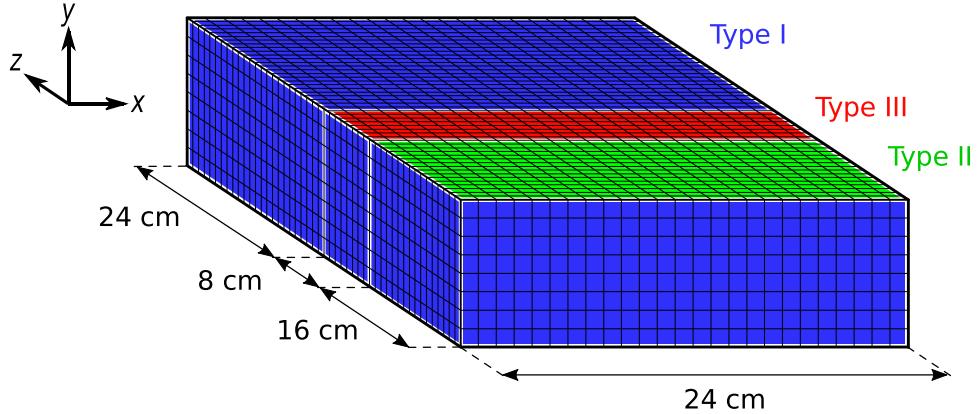


Figure 5.7: Distribution of the three types of MPPCs around the six faces of the SuperFGD Prototype. There are 1,152 Type I (blue) MPPCs, 384 Type II (green) MPPCs and 192 Type III (red) MPPCs. The colours do not represent the position of the MPPCs themselves, only the positions of the fibres that they are reading out. The MPPCs themselves may be on the opposite side of the detector to the displayed colours. The z -axis points in the direction of the particle beam for the majority of runs in the particle beam tests.

installed at the downstream ends of these faces, where we expected the stopping point of $0.8 \text{ GeV}/c$ protons to be.

Once all MPPCs were installed onto the prototype, they were connected to the readout electronics. A photo of the fully instrumented SuperFGD Prototype, with MPPCs connected to readout electronics housed in minicrates mounted to the aluminium bars of the steel platform, is shown in Fig. 5.8.

Each MPPC had a unique operating voltage specified by the manufacturer. The range this operating voltage could be varied depending on the MPPC type (see Tab. 5.1). A common voltage was supplied to all MPPCs connected to the same CITIROC. The voltage supplied to the individual MPPCs was then tuned for each MPPC using CITIROC embedded 8-bit 4.5 V DACs that were connected to each MPPC anode. To ensure that the required operating voltages could be supplied to all MPPCs with this method, MPPCs with similar operating voltages were connected to the same CITIROC. As such, the range of the operating voltages of MPPCs in each CITIROC was no greater than 1 V.

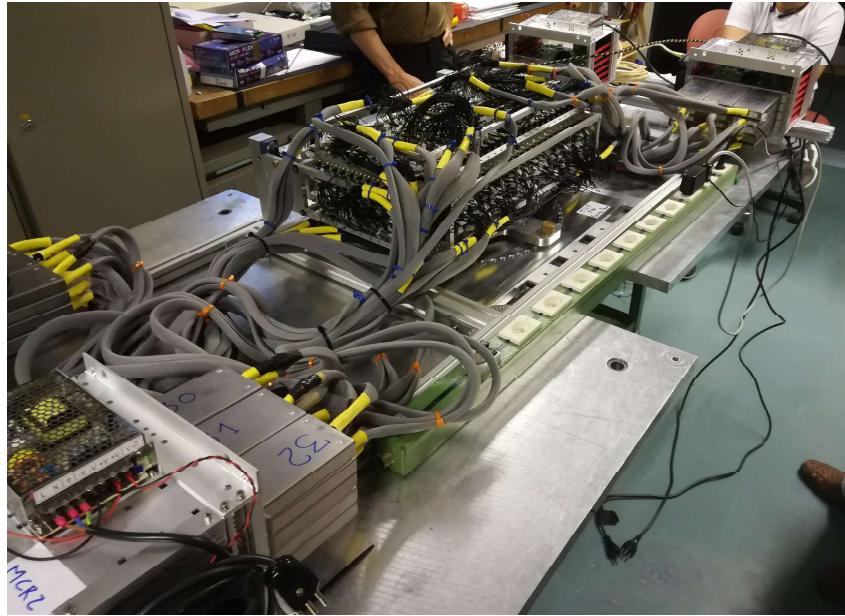


Figure 5.8: The SuperFGD Prototype with all readout electronics attached. The minicrates housing the front end electronics can be seen on each of the four corners of the steel platform the SuperFGD Prototype is mounted on.

5.5 Calibration

Once the SuperFGD Prototype had been fully assembled, the various readout channels for each MPPC were calibrated. The HG and LG channels needed to be calibrated in order to convert the ADC units to the more physically understandable unit of photoelectrons (p.e.), which corresponds to the number of incident photons on the MPPC. The ToT channel instead requires conversion from units of time to p.e. The methods used to calibrate each channel are as follows:

- HG Calibration - find the ADC/p.e. ratio by analysing finger plots produced from MPPC dark counts (only viable for Type III MPPCs, which have a high dark count) or from calibration measurements performed by using a custom LED set-up to get output from the MPPCs [77]. An example of a finger plot is shown in the left plot of Fig. 5.9. The peaks seen in the plot correspond to the number of ADC counts seen when one, two, three or four photons are incident on the MPPC. Measuring the difference in ADC units between neighbouring peaks gives the ADC/p.e. ratio.
- LG Calibration - the only difference between the LG and HG channels are the

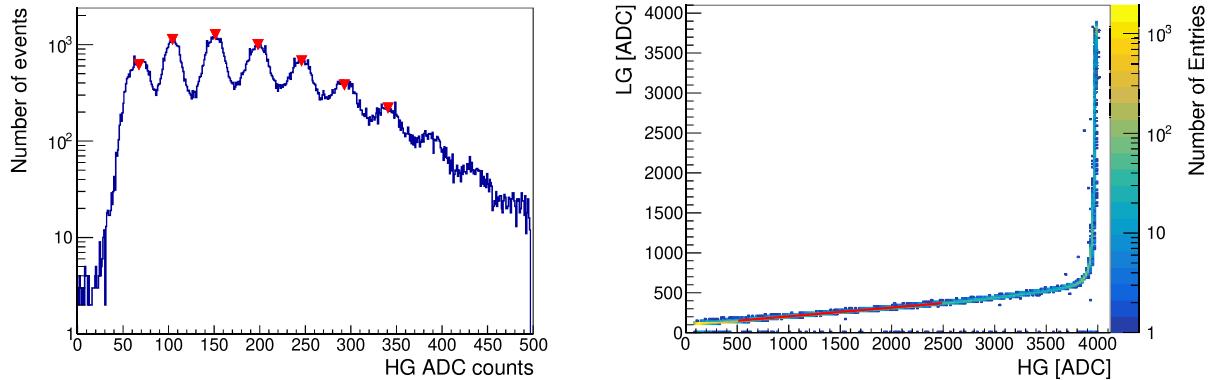


Figure 5.9: Left: typical calibration plot used to extract the ADC/p.e. ratio for the HG signal path. Right: The LG data plotted against the HG data. A linear fit is shown in red, the gradient of which can be used to scale the ADC/p.e. ratio of the HG channel to correspond to the LG channel.

different gains of the pre-amplifiers. As such, the ADC/p.e. ratio of the LG channel can be found with a linear fit of the LG readings to the HG readings and scaling the HG ADC/p.e. ratio appropriately. An example of one of these fits for a single channel is seen in the right plot of Fig. 5.9.

- ToT Calibration - the time over threshold can be fit to HG and LG readings to find the relationship between the time over threshold and the signal amplitude. Specifically, we fit the ToT signal to the HG readings of 100 p.e. or below and to LG readings of 100 p.e. or above. We performed the fits using a fifth-order polynomial. Examples of the fits from a single channel are shown in Fig. 5.10.

The calibration at the CERN beam-test was the first ever calibration of the SuperFGD Prototype. As there was no prior knowledge of the detector response, three different gains settings were used during the calibration process, resulting in ADC/p.e. ratios of 42, 29 and 23 ADC/p.e. The reason for testing different gains settings was to try and maximise the signal produced by a muon travelling through the SuperFGD Prototype, without saturating the HG channel.

As seen in Tab. 5.1, the different types of MPPCs used for the prototype have different gains. As such, each MPPC type had custom pre-amplifier settings in an attempt to match the ADC/p.e. ratios of each MPPC type. Fig. 5.11 shows the distribution of the ADC/p.e. ratios for the HG channel's 29 ADC/p.e. setting. There is a natural spread in the ratios for each MPPC type, due to the uncertainties in gains and voltages. The

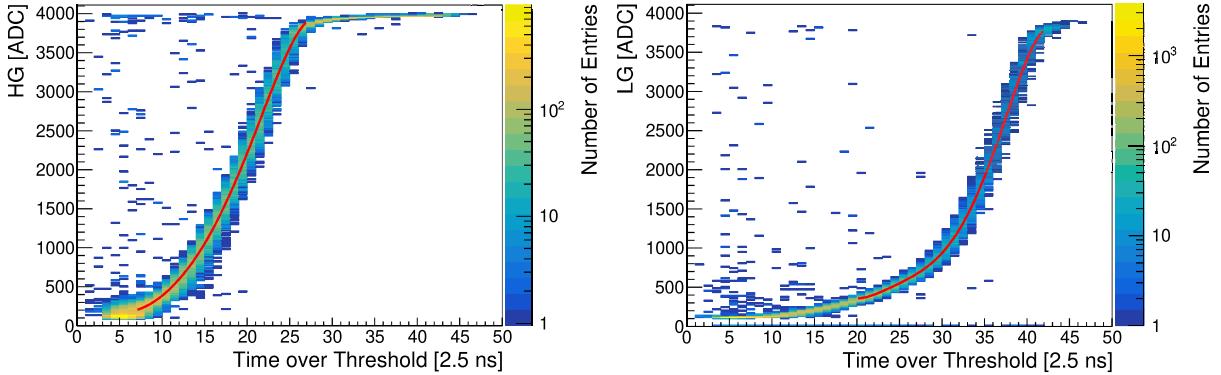


Figure 5.10: Examples of the fifth-order polynomial fits (shown in red) of ToT data to HG (left) and LG (right) data. The data points appearing as short lines which are significantly off the S-shaped curves are due to a mis-tagging between the HG or LG signal and its corresponding ToT signal.

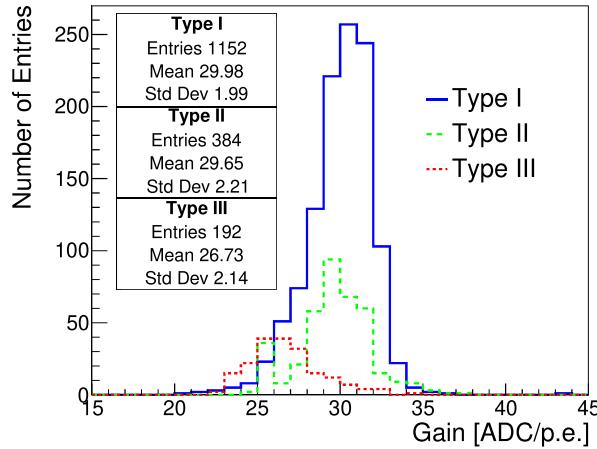


Figure 5.11: ADC/p.e. ratio distributions for all Type I (blue solid), II (green dashed) and III (red dotted) MPPCs with CITIROC HG pre-amplifier settings of 50, 43 and 45 respectively.

figure also shows the result of trying to match the ADC/p.e. ratios of the different MPPC types. Types I and II match relatively well, with Type III having a slightly lower ratio on average. It will become apparent to the reader throughout this section how difficult it was to match the behaviour of the Type III MPPCs to the other types, as the Type III properties—especially their dark count—are significantly different to the other MPPCs.

At the CERN and LANL beam tests, the HG ADC/p.e. ratio measurements were carried out at the beam-line location. The LG and ToT calibrations were performed using beam data. The MPPC gains have a significant temperature dependence, with

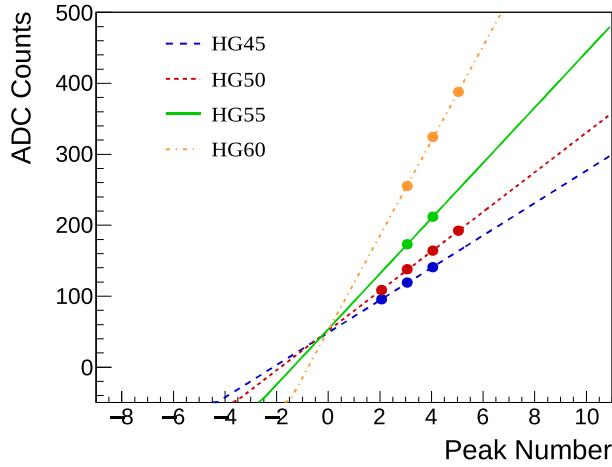


Figure 5.12: Pedestal finding method for a single channel with four different HG values. The ADC values of each photoelectron peak are plotted and extrapolated. The point of intersection is taken as the pedestal position [76].

roughly a 10% change in ADC/p.e. ratio with a change in temperature of 1°C. At CERN, the temperature around the prototype was not monitored and was assumed to be constant. The test area at CERN was well-insulated and air-conditioned, so this was not an unreasonable assumption. At LANL, the test area had a much more varied temperature, varying throughout the day and from day-to-day. As such, the temperature near the detector was monitored throughout the data-taking period, so that calibration constants could be adjusted to correspond to the temperature conditions of each data run.

There were also dedicated calibration runs to identify the pedestal value of each MPPC. The pedestal corresponds to a peak in an MPPC's finger plot caused by electronics noise, which has the lowest ADC value of all the observable peaks. The ADC value of the pedestal varied significantly between MPPCs, hence the need to measure it for each individual MPPC. Knowing the value of the pedestal is important when taking low amplitude measurements, although for the CERN and LANL beam tests the amplitude threshold was set above 0.5 p.e., so the electronic baseline was not observed. Fig. 5.12 shows the method used to find the pedestal value for a single MPPC. For several pre-amplifier gain settings, the ADC values of several peaks were identified from finger plots of calibration data. Linear fits were performed on the peak number vs ADC count data from each gain setting and extrapolated to lower ADC values. The point at which the linear fits intersected was taken to be the pedestal value.

It could reasonably be questioned, seeing as three different methods (HG, LG and ToT)

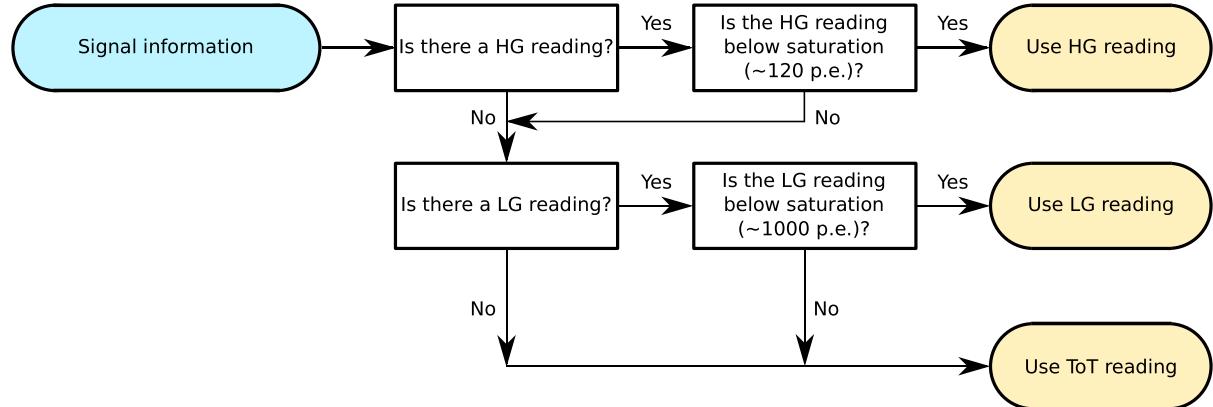


Figure 5.13: The algorithm used to decide which output channel was used to represent the signal amplitude.

are used to measure the amplitude of MPPC signal, “which amplitude measurement is used in the data analysis?”. The answer is that, for most analyses, a mixture of HG, LG and ToT signals are used. For each measured amplitude, a simple algorithm is used to decide which channel is used to represent the signal amplitude. This algorithm is shown in Fig. 5.13. The HG reading is taken if it exists and is below the HG channel’s saturation level, which is around 120 p.e. depending on the gain settings used. If one of these conditions are not met, the LG signal is used instead, unless it is above the LG channel’s saturation level of around 1000 p.e. or does not exist. Otherwise, the ToT channel is used to estimate an amplitude value. This algorithm allows amplitude readings to be made for every signal detected regardless of signal amplitude or electronics dead-time. The output of the algorithm is referred to as the “PE” output. In future sections, it will be made clear if an analysis utilises the HG, LG, ToT or PE signal.

5.6 The CERN Particle Beam Test

Data was collected using the SuperFGD Prototype at the CERN Proton Synchrotron (PS) T9 beam-line in 2018 over two separate data-taking periods: one in June-July and one in August-September. The PS delivers semi-regular spills of protons to the T9 beam-line, where they impinge on a target to produce a mixture of hadrons, electrons and muons [78]. The spills are delivered once or twice every 15 s and provide a steady flux of protons over 0.4 s. The target can be changed to alter the beam composition. For the SuperFGD

Prototype particle beam test, an aluminium target was used. The target was 55 m downstream of the SuperFGD Prototype’s position in the T9 experimental area, as such the only hadrons with decay times long enough to reach the prototype were protons and pions.

The momentum range of particles in the beam downstream of the aluminium target was between 0.5–10 GeV/ c . Beam-line magnets were used to adjust the beam momenta, so that only particles of a narrow momenta range were incident on the SuperFGD Prototype. A beam of positively charged or negatively charged particles could be chosen by reversing the magnetic fields. For some data runs, we moved a thick (50–100 mm) lead beam stopper into the beamline, which completely suppressed the hadronic component of the beam and partially suppressed the electron component, leaving mainly muons.

The SuperFGD Prototype was assembled for the first time immediately before the first data-taking period at CERN. As such, this first data-taking period was very much a case of getting used to using the detector and making sure everything functioned as expected. The second data-taking period was able to focus more on defining the properties of the detector.

For both data-taking periods at CERN, a TPC detector was installed upstream of the SuperFGD Prototype. This TPC was not part of the SuperFGD Prototype beam test; it was a separate experiment that we shared the beam-line with.

Several scintillator detectors and a Cherenkov radiation detector were also installed upstream of the SuperFGD Prototype. These were part of our experiment and were used as triggers to inform us of incoming particles. The scintillator detectors were used for time-of-flight information, allowing identification of particles. The Cherenkov detector was used to identify electrons, as the gas pressure (< 3 bar) only allowed the detection of electrons for the range of momenta used. Each detector was directly connected to the SuperFGD Prototype via a dedicated FEB. Triggers on specific particles or groups of particles could be made by triggering on a combination of signals from the scintillator and Cherenkov detectors. The trigger combinations and the corresponding particles, as well as the resulting purity, are shown in Tab. 5.2. Some samples had a relatively high contamination from unwanted electrons, as not all electrons were detected in the Cherenkov detector.

The SuperFGD Prototype itself was placed in a magnet—named MNP17—whose

Beam Mode	Desired Particles	Trigger Setup	Purity
Hadrons	All	$S_{2L} + S_3 + S_1$	N/A
Hadrons	p	$S_{2S} + S_3 + S_1$	> 90%
Hadrons	e	$S_{2L} + S_3 + C_2$	> 90%
Hadrons	π/μ	All (no e or p signal)	[e (40–50%)]
Muons	μ	All (no e or p signal)	[e (10–20%)]

Table 5.2: The trigger combinations used to trigger on specific particles and the purity of the resulting samples. The beam mode corresponds to when the beam stopper was removed (hadrons) or present (muons) in the beam-line. S_{2L} and S_{2S} refer to the temporal length of the signal output by the S_2 scintillator detector (long and short respectively). The π/μ and μ samples triggered on all trigger signals unless they corresponded to e or p signals. Square parentheses indicate the contamination from unwanted particles in the sample.

magnetic field could be varied or switched off entirely. Another variable magnet—named MDX—was fixed upstream of the SuperFGD Prototype.

The positions of all components of the beam test in the T9 beam-line are shown in Fig. 5.14.

One last thing before me move on to reviewing the beam test data: throughout our discussions we use the term “hit” to describe a signal from a single readout channel, be the hit from scintillation caused by a particle passing through the detector or from some other source. We also use the nomenclature “event” to represent the space in time when a particle passed through the detector and was detected.

5.7 CERN Beam Test Data Analysis: Basic Detector Properties

The CERN particle beam test was an excellent opportunity to characterise the inherent properties of the SuperFGD Prototype. Before the CERN particle beam test, the SuperFGD had only ever been a conceptual idea, with the exception of a small $5 \times 5 \times 5$ cube prototype that was mainly used to test the readout electronics. With the larger SuperFGD Prototype, detailed studies could be carried out using a large number of channels and

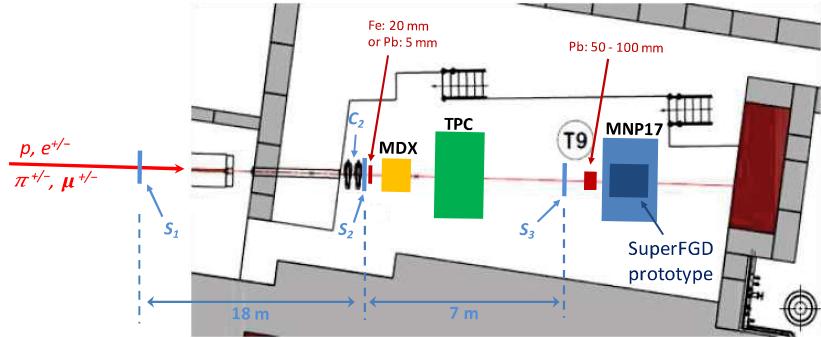


Figure 5.14: The set-up of the CERN particle beam test in the T9 beam-line. The positions of the S_1 , S_2 and S_3 scintillator detectors are shown, as well as the positions of the MDX and MNP17 magnets, the upstream TPC, the C_2 Cherenkov detector, the beam stoppers and the SuperFGD Prototype itself.

long particle trajectories, a much more comparable set-up to the final SuperFGD. This section will detail the analyses performed using the data collected with the SuperFGD Prototype at the CERN particle beam test, including particle events, hit time structure, hit amplitude threshold, channel response, optical crosstalk between cubes, light attenuation, cube response and the time resolution of the readout electronics. The results of these studies have gone on to inform the construction of the final SuperFGD design and the computer simulations of the SuperFGD and its electronics.

5.7.1 Particle Events

Despite there being many different particles in the particle beam, we will mainly focus on protons and muons within our analyses of the CERN particle beam data. These were by far the most commonly detected particles in the SuperFGD Prototype and were easily differentiable.

Fig. 5.15 shows 2D event displays of a muon event in the SuperFGD Prototype detected during a 2 GeV/c muon data run at the CERN particle beam test. The different plots show the charge amplitude and location of hits detected in each plane of the detector. These plots could be combined to make a 3D plot of the particle event, but 2D plots are easier to read in this format.

The muon in Fig. 5.15 was travelling parallel to the z -axis. We would expect to observe

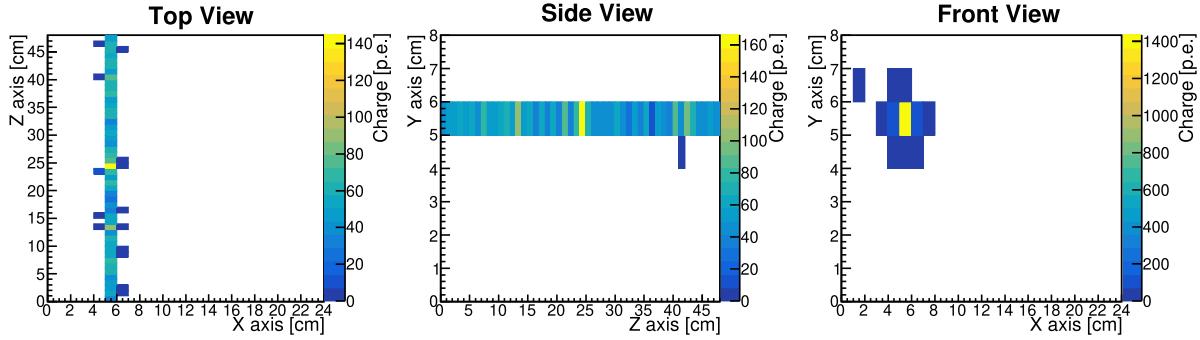


Figure 5.15: 2D displays of hits detected in a muon event at the SuperFGD Prototype. Hits from the three different planes of detection are show, with their colour representing the detected charge from the hit. No charge cuts are applied to the data. A narrow time window is applied to remove hits not caused by the muon.

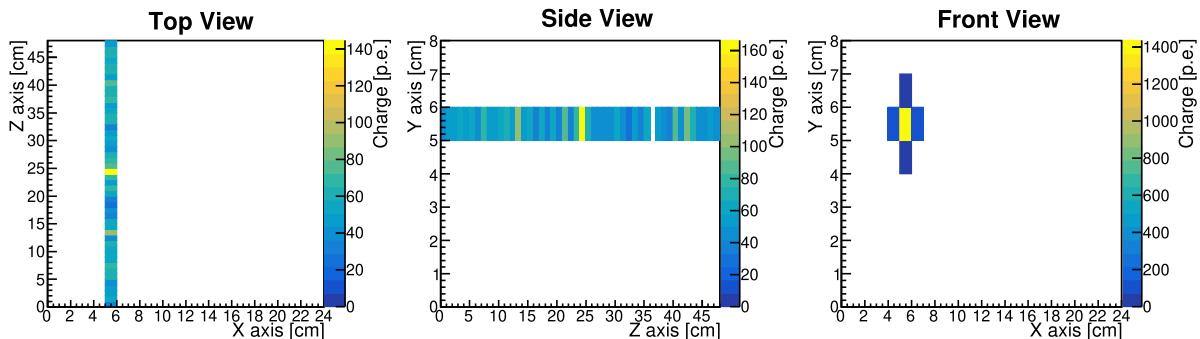


Figure 5.16: 2D displays of hits detected in a muon event at the SuperFGD Prototype. Hits from the three different planes of detection are show, with their colour representing the detected charge from the hit. A charge cut removing hits below 20 p.e. is applied to remove hits caused by light leaking between cubes and a narrow time window is applied to remove hits not caused by the muon.

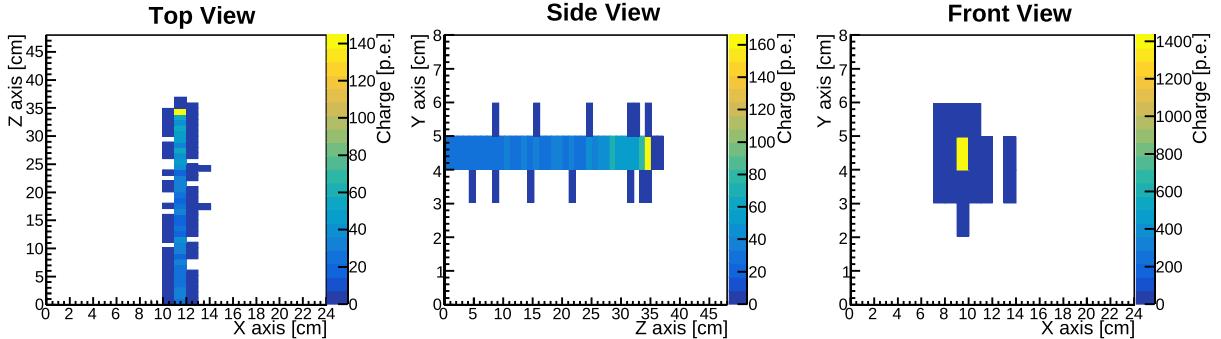


Figure 5.17: 2D displays of hits detected in a stopping proton event at the SuperFGD Prototype. Hits from the three different planes of detection are shown, with their colour representing the detected charge from the hit. No charge cuts are applied to the data. A narrow time window is applied to remove hits not caused by the proton.

a single track from this event, but we see that several hits are adjacent to the main trajectory of the muon, especially in the x - y and x - z views. These extraneous hits are of low amplitude and can be removed by placing a cut of 20 p.e. on the event, so that only hits with amplitude above 20 p.e. are kept. This is what is shown in the 2D displays of Fig. 5.16, where the muon’s trajectory has been isolated. These low charge hits seen in cubes the muon did not pass through are caused by light leaking between cubes. The properties and uses of this leaked light are explored in Sec. 5.7.5.

In Fig. 5.17, 2D displays of a proton stopping inside the SuperFGD Prototype are shown. This event was detected during a 800 MeV/c hadron beam run and was tagged with a proton particle trigger. No charge cut has been applied to the event, hence light leaking into neighbouring cubes is still visible. One interesting feature of these stopping proton events is the high charge deposit near the end of the track where the proton is halted. This high signal amplitude is useful for analysing the SuperFGD Prototype’s response to hits of high amplitude.

5.7.2 Hit Time Structure

Before we show any statistical analysis of the CERN particle beam data, we will show the structure of hits seen in the SuperFGD Prototype over time. This should give a general idea of how the detector responded to the beam and the time-frames in which events occur.

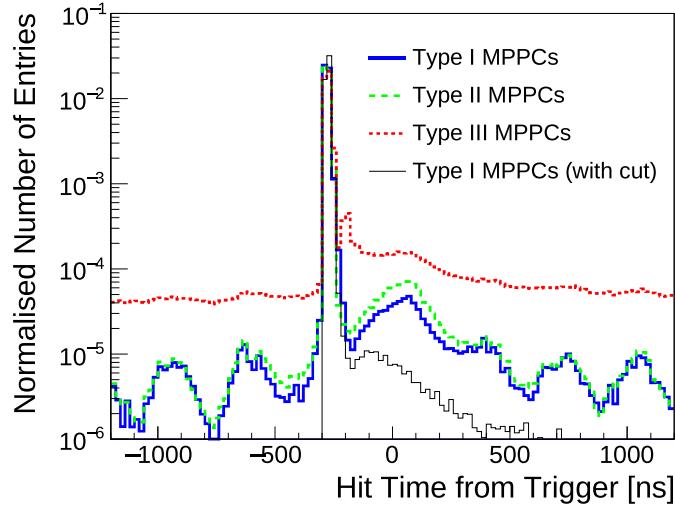


Figure 5.18: The distribution of hits over time with respect to the particle trigger time. Data for all three MPPC types are shown separately, along with Type I MPPC data with cuts applied to remove beam structure [76].

Fig. 5.18 shows the number of hits tagged with a muon trigger plotted against the difference in the hit time in the detector and the measured time of the associated trigger, labelled the “hit time from trigger”. This data was collected over a data run using a 2 GeV/c muon beam (i.e. the lead stopper was in the beam-line suppressing the hadronic component of the beam). The largest peak in the plot, which we can assume corresponds to the particles that caused the trigger, occurs at a “hit time from trigger” value around -250 ns. In theory, the hit time for the particle event should occur after the trigger and have a positive “hit time from trigger” value. However, due to signal travel time over cables, hits from particles would occur in the detector before the particle trigger signal reached the readout electronics and so the particle hits are recorded earlier than the particle trigger.

Other than the main peak seen at -250 ns corresponding to the main particle event, a repeating structure is seen in the Type I and II MPPC signals with peaks roughly every 350 ns. These extra peaks are caused by particles that did not trigger the particle trigger and so were tagged by a trigger signal caused by another particle. The repeating structure of these hits corresponds with the sub-structure in the PS’s proton delivery to the T9 beam-line. On larger time scales than those shown in Fig. 5.18, larger peaks can also be seen every 2 μ s that correspond to the 477 kHz PS revolution frequency.

The peaked structure is not seen for Type III MPPCs due to the relatively high dark

MPPC	Afterpulse/muon hit ratio	Afterpulse lag [ns]	Mean charge [p.e.]
Type I	0.0029	250	4.41 ± 1.42
Type II	0.0051	255	4.67 ± 1.36
Type III	0.1127	125	8.93 ± 2.60

Table 5.3: The properties of afterpulse hits for the different MPPC types.

count rate of those MPPCs drowning out the signal. However, a secondary peak about 100 ns after the main peak can be seen in the Type III MPPC hit time structure shown in Fig. 5.18. The Type I and II MPPCs also see a secondary peak but it is much broader and peaks later in time. To better visualise the secondary peak, cuts were applied to the Type I hits to remove the beam structure. The black line in Fig. 5.18 represents the Type I MPPC hits, but with the beam structure removed. The beam structure was removed by cutting all hits that appeared to originate from muons outside of the expected event time. More specifically, all hits with an amplitude above 20 p.e. with a hit time from trigger below -287.5 ns or above -225 ns were removed. The black line shows a clear “afterpulse” of hits after the main muon peak.

The exact cause of this afterpulse is not clear. The timescale of the afterpulse relative to the main muon event indicates it’s not due to light crossing between cubes and being detected later by neighbouring MPPCs (see Sec. 5.7.5). The travel and detection time of this optical crosstalk occurs at timescales of around 1 ns, much smaller than the 250 ns delay between the main muon hits and the afterpulse peak. The origin of the afterpulse is possibly hinted at by the differences in the afterpulse properties between MPPC type. These different properties are shown in Tab. 5.3. In the table we show the mean hit charge of the after pulse hits for each MPPC type, the time difference between the peak of the main muon event and the peak of the afterpulse (the afterpulse lag) and the ratio between afterpulse hits and hits from the main muon event. Afterpulse hits were taken as all hits with amplitudes smaller than 20 p.e. and main muon hits were hits above 20 p.e. with hit times from trigger between -287.5 ns and -225 ns. The properties are very similar for the Type I and II MPPCs, but Type III MPPCs have a ratio of afterpulse hits to muon hits 100 times higher, plus the afterpulse lag is 125 ns faster and the mean charge twice as large. This seems to indicate the afterpulse is an artefact of the electronics.

Investigations are still ongoing as to the exact origin of the afterpulse. For now, the

afterpulse hits are removed from analyses by applying a 20 p.e. cut to all hits as well as a time window corresponding to the event in question.

5.7.3 Hit Amplitude Thresholds

The minimum hit amplitude—i.e. the hit amplitude threshold—that can be detected and output by the readout electronics of the SuperFGD Prototype is dictated by the threshold discriminator of the CITIROC and the zero-suppression of the FPGA. The threshold discriminator value was universally set across all channels, but the FPGA threshold was set for each CITIROC. These values of the thresholds can be adjusted, but the resulting hit amplitude threshold is difficult to judge. We used the CERN particle beam data to estimate the hit amplitude threshold for every HG channel in the detector with the threshold settings used during the CERN particle beam test.

Estimates on the hit amplitude thresholds were made with two different methods. Method 1 was to find the lowest HG hit amplitude recorded for each data channel and use that as the hit amplitude threshold. Method 2 instead used the lowest HG hit amplitude for each channel for every spill in a data run. The average HG hit amplitude threshold per spill could then be estimated. The advantage of the latter method was that the change in hit amplitude threshold over time could be accounted for. With variations in temperature, we would expect the hit amplitude threshold to vary. The method does require, however, a significant number of hits to occur in a channel in order to calculate a good estimate of the mean. The former method does not have this drawback, but does not account for any variation in the hit amplitude threshold.

The results of the measurement of the hit amplitude threshold for every channel in the detector are shown in Fig. 5.19. The results of both methods are given. The MPPC type and the corresponding FEB number for each channel is also shown. The mapping of FEB numbers to the areas of the detector they readout is shown in Appendix A. The channel numbers can also be interpreted as positions on the detector surface by dividing the channel number by 32 to get the number of the cable bundle it is connected to, then taking the remainder and using the channel mapping in Appendix A to find the exact position of the channel.

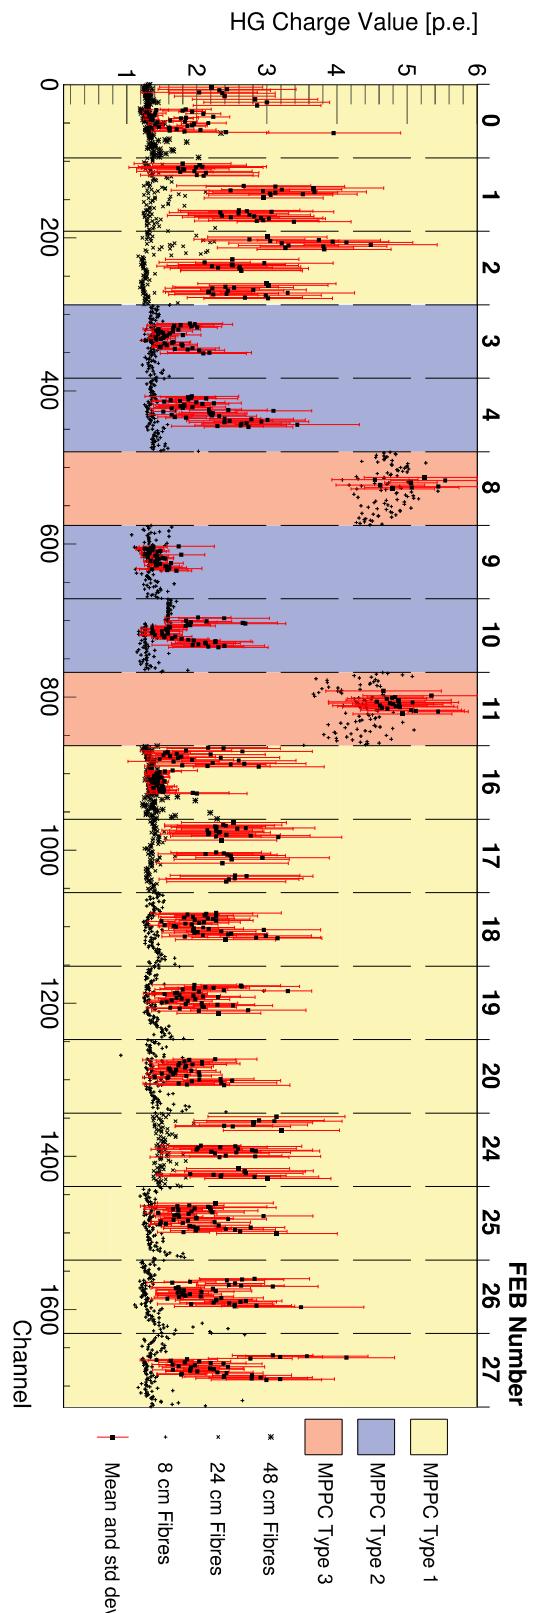


Figure 5.19: Measured values of the hit amplitude thresholds for each channel. Colours represent different MPPC types and the fibre length of the channel is represented by the marker style. Two different representations of the hit amplitude threshold are shown: the minimum HG value measured over the whole data run (Method 1 in the text) is shown for every channel by either a plus, cross or star marker, with no error bars. The mean value of the minimum HG values measured for each spill (Method 2 in the text) is represented by a solid black square for channels with sufficient statistics. The standard deviation of the mean value is shown by the red error bars. Data points with standard deviations above 1 p.e. are not shown.

We can see from Fig. 5.19 that the minimum hit amplitudes (the results of the first method) for Type I and II MPPCs are very similar, with values around 1.2 p.e. The Type III MPPCs have higher values around 4.8 p.e. This is not an unexpected result, as the CITIROC thresholds were set differently for each MPPC type. The thresholds for the MPPC Type III were set deliberately high to reduce dark count hits, given this type of MPPC had a crosstalk probability 10 times higher than the other MPPCs.

The mean value and standard deviation of the lowest hit charge per spill (the results of method two) differ significantly between channels. Such variations can be attributed to the different number of hits seen by each channel per spill. Channels that saw many hits—for example some channels in FEBs 0, 3, 9 and 16—have very small standard deviations and tend to have mean values around 1.4 p.e. Channels with lower statistics varied between 2–3 p.e. Channels that had a standard deviation over 1 p.e. were ignored, as it was assumed the number of hits seen in these channels was too small to calculate an accurate mean. The Type III MPPCs again gave different results to the other MPPCs, with average values around 5 p.e.

5.7.4 Channel Response

There are several physical differences between the different readout channels in the detector. There are three different types of MPPC reading off the channels and three different lengths of WLS fibre. Faults or inconsistencies within individual pieces of instrumentation could also cause some channels to respond unexpectedly. In this section, we will investigate the uniformity of channels across the detector, looking at the differences in signal caused by the different instrumentation of some channels. We will also look for individual channels that display strange behaviour compared to other channels. This is also a good opportunity to compare the different signal paths for each channel: HG, LG, ToT and PE.

The left plot of Fig. 5.20 shows the amplitude distribution for one readout channel with a Type I MPPC and a 24 cm WLS fibre. Distributions are shown for each signal path. The data used was collected from a single data run with 2 GeV/c muons. 385,925 events were recorded in the data run. Muon events are ideal for channel response studies as muons are minimally ionising and produce a consistent signal as they propagate through the prototype. For the data shown in the plot, a cut was applied so that a hit was ignored

MPPC	Fibre	HG [p.e.]	LG [p.e.]	ToT [p.e.]	PE [p.e.]
Type I	48 cm	78.45 ± 3.07	109.74 ± 5.33	104.10 ± 5.90	110.37 ± 5.22
Type I	24 cm	56.30 ± 4.29	56.14 ± 4.38	54.86 ± 4.42	55.88 ± 4.34
Type I	8 cm	62.30 ± 5.95	62.67 ± 6.00	61.50 ± 5.96	62.42 ± 5.92
Type II	8 cm	63.79 ± 5.35	64.21 ± 5.57	62.65 ± 5.63	63.94 ± 5.55
Type III	8 cm	53.97 ± 6.91	53.87 ± 6.81	52.26 ± 6.93	53.49 ± 6.96

Table 5.4: Mean and standard deviation of the light yields obtained from the different readout signal paths for the different fibre lengths and MPPC types instrumenting the SuperFGD Prototype.

if its PE signal was below 20 p.e. This cut was to remove the afterpulse hits as mentioned in Sec. 5.7.2. The cut should not affect the number of muon hits, which resulted in hit amplitudes of at least 50 p.e.

The distributions of the different signal paths seen in Fig. 5.20 match very well, with similar peak positions and variations. It should be noted that, unlike the other signal paths, the ToT signal is discrete in nature as it utilises timings measured in discrete units of 2.5 ns.

Shown in the right plot of Fig. 5.20 is the distribution of the signal amplitude means for each of the 384 channels with Type I MPPCs and 24 cm WLS fibres over the whole 2 GeV/c muon data run. The distribution of means for each signal path is shown. The same 20 p.e. cut to remove afterpulse hits is applied here. The plot shows the means for the different signal paths match well, indicating the calibration process for these channels was successful. The standard deviation in the means for each signal path is around 4 p.e. for the normalised distributions. This is an acceptable response considering no corrections for time-walk or light attenuation are being applied at this stage.

Figs. 5.21, 5.22, 5.23 and 5.24 are the same distributions as shown in Fig. 5.20 but are taken from readout channels with different combinations of MPPC type and fibre length. Tab. 5.4 shows the means and standard deviation from each plot. Comparing these figures shows how the instrumentation of the readout channels affects the channel response. Each figure shows that the different signal paths are calibrated well, with similar amplitude distributions to other signal paths with the same channel instrumentation.

Fig. 5.21 displays some interesting behaviour due to the fact that the 48 cm WLS fibres

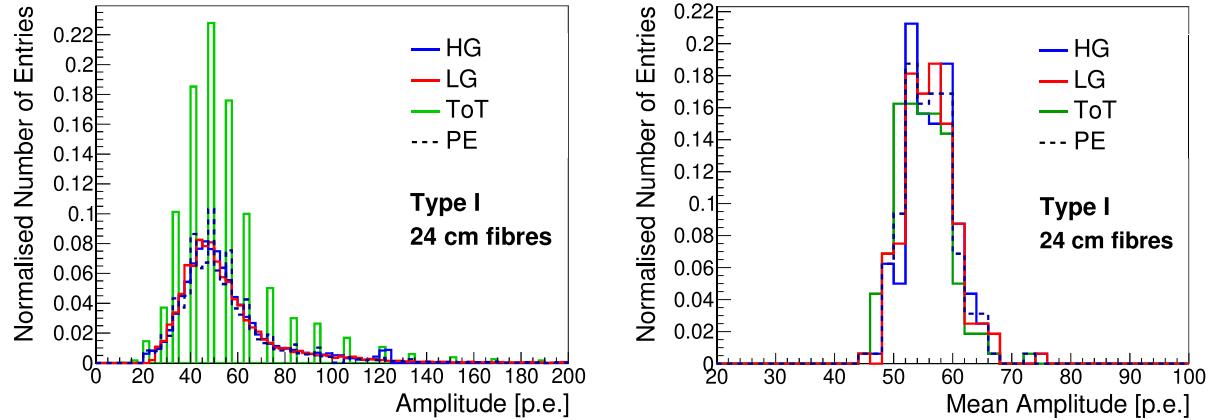


Figure 5.20: Left: Response of one readout channel to a muon particle run. All signal paths are shown. The readout channel uses a Type I MPPC and reads off a 24 cm fibre. Right: Mean signal amplitudes for the 384 readout channels connected to 24 cm fibres and Type I MPPCs. The data run is the same muon run used in the left plot.

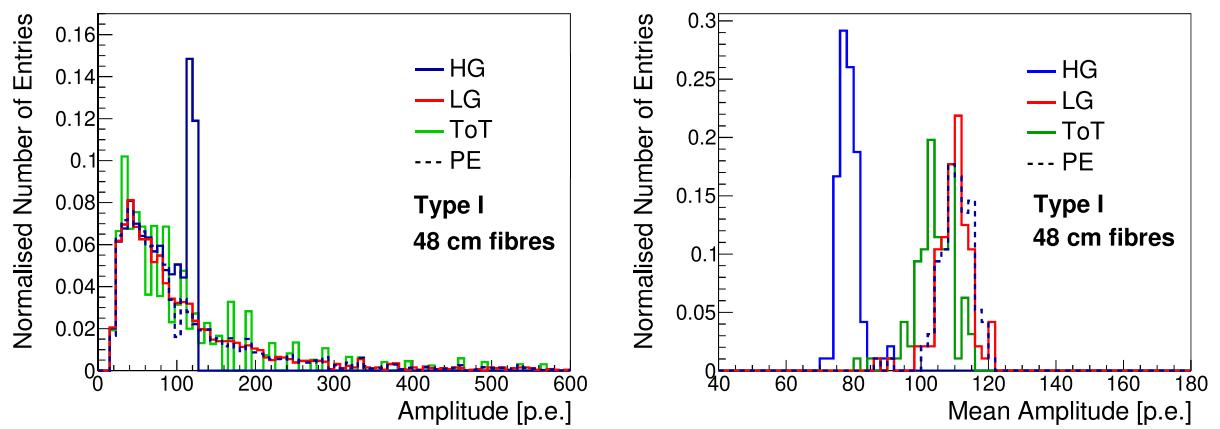


Figure 5.21: Left: Response of one readout channel to a muon particle run. All signal paths are shown. The readout channel uses a Type I MPPC and reads off a 48 cm fibre. Right: Mean signal amplitudes for the 192 readout channels connected to 48 cm fibres and Type I MPPCs. The data run is the same muon run used in the left plot. All signal paths are shown.

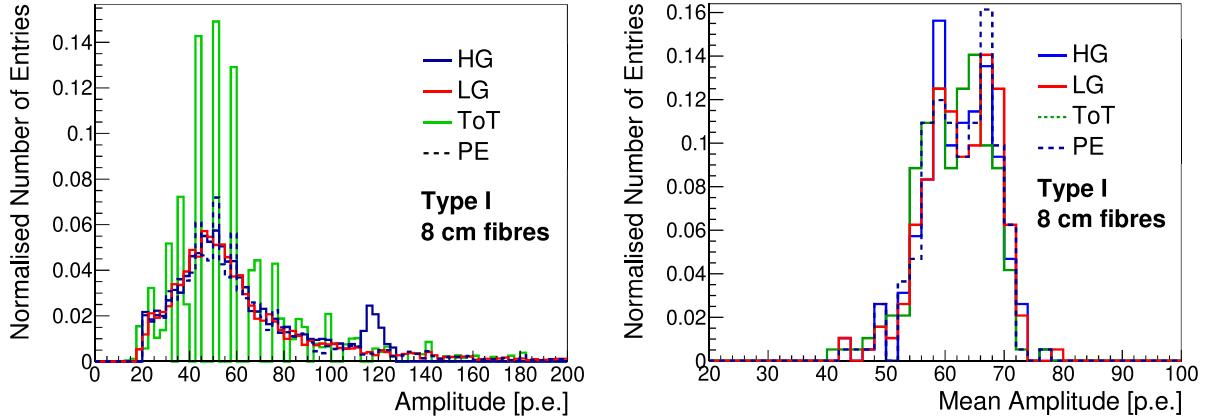


Figure 5.22: Left: Response of one readout channel to a muon particle run. All signal paths are shown. The readout channel uses a Type I MPPC and reads off a 8 cm fibre. Right: Mean signal amplitudes for the 576 readout channels connected to 8 cm fibres and Type I MPPCs. The data run is the same muon run used in the left plot. All signal paths are shown.

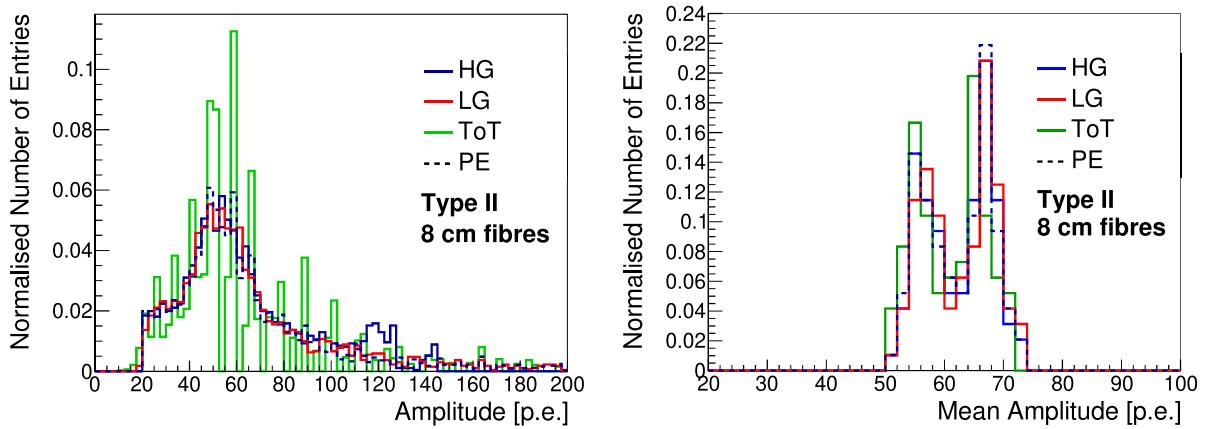


Figure 5.23: Left: Response of one readout channel to a muon particle run. All signal paths are shown. The readout channel uses a Type II MPPC and reads off a 8 cm fibre. Right: Mean signal amplitudes for the 384 readout channels connected to 8 cm fibres and Type II MPPCs. The data run is the same muon run used in the left plot. All signal paths are shown.

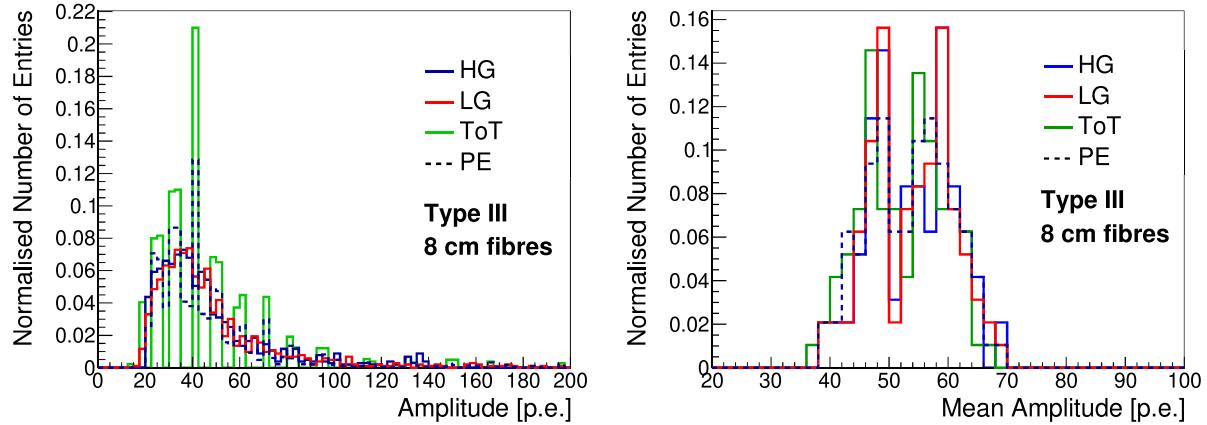


Figure 5.24: Left: Response of one readout channel to a muon particle run. All signal paths are shown. The readout channel uses a Type III MPPC and reads off a 8 cm fibre. Right: Mean signal amplitudes for the 192 readout channels connected to 8 cm fibres and Type III MPPCs. The data run is the same muon run used in the left plot. All signal paths are shown.

are uniquely parallel to the particle beam. As such, channels reading signals from 48 cm WLS fibres receive light accumulated from several cubes along the particle trajectory. As a result, we see a much different distribution compared to the other channels. The average amplitude and the spread in amplitude values is much larger. We can also clearly see the saturation peak of the HG signal path at around 120 p.e. The mean amplitudes for the HG signal path do not match up with the mean amplitudes of the other signal paths in the right plot in the figure due to the saturation of the HG signal.

There is some tension between the average channel amplitudes for 24 cm fibres and 8 cm fibres. This is due to light attenuation along the fibre (see Sec. 5.7.6). The muon beam was centred on the detector's x - y plane, so signals along the 24 cm fibres had to travel three times as far on average compared to the signals in the 8 cm fibres, leading to smaller amplitudes detected in the 24 cm fibres.

The figures showing hit amplitude distributions for 8 cm fibre channels are an excellent resource to investigate differences between the different MPPC types as the 8 cm fibres are the only ones with more than one MPPC reading out the signals. The main difference we see in the response from each MPPC type is that the mean hit amplitude is noticeably smaller for Type III MPPCs compared to Types I and II. This is possibly due to the smaller active surface of the Type III MPPCs, leading to more light loss between the

WLS fibre and the MPPC.

It is also interesting to note that the mean amplitudes for channels reading off 8 cm WLS fibres seem to have mean amplitude values distributed in a double peak shape, with the peaks separated by about 10 p.e. For the other fibre lengths, there was only one noticeable peak in the mean amplitude distribution. It is unlikely that the length of the fibre is the cause of this double peak, as we would expect shorter fibre lengths to result in more consistent amplitude values. The cause may actually be the Tyvek sheets placed between each x - z plane of cubes. The Tyvek sheets affect the leakage of light between cubes, which we investigate and explain fully in Sec. 5.7.5. The sheets make it much less likely for light to leak to neighbouring cubes in the y directions. As such, the WLS fibres parallel to the y -axis—the 8 cm fibres—will detect less light originating from a hit if there is light leakage in the x and z directions compared to if there is no light leakage. This is because the 8 cm fibre passing through the hit cube will not detect the leaked light. If there is no light leakage, then no light is lost to the neighbouring cubes, leading to more light entering the 8 cm fibre. This is the possible cause of the double peak distribution seen in the 8 cm fibres, the lower peak corresponding to light leakage occurring and the higher peak corresponding to no light leakage. We do not see the same pattern for the 24 cm fibres, which are parallel to the x axis, because these fibres will still pick up light that has leaked to neighbouring cubes along the x -axis.

5.7.5 Optical Crosstalk Between Adjacent Cubes

We have mentioned previously that light produced within a single cube can potentially pass from that cube into neighbouring cubes, despite the reflective materials etched onto the surface of each cube. From our investigations into channel responses, we've already seen how the Tyvek sheets can affect this optical crosstalk, as we will call it from here on.

This optical crosstalk can have an important impact on event and track reconstruction. Hence, we have investigated the properties of the optical crosstalk hits and their frequency.

Possibly one of the most pressing questions arising from optical crosstalk is how to measure the total amount of light emitted by a particle passing through a cube. Without optical crosstalk, we would expect the light to be distributed roughly evenly between the three orthogonal WLS fibres within the cube. Yet, with the phenomenon of optical

crosstalk, light is also collected by fibres passing through neighbouring cubes. The fibres in the cube the light originated from are also likely to pick up light that leaked through from neighbouring cubes. Recombining leaked light with the main light signal, as well as removing unwanted light signals, presents a difficult task.

In our study looking into optical crosstalk, we have attempted to model the distribution of light about the cube and how it is picked up by the surrounding WLS fibres. Fig. 5.25 shows the expected distribution of light produced by a particle travelling parallel to the z -axis about a single cube in the x - y plane. The cube the light originated in is shown in orange. The cubes the light leaks into are shown in blue. The WLS fibres, shown in green, are labelled with the light amplitude that they pick up. $M_{\text{main},i}$ is the light picked up by fibres parallel to the i -axis passing through the cube where the light originated (the “main” cube) and $M_{\text{xtalk},i}$ is the light picked up by fibres parallel to the i -axis passing through the neighbouring cubes containing leaked light (the “crosstalk” cubes).

In our studies of optical crosstalk, we used hadron beam data with either a proton or μ/π trigger to accumulate samples of proton events and μ/π events respectively. Cuts to the data were applied such that only events where the particle travelled perfectly parallel to the detector were used. The cuts made it simple to separate main hits and crosstalk hits as the main hits would constitute a straight track through the detector, all other hits would be crosstalk hits. Additional cuts were applied to the proton data sample so that only events where protons stopped in the SuperFGD Prototype detector remained in the sample.

Fig. 5.26 shows the hit amplitude distributions for both main hits and crosstalk hits. The top two plots show main hit amplitudes ($M_{\text{main},i}$) and the bottom two show crosstalk hit amplitudes ($M_{\text{xtalk},i}$). The left plots contain data from WLS fibres parallel to the y -axis and the right plots contain x -fibre data. Each plot separates the data depending on if the event was a μ/π event or a stopping proton event. For the stopping proton data, hits were sampled depending on the distance d of the hit from the proton stopping point.

The top two plots of Fig. 5.26 fit well with the event display of a stopping proton event shown in Fig. 5.17. Main hits detected at the stopping point of the proton have, on average, a larger light yield than hits detected upstream of the stopping point. The further from the stopping point the hits were, the lower the light yield. μ/π hits give the lowest light yields of around 50 p.e., which is consistent with the amplitude distributions

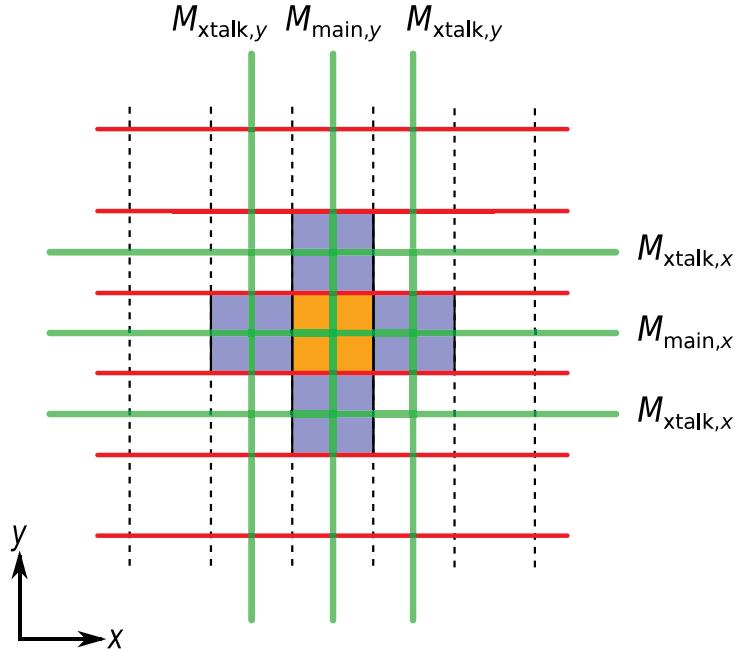


Figure 5.25: A diagram showing the distribution of light about the cube from which the light originated. The light was produced by a particle travelling into the page (along the z -axis). The WLS fibres that carry the light signals to the MPPCs are also shown and are labelled according to the expected light amplitude within the WLS fibres. The cube in orange is where the light signal originated and the blue cubes are where we expect optical crosstalk light to appear. The WLS fibres are shown in green and Tyvek sheets are shown in red. Only cubes in the x - y plane are shown, but we would also expect light to leak to and from cubes along the z -axis.

shown in Sec. 5.7.4. The distributions of main hit light yields differ slightly between x - and y -fibres. The distributions peak at main hit light yields around 10% higher for x -fibres relative to y -fibres. This lines up with the theory outlined in Sec. 5.7.4 that the x -fibres pick up more light due to the optical crosstalk being more likely to occur between cubes along the x -axis as opposed to cubes along the y -axis due to the Tyvek sheets.

The lower two plots of Fig. 5.26 showing the distributions of crosstalk hits give some insight as to the effectiveness of the Tyvek sheets in preventing light leakage between cubes. We see for the $d = 0$ cm and $d = 1$ cm samples that the crosstalk hit light yields are significantly higher on average in the y -fibres. The y -fibres measure light that has leaked into neighbouring cubes along the x -axis, so this indicates that the Tyvek sheets block a significant portion of the light leaking to cubes along the y -axis, which is measured by the x -fibres. For the $d = 24$ cm sample, there is little difference between the crosstalk

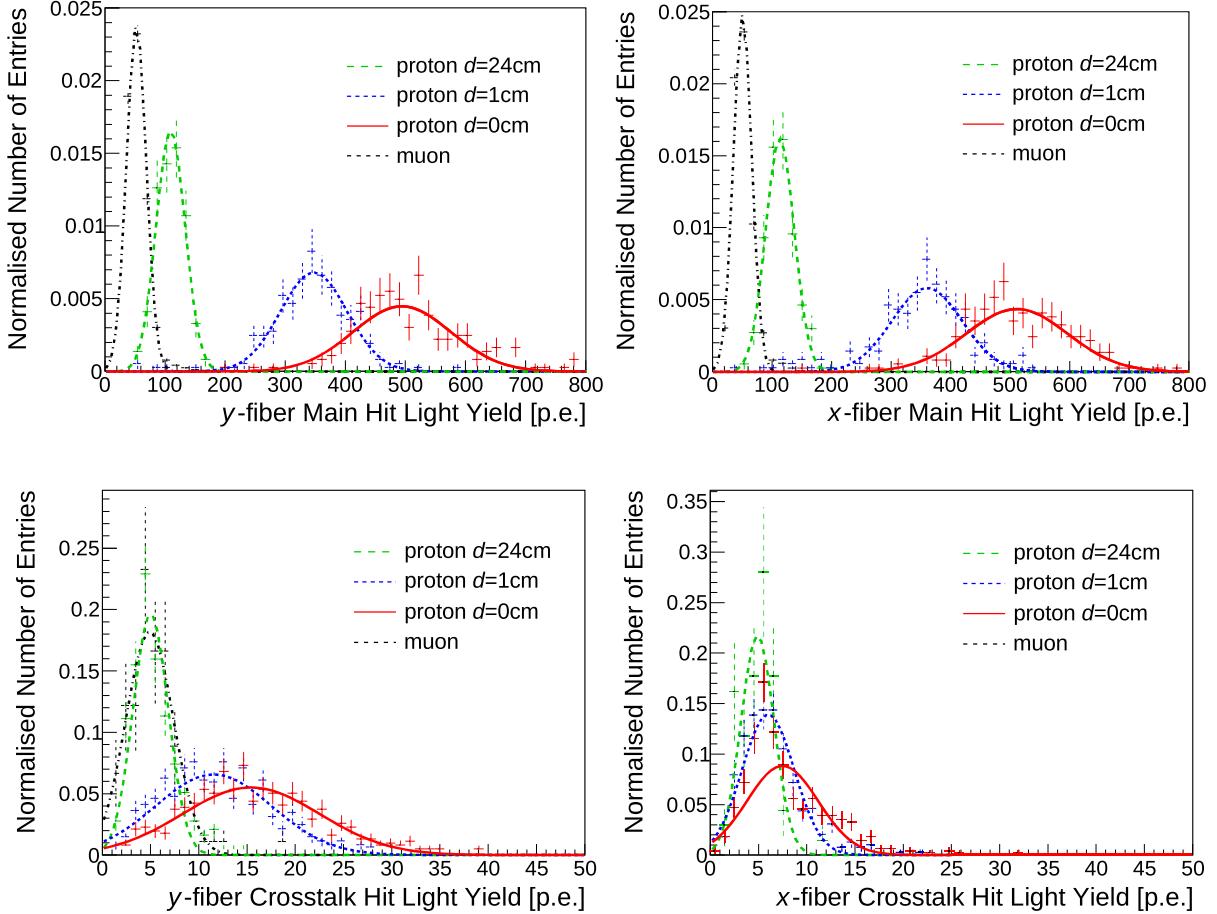


Figure 5.26: Light yield measurements from main hits and crosstalk hits using fibres perpendicular to the particle's direction of travel. Distributions from both the proton and μ/π triggers are shown. The selected protons stopped within the prototype volume, leaving deposits of different energies along their tracks. The light yield for these protons is measured at three different distances d from the stopping point: at $d = 0$ cm, $d = 1$ cm and $d = 24$ cm. The bottom right plot does not include a μ/π crosstalk p.d.f. since very few crosstalk hits were recorded for this sample.

hit amplitude distributions, which could be because the crosstalk amplitudes are so close to the hit amplitude threshold.

To quantitatively determine the level of light leakage for each light deposit within a cube, we used the equation

$$\kappa_i = \frac{M_{\text{xtalk},i}}{M_{\text{total},i}} = \frac{M_{\text{xtalk},i}}{M_{\text{main},i} + 2M_{\text{xtalk},i}}, \quad (5.1)$$

where κ_i is the ratio $M_{\text{xtalk},i}$ to the total amount of deposited light within a cube that was detected in fibres parallel to the i -axis ($M_{\text{total},i}$). In Eq. 5.1, we have calculated $M_{\text{total},i}$ by summing up signals from several fibres. By summing signals, we have accounted for light leakage into each neighbouring cube and the extra light detected in the main cube leaking from neighbouring cubes on the z -axis. This extra light originates from other cubes the particle passed through on the z -axis. We assume that the crosstalk signals are symmetric about the main cube and that the crosstalk signals between cubes along the x -axis are equal to crosstalk signals between cubes along the z -axis. Hence, we would expect $M_{\text{main},i}$ to contain signals corresponding to the light remaining in the main cube from the initial light deposit, the leaked light from neighbouring z -cubes—equal to $2M_{\text{xtalk},x}$ —and the leaked light present in neighbouring cubes along the i -axis—equal to $2M_{\text{xtalk},j}$, where j is the opposite axis to i (e.g. if $i = y$ then $j = x$). Because $M_{\text{main},i}$ already contains light signals equivalent to the light signals that leaked into neighbouring j - and z -cubes, all we need to add to equate to $M_{\text{total},i}$ is $2M_{\text{xtalk},i}$, which accounts for the leaked light along the j -axis that wasn't picked up by fibres passing through the main cube. These summations result in Eq. 5.1. This equation also accounts for the different amplitudes of crosstalk signals seen between cubes along the x -axis and cubes along the y -axis due to the Tyvek sheets.

Notice that when comparing crosstalk amplitudes to total light deposit, we only use light signals detected from a single plane and do not sum light signals from different planes in order to calculate the total light deposit from the particle scintillation. We do this assuming that light within a cube is shared equally between the three orthogonal fibres within the cube, hence the ratios of crosstalk signals to total charge deposit should be the same no matter if the charges are summed or taken from one plane.

As seen in Fig. 5.26 with the low crosstalk light yields for x -fibres, we must be careful when measuring κ_x as it may be influenced by the hit amplitude thresholds of the electronics discussed in Sec. 5.7.3. To avoid crosstalk signals close to the hit amplitude thresholds, we can instead use signal measurements from the z -fibres to estimate the signal percentage leaking into neighbouring y -cubes. z -fibres are parallel to the particle trajectory and hence pass through many cubes where light was produced, yielding much higher light amplitudes for both crosstalk and main hits. Fig. 5.27 shows the crosstalk and main hit light yields as measured along z -fibres. It is plain to see the hit amplitudes are much larger compared to x - and y -fibres. The plots displaying z -fibre crosstalk data

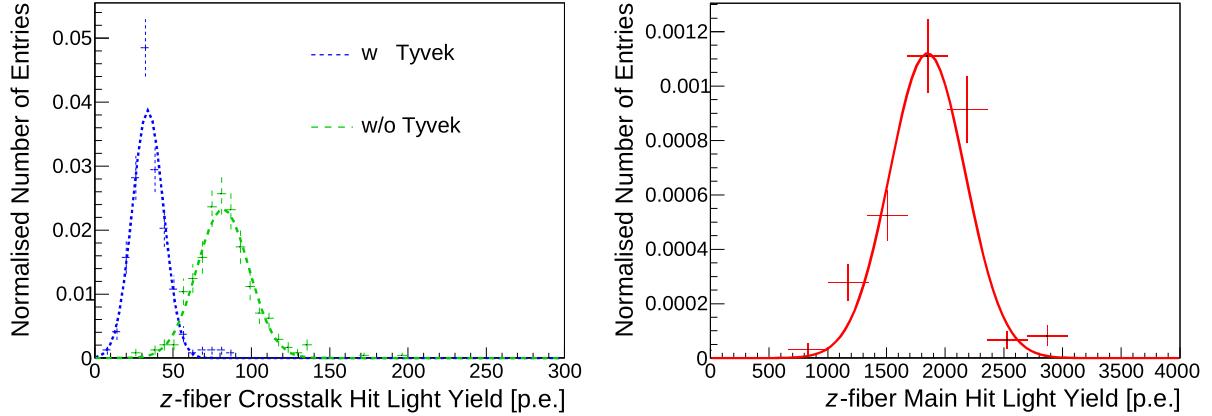


Figure 5.27: Light yield amplitudes for crosstalk (left) and main (right) hits detected with z -fibres with 2.0 GeV/c muon data and the μ/π trigger. The crosstalk hits have been separated into samples of hits that were affected by Tyvek (blue dotted) and hits that weren't affected by Tyvek (green dashed).

show distributions for crosstalk hits unaffected by Tyvek (i.e. from neighbouring x -cubes to the main cube) and crosstalk hits that are affected by Tyvek (i.e. from neighbouring y -cubes to the main cube). This is another clear demonstration of the lower crosstalk hit amplitudes seen in cubes shielded by Tyvek.

Fig. 5.28 shows both the κ_y measurements (left plot) and κ measurements taken using the z -fibres (right plot). The κ_y plot was compiled using data from a stopping proton sample using only hits at the end point of the proton track. This sampling ensured the crosstalk hits would be large in amplitude and unaffected by hit amplitude thresholds. The plot itself shows the distributions of κ_y values taken for each main cube hit in the data run. The average value of κ_y is calculated as $2.94 \pm 0.05\%$. In other words, nearly 3% of light produced in a cube will leak into each neighbouring cube without Tyvek protection, making 12% in total. The right plot of Fig. 5.28 showing the z -fibre measurements splits the data into crosstalk hits taken from y -cubes neighbouring the main cube (“w Tyvek”) and crosstalk hits taken from x cubes neighbouring the main cube (“w/o Tyvek”). The fact that the average κ value of data taken from the neighbouring x -cubes ($4.11 \pm 0.06\%$) does not match the average κ_y value shows this method of using z -fibres may be flawed. This is possibly because the method assumes crosstalk signals are produced for every main hit, but looking at the event displays in Fig. 5.17 this is not true in the x - z or the z - y planes. This would result in a lower κ value when using z -fibres, as we observe. However, this is the best evaluation of the κ_x value we have while avoiding hit amplitude thresholds.

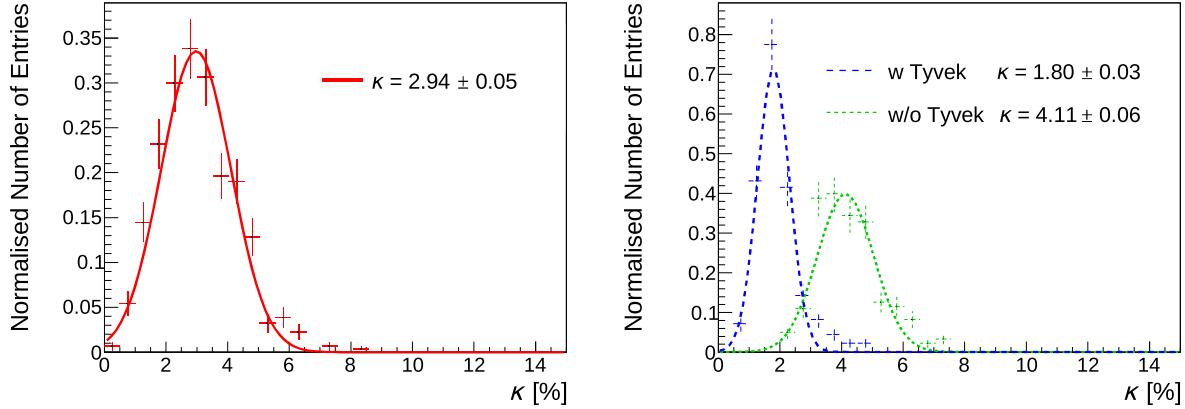


Figure 5.28: Measurement of the percentage fraction of light κ_y that flows from main cubes to neighbouring x -cubes and then enters WLS fibres parallel to the y -axis. The left plot corresponds to the best estimate of $\kappa_y = 2.94 \pm 0.05$ made using the stopping point of proton events. The right plot shows similar measurements but using z -fibre signals, which allow measurements of crosstalk light levels in cubes shielded by Tyvek sheets.

The value of κ_x we measure is equal to $1.80 \pm 0.03\%$. In future, this measurement could be improved by accounting for main hits that do not produce crosstalk hits by weighting κ measurements by the ratio of main hits to crosstalk hits.

5.7.6 Light Attenuation in WLS Fibre

There is some light loss as photons propagate along the WLS fibres in the SuperFGD Prototype. As such, the light amplitudes detected by the MPPCs are dependent on how far the signal had to travel along the WLS fibre to reach the photosensor. We can use the attenuation equation

$$y(d) = Y_0 \left(\alpha e^{\frac{-d}{L_S}} + (1 - \alpha) e^{\frac{-d}{L_L}} \right), \quad (5.2)$$

to estimate the light attenuation along a fibre. $y(d)$ is the amplitude of attenuated light after travelling a distance d along the fibre, Y_0 is the unattenuated light yield, α is a weighting factor and L_S and L_L are respectively short and long attenuation constants [63]. L_L is known to be 4 m from the manufacturer specifications [71]. The other constants are dependent on WLS fibre length and had to be measured.

The measurements of the light attenuation in the SuperFGD Prototype's 24 cm (x)

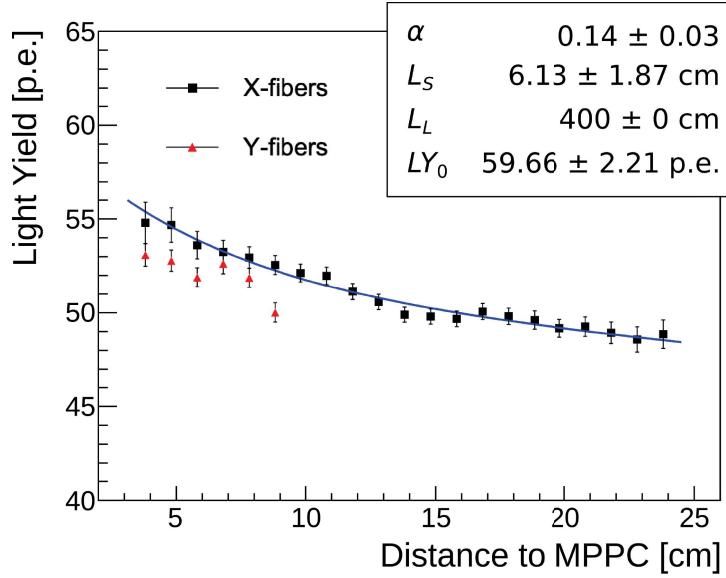


Figure 5.29: Light yields as a function of signal propagation distance in x -fibres (black squares) and y -fibres (red triangles). The x -fibre data is fit using Eq. 5.2 (blue line). The L_L parameter is fixed to 4 m according to manufacturer specifications [76].

and 8 cm (y) long fibres are presented in Fig. 5.29. A 2.0 GeV/c muon data run with the μ/π trigger was used. The data in Fig. 5.29 are obtained by selecting straight tracks parallel to z -fibres so that the x and y coordinates are known for all hits. The light yields are then obtained by averaging the data of all channels in the x - z (y - z) plane for each specific y -distance (x -distance). y -fibres show a slightly smaller light yield (3%) than that of x -fibres while showing a similar attenuation trend. As discussed in Sec. 5.7.4, this difference in light yields is likely due to the Tyvek shielding between cubes along the y -axis. Eq. 5.2 was used to fit the results for the x -fibres. The fit results show the best parameters to be $Y_0 = 59.66 \pm 2.21$ p.e., $L_S = 6.31 \pm 1.87$ cm and $\alpha = 0.14 \pm 0.03$.

A fit to the y -fibre data is difficult due to the low number of data points. For the data analysis, we assumed the y -fibres had equal attenuation constants to the x -fibres, which is not unlikely given their similar attenuation slopes in Fig. 5.29. In future, it would be valuable to take dedicated measurements of the attenuation constants of the WLS fibre lengths to be used in the final SuperFGD design, so that attenuation can be accounted for with more accuracy.

5.7.7 Cube Response

A study on the CERN beam test data was performed to check the individual response of cubes. Potential failures in the WLS fibre and SiPM or presence of impurities in the scintillator could introduce a source of non-uniformity across cubes. Hence, checking for inconsistent signals is an excellent screening for faulty hardware. Moreover, analysis tools used for this study are the same as those for the determination of hit position, the first step towards event reconstruction.

For this study, the cube from which a signal originated had to be identified. Studying the hit time of an individual channel in one projection is not sufficient to provide information about the cube position along the WLS fibre. There is a spread in photon arrival time at the photosensor due mostly to processes local to the energy deposition site such as the scintillation process in the cube and the wavelength conversion process in the WLS fibre. There is also a time delay due to light propagation down the WLS fibre. These time-dependent processes, along with the processing of the photosensor's electrical signal by the CITIROC preamplifier and shaper, exclude the possibility of resolving the hit position along the fibre based on the hit profile of a single channel with a resolution better than 10 cm. It is therefore impossible to resolve the main cube hit from adjacent cubes that might collect a fraction of photons from optical crosstalk. That type of distinction can only be made by studying the coincidence of hits in two or more projections.

0.8 GeV/ c muons were used in the study. Straight tracks were selected to minimize differences in track length through each cube. The average light yields for a number of cubes with the SuperFGD Prototype are shown in Fig. 5.30.

After attenuation corrections, average light value for both the horizontal and vertical fibres were 58 p.e. and 59 p.e. respectively, with standard deviations of 7 p.e and 8 p.e. No corrections were made to account for the crosstalk picked up by the fibres or the effect of the Tyvek sheets between horizontal cube layers. As such, we would expect the average value of the corrected light yields for cube signals read out by vertical fibres to be slightly lower than those read by horizontal fibres, as the Tyvek sheets will block crosstalk passing into other cubes along the vertical fibre. We do not see this with our fitted means, though there are some overly populated bins for the vertical fibre histograms between 50–60 p.e., which are not accounted for in the Gaussian fit.

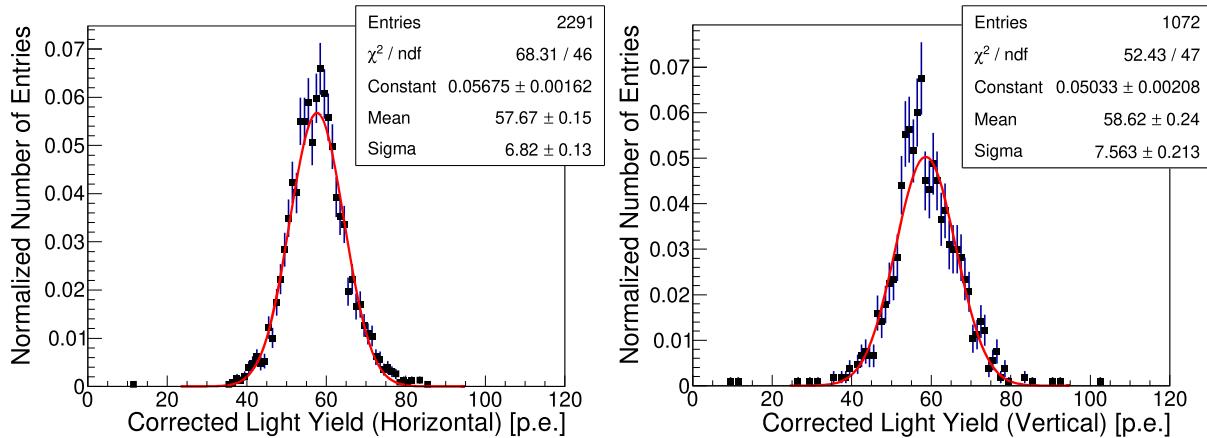


Figure 5.30: Left: cube light yields for 2,291 cubes read out by horizontal fibres. Right: cube light yields for 1,072 cubes read out by vertical fibres. Signal attenuation is corrected for. The mean light yields along vertical and horizontal fibres are within 1 p.e. after corrections but are separated by more than one standard deviation.

5.7.8 Signal Travel-Time and Time-Walk Corrections

Up to now there has been little discussion of the timing measurements gained from the CERN beam test. This will be rectified in this section, where we look at how to correct timing measurements to account for time-walk and signal travel time. These corrections will be utilised in the time resolutions analysis of the SuperFGD Prototype, where accurate and consistent timing measurements are key.

Signal travel time is simply the time taken for light to travel along the WLS fibre to the MPPC at the end of the fibre. This will differ depending on how far the signal had to travel along the fibre. The core of the WLS fibres is composed of polystyrene with a refractive index of 1.59. The speed of light in the fibre will thus be around $1.89 \times 10^8 \text{ ms}^{-1}$.

The light speed in the WLS fibres is known, so all we must do is work out how far along the WLS fibre each signal has travelled before it reached the MPPC. This was relatively simple in the CERN beam test where most particle tracks were parallel to the z -axis of the detector. Hence the x coordinate of a hit in the z - y plane can easily be inferred by looking for this single hit in the x - z plane with the same z value. A similar method can be utilised to find the y coordinate of hits in the x - z plane. Then, after accounting for which side of the detector the MPPC is on, one can calculate the distance of the hit to

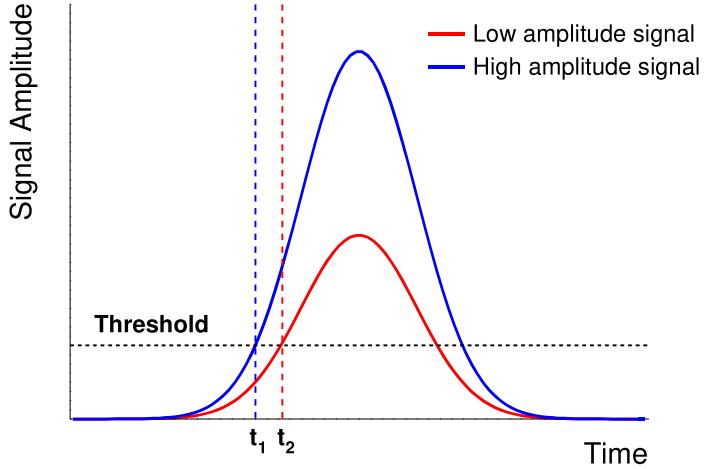


Figure 5.31: A diagram showing how signal amplitude can affect the rising edge timing. Both signals in the diagram have the same width and mean, but differ in amplitude. The higher amplitude signal crosses the discrimination threshold first at time t_1 , despite both signals beginning at the same point. This phenomenon is known as time-walk.

the corresponding MPPC by finding the distance from the hit to the side of the detector the MPPC is on. For all hits, we took the distance from the centre of the cube the hit originated in to the side of the detector. We also added on an extra 2.3 mm to account for the separation of the MPPC from the side of the detector.

Time-walk is the phenomenon whereby the amplitude of a signal can affect its timing measurement. For the SuperFGD Prototype, timestamps are created when a signal rises above the discrimination threshold (rising edge) and when it drops below (falling edge). It is the rising edge time we use to signify the hit time of a signal. The amplitude of a signal dictates how quickly it rises above the discrimination threshold, with larger amplitudes rising faster (see Fig. 5.31).

To correct for time-walk, we model amplitude-time distributions of a reference data set and apply the model to other data sets. This allows us to account for the time taken for the signal to reach the discriminator threshold. The model used was a combination of a constant, an exponential function and an inverse function:

$$t = a - be^{-cx} - \frac{d}{x + e}, \quad (5.3)$$

where a , b , c , d and e are constants, t is the hit time and x is the signal amplitude. The

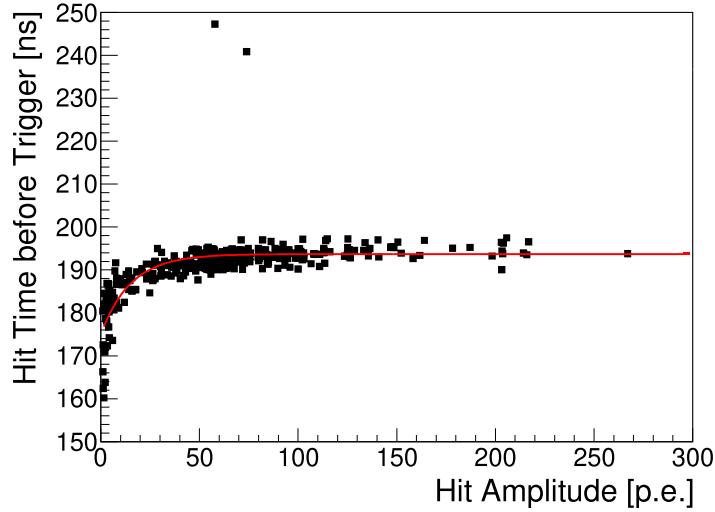


Figure 5.32: The amplitude-time distribution of a single readout channel. The channel is from FEB 19, channel number 69, with MPPC Type I. The red line is a fit to the time-walk model. Hit times are relative to the trigger time.

reference data set was with a 0.8 GeV/c hadron beam data run with an electron trigger.

Fits using Eq. 5.3 were performed for every channel on the SuperFGD Prototype, so that each channel had unique values of a , b , c , d and e to account for any differences in channel response. An example of the amplitude-time distribution for a single channel, with a fit following Eq. 5.3, is shown in Fig. 5.32. The hit times are taken relative to the trigger time for the event. The fit matches the distribution of points well. There is some width to the distribution because of the distribution in hit times over an event.

To summarise, to correct rising edge timing measurements in the data for signal travel time and time-walk, the equation

$$t_{\text{corr}} = t_{\text{meas}} - \frac{D}{v} + be^{-cx} + \frac{d}{x + e} \quad (5.4)$$

was used, where t_{corr} is the corrected rising edge time, t_{meas} is the measured rising edge time, D is the distance the signal travelled along the WLS fibre, v is the speed of the signal in the fibre and b , c , d and e are constants from the fit of Eq. 5.3.

It became apparent after fitting each channel that the fits for Type III MPPCs were not matching the data. For these Type III MPPCs, the exponential component was removed

from the fit, which resulted in a much better match to the data.

5.7.9 Time Resolution

The main contributions towards the time resolution of the SuperFGD Prototype’s channels are the timing uncertainties in the scintillation and wavelength shifting processes and the readout electronics response. Previous unpublished measurements of the channel time resolutions of a small $5 \times 5 \times 5$ cubed scintillator detector prototype with a 5 GHz sampler gave an average value of 0.95 ns. Due to the fast sampler used and the absence of the SuperFGD readout electronics, this value is taken to represent the contribution to the time resolution from scintillation and wavelength shifting. The readout electronics are predicted to contribute a 0.7 ns uncertainty to the timing measurement [76]. Adding the two contributions in quadrature, we would expect the SuperFGD Prototype to have an average channel time resolution of roughly 1.2 ns.

For the CERN beam test analysis, the SuperFGD Prototype’s time resolution was measured by comparing hit times to a consistent timing within an event. This was done for each readout channel, building up a distribution of hit time differences relative to the consistent time reference. The variation in the hit time difference was used to calculate the channel’s time resolution.

A typical time reference to use when measuring a detector’s channel time resolution is the trigger time for the event. However, the CERN beam test’s trigger showed significant uncertainty relative to the expected time resolution of the prototype. The uncertainty originated from the long lengths of cabling connecting the trigger system to the readout electronics. As such, instead of using the event trigger, the time reference used for each event was the hit time in the 24th z -layer. This assumed straight tracks through the detector for each event and consistent particle speeds between each event. A sample of 2 GeV muon events travelling parallel to the z -axis was used for the time resolution analysis. Given the low ionisation rate of muons, the assumption of consistent speeds between each event is justified.

The method to measure each channel’s time resolution was as follows:

- the hit times within selected events were corrected for time-walk and WLS fibre

travel time;

- a dataset of $(T_{\text{ref}} - T_{\text{hit}})$, where T_{ref} is the reference time and T_{hit} is the hit time measured in the channel being analysed, was collected for each channel using all the selected events;
- a Gaussian distribution was fitted to the dataset if there were more than 250 entries;
- the combined resolution of the channel and the time reference was taken as the standard deviation of the fitted Gaussian distribution;
- considering the time reference was also a simple channel hit time, it was assumed the reference and the channel had the same time resolution, so the standard deviation of the Gaussian was divided by $\sqrt{2}$ to deconvolve the two resolutions and extract the channel time resolution.

Using the above method, time resolution values for 854 of 1,728 channels were measured. Not all channels were able to have channel time resolution measurements as the muon beam was centred on the detector, meaning channels near the edge of the detector saw very few if any particle hits. The channel time resolution results are shown mapped to their detector coordinates in Fig. 5.33. The channel map shows that there is some structure to the x - z map between $25 < z < 32$, with the rest of the values appearing to be consistent across the detector. The structured region corresponds to channels readout by FEB 18. The time resolutions of FEB 18 channels have time resolutions larger than the average by about 0.2 ns. The cause of this is as yet unknown, as data from FEB 18 looks otherwise normal. The most likely cause is a different response of FEB 18 to the clock and synchronization signals sent by the Master Clock Board, although the cause of this difference in time resolution measurement is still under investigation. The channels from FEB 18 are ignored for the remainder of this section.

The distribution of the time resolution values for each channel is shown in Fig. 5.34. There are three sub-plots in the figure with different segregations of channels to demonstrate differences seen due to channel properties. Fig. 5.34a shows the distribution for all the measured channels, except FEB 18 channels. The distribution has a mean value of 1.14 ns and a standard deviation of 0.06 ns. This shows the average channel time resolution is within one standard deviation of the expected value of 1.2 ns.

Figs. 5.34b and 5.34c show channel distributions for channels with different WLS fibre

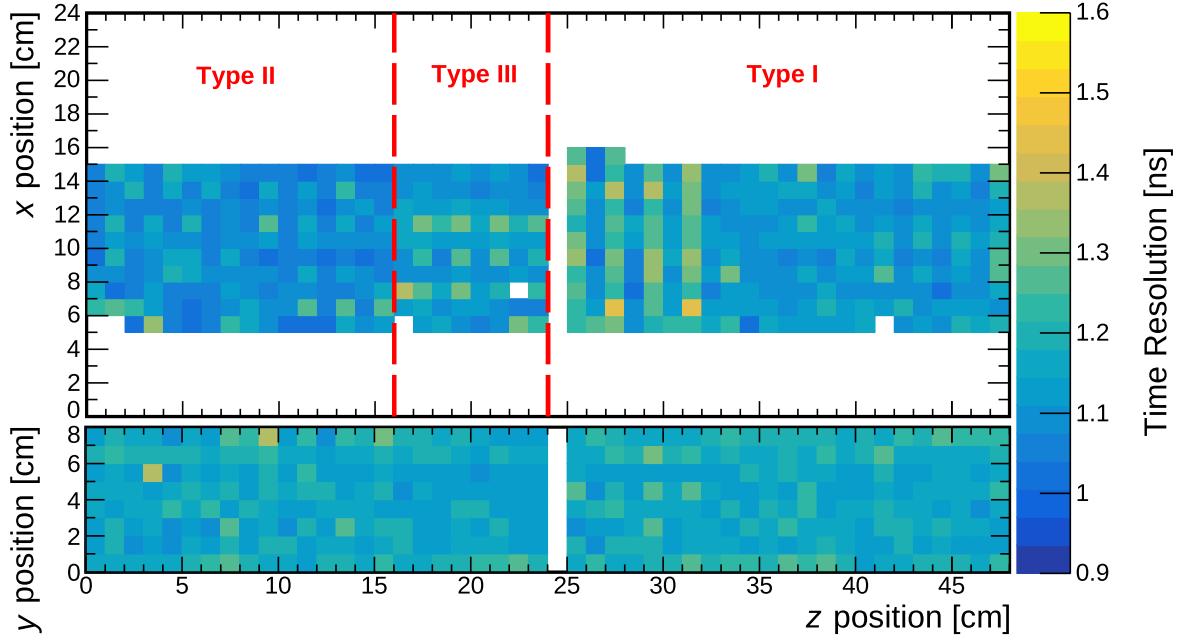


Figure 5.33: A map of readout channels and their measured time resolution values. Colour represents the channel time resolution value. The upper plot shows the x - z plane and the lower plot shows the z - y plane. Dashed red lines indicate the different MPPC types. The z - y plane channels all use Type I MPPCs.

lengths and different MPPC types respectively. Fig. 5.34b displays histograms of the time resolutions for channels that use Type I MPPCs with 8 cm fibres and 24 cm fibres. There is a clear offset between the two distributions, despite fibre length being corrected for in the analysis. A possible cause of this observed offset are the Tyvek layers between horizontal cube layers.

The three different MPPC types are isolated in Fig. 5.34c. Only channels with 8 cm fibres are included to remove any effects from fibre length. The peaks of the three distributions are in similar positions, with the peak of the Type II MPPCs' about 0.04 ns lower on the time resolution axis compared to the other MPPCs' peaks. There is also a small tail for each MPPC distribution at larger time resolution values. This may be an artifact of the method used to measure the time resolution values, as the fitted Gaussian to the channel time difference distribution was not always accurate. These distributions could potentially be improved in their representation of the actual time resolution by cutting out channels with poor Gaussian fits to the distribution.

The average time resolution measured here and the previous analyses shown resulting

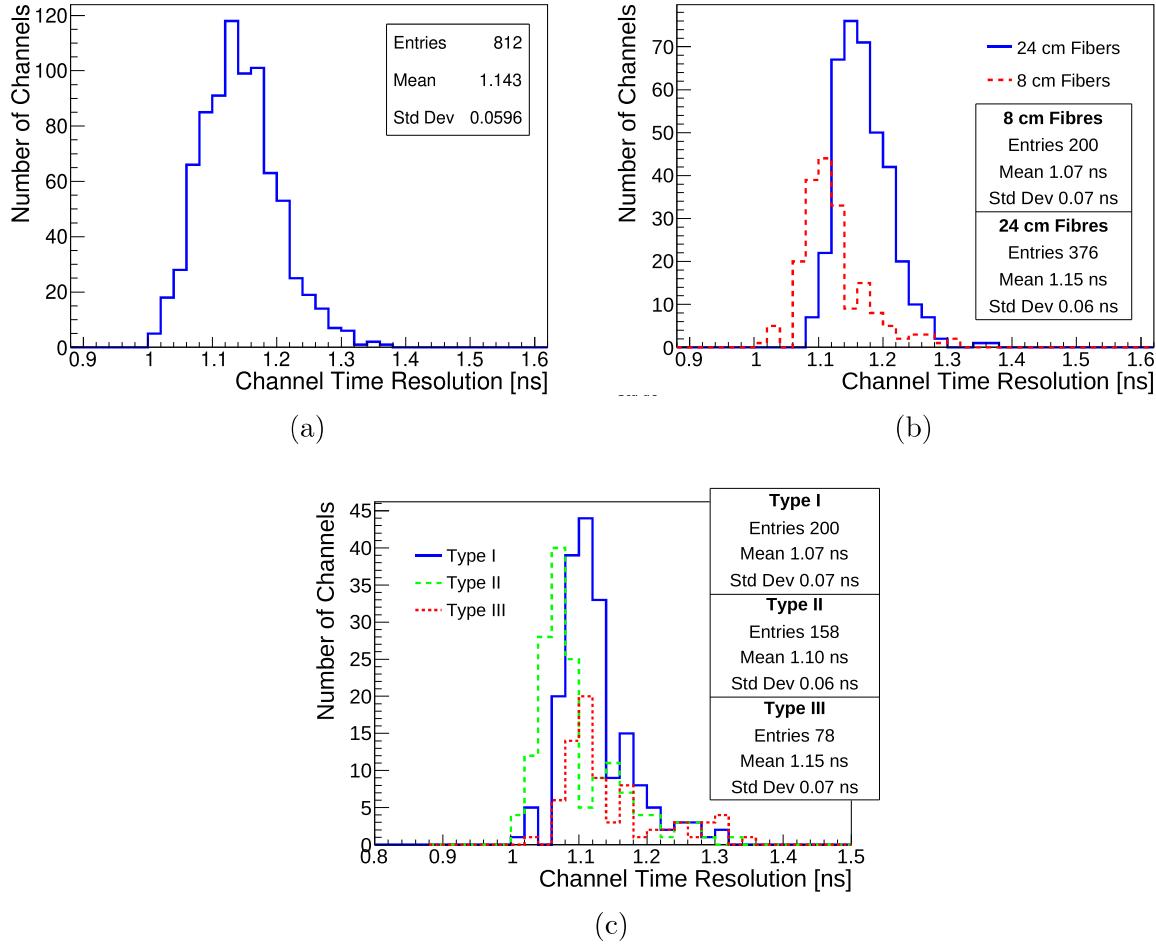


Figure 5.34: The distribution of the measured channel time resolutions. Channels read out by FEB 18 are excluded. (a) shows all the channels with a measured time resolution. (b) displays two distributions, one corresponding to channels with 8 cm long fibres (red dashed) and the other to channels with 24 cm long fibres (blue solid). Only channels with Type I MPPCs are shown. (c) shows three distributions corresponding to channel MPPC type. Type I (blue solid), Type II (green dashed) and Type III (red dotted) MPPCs are all shown, but only channels with 8 cm fibres are included.

from the CERN beam test have been invaluable in the simulation work being performed for the upcoming SuperFGD to be installed in ND280. More questions still remain regarding the SuperFGD's response to low energy nucleons, which is examined using the LANL beam test data.

5.8 The LANL Particle Beam Test

In December of 2019, another particle beam test of the SuperFGD Prototype was carried out at LANL. This beam test allowed us to build upon what we learned at CERN and explore the response of the prototype to neutrons. Given that the detector properties had been measured with the CERN data, there was much more of a focus on physics event analyses with the LANL data. I acted as a Run-Coordinator at the 2019 LANL beam test and had many responsibilities including deciding run configurations, carrying out data analysis and organising and training beam test workers.

The LANL data is still being analysed more than a year after it was first taken. Much of the analyses are still incomplete, but we will do our best to show the main results here with warning that some of them are still preliminary results. In December 2020, another SuperFGD Prototype beam test took place at LANL, with the results being analysed at the time of writing. Future publications will address the results of the 2020 LANL beam test.

For the 2019 LANL beam test, we were able to use the Los Alamos Neutron Science Center (LANSCE) facilities to take data with a 0-800 MeV neutron beam. The neutron beam is created by firing protons accelerated by LANSCE’s linear accelerator (LINAC) into a spallation target. Gamma rays and charged particles are also produced. The charged particles are removed from the beam using magnets. Unlike the CERN particle beam, there is no way to tune the energy of particles within the beam.

The beam test was conducted within the Target 4 Flight Path 15L (4FP15L) beamline at the Weapons Neutron Research Facility. Target 4 is a spallation target comprising tungsten, with the 4FP15L flight path being 15° to the left of the initial proton beam.

The LINAC proton beam had a specific timing structure that we were able to utilise in our triggering mechanisms. One can split the beam structure into long “macropulses” containing many smaller “micropulses”. Macropulses of protons were delivered every 8.3 ms, lasting for 625 μ s. Every sixth macropulse is directed to a different target other than target 4, hence no particles are produced at the 4FP15L beamline for every sixth macropulse. Within each macropulse are hundred of micropulses, each of which are separated by 1.8 μ s in time. There are roughly seven million protons within each micropulse.

There are two possible experimental locations along the 4FP15L beamline, one being 20 m from the target and the other 90 m away. The SuperFGD Prototype was initially placed in the 90 m location, before being moved to 20 m from the beam source after two weeks of data taking. The main difference in the beam properties between the two locations is the spread in arrival time between low-energy neutrons and high-energy neutrons. At 90 m, low-energy neutrons take so long to reach the detector that high-energy neutrons from the following micropulse may overtake them. We call this mixing of neutrons between separate micropulses “wrap-around”. Wrap-around can introduce difficulties when using time-of-flight to infer the energy of a detected neutron. Wrap-around can also occur at the 20 m location, but the maximum energy of wrap-around neutrons is smaller at the 20 m location. Hence, moving to the 20 m location allowed selection of neutrons with lower energies. However, use of the 20 m location data is currently limited in this analysis, due to the increased neutron flux causing as yet unsolved triggering issues. Efforts were made to address this issue at the 2020 LANL beam test, the results of which are yet to be seen.

For the SuperFGD Prototype, we were able to trigger on every micropulse thanks to a wired signal produced whenever protons struck the target. The wired signal was connected to a dedicated FEB on the SuperFGD Prototype.

5.9 Temperature Calibration at LANL

At the 90 m location of 4FP15L, the experimental area was in a poorly insulated building. Exterior temperatures ranged significantly over a 24 hour period, meaning the temperature dependence of MPPC gains could have a significant effect on the data. For this reason, a temperature sensor was installed as close to the SuperFGD Prototype as possible, which ended up being directly on top of the detector. The temperature sensor was a Raspberry Pi (RPi) 3 with a Sense HAT attachment [79]. The sensor had to be calibrated to account for the CPU temperature of the RPi. All temperature readings were transmitted to an online data storage site via a Wi-Fi connection.

The sensory data from the RPi had to be calibrated to account for the CPU temperature of the RPi itself, which could get hot enough to affect the immediate air temperature around it. The CPU temperature, T_{CPU} , was measured and logged directly by the RPi

and used to calibrate the measured temperature, T_{meas} , using the formula

$$T_{\text{calib}} = T_{\text{meas}} - (T_{\text{CPU}} - T_{\text{meas}})/x, \quad (5.5)$$

where T_{calib} is the calibrated temperature and x is a calibration factor. This calibration factor was measured using an accurate temperature probe to take measurements of the air temperature near the SuperFGD in the 90 m location. The calibration factor was then calculated such that the calibrated temperature output from the RPi matched the temperature probe's readings. This calibration factor was updated every few hours as it was noticed that the factor did not remain constant over time. The need to update x constantly was time consuming and likely negatively affected the accuracy of the temperature measurements. Ideally, this RPi setup will not be used for the final SuperFGD design, but the variation in x was small enough that the readings give a rough estimation of the temperature variations.

When we calibrated the readout electronics at LANL with the same method as outlined in Sec. 5.5, we monitored the temperature near the SuperFGD Prototype with the RPi. The average temperature was 23.4°C. Comparing the temperature at the time of calibration with the average temperature at run time, the gains of each MPPC can be scaled as needed. The required scaling was investigated by using data runs from the beam test. Light deposits in the neutron beam test were large and frequent enough that finger plots could be compiled from the beam data. By measuring MPPC gain values over several different temperatures, we found that the gain decreases by 1.2% for an increase in temperature of 1°C. This relationship is only linear over a narrow range of temperatures, but fortunately the air temperature around the SuperFGD did not vary by more than 10°C over all data runs.

Fig. 5.35 shows the calibrated temperatures using the RPi sensor HAT near the SuperFGD Prototype whilst at the 90 m location over a 24 hour period, starting on December 14th 2019. Also shown are the uncalibrated temperature readings and the CPU temperature readings. The figure shows a variation of roughly 5°C during the 24 hour period. Considering typical data runs with the SuperFGD only lasted around 20 minutes, the temperature variation was not significant enough to account for temperature change during data runs. Current data analysis has focussed on analysing single data runs, but when several data files are combined it will be important to account for differences in

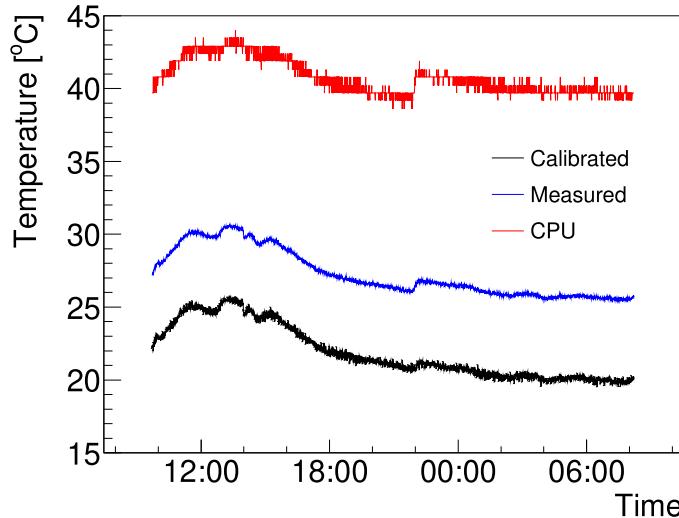


Figure 5.35: Temperature readings taken near the SuperFGD Prototype while at the 90 m location over a 24 hour period starting on December 14th 2019. The measurements taken by the RPi sensor (blue) were affected by the temperature of the RPi’s CPU (red). Hence, the measured temperature is calibrated to account for the CPU temperature (black).

temperature.

5.10 LANL Beam Test Data Analysis: Neutron Response

In our analysis of the LANL beam test data, we will focus on the response of the SuperFGD Prototype to neutrons. Specifically, what events look like when neutrons enter the detector, the defining properties of neutron interactions and the cross-section of neutrons with the SuperFGD Prototype scintillator material.

5.10.1 Observed Beam Structure and Energy Calculation

In Fig. 5.36, we show the beam structure of the LANL neutron beam as observed with the SuperFGD Prototype at the 90 m location. The left plot shows hits accumulated over a narrow time window of the hit time since the macropulse trigger for FEB 19. The hits

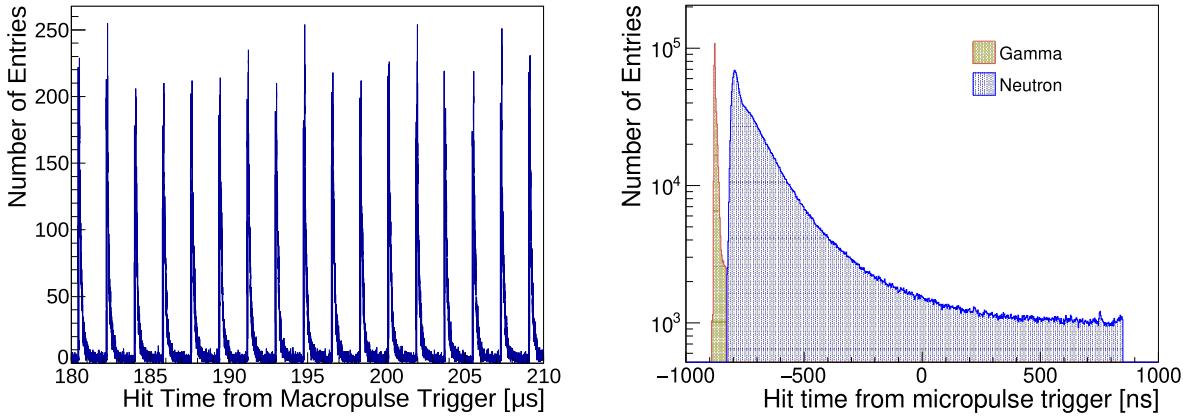


Figure 5.36: The beam structure of the LANL neutron beam as observed by the SuperFGD Prototype. The left plot shows hits detected by FEB 19 accumulated over a single data run. Hit times are taken from a specific time window after the macropulse trigger. The peaks correspond to micropulses spaced by 1.8 μ s. The right plot is the micropulse structure, where all hits for FEB 19 are plotted from a single data run, binned using the hit time from the micropulse trigger. There are two noticeable peaks, the earlier one corresponding to hits in the detector from gamma particles in the beam (filled yellow). The second peak is due to neutrons in the beam (dotted blue).

are not from a single macropulse, but are an accumulation of all hits over the data run within the displayed time window. One can clearly see the micropulse structure of the beam within the plot, with each micropulse evenly spaced by roughly 1.8 μ s.

The right plot of Fig. 5.36 shows a more detailed look at the micropulse structure. All FEB 19 hits are plotted from a single data run, binned using the hit time since the micropulse trigger time. Note that some hit times are negative due to the travel time of the trigger signal being longer than the travel time of some of the beam particles. There are two clear peaks in the micropulse structure. The earlier peak corresponds to hits from gamma particles and the later peak is from neutron hits. Because gamma particles travel at the speed of light, they will traverse the 90 m distance from the target to the prototype before the neutron particles, which are slower than the speed of light. The neutron peak also has a much longer tail than the gamma peak, due to the range of neutron energies and the resulting spread in travel time.

This right-hand plot also demonstrates how particle times-of-flight can be calculated and used to infer particle energy. We know the gamma particle speed is equal to the speed of light, c , and we know they had to travel 90 ± 1 m to reach the SuperFGD Prototype.

Hence, we can calculate the time they took to traverse the 90 m distance. If we assume the detected neutrons were produced at exactly the same time as the gamma particles, we can then use the time difference between a neutron hit and the gamma peak to calculate the time-of-flight of the neutron. Using special relativity, as the speed of the neutron is very close to the speed of light, we can then calculate the energy of the neutron, E_n :

$$E_n = \frac{m_n}{\sqrt{1 - (v_n/c)^2}}, \quad (5.6)$$

where m_n is the neutron mass and v_n , the speed of the neutron, is calculated using

$$v_n = \frac{D}{D/c + t_n - t_\gamma}. \quad (5.7)$$

In Eq. 5.7, D is the distance from the target to the SuperFGD Prototype, t_n is the neutron hit time with respect to the micropulse trigger and t_γ is the time with respect to the micropulse trigger of the gamma peak, measured as -880.0 ± 2.5 ns at the 90 m location and -712.5 ± 2.5 ns at the 20 m location.

5.10.2 Particle Events and Crosstalk Identification

Neutron interactions in the SuperFGD Prototype came in one of two forms. The simplest was if the neutron entered the detector, interacted with the material in a single cube and all resulting particles—except maybe some photons—of the interaction remained in said cube. This resulted in a single channel in each plane of the detector firing with a relatively large charge deposit. Such an event is shown in Fig. 5.37. The second type of event is where child particles of the neutron interaction exit the cube where the interaction occurred, causing more than one high charge signal to occur. Fig. 5.38 shows one such example.

In Figs. 5.37 and 5.38, no charge cuts have been applied and all hits measured with the micropulse time window are shown. As such, there are still visible crosstalk hits, as well as some dark count hits originating from the Type III MPPCs, as seen in the “Top View” plot of Fig. 5.37. These dark counts are easily removed with a charge cut, as they rarely have charge amplitudes larger than 10 p.e. However, for the crosstalk hits, a charge cut is less ideal in order to remove them. Unlike in the CERN beam test, there are several

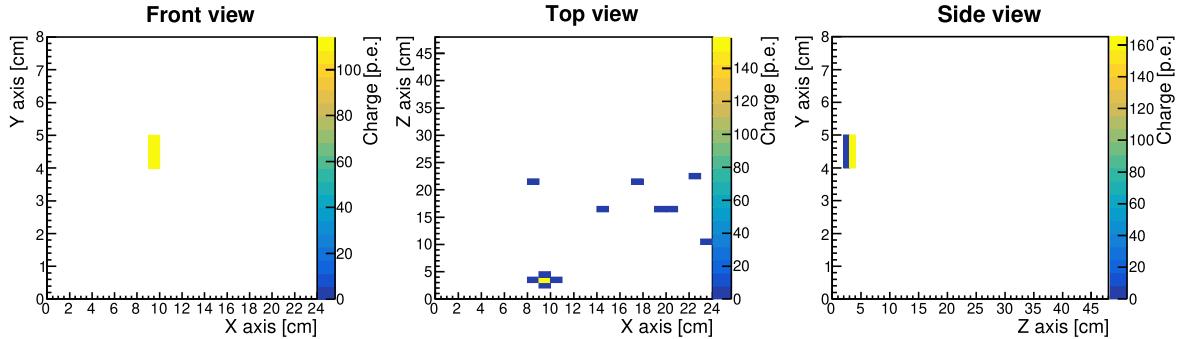


Figure 5.37: A neutron event in the SuperFGD Prototype where no charged particles left the cube in which the neutron interacted. No cuts are applied to remove crosstalk or dark counts. The dispersed hits seen around the central z -values of the “Top View” plot are due to dark counts from the Type III MPPCs.

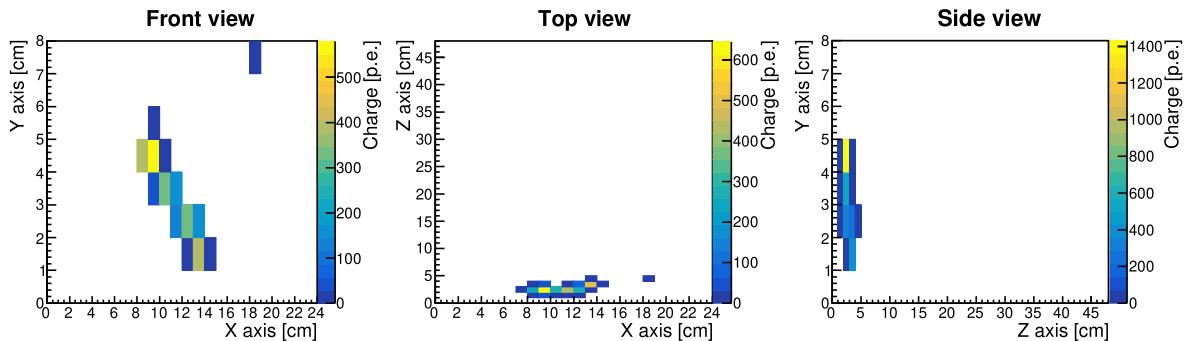


Figure 5.38: A neutron event in the SuperFGD Prototype where a charged particle was not contained within the cube where the neutron interacted. No cuts are applied to remove crosstalk or dark counts.

angled tracks in the detector, as a result of neutrons inducing charged particle movement in the detector. This results in tracks where charged particles only pass through a small section of the corner of a cube, rather than directly through its centre. In such cubes, significantly less scintillation light is produced relative to cubes where the particle passed directly through the centre. As such, the LANL data contains several hits directly caused by charged particles that have low signal amplitudes. Signal amplitudes that are so low they are of the same order as crosstalk hit amplitudes. Hence, a simple charge cut would remove many real hits as well as crosstalk hits. The left plot of Fig. 5.39 demonstrates the issue of applying a single charge cut to remove crosstalk. In the plot, Geant4 particle simulation software has been used to generate Monte Carlo (MC) simulations of the LANL neutron beam in the SuperFGD Prototype. Each hit of the MC is coloured according to

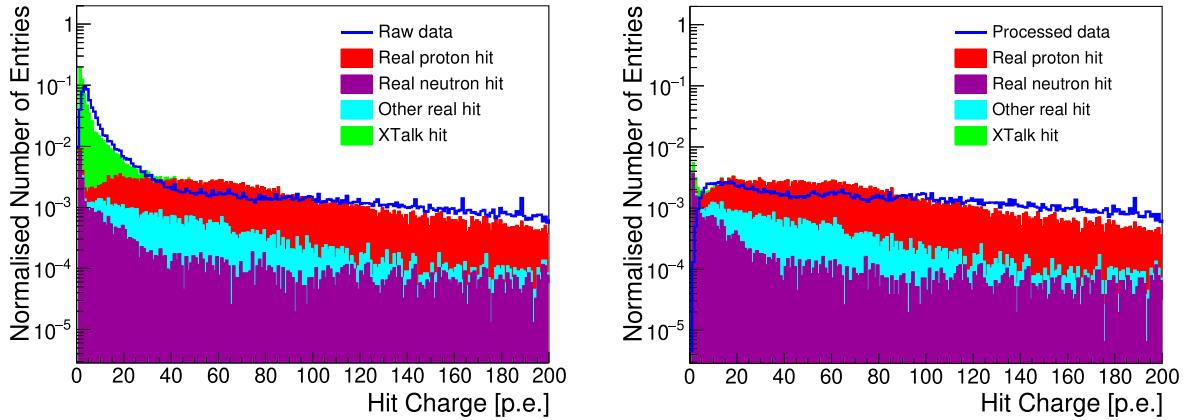


Figure 5.39: MC (filled coloured histograms) and real data (empty blue histogram) before (left) and after (right) the crosstalk algorithm is used to cut out hits that did not originate from the cube of a main charge deposit. The MC in the left plot has a real hit purity of 11.8%. The right plot’s MC has a real hit purity of 98.3% with an efficiency of 95.6%.

the cause of the hit, be it a certain particle or from crosstalk. MPPC dark counts were not included in the simulation. The optical crosstalk simulations are heavily based on the results from the CERN beam test in Sec. 5.7.5. The plot clearly shows that applying a charge cut at 40 p.e. to remove all crosstalk hits would also remove a significant number of hits originating from particles. The raw data from a data run in the LANL beam test is also shown in the plot for comparison to the MC.

For the final SuperFGD, it is planned to incorporate pattern recognition software to the particle tracks in order to identify particle events. Removing crosstalk hits would be extremely useful in order for said pattern recognition software to work reliably, as crosstalk hits can add a degree of ambiguity to the location of interaction vertices. An algorithm has been developed to tackle the issue of crosstalk identification, utilising not only hit amplitude information from a single hit, but from surrounding hits and hits in multiple planes.

The crosstalk identification algorithm is shown in Fig. 5.40. Hits that have not yet been labelled as “real” or “crosstalk” by the algorithm are fed into the algorithm one-by-one. When a hit enters the algorithm—which we will call the “initial hit”—the hit and all corresponding hits in other planes of detection are immediately identified as real hits if the initial hit has a charge amplitude larger than 50 p.e. This is because we would not expect to see any crosstalk hits of such large amplitude. By “corresponding hits” we mean

hits that originated from the same cube as the initial hit.

To cut out dark count hits, if a hit occurs outside a time window of ± 12.5 ns about the highest amplitude hit in the event, then the initial hit and corresponding hits are labelled as crosstalk. Note we are using the term “crosstalk” to encompass dark counts too.

The algorithm goes on to check the signals from neighbouring channels to the channel in question where the initial hit was seen. If these neighbouring channels have charge amplitudes, A_N , larger than a calculated threshold to suggest the current channel’s signal is caused by crosstalk from said neighbouring channels, then the initial hit is labelled as crosstalk. The exact condition to label the initial hit as crosstalk is

$$A_N > A_I \left(\frac{1 - 2(\kappa_i + \sigma_{\kappa,i})}{\kappa_i + \sigma_{\kappa,i}} \right), \quad (5.8)$$

where A_I is the charge amplitude of the initial hit, κ_i is the ratio of light leaking into a single neighbouring cube as detected by a WLS fibre parallel to the i -axis, as discussed in Sec. 5.7.5 and $\sigma_{\kappa,i}$ is the standard deviation in κ_i . In other words, Eq. (5.8) calculates a tentative upper limit of the light yield we would expect to see in neighbouring crosstalk cubes relative to the light signal measured along the WLS fibre through the main cube and uses this light yield as a threshold. The expected light yield calculation is found by rearranging Eq. (5.1) to find $M_{\text{xtalk},i}$, which we assume to be A_I , then substituting κ_i for $\kappa_i + \sigma_{\kappa,i}$. We make this substitution to encapsulate as many of the possible crosstalk amplitude values as possible, without making the threshold so large as to misidentify real hits as crosstalk hits.

The effectiveness of the crosstalk identification algorithm is demonstrated in Fig. 5.39. The left plot of the figure shows MC and real LANL data with no cuts applied. The right plot shows the same data sets but with crosstalk hits identified and removed using the crosstalk algorithm. The crosstalk algorithm has removed all but 2% of the crosstalk hits and has an efficiency of 95.6% for real hit retention. Conversely, if a 40 p.e. charge cut is imposed to cut out crosstalk hits, this results in a purity of 99.5%, but an efficiency of 71.8%. Clearly, when looking at MC data, the crosstalk algorithm is a much better identifier of crosstalk and dark hits as opposed to a simple charge cut. However, the MC may not accurately represent the real data. Indeed, we do see some discrepancies between MC and real data at low energies. In the right plot of Fig. 5.39, there are many MC hits

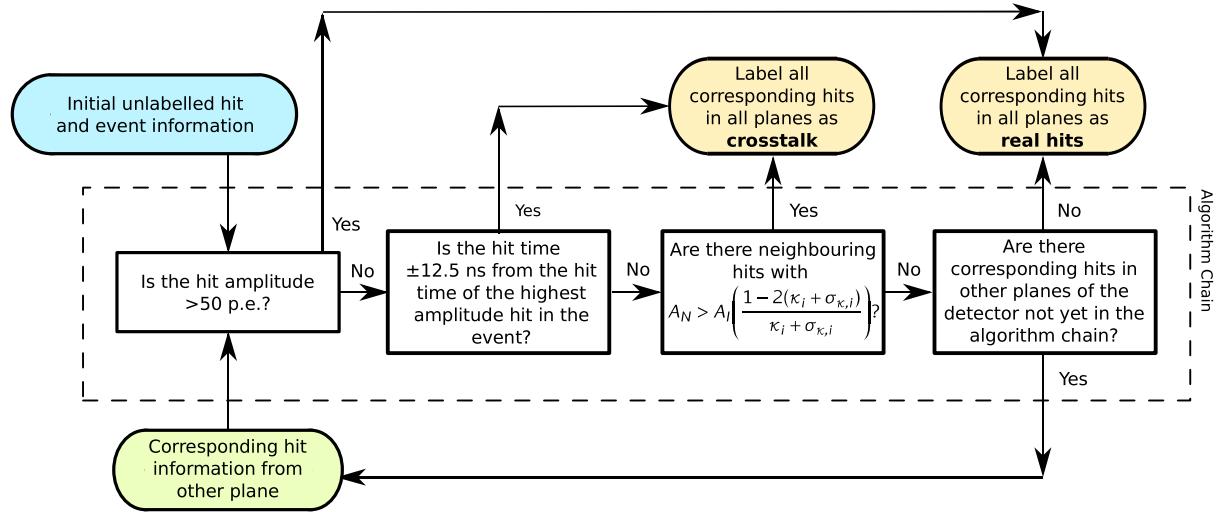


Figure 5.40: The algorithm to determine if a light signal originates from the main charge deposit of a charged particle or from the optical crosstalk of said light deposit into neighbouring cubes. The algorithm also identifies dark counts, but does not discriminate between crosstalk hits and dark count hits. The initial hit enters the algorithm chain where it must pass several conditions. If an initial hit passes all the conditions, then other hits originating from the parent cube of the initial hit, i.e. the hits from other planes, are passed into the algorithm chain one-by-one. If they all pass, they are all labelled as “real hits”. Otherwise, they are all labelled as crosstalk hits.

below around 5 p.e. that were labelled as real hits, but in the data there are extremely few real hits at these energies. A likely explanation for this is that no hit amplitude thresholds are implemented in the MC data. The simulation software is still being developed and the hit amplitude thresholds will be added in time for the analysis of the 2020 LANL data.

For a more visual demonstration of the effectiveness of the crosstalk algorithm, Fig. 5.41 and Fig. 5.42 show the neutron event displays previously seen in Figs. 5.37 and 5.38 but with crosstalk and dark hits identified and removed using the crosstalk algorithm.

5.10.3 Event Reconstruction and Neutron Selection

Unlike the CERN beam test analysis, there has been a significant focus on the reconstruction of events with the LANL beam test data. This is mainly because neutron interactions within the detector are of major interest, given the SuperFGD is being built partly to improve neutron detection.

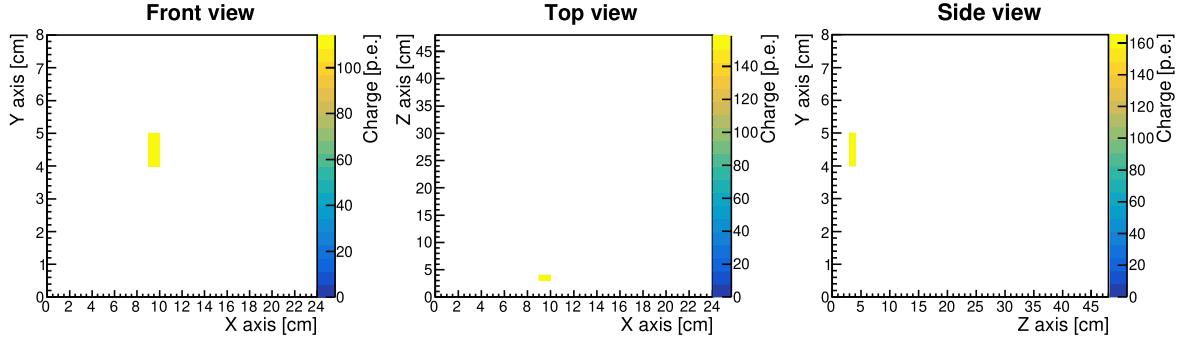


Figure 5.41: A neutron event in the SuperFGD Prototype where no charged particles left the cube in which the neutron interacted. Cuts are applied to remove crosstalk and dark counts using the crosstalk identification algorithm.

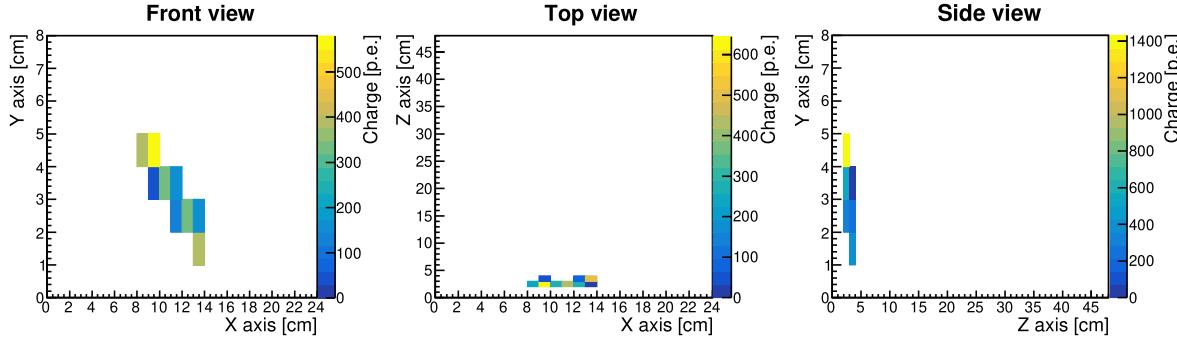


Figure 5.42: A neutron event in the SuperFGD Prototype where a charged particle was not contained within the cube where the neutron interacted. Cuts are applied to remove crosstalk and dark counts using the crosstalk identification algorithm.

Before any reconstruction algorithms are run, hits within micropulses are organised into time clusters. These clusters are defined as groups of hits where the earliest hit time is at least 17.5 ns later than the latest hit time of another cluster. This separation in time is used to separate sub-events within micropulses, as it is possible for multiple events to occur within a single micropulse.

Once hits are grouped into sub-events, the hit information from the three two-dimensional readout planes are combined to create three-dimensional hit voxels. A voxel is created if hits from all readout planes originate from the same x -, y - and z -values. The co-ordinates of each voxel are the common x , y and z co-ordinates of the 2D hits. Timing and signal amplitude information is also stored in each voxel. The voxel time is recorded as the earliest of the hit times from the 2D hits used to create the voxel. The signal amplitude of the voxel is taken as the sum of hit amplitudes from the voxel's

constituent 2D hits.

Now using the 3D voxel information, the voxels are grouped together according to spatial clustering. This is one more step to separate sub-events within a micropulse, as separate interactions may occur at similar times, meaning they are not parsed by the time clustering step. Voxel spatial clusters are taken to be groups of voxels that are all at most 1.8 cm from another voxel in the same spatial cluster. This spatial clustering is heavily based on the DBSCAN (density-based spatial clustering of applications with noise) algorithm.

With voxels detected within a micropulse grouped into separate events, various geometric properties of the voxel clusters are measured. Linearity, planarity and sphericity are the measured geometric properties. For the LANL beam test event selection, events are only selected if the cluster linearity is high enough. The largest distance between two adjacent voxels in a cluster is also calculated. If this distance is larger than 2 cm, implying a large gap in the reconstructed track, the event is left out of the event selection.

The tracks in the event are reconstructed by finding the principal vectors of the voxel cluster. These principle vectors are found by calculating the covariance matrix of the cluster voxels and finding the matrix's eigenvectors.

Finally, the event vertex is identified by finding the voxel with the earliest hit time and lowest z -value.

5.10.4 Neutron Energy and Cross-section Measurements

Using the neutron energy measurement technique outlined in Sec. 5.10.1, a histogram of neutron energies could be constructed using LANL data. This was a useful check to corroborate the energy reconstruction by comparison to simulated data where the particle energy was available.

In Fig. 5.43, neutron energy spectra are shown using measurements from real data and from MC simulations. Both data sets correspond to the SuperFGD Prototype being situated 90 m from the target. The real data spectrum utilised the time-of-flight energy measurement technique and the neutron event selection algorithm. The neutron energy was measured using the voxel information at the identified vertex of the event. The

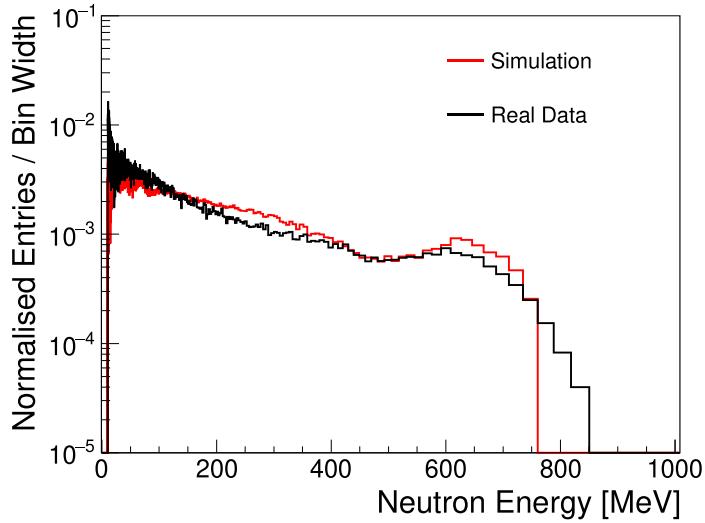


Figure 5.43: Neutron energy profiles measured with real data (black) and extracted from MC simulation truth information (red).

MC simulations used the truth information provided by the MC to return the neutron energy profile for neutrons that interacted in the simulated detector. The simulation was informed using the neutron energy flux provided by LANL, using Geant4 to propagate the neutrons from the target to the simulated detector. The result is a decently matched shape between the two spectra. Most notable is the shared peak in neutron energy around 600 MeV. The bin width is varied due to the discrete nature of the time measurements from the SuperFGD Prototype, which measures time in 2.5 ns increments. Transforming time to energy using the time-of-flight technique leads to improved energy resolution for low energy (large time-of-flight) measurements. The same binning is used for the MC data for comparison with the real data. Also, the MC energy spectra stops at around 750 MeV, but the measured spectra goes on above 800 MeV. This is due to the provided neutron flux from LANL having a maximum value of 763 MeV, although neutrons above this energy are expected.

The sharp cut off in neutron energies for both datasets around 10 MeV is due to the maximum time-of-flight possible within a micropulse. Neutrons with energies below 10 MeV will arrive at the same time as neutrons from the next micropulse, meaning they could be misidentified as high energy neutrons.

Knowing the energy threshold of neutron interactions with the SuperFGD's plastic

scintillator is useful, but we can go one step further and measure the cross-section of neutrons with the scintillator. If we assume the neutron beam follows a general attenuation of flux as it passes through the scintillator, we can model the flux, Φ , as

$$\Phi(z) = \Phi_0 e^{-n\sigma z}. \quad (5.9)$$

In this equation, z is the distance the neutron beam has travelled through the detector, Φ_0 is the flux at $z = 0$, σ is the interaction cross-section and n is the number density of scintillator molecules. This is also known as the Beer-Lambert Law, but mainly when referring to the attenuation of light particles [80]. If we take the front-most hit in each event as the position where the neutron interacted with the detector, we can build up a dataset of the number of neutron interactions in the SuperFGD Prototype as a function of z . The number of neutron interactions will be proportional to flux, so the resulting dataset will follow an exponential as shown in Eq. (5.9). We can fit this exponential to find the neutron-scintillator cross-section. By building datasets for several neutron energies, values of the energy dependent cross-section can be measured.

Shown in Fig. 5.44 is an example of the number of neutron interactions per z -layer of the SuperFGD Prototype for a specific energy bin. The energy bin in this case is 0–40 MeV. An exponential fit to the data is also shown. For now, the error bars are purely statistical. In the future, we are looking to incorporate systematic uncertainties into our calculations.

Fig. 5.45 shows the culmination of fitting the extinction profile for all specified energy bins between 0 and 800 MeV. The result is a plot of the neutron-scintillator cross-section as a function of energy. We have shown the results when using the beam test data and the simulated data. The simulated data should represent the result of the cross-section model calculations within the simulation, so comparison of the beam test data result to the simulation results should signify how accurate the beam test results are. The vertical error bars in the plot purely correspond to statistical uncertainties. See Sec. 5.11 for details on how these uncertainties could be better estimated.

The general trend of the beam test data cross-section results matches the simulation cross-section results very well. The main difference is in the tail of the distribution where the smaller number of events in the beam test data leads to larger uncertainties, although these uncertainties do not seem well represented by the vertical error bars.

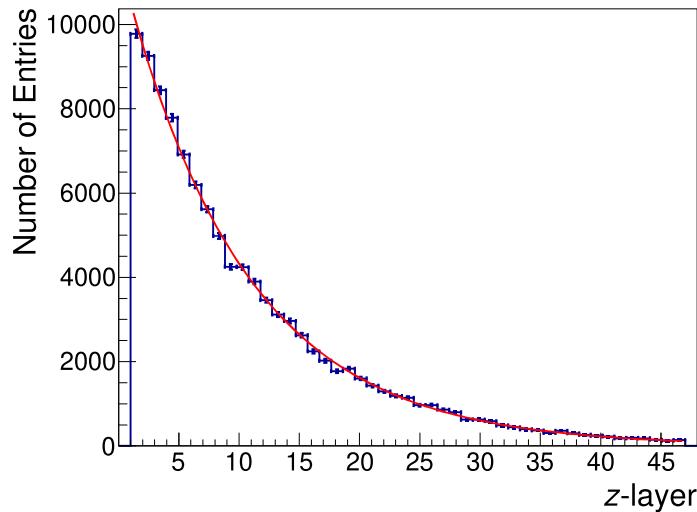


Figure 5.44: The neutron beam attenuation as it travels through the SuperFGD Prototype for neutrons with kinetic energies below 40 MeV. The number of neutron interactions at each z -layer (blue) is fit with an exponential function (red) to extract the neutron-scintillator cross-section for low energy neutrons.

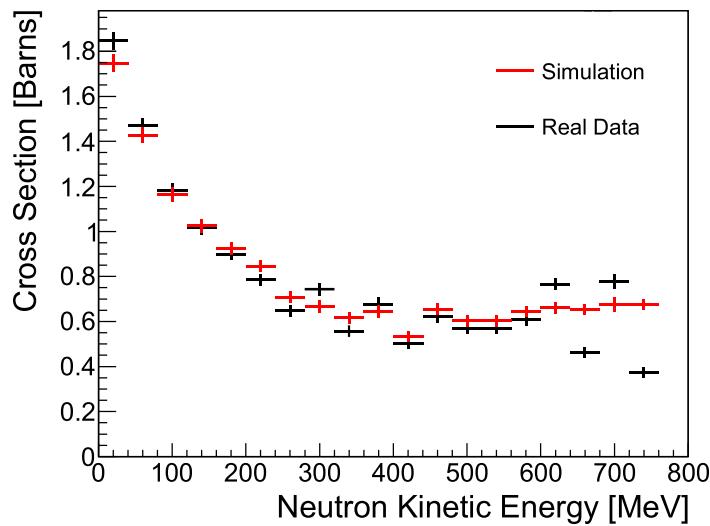


Figure 5.45: The energy dependent cross-section for all neutron-scintillator interactions as measured using the LANL beam test data (black) and the simulation data (red). Horizontal energy bars correspond to the neutron energy bins. Vertical error bars are derived from statistical uncertainties from the exponential fit.

5.11 Further Analysis

We have shown detailed yet still preliminary results for the LANL beam test with the SuperFGD Prototype. There are still plans for further work with both the 2019 and 2020 data.

The cross-section analysis detailed in Sec. 5.10.4 gives a basic idea of the total cross-section for neutrons incident on the SuperFGD Prototype scintillator material. However, there are several aspects to the analysis that could be improved given more time. One improvement could be in correcting for interactions that were either not detected or not selected during the analysis. These could be accounted for by including a layer dependent correction factor representing the ratio between single-track events and other-than single-track events. This correction factor could be estimated by studying MC. For such a study, though, the current MC would need to be improved to guarantee accuracy.

One way to improve the LANL beam test MC is to weight the transverse spread of the simulated neutron beam according to recorded data. The transverse spread of interactions increases along the beamline, most likely due to invisible scatterings of neutrons. This invisible scattering is not currently tuned to the SuperFGD geometry in the LANL beam test MC. Initial studies show adding a weight to invisible scattering interactions can match the MC to data. The effect on the transverse spread of the beam from applying different weights in the MC is shown in Fig. 5.46 and compared with the data. From the figure, it appears that applying a weight of 0.5 to the invisible scattering matches the transverse spread of the simulated beam to the real data well. This is most likely the weighting that will be applied in future versions of the MC and used to extract correction factors for the cross-section.

Another contribution to the uncertainties in the LANL analyses that is not yet accounted for is the uncertainty in measured energies. Energy measurements heavily rely on timing information, so the time resolution can be translated to an uncertainty in the measured energy of an interacting neutron. Using the previously measured time resolution of 1.14 ns (see Sec. 5.7.9), uncertainties on the measured neutron energies can be calculated and propagated to uncertainties in the cross-section calculation, for example.

Finally, an aspect to the LANL beam test that has not been addressed in enough detail as of yet is the wrap-around events that can occur between micropulses. This is when low

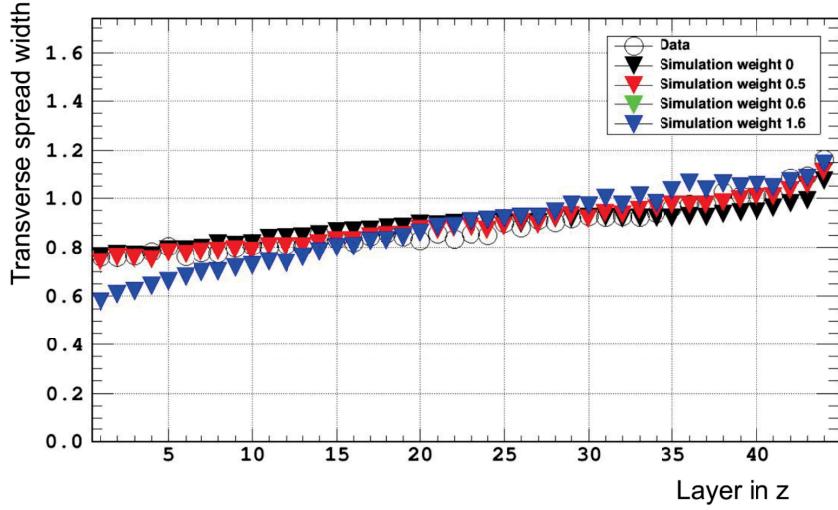


Figure 5.46: The transverse spread of the simulated beam along the SuperFGD prototype with different invisible scattering weights. The transverse spread of the beam test data is also shown for comparison. A weighting of 0.5 appears to give a similar distribution to real data.

energy neutrons travelling relatively slowly arrive at the detector at the same time as high energy neutrons travelling quickly from a later micropulse. Due to the $1.8 \mu\text{s}$ spacing of the micropulses, at the 90 m location, neutrons with kinetic energies of less than 10 MeV can wrap-around and contaminate samples of high energy neutrons. In the 2020 LANL beam test, a special micropulse spacing of 3.6 ns was used for a limited time. Such a spacing will potentially allow the observation of neutrons with a minimum energy of 2 MeV.

With the SuperFGD set to be installed in ND280 by 2022, the analyses shown here and the analyses currently in progress will be invaluable resources to help inform users of the performance and use of the brand new detector technology. With our results on time resolution, neutron detection and energy thresholds, we have demonstrated the benefits of the SuperFGD over a typical FGD design and how it can potentially reduce systematic errors at T2K. The small time resolution, combined with the time of flight detectors to be installed with the SuperFGD, will allow backwards-going particles to be more efficiently identified. Particle identification will also be more efficient with the lower energy thresholds of the SuperFGD, including neutron detection for anti-neutrino interactions. As outlined in Sec. 3.4, these properties will lead to improvements over the current ND280 and a reduction in systematics, leading to more precise δ_{CP} measurements.

As part of a study into global fits of neutrino data, we can demonstrate the improved sensitivity of the T2K experiment thanks to the reduced systematics and increased POT. In the following chapter, we will detail the three-neutrino global fit implementation within the global fitting software GAMBIT and show how it can be used to estimate the improved constraints of the T2K experiment after the ND280 Upgrade.

Chapter 6

Neutrino Oscillation T2K Fit with GAMBIT

T2K is one experiment of many trying to measure the oscillation parameters of neutrinos. In the past, a number of groups have combined the results of several neutrino oscillation experiments into global fits of the data, some of the most recent being Global Fit [81] and NuFit [82]. To no fault of the authors of such global fits, the predictions of experimental data within said global fits often lack experimental nuances and are often forced to source information about experiments from uncorroborated sources such as PhD theses and presentations.

Global fits can be an invaluable resource when studying the field of neutrino oscillations. Because different experiments are sensitive to different areas of the parameter space, combining the results of several experiments can reveal interesting and informative parameter space constraints. That being said, carrying out accurate and valid global fits is a challenging task, as it requires simulating several different experiments and sourcing information that may not be publicly available. As part of their goal to fit all BSM phenomena across many different areas of physics, the global fitting software GAMBIT (Global And Modular BSM Inference Tool) is aiming to produce a global fit of three-neutrino oscillation data using advice and input from neutrino experimentalists that have the knowledge to implement experiment specific nuances and retrieve the required information from reliable sources. This work is still in progress, but the T2K likelihood calculation has been implemented and will be shown here as an example of how the GAMBIT global fit can lead to improvements over past global fits.

6.1 The GAMBIT Software

First, we will go into more detail on what the GAMBIT software is and what results have been published using GAMBIT. The details here only scratch the surface of GAMBIT’s functionalities and capabilities, but should give the reader enough information to understand how GAMBIT works in the context of the neutrino oscillation global fit. For more in-depth writing on GAMBIT, see the published literature [83, 84].

The GAMBIT collaboration have been publishing papers on global fits using their software since 2017. They have several notable results regarding dark matter and collider physics [85–87]. Recently, they have been delving into neutrino physics with a study into the extension of the Standard Model by adding three right-handed neutrinos, plus they have published a bound on the mass of the lightest neutrino using both terrestrial and cosmological experimental data [88, 89]. This has led up to the development of a global fit for neutrino oscillations with three-neutrino flavours.

For a typical user of GAMBIT that wants to run a scan over a particular parameter space to produce constraints, they will have to define a model (i.e. a collection of parameters), write functions within GAMBIT to calculate observables and compare these calculated observables to experimental data in order to produce a likelihood value. Models in GAMBIT have a hierarchical structure, with parent and children models. Children models contain a subspace of their parent model’s parameter space. Models can be used simultaneously, for instance SM parameters can be included as nuisance parameters while testing a BSM model.

The GAMBIT software itself is impressively vast in the number of options it provides when scanning over a parameter space. These can be options for how to explore the parameter space, with several algorithms available for either frequentist or Bayesian studies, or how to distribute the computational resources, through a two-level parallelisation system in which the parameter exploration and per-point physics computations are parallelised using MPI and OpenMP, respectively. As is stands, the software contains an extensive number of implemented models and experimental datasets, but new models and datasets can easily be added.

The modular aspect of GAMBIT splits the software into different areas of physics. For example, dark matter physics, flavour physics, neutrino physics and so on. Each

module contains “module functions” that have specified capabilities, such as a likelihood calculation. These module functions can also have dependencies on other module functions. If a model is used in a scan, a module function must be specified that can calculate the relevant likelihoods and satisfy the dependencies of other module functions. It is in these specifications that different calculation methods can be requested when performing a scan, removing or adding contributions to the total likelihood, for example.

6.2 NuFit’s Global Fit of Neutrino Oscillation Data

The fit of neutrino oscillation data performed by NuFit is impressive in scope. The group have combined data from many experiments, including experiments using neutrinos from accelerators, from nuclear reactors and from atmospheric and solar production. Their most recent analysis from 2020 shows some interesting results considering the possible tensions between the favoured regions for T2K and NOvA on the δ_{CP} value [82].

We have identified some areas of the NuFit analysis where we believe improvements can be made and have explored the impact of implementing said improvements. Looking at the calculation of the T2K likelihood alone, there are some assumptions made that reduce the accuracy of the calculation. NuFit predict the expected event spectra in T2K’s far detector by using T2K’s neutrino flux and cross-sections. They also account for energy reconstruction within SK. However, the energy reconstruction is a simple gaussian approximation and the adjustments to the flux and cross-section values performed by T2K using fits to near-detector data are not accounted for, except for an adjustment of “efficiencies” to match the NuFit event predictions to some MC of T2K’s events at SK shown in a public presentation. This basic approach to sources of systematic error can produce inaccuracies in the parameter space constraints.

6.3 GAMBIT’s Global Fit of Neutrino Oscillation Data

GAMBIT aims to improve upon NuFit’s approach to neutrino oscillation global fits by, firstly, sourcing as much information as possible from official sources. NuFit undoubtedly

took the same approach regarding information sourcing, but having experimentalists from several neutrino oscillation experiments in the GAMBIT collaboration may allow for some additional insight when making use of official information. GAMBIT also aim to provide more accurate energy reconstruction simulations for each likelihood calculation. T2K provide plots showing how reconstructed energies are distributed for events with a single true energy, which is informative and will allow GAMBIT to improve upon the simple Gaussian approach. Lastly, for the T2K likelihood calculation, GAMBIT will implement the near-detector fit parameters and profile over these parameters in scans across the parameters space. This may not be applicable for all likelihood calculations, but T2K provide some of the most significant constraints on the neutrino oscillation model, so committing to improving the calculation of T2K’s likelihood will potentially have a large impact.

6.4 Likelihood Calculation for T2K

For the implementation of the T2K experiment within GAMBIT, the basis of the likelihood calculation is similar to NuFit’s but we have expanded upon the calculation by introducing weights for the near-detector fits and a more accurate treatment of energy reconstruction [90]. At SK, the reconstructed energy spectra for neutrino interactions (i.e. the number of detected neutrino events at SK as a function of reconstructed energy) are predicted using information known about the T2K experiment and its far detector. For reconstructed energy bin i with minimum reconstructed energy E_i and maximum reconstructed energy E_{i+1} , for a specific interaction channel a with produced neutrino flavour and sign β and observed neutrino flavour and sign γ , the number of events expected to occur in said reconstructed energy bin, $N_i^{a,\beta,\gamma}$, is calculated using

$$N_i^{a,\beta,\gamma} = \int_{E_i}^{E_{i+1}} dE_{\text{rec}} \int_0^\infty dE_\nu R_{a,\gamma}(E_{\text{rec}}|E_\nu) \frac{d\Phi_\beta}{dE_\nu} \sigma_{a,\gamma}(E_\nu) \epsilon_\gamma(E_\nu) P_{\nu_\beta \rightarrow \nu_\gamma}(E_\nu). \quad (6.1)$$

The constants and parameters in this equation are as follows: E_{rec} is the reconstructed energy of the parent neutrino, $R_{a,\gamma}$ is the reconstruction function for interaction channel a and observed neutrino ν_γ , E_ν is the true energy of the neutrino, Φ_β is the ν_β flux at SK assuming no neutrino oscillation, $\sigma_{a,\gamma}$ is the cross-section for interaction channel a of ν_γ particles on water, ϵ_γ is the ν_γ detection efficiency at SK and $P_{\nu_\beta \rightarrow \nu_\gamma}$ is the probability of

a ν_β produced at J-PARC oscillating to a ν_γ at SK.

T2K present their data as events binned in reconstructed energy bins split over several data samples. These data samples are defined by the cuts applied to the data to select specific neutrino interaction channels. There are five data samples in total, corresponding to CCQE interactions of ν_μ particles in FHC mode, CCQE interactions of $\bar{\nu}_\mu$ particles in RHC mode, CCQE interactions of ν_e particles in FHC mode, CCQE interactions of $\bar{\nu}_e$ particles in RHC mode and CC1 π interactions of ν_e particles in FHC mode.

In GAMBIT, to simulate the five different data samples of T2K, Eq. 6.1 was used while summing over the β parameter for muon and electron flavours of neutrinos and anti-neutrinos, which accounts for ν_e , $\bar{\nu}_e$ and $\bar{\nu}_\mu$ (ν_μ) contamination of the J-PARC FHC (RHC) neutrino beam. SK is unable to differentiate between leptons and their anti-particles, hence the need to sum over neutrino and anti-neutrinos. Background events were also accounted for in each sample. Many neutrino interactions can potentially pass the selection cuts applied to remove said interactions, albeit the number of these events will be small for most interaction channels. The exception to this is the CCQE, 2p2h and CC1 π interaction channels, which have very similar final state particles if the pions from CC1 π are misidentified or missed altogether. As such, the events from these interaction channels at SK are calculated directly in GAMBIT for each data sample, whereas all other interaction channels are regarded as background events. These background events are difficult to simulate as their numbers rely heavily on selection cuts. Given this difficulty, GAMBIT uses extracted values of the background events from T2K MC.

Neutrino flux values as a function of true neutrino energy at SK, $\frac{d\Phi_\beta}{dE_\nu}$, are provided by T2K and used in the GAMBIT calculation, specifically the 2016 flux release [91]. T2K derive their flux values from detailed beam simulations. Interpolation of the flux values is used in GAMBIT considering the flux information from T2K is only provided via data points. Also, the flux provided by T2K is given in units of $/\text{cm}^2/50\text{MeV}/10^{21}\text{POT}$, where POT represents the Protons On Target in the beam production. To scale the flux up to any desired period in time, one must include the corresponding POT factor. For the run1–10 data used in GAMBIT, the corresponding POT is 1.9664 for FHC and 1.63456 for RHC.

T2K also tune their flux predictions using fits to the near-detector data, called the BANFF (Beam And ND280 Flux extrapolation task Force) fit [92]. These fits are used to

produce weights that are applied to bins across the energy spectra. The BANFF weights are made public by T2K and are used by GAMBIT in their energy spectra predictions. This is one step beyond the NuFit analysis which did not account for near-detector fits.

The neutrino interaction cross-sections, $\sigma_{\alpha,\gamma}$, are extracted from the neutrino interaction generator software NEUT [93, 94]. NEUT combines theoretical models with experimental data to give accurate interaction cross-sections as a function of energy. These cross-section values are stored in GAMBIT and interpolated at run-time.

Neutrino oscillation probabilities are calculated directly using the open source software Prob3++ [95]. This software is also used by T2K in their oscillation analyses. Prob3++ uses no approximations in its probability calculations and also accounts for traversal through matter. In GAMBIT, Prob3++ has been incorporated as a plug-in (“backend” in GAMBIT terminology), allowing efficient and disk-less communication of oscillation parameters and probabilities between GAMBIT and Prob3++ during a scan.

The integration over true neutrino energy and the reconstruction function, R , are used in Eq. (6.1) to simulate the energy reconstruction of SK. As stated previously, NuFit used Gaussian reconstruction functions for their global fit. GAMBIT are using a mixture of Gaussian and log-normal distributions for their reconstruction functions. This is based on T2K MC showing the reconstructed energy distribution for a single true neutrino energy, as shown in Fig. 6.1. For the $2p2h$ interaction channel, two log-normal distributions were included to simulate both Δ and non- Δ enhanced events.

Each reconstruction function has several independent parameters, the width of the function and the offset from the true energy, where the true energy is defined as the mode of the R function. Each of these parameters is fitted using a χ^2 fit to MC provided by T2K broken down by sample and interaction channel. MC for five different points from the oscillation parameter space was provided and the χ^2 fit used all points in the fit to ensure a local minima was not found from using a single parameter space point. The values of these parameter space points are shown in Tabs. 6.1 and 6.2. The detector efficiency, ϵ_γ , is also fitted to this MC. These parameters and their nominal values are shown in Appendix C.

GAMBIT’s predicted energy spectra for T2K’s run1-10 data are shown in Figs. 6.2, 6.3, 6.4, 6.5 and 6.6. The energy spectra in Fig. 6.2 correspond to a point in the oscillation parameter space indicated by Tab. 6.1, named Asimov A by T2K [97]. Each

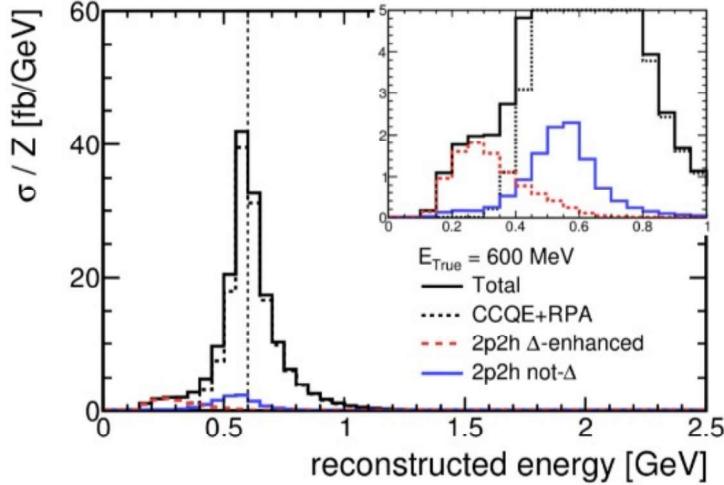


Figure 6.1: Simulated data showing the distribution of reconstructed energies for neutrino events of a single true energy ($E_\nu = 600$ MeV). The reconstructed energy assumes all events are from the CCQE interaction channel. As well as the sum of events from all interaction channels, the individual interaction channels, CCQE, $2p2h$ Δ -enhanced and $2p2h$ not- Δ are also shown [96].

plot in the figure depicts the energy spectra for a different data sample. The coloured histograms represent events from the different interaction channels, which were calculated with Eq. (6.1), except for the NC+CCother channel, which is taken as background and extracted directly from T2K's MC. The T2K MC, produced using the experiment's MaCh3 (Markov Chain 3) software, is also displayed in each plot for comparison to GAMBIT's predictions. This MC has not been approved by T2K for external use, but is used in this thesis to demonstrate how it could be used if approved in the future. If not approved, GAMBIT will instead source reconstruction function parameters from publicly available MC plots, which only show energy spectra from one point on the parameter space and usually don't break down the interactions in an ideal way to extract reconstruction function parameters.

The provided MC was broken down per interaction channel, but the displayed MC here is only shown as a sum over interaction channels. In each data sample, the MC matches the GAMBIT predictions very well, especially considering how the spectra is not directly tuned to fit the MC data. Figs. 6.3, 6.4, 6.5 and 6.6 contain energy spectra corresponding to set 1, 2, 3 and 4 of the oscillation parameters respectively, as shown in Tab. 6.2. In these figures, the trends of the GAMBIT predicted energy spectra follow those of T2K's MC relatively well, although there is some divergence between the GAMBIT predictions

Parameter	Asimov A
$\sin^2(\theta_{23})$	0.52800
$\sin^2(\theta_{12})$	0.30700
$\sin^2(\theta_{13})$	0.02180
$\Delta m_{32}^2 \approx \Delta m_{31}^2$	0.00251 eV ²
Δm_{21}^2	0.00008 eV ²
δ_{CP}	-1.60100

Table 6.1: The Asimov A neutrino oscillation parameters.

Parameter	Set 1	Set 2	Set 3	Set 4
$\sin^2(\theta_{23})$	0.45	0.6	0.45	0.6
$\Delta m_{32}^2 \approx \Delta m_{31}^2$	0.00235 eV ²	0.00235 eV ²	0.0026 eV ²	0.0026 eV ²

Table 6.2: The neutrino oscillation parameters used in addition to the Asimov A point when comparing GAMBIT’s calculations to T2K’s MC. Any parameters not shown are set to the same values as Asimov A.

and T2K MC in the first peak of the ν_μ and $\bar{\nu}_\mu$ samples of set 3 and 4. This highlights the difficulty to match the GAMBIT predictions and T2K MC across the whole parameter space, even with complex reconstruction functions. However, the match is good enough that there should not be significant effects on the resulting parameter space constraints.

The T2K likelihood itself is calculated using a Poisson likelihood, the same as T2K use for their oscillation analysis [98]. To avoid the calculation of a log-factorial, and to calculate a distribution that is related to the χ^2 distribution, the poisson maximum log-likelihood ratio is calculated instead of the absolute poisson log-likelihood. As such, the maximum log-likelihood ratio for sample j , \mathcal{L}_j , is calculated by using the observed events in the i^{th} energy bin from T2K, n_i , and the predicted events in the i^{th} energy bin from GAMBIT (including background), s_i , with the following formula:

$$\mathcal{L}_j = \sum_{i=0}^N n_i - s_i - n_i \ln(n_i/s_i), \quad (6.2)$$

where N is the total number of energy bins for the sample. The value of N was chosen based on the values used by T2K in their oscillation analyses and is equivalent to the number of energy bins for each sample in Fig. 6.2 [98]. When using the T2K likelihood

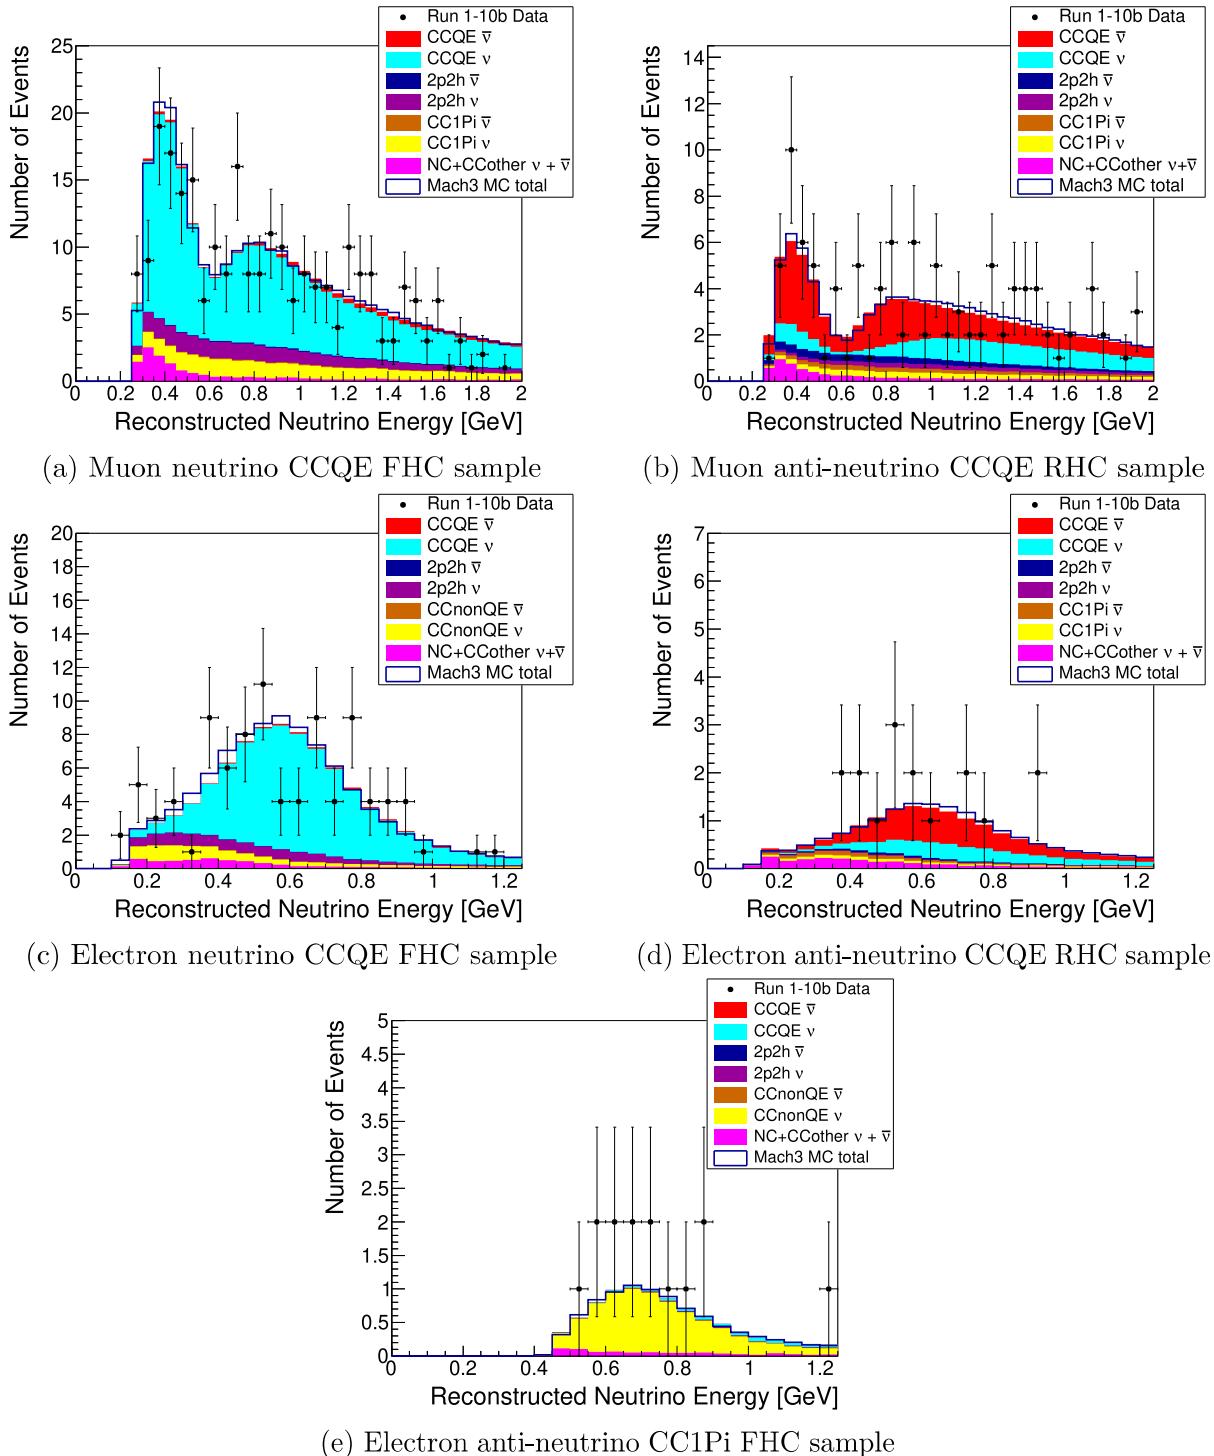


Figure 6.2: SK energy spectra of neutrino events as predicted by GAMBIT for the Asimov A parameter space point, broken down by interaction channel (coloured histograms). The run1-10 data from T2K (data points) and MC predictions from T2K's MaCh3 fitting software summed across all interaction channels (unfilled blue histogram) are also shown.

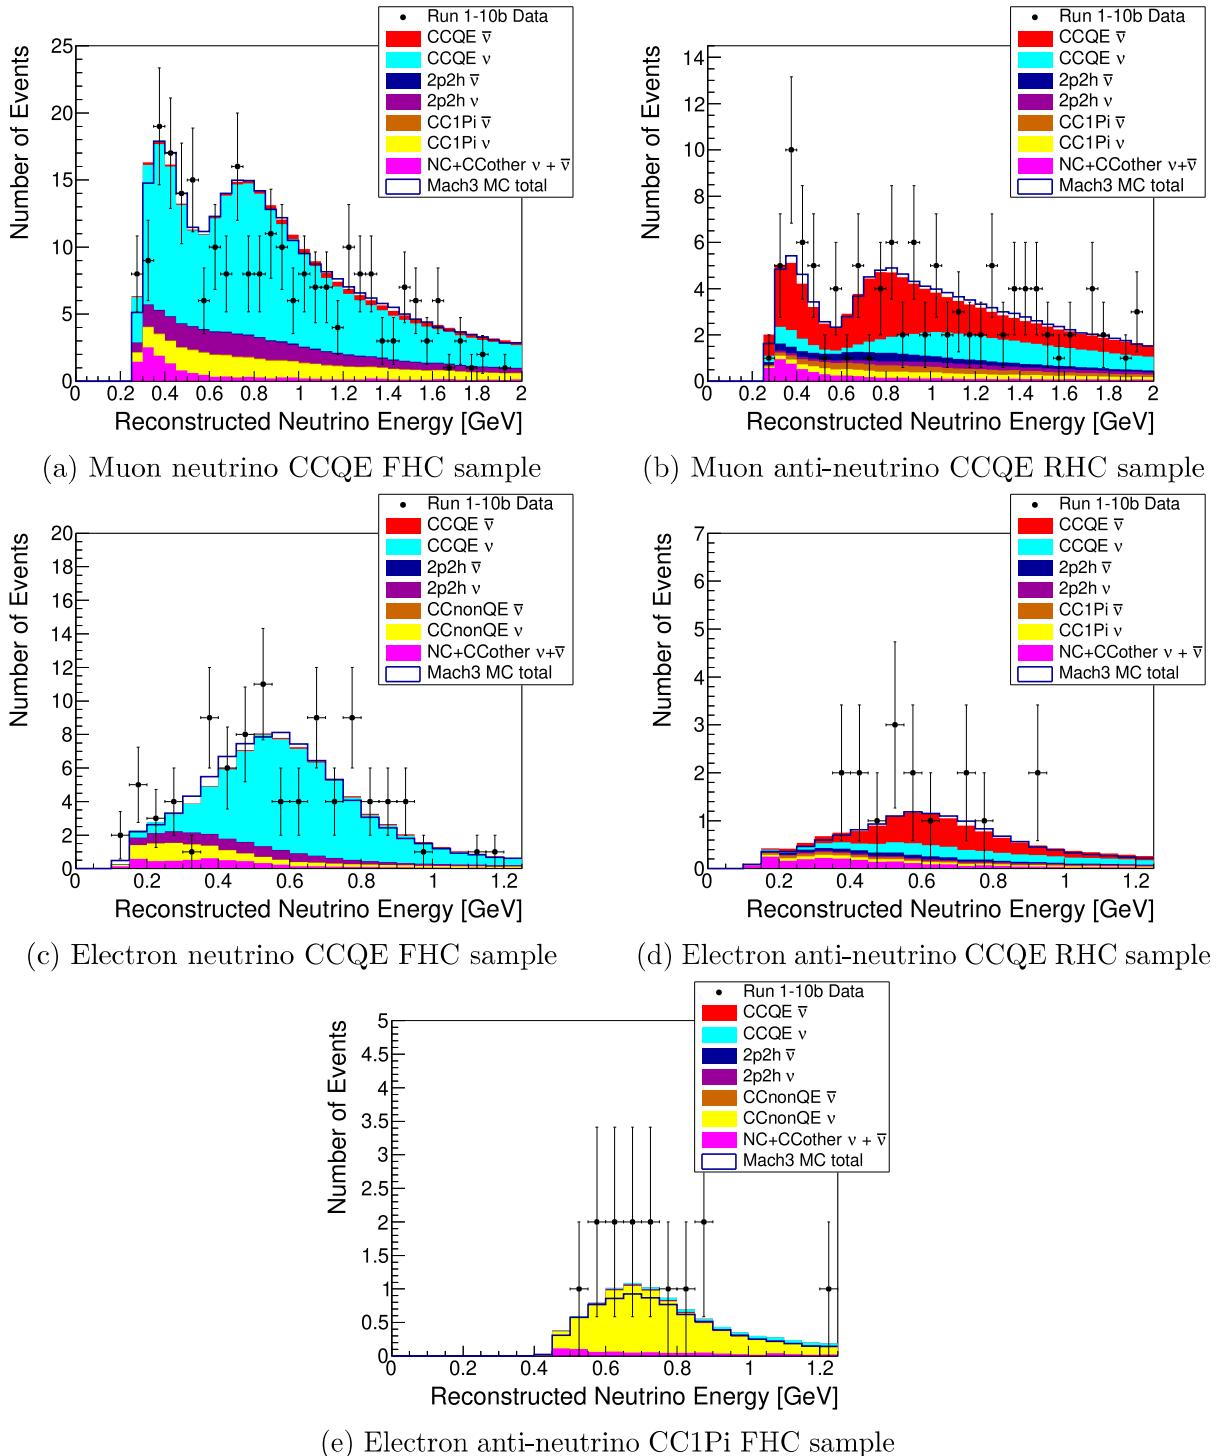


Figure 6.3: SK energy spectra of neutrino events as predicted by GAMBIT for the parameter set 1, broken down by interaction channel (coloured histograms). The run1-10 data from T2K (data points) and MC predictions from T2K's MaCh3 fitting software summed across all interaction channels (unfilled blue histogram) are also shown.

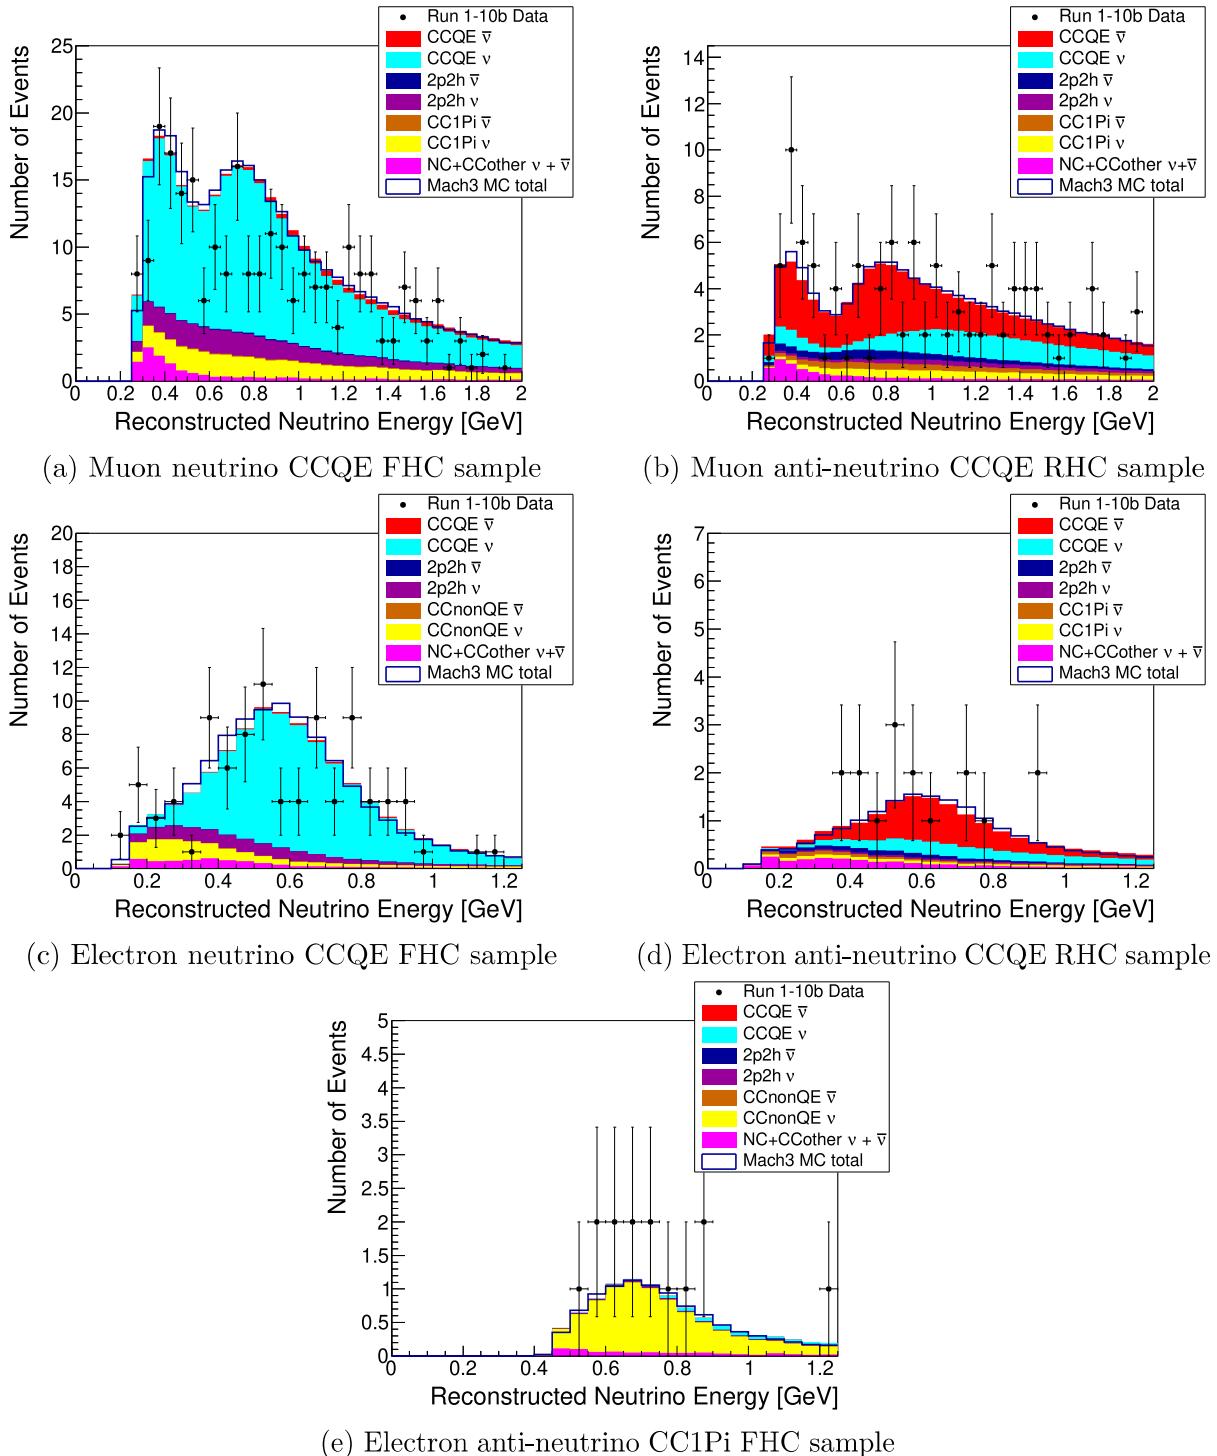


Figure 6.4: SK energy spectra of neutrino events as predicted by GAMBIT for the parameter set 2, broken down by interaction channel (coloured histograms). The run1-10 data from T2K (data points) and MC predictions from T2K's MaCh3 fitting software summed across all interaction channels (unfilled blue histogram) are also shown.

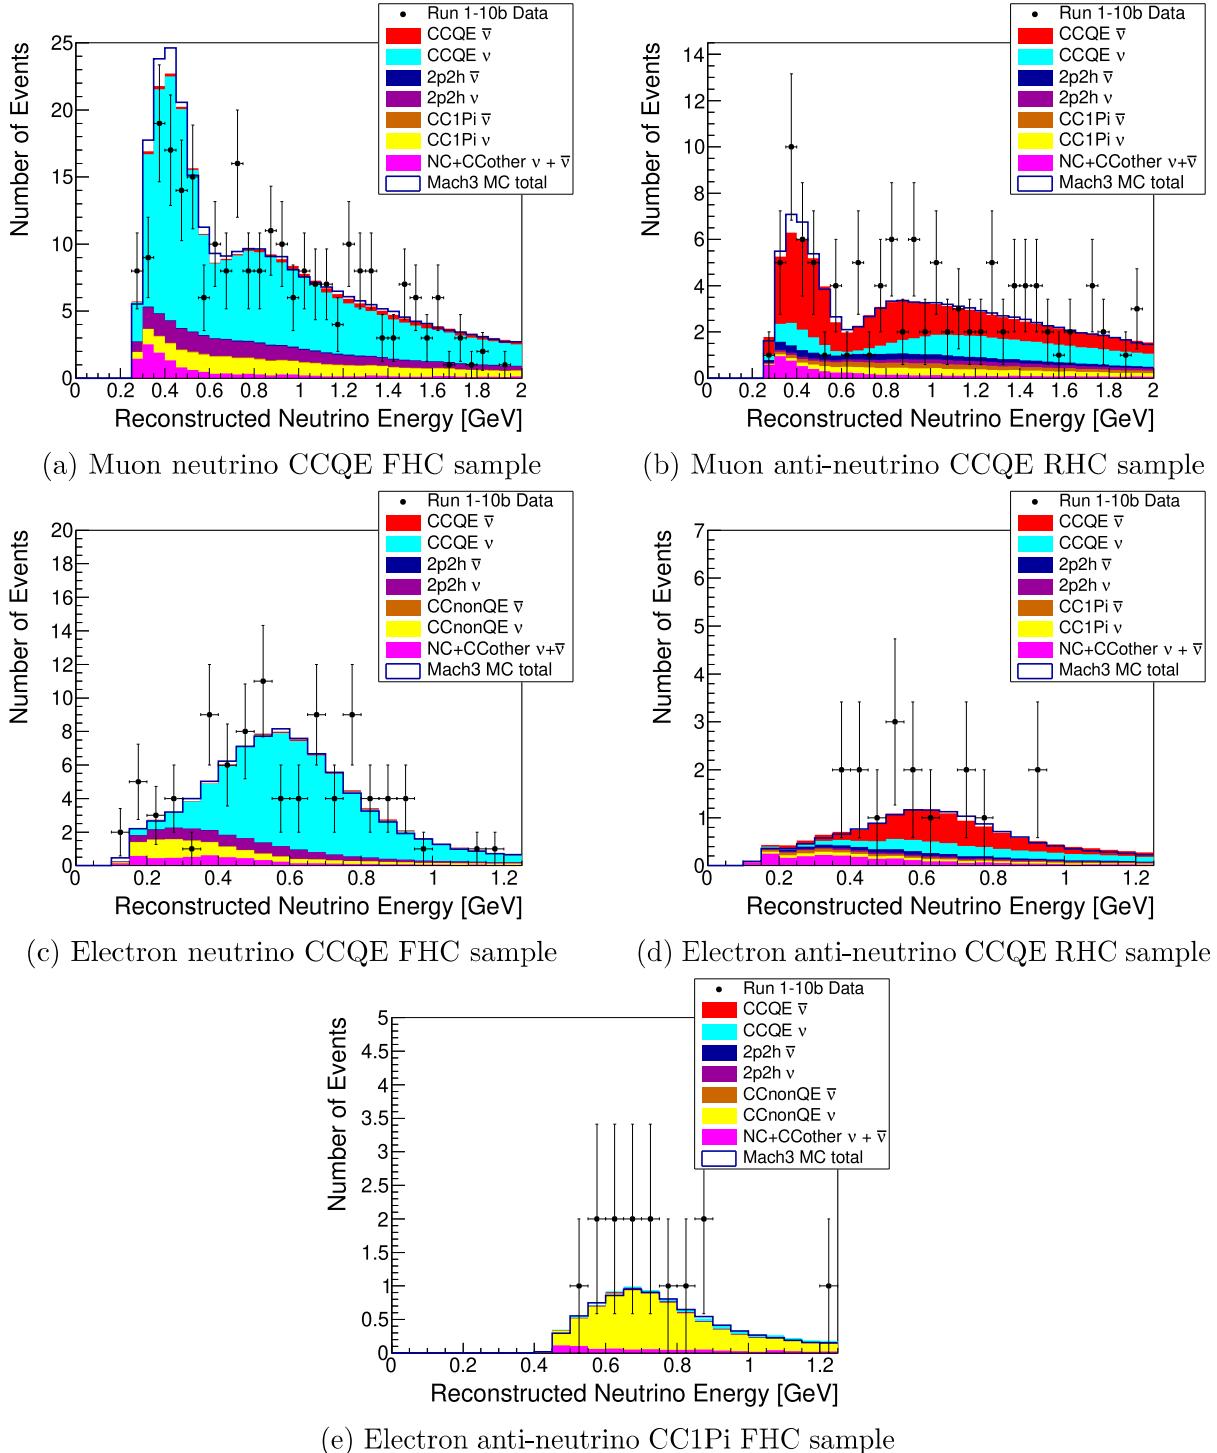


Figure 6.5: SK energy spectra of neutrino events as predicted by GAMBIT for the parameter set 3, broken down by interaction channel (coloured histograms). The run1-10 data from T2K (data points) and MC predictions from T2K's MaCh3 fitting software summed across all interaction channels (unfilled blue histogram) are also shown.

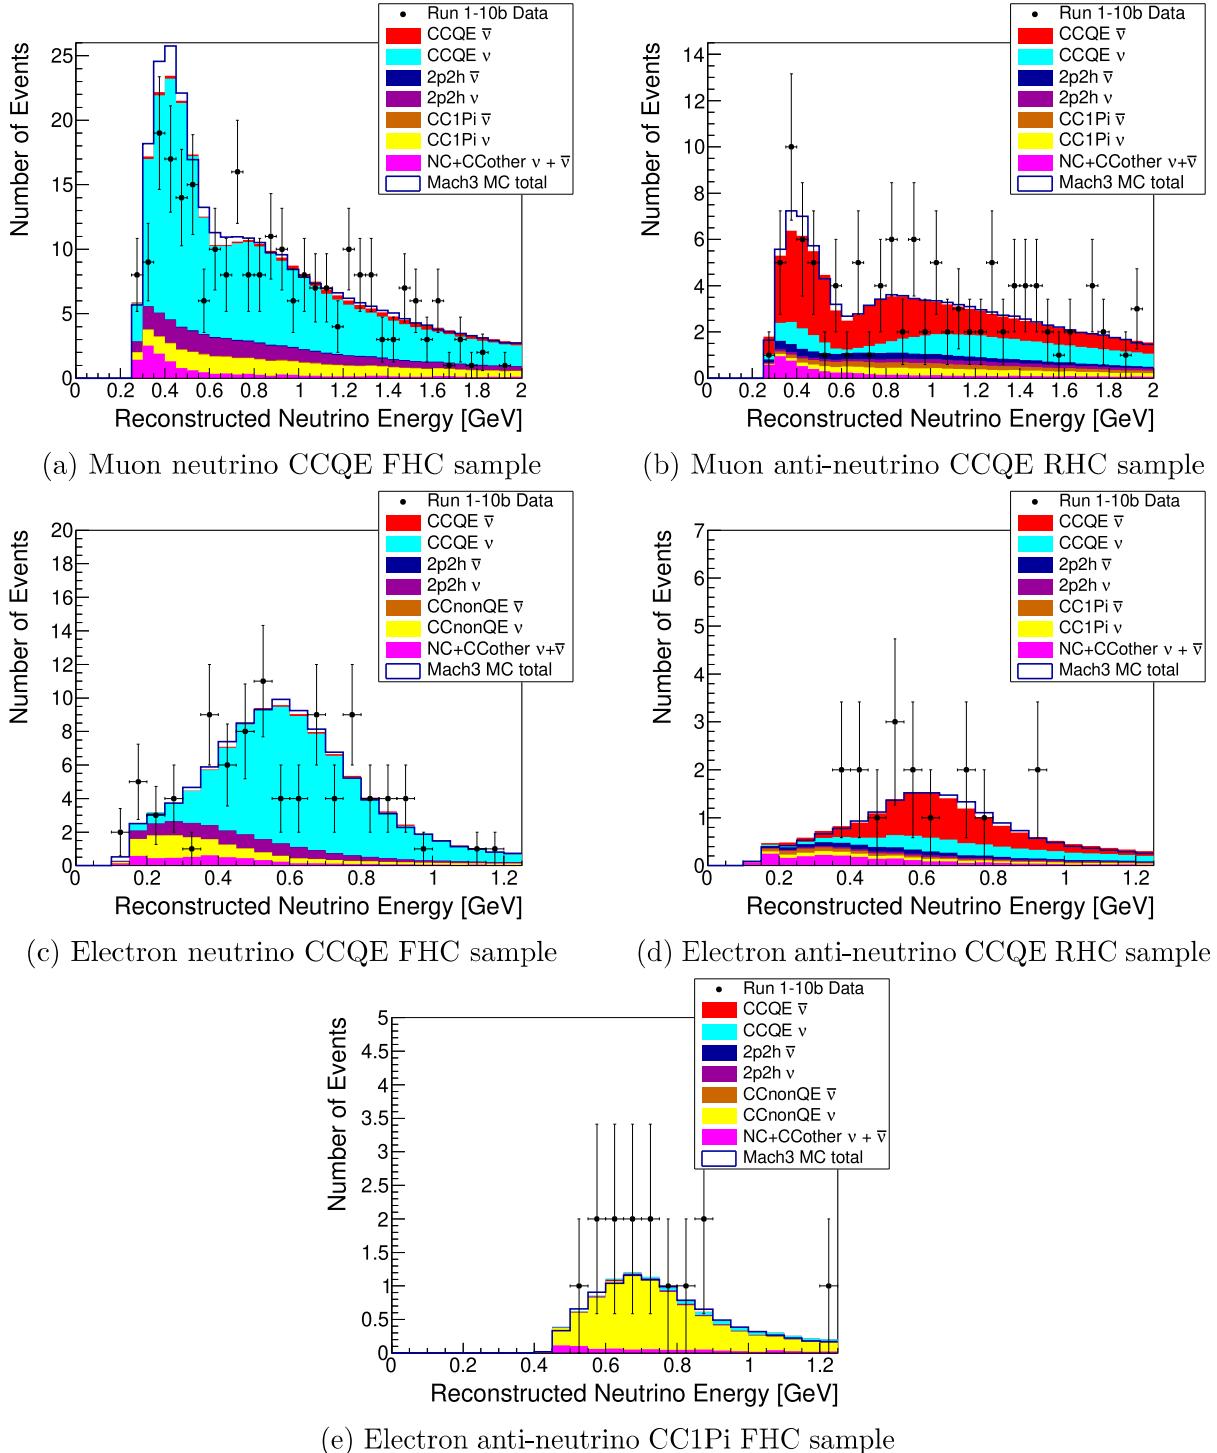


Figure 6.6: SK energy spectra of neutrino events as predicted by GAMBIT for the parameter set 4, broken down by interaction channel (coloured histograms). The run1-10 data from T2K (data points) and MC predictions from T2K's MaCh3 fitting software summed across all interaction channels (unfilled blue histogram) are also shown.

calculation to constrain parameters space, the values of \mathcal{L}_j can simply be summed—thanks to their logarithmic format—to give the total likelihood for the T2K experiment.

When developing the likelihood calculation for the T2K experiment within GAMBIT, several T2K resources were used. Given some developers of the GAMBIT code had access to internal T2K information, efforts were made to ensure only publicly available information was utilised in the GAMBIT code. Any internal information used was included with the explicit permission of the T2K Collaboration. All T2K resources used in GAMBIT and their sources are listed in Appendix C.1.

6.5 Including Systematic Uncertainties

T2K puts rigorous effort into accounting for all their experimental systematic uncertainties when analysing their data. These systematics can be categorised into flux, cross-section and detector systematics. The flux systematics stem from uncertainties in the flux models and simulations. The cross-section systematics are uncertainties both in the cross-section model parameters and the cross-section model itself, as there is no single agreed-upon model for neutrino cross-sections. The detector systematics include uncertainties in energy reconstruction, particle identification and momenta/energy measurements. T2K produce a covariance matrix containing all the systematics and their correlations between each other.

Because GAMBIT are calculating event predictions differently to T2K, it is difficult to incorporate some of these systematic values into the GAMBIT event predictions. Flux systematics are relatively simple, as they can be taken as the uncertainties in the BANFF flux weights found through the near-detector fit. Cross-section uncertainties are more troublesome, as the BANFF weights for cross-sections are applied directly to NEUT events rather than to the energy spectra itself. The cross-section BANFF weights simulate adjusting cross-section model parameters without having to rerun the whole MC. T2K implement these weights using internal software, which is not available to GAMBIT. Given time, GAMBIT may be able to build their own software that could apply the cross-section BANFF weights to their predictions. As of the time of writing, however, there is no such software written by GAMBIT, and so the cross-section uncertainties are not included in GAMBIT’s results. As seen from Tab. 3.1, the BANFF cross-section

systematic uncertainties are of the order of 10% when ND280 constraints are not included, so it is important to account for them in some way. GAMBIT have addressed this by adding normalisation uncertainties corresponding to the cross-section uncertainties without ND280 constraints in Tab. 3.1 to their energy spectra predictions in the T2K likelihood calculation.

Given GAMBIT’s method of simulating the neutrino energy reconstruction at SK, it is not possible to implement the detector systematics in GAMBIT the same as T2K do. However, an equivalent approach would be to treat the reconstruction function parameters used in GAMBIT as nuisance parameters, incorporating the uncertainty in the energy reconstruction values directly into GAMBIT results.

To summarise, GAMBIT have implemented flux, cross-section and detector systematic uncertainties into their T2K likelihood calculation. The cross-section uncertainties are included as normalisation uncertainties. The flux and detector systematics are incorporated by treating the flux BANFF weights and the reconstruction function parameters as nuisance parameters, which are varied about their nominal parameters when scanning across the oscillation parameter space. The nuisance parameters are constrained using Gaussian log-likelihood terms centred on the nominal values of the parameters. These log-likelihood terms are added to the final log-likelihood from the T2K likelihood calculation. Constraining the nuisance parameters allows penalisation of the parameters when they stray too far from their nominal values. Without these constraints, the scanning algorithms used by GAMBIT may acquire unphysical or unreasonable values of these parameters when scanning across the parameter space.

6.6 The T2K Likelihood Scanning Algorithm: Differential Evolution

GAMBIT includes a multitude of scanners capable of calculating likelihood values for many points across parameter space and converging to the point of highest likelihood [99]. For the T2K likelihood, a differential evolution scanner was used [100–103]. The scanner was implemented using a plugin named Diver. The specific algorithm used to create the results shown later in this thesis was a self-adaptive differential evolution algorithm,

named jDE, using rand/1/bin evolution [104].

We will not detail the rand/1/bin algorithm here, as it is already written up clearly by Martinez et al. [99]. It would be informative, however, to summarise the basic process of differential evolution. There are three main steps to differential evolution: mutation, crossover and selection. A set of vectors, the number of which is decided by the user, is randomly chosen inside the parameter space to be scanned. These target vectors are then shifted, or “mutated”, to produce a set of base vectors. These base vectors are derived from difference vectors between two target vectors and have a scale factor applied to them, which is chosen by the user. The mutation is based on the difference vectors between vectors within the target set. The crossover step then selects components of the base vectors and target vectors to produce trial vectors. The component selection is based on another user set value called the crossover rate. A selection is then made between the trial vectors and their corresponding target vector based on the likelihood value of the vectors. Whichever vector has the more desirable likelihood is chosen and used in the next generation of points.

Differential evolution is an ideal algorithm for multi-dimensional problems. This multi-dimensionality, as well as the improved mapping of likelihood contours over other algorithms, is why GAMBIT chose to use this for their neutrino global fit. A step further could have been taken by using a rand-to-best/1/bin evolution for quicker convergence, but in order to sufficiently sample the parameter space only the rand/1/bin evolution was used [99].

Based on the basic outline of differential evolution, one can comprehend that the scaling factors of the base vectors and the crossover rate are the two main parameters that dictate convergence speed and area exploration. For GAMBIT’s likelihood scans, we aim to produce contour plots for the oscillation parameter space, so the location of the global minima is not our only goal. Hence, parameter values are set such that the parameter space is explored sufficiently. Although, the jDE self-adaptive algorithm will evolve the values of the scaling factor and crossover rate with each generation of target vectors, but this still allows for suitable parameter space exploration.

Parameter	Range	Prior
$\sin^2(\theta_{13})$ w/ RC	[0.0173, 0.0270]	Gaussian(0.0218, 0.007)
$\sin^2(\theta_{13})$ w/o RC	[0, 0.0700]	Uniform
$\sin^2(\theta_{12})$	[0.282,0.338]	Gaussian(0.307, 0.013)
$\sin^2(\theta_{23})$	[0.360,0.640]	Uniform
Δm_{21}^2 [eV ²]	[7.25,7.81] $\times 10^{-5}$	Gaussian(7.53×10^{-5} , 0.18×10^{-5})
Δm_{32}^2 NO [eV ²]	[2.25,2.71] $\times 10^{-3}$	Uniform
Δm_{32}^2 IO [eV ²]	[-2.71,-2.25] $\times 10^{-3}$	Uniform
δ_{CP}	[- π , π]	Uniform

Table 6.3: The scanning ranges and prior constraints on the neutrino oscillation parameters used in the GAMBIT global fit. The numbers shown for Gaussian priors show the mean of the Gaussian and the standard deviation respectively.

6.7 Constraints on the Oscillation Parameter Space

Using the T2K likelihood calculation as described above and the diver scanning plugin, we can map the likelihood function of the T2K results across the oscillation parameter space. Such a scan reveals areas of the parameter space supported by T2K’s results. In our scans using GAMBIT, we have performed several scans with different configurations. In each configuration, we show results assuming NO, results assuming IO, results with the reactor constraints on θ_{13} and results without the reactor constraints. Normalisation uncertainties to account for cross-section systematics are also included in all the configurations. The different configurations are four-fold, one with no variation in systematic nuisance parameters such as flux weights and reconstruction function parameters (a similar approach to NuFit), one with a variation in flux systematics only, one with a variation in reconstruction parameters only and one with variations in flux systematics and reconstruction function parameters.

The ranges of oscillation parameters scanned over for each configuration are shown in Tab. 6.3. Also shown are the prior constraints applied to the likelihood calculation according to prior knowledge, for example from reactor constraints on θ_{13} .

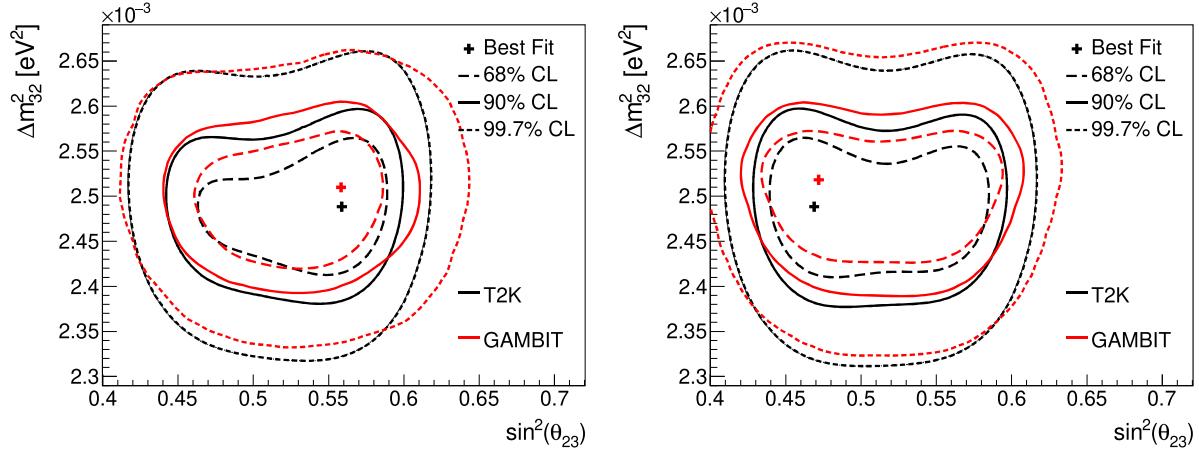


Figure 6.7: The run1-10 T2K oscillation analysis Δm_{32}^2 - $\sin^2(\theta_{23})$ contours compared to the GAMBIT contours when nuisance parameters are kept fixed. The NO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

6.7.1 No Variation of Systematic Nuisance Parameters

Figs. 6.7, 6.8, 6.9 and 6.10 show the resulting constraints on the oscillation parameter space from the GAMBIT T2K likelihood calculations when no variation of nuisance parameters is included. For comparison, the T2K constraints from their 2020 oscillation analysis using T2K runs 1-10 are also shown in each plot. The constraints are Frequentist and are shown in the form of confidence limits with regards to the true point on the parameter space. The plots show constraints for both the Δm_{32}^2 - $\sin^2(\theta_{23})$ parameter space and the δ_{CP} - $\sin^2(\theta_{13})$ parameter space, for both the NO and IO mass hierarchies and with and without the reactor constraint on θ_{13} . The reactor constraint is the same as used in the T2K 2020 analyses, which was sourced from [44].

The comparison between the GAMBIT constraints and the T2K constraints is satisfactory. The shape and magnitude of the contours for both GAMBIT and T2K are very similar, which attests to the improvements made in energy reconstruction and BANFF weight implementation. For the Δm_{32}^2 - $\sin^2(\theta_{23})$ results (Figs. 6.7 and 6.9), there is some offset or over-constraint with respect to the Δm_{32}^2 value for the GAMBIT results. This is most likely due to the energy reconstruction at SK still not being implemented to perfect accuracy, leading to mismatches between the GAMBIT energy spectra and T2K's MC. Such mismatches can be seen when comparing the GAMBIT energy spectra to T2K's MC

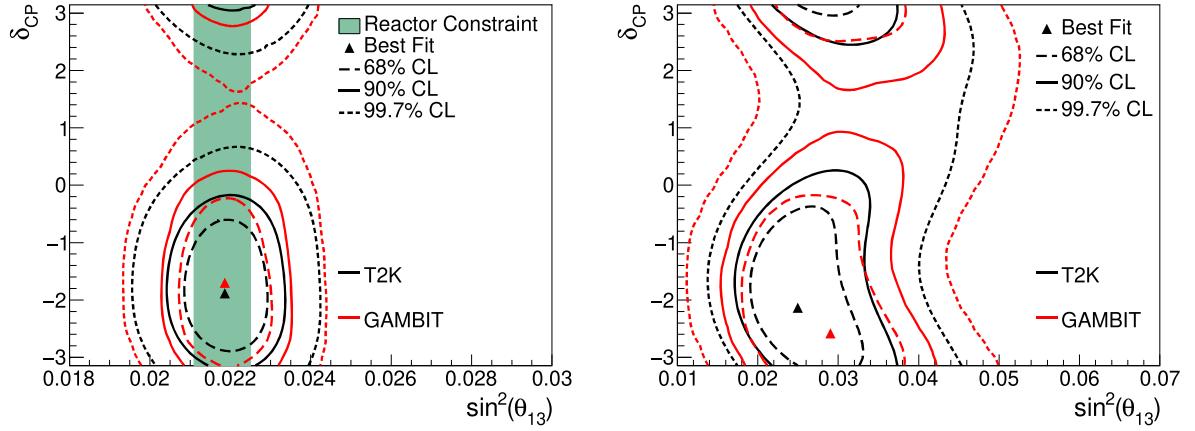


Figure 6.8: The run1-10 T2K oscillation analysis δ_{CP} - $\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when nuisance parameters are kept fixed. The NO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

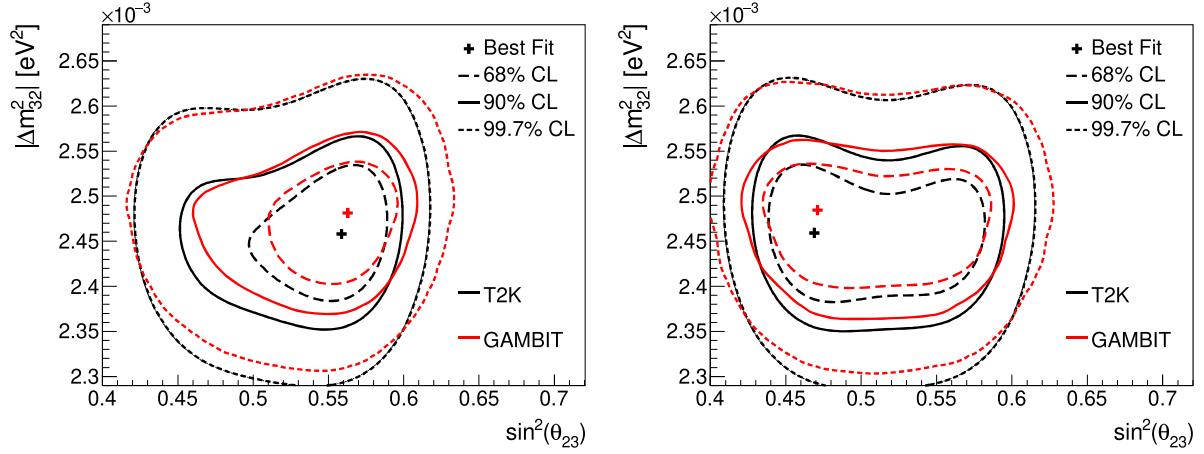


Figure 6.9: The run1-10 T2K oscillation analysis Δm_{32}^2 - $\sin^2(\theta_{23})$ contours compared to the GAMBIT contours when nuisance parameters are kept fixed. The IO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

in Figs. 6.5a and 6.6a, for example. In theory, these over-constraints seen in the GAMBIT contours may be alleviated when reconstruction function nuisance parameters are varied.

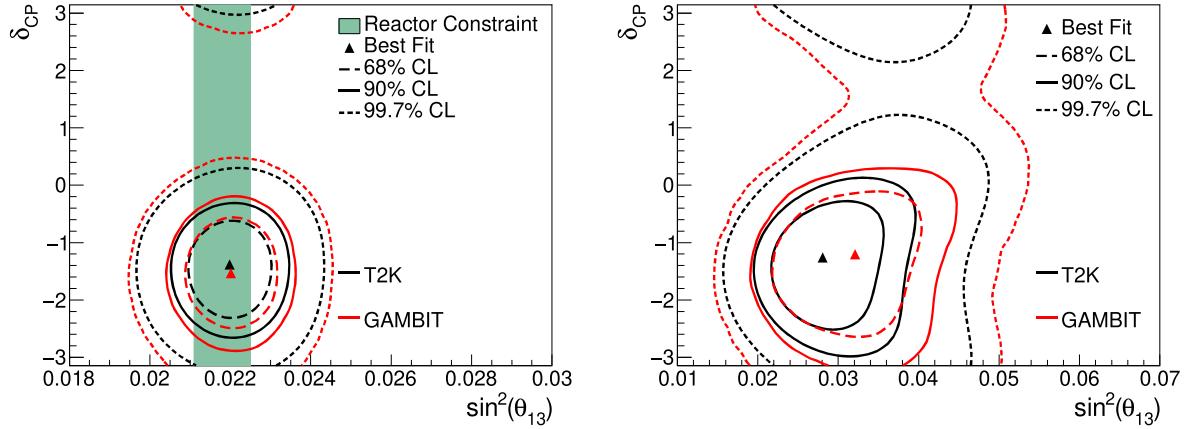


Figure 6.10: The run1-10 T2K oscillation analysis δ_{CP} - $\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when nuisance parameters are kept fixed. The IO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

6.7.2 Flux Parameters Variation

The flux BANFF uncertainties are included by adding BANFF weights as nuisance parameters in a parameter space scan and varying them across their uncertainty ranges. Ideally, these parameters would be varied independently, but there are 50 flux parameters in total, leading to increased run times for parameter space scans. The constraints shown in Figs. 6.11, 6.12, 6.13 and 6.14 were made utilising parameter space scans where all flux nuisance parameters were varied according to a single common scaling factor. This factor dictated the size of the uncertainty added to the nominal value of each parameter. For each flux parameter, the product of the scaling factor and the parameter's uncertainty value was taken and added to the nominal value. This scaling factor was varied between -1 and +1. This method of flux parameter variation only samples a small area of the flux parameter space. See Sec. 6.8 for potential improvements to this method in the future.

As seen from the figures, including the flux systematics causes the confidence limits to become more constrained compared to when all systematic nuisance parameters are fixed. This is due to improvements in the likelihood near the best fit point of the parameter space, but less significant improvements elsewhere, causing the likelihood ratio to decrease away from the best fit point. Usually when including systematic parameters, confidence limits become less constrained due to the extra freedom given when varying parameters.

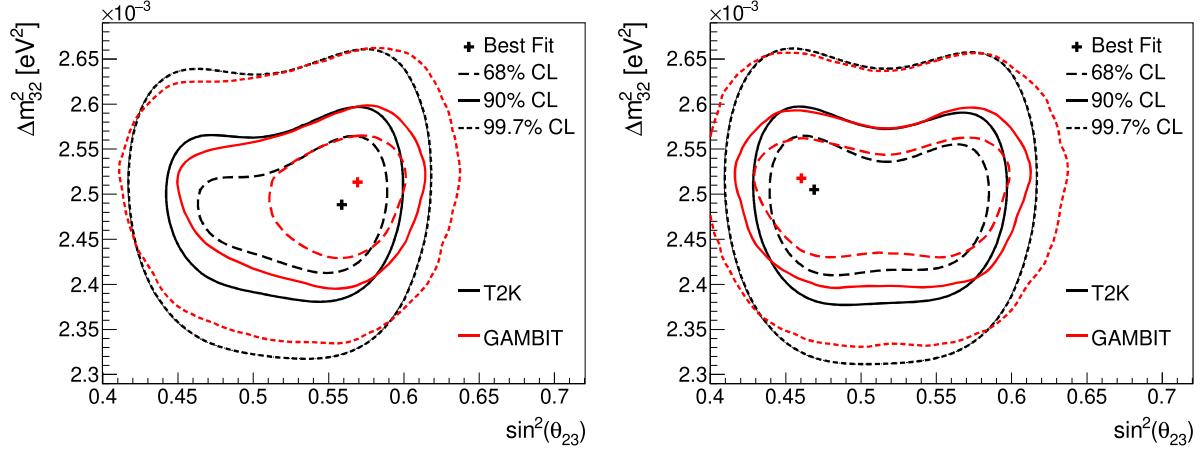


Figure 6.11: The run1-10 T2K oscillation analysis Δm_{32}^2 - $\sin^2(\theta_{23})$ contours compared to the GAMBIT contours when flux nuisance parameters are profiled over. The NO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

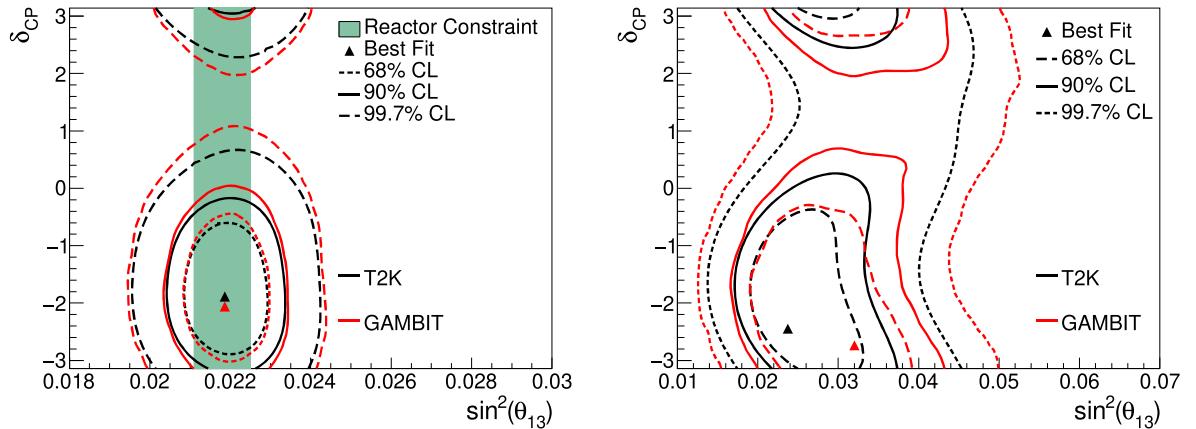


Figure 6.12: The run1-10 T2K oscillation analysis δ_{CP} - $\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when flux nuisance parameters are profiled over. The NO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

The observed behaviour could be blamed on the common scaling factor used to vary the flux parameters. This will be tested in future scans with alternative flux variation methods.

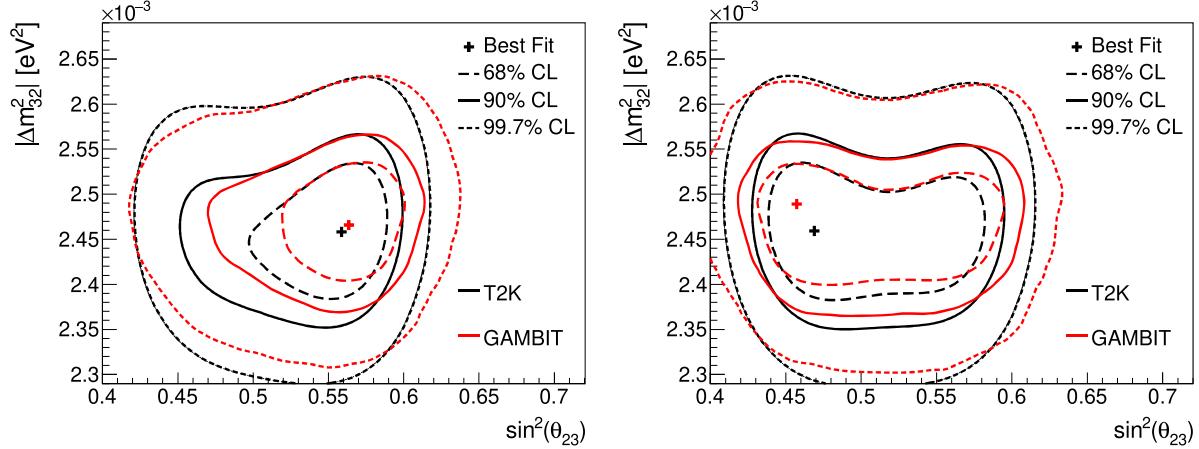


Figure 6.13: The run1-10 T2K oscillation analysis Δm_{32}^2 - $\sin^2(\theta_{23})$ contours compared to the GAMBIT contours when flux nuisance parameters are profiled over. The IO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

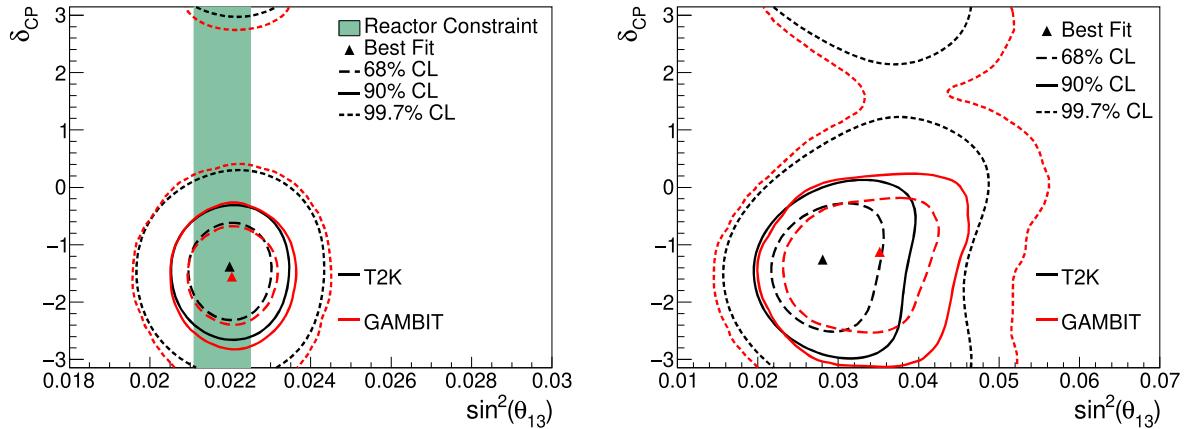


Figure 6.14: The run1-10 T2K oscillation analysis δ_{CP} - $\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when flux nuisance parameters are profiled over. The IO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

6.7.3 Reconstruction Function Parameters Variation

Parameter space constraints when reconstruction function parameters are varied as nuisance parameters are shown in Figs. 6.15, 6.16, 6.17 and 6.18. As with the flux parameters, the reconstruction function parameters were varied with a common scalar method. Considering there were almost 100 reconstruction function parameters in total,

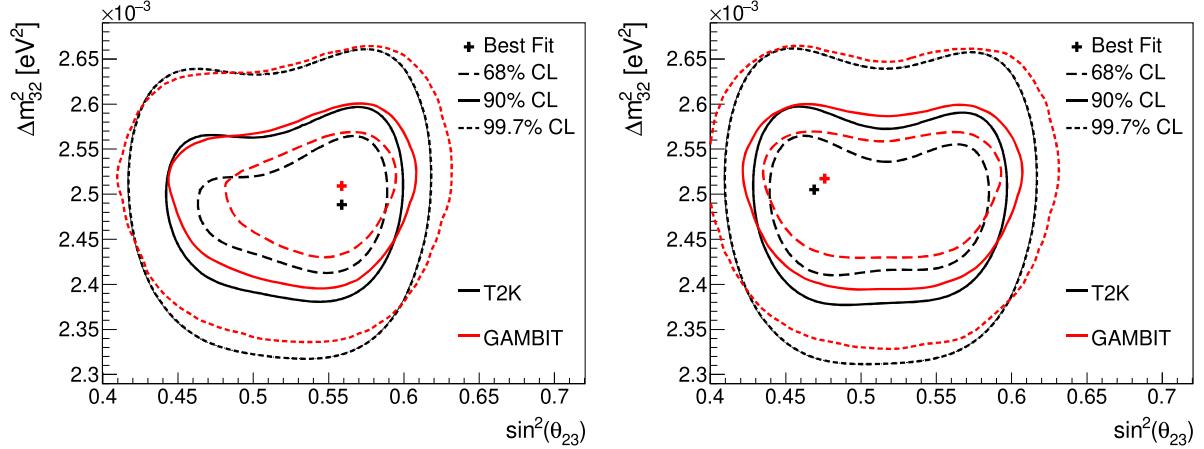


Figure 6.15: The run1-10 T2K oscillation analysis Δm_{32}^2 - $\sin^2(\theta_{23})$ contours compared to the GAMBIT contours when reconstruction function nuisance parameters are profiled over. The NO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

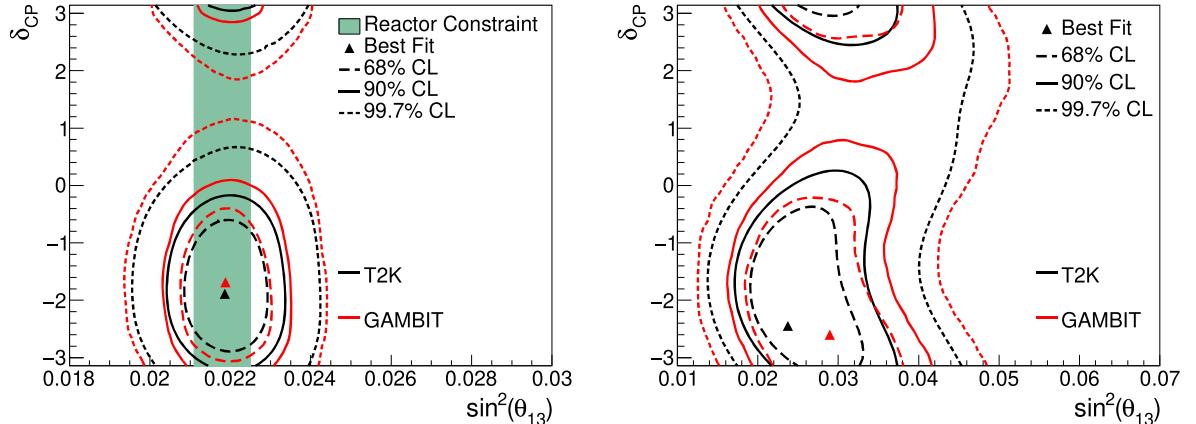


Figure 6.16: The run1-10 T2K oscillation analysis δ_{4p} - $\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when reconstruction function nuisance parameters are profiled over. The NO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

the common scaling method was necessary for convergence to be achieved in parameter space scans.

The results shown in the figures are more constrained compared to when there is no nuisance parameter variation. The constraint is not as significant as when flux parameters are the only varied nuisance parameter. Again, this behaviour could be connected to the

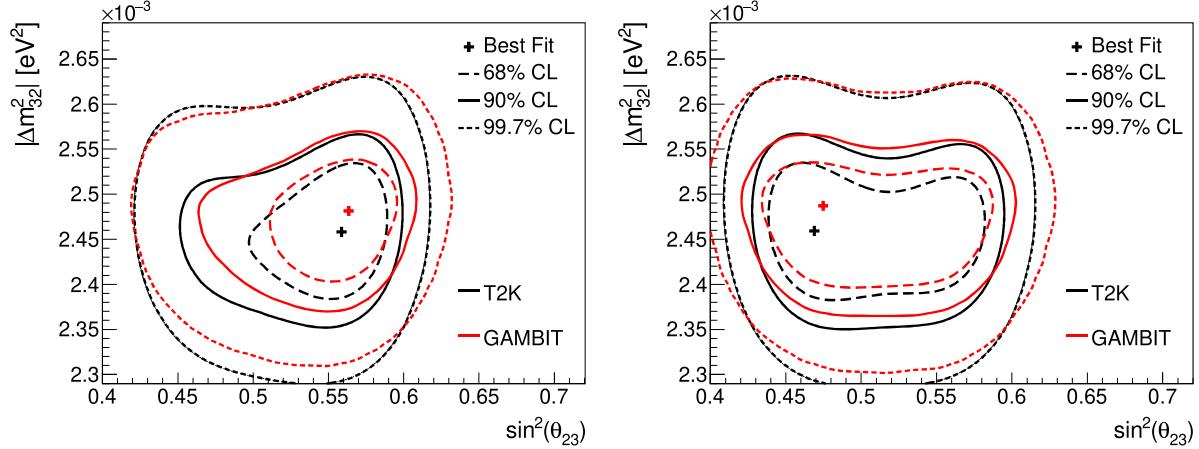


Figure 6.17: The run1-10 T2K oscillation analysis Δm_{32}^2 - $\sin^2(\theta_{23})$ contours compared to the GAMBIT contours when reconstruction function nuisance parameters are profiled over. The IO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

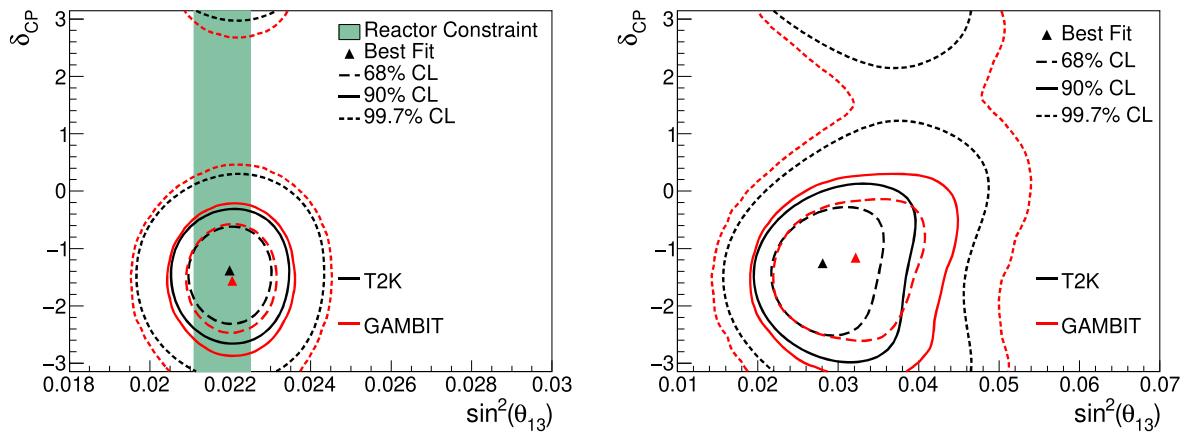


Figure 6.18: The run1-10 T2K oscillation analysis δ_{4p} - $\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when reconstruction function nuisance parameters are profiled over. The IO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

common scaling method used for the parameter variation.

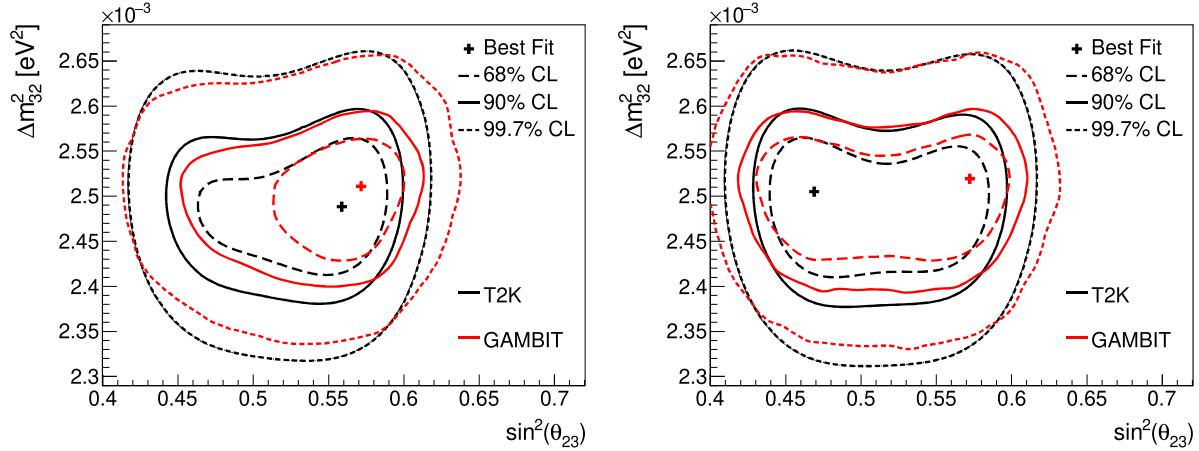


Figure 6.19: The run1-10 T2K oscillation analysis Δm_{32}^2 - $\sin^2(\theta_{23})$ contours compared to the GAMBIT contours when flux and reconstruction function nuisance parameters are profiled over. The NO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

6.7.4 Flux and Reconstruction Function Parameters Variation

Figs. 6.19, 6.20, 6.21 and 6.22 show parameter space constraints when all systematic uncertainties are included. That is, cross-section uncertainties, flux uncertainties and energy reconstruction uncertainties. The latter two uncertainties are included by common scaling of the nuisance parameters, with independent scaling factors for the flux and reconstruction function parameters.

Unsurprisingly, including both flux and reconstruction function uncertainties gives similar results to when the uncertainties are implemented individually. The parameter space constraints are narrowed, especially the 68% confidence limit for the Δm_{32}^2 - $\sin^2(\theta_{23})$ parameter space scans.

A comparison of the GAMBIT calculated contours when systematic uncertainties are kept constant and when both flux and reconstruction function parameters are varied is shown in Fig. 6.23. The contours shown use reactor constraints on θ_{13} and assume the NO mass hierarchy. The figure highlights that adding systematic uncertainties into the parameter space scan appears to further constrain the scan results.

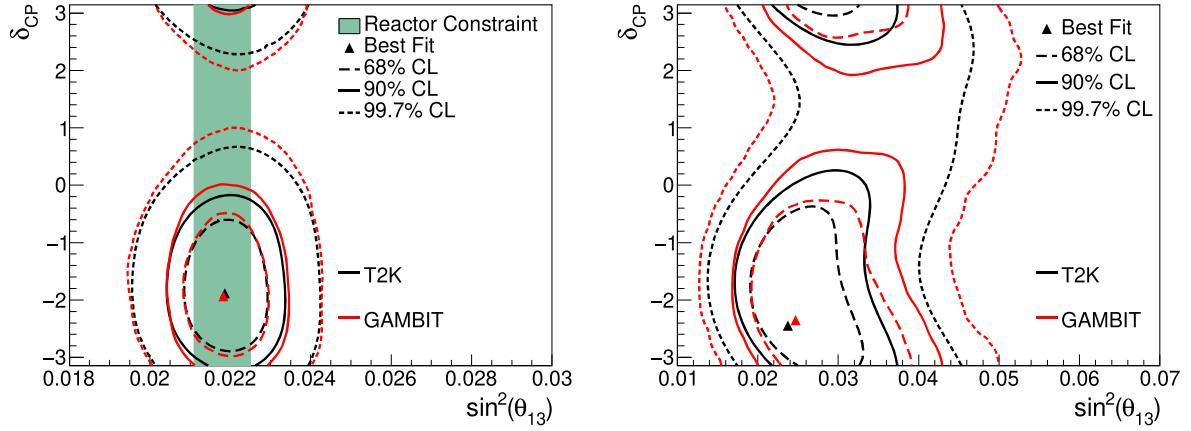


Figure 6.20: The run1-10 T2K oscillation analysis δ_{CP} - $\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when flux and reconstruction function nuisance parameters are profiled over. The NO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

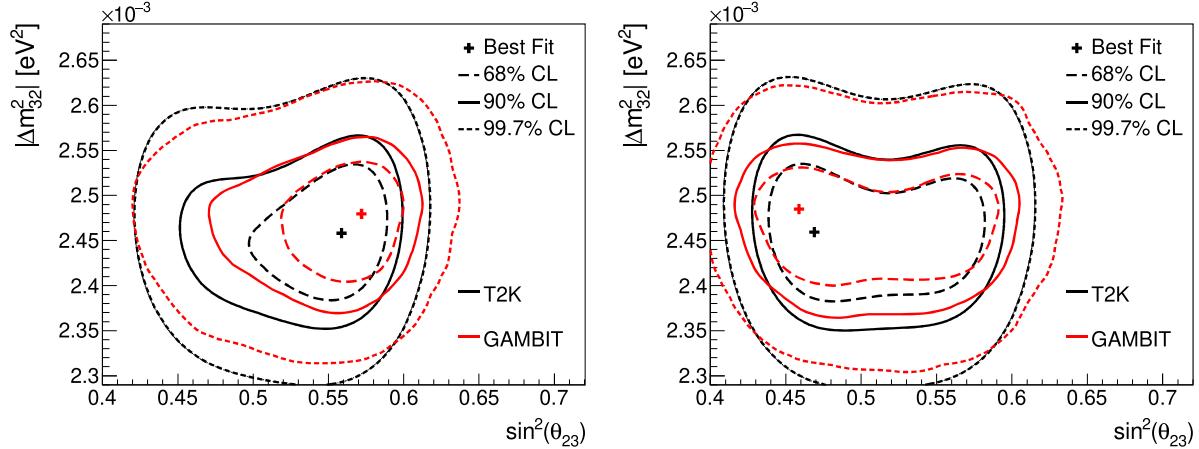


Figure 6.21: The run1-10 T2K oscillation analysis Δm_{32}^2 - $\sin^2(\theta_{23})$ contours compared to the GAMBIT contours when flux and reconstruction function nuisance parameters are profiled over. The IO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

6.8 Further Work

As seen from the previous sections, the T2K likelihood calculation is close to completion. Some improvements could be made in the implementation of systematic uncertainties. Currently, cross-section uncertainties are implemented only through normalisation uncertainties in predicted energy spectra. Ideally, these would be implemented directly to the

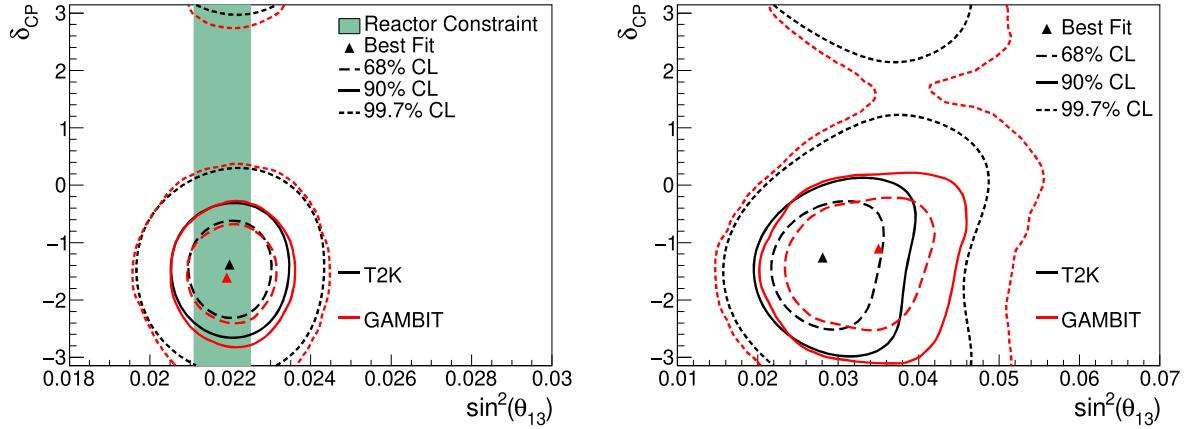


Figure 6.22: The run1-10 T2K oscillation analysis δ_{CP} - $\sin^2(\theta_{13})$ contours compared to the GAMBIT contours when flux and reconstruction function nuisance parameters are profiled over. The IO mass hierarchy is used. The separate plots show results with the reactor constraints to θ_{13} applied (left) and not applied (right).

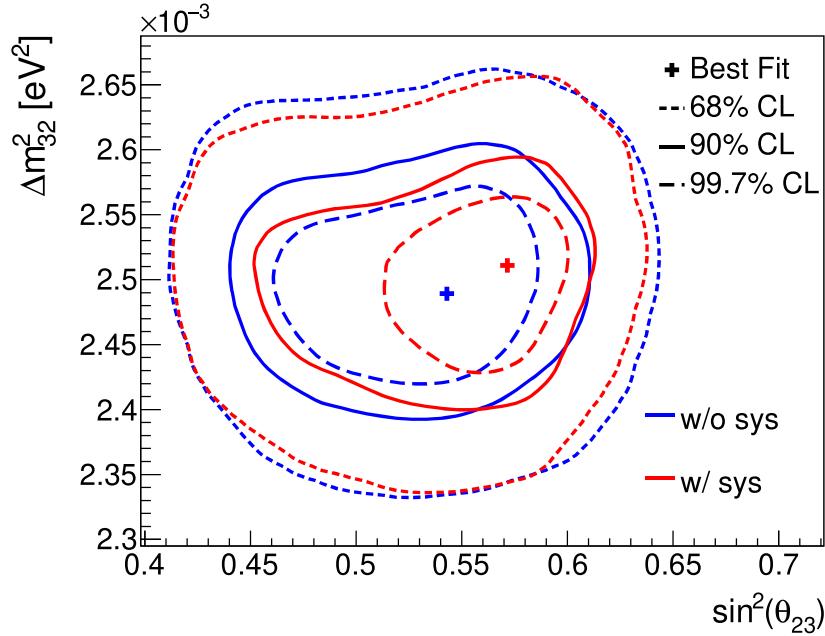


Figure 6.23: A comparison of the GAMBIT calculated contours when systematic uncertainties are left unvaried (blue) and when they are varied (red). Contours are for the NO mass hierarchy and use reactor constraints.

cross-sections by varying the cross-section model parameters. This is one possible goal for the future of the T2K likelihood calculation: for GAMBIT to find a way to re-weight the NEUT cross-sections using the BANFF parameters from the near-detector fit. Such

a task is ambitious and would require significant input from T2K collaborators. Yet, it would benefit GAMBIT’s global fit in the long run as T2K reduce uncertainties from their detectors and cross-section uncertainties become more important relative to other uncertainties.

All the flux BANFF weights and the reconstruction function parameters can be individually varied in parameter space scans as nuisance parameters according to their uncertainties. However, the large number of these parameters has led to the use of a common scaler to be varied rather than each individual parameter. This leaves a large area of the nuisance parameter space left unexplored. For the flux parameters, a potential improvement is to use a principal component analysis method. This would involve decomposing the BANFF covariance matrix into uncorrelated eigenbasis parameters. These eigenbasis parameters can be varied in scans and used to reconstruct the flux parameters. The impact of the eigenbasis parameters on the variation in BANFF parameters can be estimated by the relative size of their corresponding eigenvalues. As such, only a few of these eigenbasis parameters with large eigenvalues need be included in a scan of the parameter space. Including just 13 of the 97 BANFF parameters would cover 99% of the BANFF parameter space. This method of parameter variation has been implemented within GAMBIT. Unfortunately, the eigen libraries used in the method were returning inconsistent results when used on the Joliot Curie cluster used to calculate the parameter space scans shown here. The eigen calculations within GAMBIT have been shown to work on other systems, but these systems lack the computing power to return converged scans of the parameter space. Work is ongoing to return converged scans using the principal analysis method. With the increased sampling of the nuisance parameter space, it could significantly improve results with respect to how they compare to T2K’s oscillation analysis contours.

Work on likelihood calculations within GAMBIT for other experiments is on-going. Currently, the experiments being added are NO ν A, MINOS, Daya Bay, KamLAND and IceCube Deep Core. The full list of experiments to be included will reach far beyond these first few experiments, but the final list is not yet set in stone. For now, GAMBIT’s goal is to include likelihood contributions from at least these additional experiments: SNO, SK, Reno, Double Chooz and K2K. It is highly likely that GAMBIT will publish initial results with just a few experiments, for example just T2K and NO ν A.

The T2K calculation is being regarded as a template for the other calculations to follow.

Some hurdles are somewhat more difficult to overcome for other experiments, however, as sourcing information like fluxes and cross-sections can be somewhat more obtuse to achieve. MINOS, for example, do not explicitly publish fluxes for their far-detector, instead relying on near-detector fluxes being converted to a far-detector flux with a conversion factor. The NO ν A experiment have a special treatment of cross-sections, with a custom adjustment of the momentum transfer parameters for the $2p2h$ cross-section to better fit their MC to their observed data. The exact adjustments to the parameters aren't available publicly and we are liaising with NO ν A to see if we can get the required information to reproduce their cross-sections.

Once enough experiments are implemented into GAMBIT's global fit, the framework can easily be used to scan over the parameter spaces of BSM models other than the three-neutrino model. One example could be a sterile neutrino model such as the see-saw mechanism. Another goal once several experiments are complete is to establish a framework for a Bayesian analysis as opposed to the Frequentist studies presented in this thesis.

Of course, GAMBIT will also aim to update their neutrino global fit as experiments release new data and change their configuration over time, such as T2K, who will be releasing data from their T2K-II phase in the next few years. There are also the new up-and-coming experiments such as DUNE and Hyper-Kamiokande that can be added once their data is published.

6.9 Next Generation Experiments

Given the imminent new phase of neutrino experiments, with detectors of as yet unseen volume and sensitivity and the extended running periods of current experiments, the future is bright and brings hope of significant discovery. Global fits will be an invaluable tool for combining the data of these new experiments.

We can use the T2K likelihood calculation framework currently in GAMBIT to indicate the augmented sensitivity of T2K after the T2K-II phase.

6.9.1 Impact of T2K-II

In the T2K likelihood scans shown in Sec. 6.7, flux, cross-section and detector systematics are implemented. By scaling these systematics appropriately, so that they correspond to the predicted systematics after the ND280 upgrade, we can predict the magnitude of the oscillation parameter space contours after T2K-II.

For this sensitivity study, we produced a Asimov datasets using the GAMBIT event predictions as described in Sec. 6.4.

Two datasets in total were produced, a “current” dataset and an “upgraded” dataset. The current dataset was created with current POT values. The upgraded dataset corresponded to data after T2K-II, with the predicted POT of T2K-II. The predicted POT for T2K-II is 20×10^{21} POT as a sum of FHC and RHC beams. We assumed an equal split between FHC and RHC beam time, giving 10×10^{21} POT for both beam types.

The oscillation parameters used to produce both datasets were the same as the Asimov A parameter point shown in Tab. 6.1, except the value of δ_{CP} was set to a maximally CP violating value of -90° .

When scanning over the two datasets to extract confidence limits, differences in systematic uncertainties were accounted for where possible. For both datasets, the nominal flux systematic parameters were set equal to the nominal values from T2K’s 2020 oscillation analysis [105]. The uncertainties in these values were implemented differently, however. For the current dataset, the uncertainties for the flux parameters were implemented as described for the T2K 2020 oscillation analysis. For the upgraded dataset, the uncertainties used were 77% of the ones quoted in T2K’s oscillation analysis. This was extracted from T2K’s sensitivity studies for the ND280 upgrade [59]. The cross-section and detector systematics, implemented as normalisation uncertainties and reconstruction function uncertainties in GAMBIT, were kept the same for both datasets. Both sources of uncertainty would be expected to improve in T2K-II, but due to the implementation of these in GAMBIT, it is not applicable to change these for the upgraded dataset.

Confidence limits derived from both the current and upgraded datasets can be seen in Fig. 6.24. The reactor constraint on θ_{13} is applied and the NO mass hierarchy is assumed. Not only do the confidence limits become more constrained for the upgraded dataset,

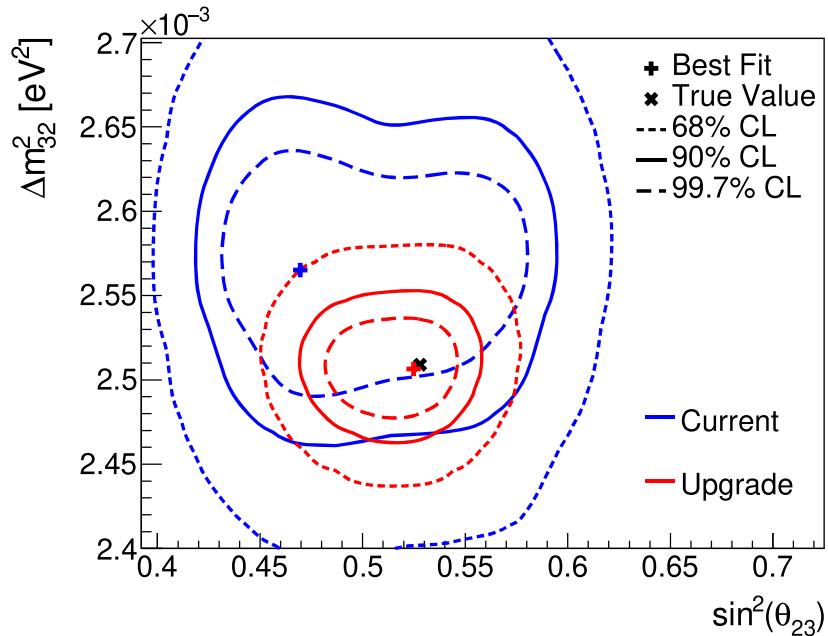


Figure 6.24: GAMBIT’s predictions for T2K’s parameter space confidence limits with current (blue) and upgraded (red) POT values and systematic uncertainties. The upgraded values correspond to values predicted at the end of T2K-II. The parameter space point used to simulate both datasets is also shown.

the best fit point is also much closer to the true value used for the mock datasets. This demonstrates the appreciable impact T2K-II could have on neutrino oscillation parameter measurements.

One of the most anticipated measurements from T2K-II is the improved constraint on δ_{CP} . In Fig. 6.25, the predicted confidence limits are shown for both the current experiment configuration and POT and the T2K-II configuration and POT. These predictions were made with the same datasets used for Fig. 6.24. The predictions for the upgraded dataset show that $\delta_{\text{CP}} = 0$ can be excluded at 3σ , if neutrino oscillation is maximally CP violating. There is also the potential for a 5σ measurement of the value of δ_{CP} , unlike with the current configuration and POT of T2K. The extension to T2K’s runtime and the upgraded detectors, including the SuperFGD, will clearly have a significant impact on the neutrino oscillation field.

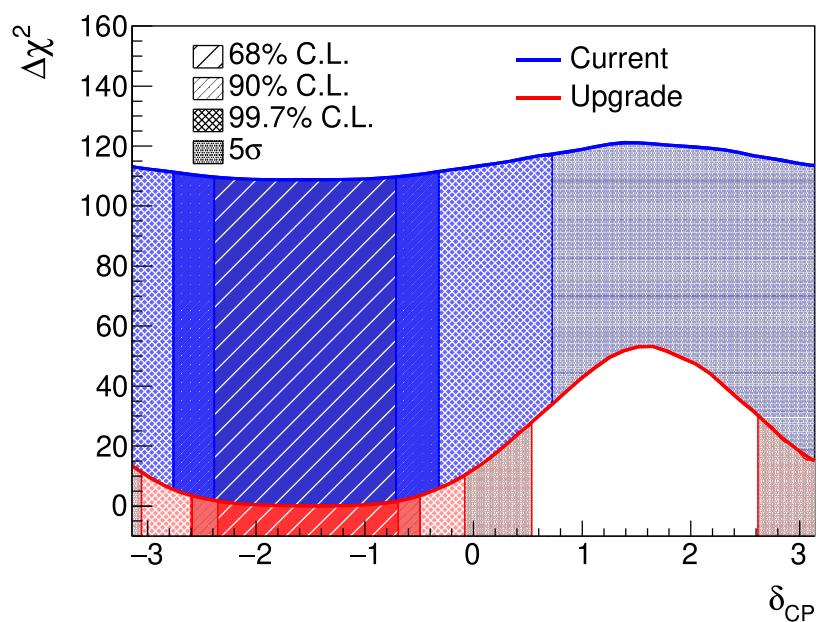


Figure 6.25: GAMBIT's confidence limits for the δ_{CP} oscillation parameter using the current (blue) and upgraded (red) POT values and systematic uncertainties. The upgraded values correspond to values predicted at the end of T2K-II.

Chapter 7

Conclusion and Remarks

In this thesis, several studies of the SuperFGD Prototype were shown. The SuperFGD is a brand new innovative design of scintillator detector able to track charged particles travelling in any direction. The testing of the prototype has given invaluable information for the construction, operation and simulation for the SuperFGD to be installed in ND280 in 2022. The installation is part of the proposed ND280 upgrade, which aims to improve systematic uncertainties in the T2K experiment from 6% to 4%. The detector upgrades are to be used in the extended run of T2K, named T2K-II. This extended run will increase the POT of T2K nearly ten-fold to 20×10^{21} POT.

The SuperFGD Prototype studies included signal threshold measurements, optical crosstalk characterisation, light attenuation modelling and time resolution measurements, to name a few. Preliminary studies of neutron data and the neutron cross-section as a function of energy were presented, with suggested improvements for further studies. The low energy threshold of the SuperFGD design was demonstrated, with neutrons with energies as low as 10 MeV producing signals in the detector. There is even the potential to observe lower energy neutrons in the analysis of the 2020 LANL beam test data.

Also shown in this thesis were the beginnings of a neutrino oscillation data global fit within the BSM global fitting software GAMBIT. The implementation for the T2K experiment within the global fit is all but complete, pending some possible minor alterations to energy spectra predictions. Likelihood calculations for other experiments are in-progress, these experiments being NO ν A, MINOS, Daya Bay and IceCube DeepCore and KamLAND.

The T2K likelihood calculation within GAMBIT seeks to surpass previous implementations of global fits by utilising official data sources rather than theses and presentations. It also benefits from the experience and knowledge of T2K members. All information used in the calculation was extracted from public sources except the reconstruction function

parameters, which were estimated using fits to internal MC energy spectra from several points on the parameter space. The use of these MC energy spectra have not been approved for use within GAMBIT, but could potentially be released publicly later in 2022 when T2K review their approach to presenting data and information. It is the author's hope that the analyses shown in this thesis using the internally sourced information will go some way to allowing the data to be used within GAMBIT and in future global fits.

Studies using GAMBIT's T2K likelihood calculation give similar oscillation parameter space constraints to T2K's oscillation analysis. There is some deviation from T2K's constraints when nuisance parameters, such as the flux corrections and reconstruction function parameters, are allowed to vary in scans across the parameter space. Further studies into why the confidence limits become more constrained when these systematic uncertainties are added would be necessary before publishing any neutrino global fits from GAMBIT including the T2K experiment.

Finally, the impact of the SuperFGD and T2K-II as a whole was examined using the T2K likelihood calculation within GAMBIT. Confidence limits on the parameter space for both the current T2K setup and beam time and the upgraded setup and beam time were shown. The benefits of the upgrade were particularly clear in the increased constraints on the $\Delta m_{32}^2 - \sin^2(\theta_{23})$ parameter space and the potential 5σ exclusion of a zero δ_{CP} value.

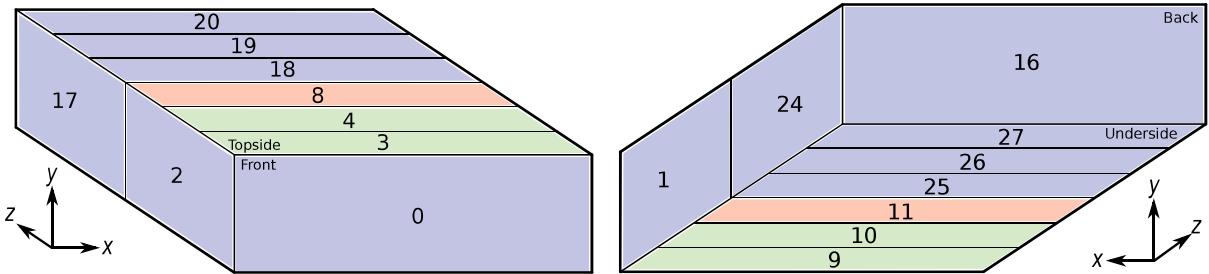


Figure A.1: The mapping of FEB numbers to mounted MPPCs. The number inside each rectangular region on the detector surface indicates the number of the FEB that reads signals from MPPCs mounted to that surface. Colours represent the MPPC type: blue for Type I, green for Type II and red for Type III. Two perspectives of the detector are shown so that all surfaces are visible. The left image shows the detector from the front with a top-down view. The right image shows the back of the detector looking upwards.

Appendix A

SuperFGD Prototype FEB and Channel Mapping

Many analyses rely on FEB and channel numbering to associate signals with an area or channel of the SuperFGD Prototype. A single FEB reads signals from 96 MPPCs distributed across the surface of the detector. The readout cabling of the 96 MPPCs are split into three cable bundles with 32 channels each. Each cable bundle is connected to one of the CITIROCs mounted in the FEB. Fig. A.1 shows mapping of FEB numbers to areas of the detector surface. The indicated surface is where all the MPPCs connected to the FEB are mounted. There were 18 FEBs dedicated to reading MPPCs, numbered 0–4, 8–11, 16–20, 24–27. The numbering is admittedly confusing, being a relic of the readout electronics development.

Fig. A.2 shows the mapping of individual channels to the detector surface. More precisely, cable bundle numbers are mapped to the surface of the detector and the channel numbering within each cable bundle is shown that users can translate to an individual

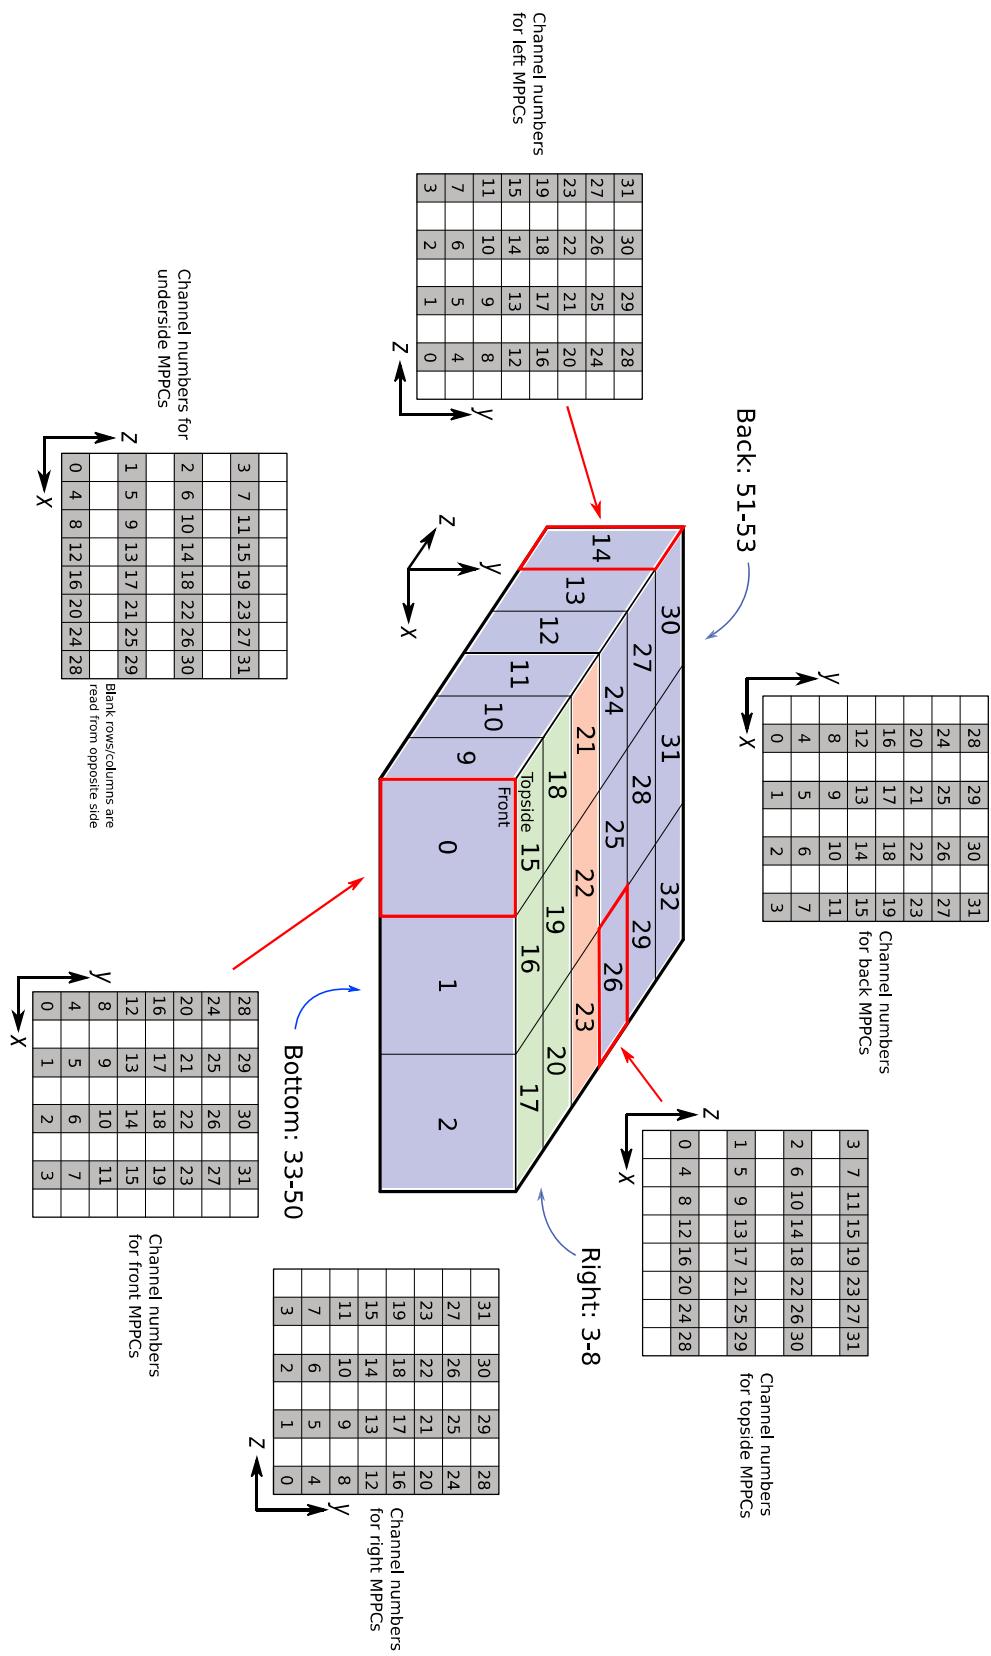


Figure A.2: The mapping of cable bundles to the detector surface and the order of channel numbers associated with each cable bundle. Numbers on the diagram of the detector indicate the cable bundle number that reads out MPPCs on that area of the detector. The grey and white squares represent cubes on the outer-surface of the detector. Grey squares indicate an MPPC is present on the surface and white squares indicate there is no MPPC, instead the MPPCs are mounted on the opposite side of the detector. The numbers on grey squares indicate the channel number within the cable bundle.

MPPC on the detector surface.

FEB	DAC10b threshold	DAC10b_t threshold	HG shaper time constant	FPGA ...ASIC 0	ADC threshold ...ASIC 1	HG for... ...ASIC 2	FPGA threshold	ADC LG	MPPC Type
0	250	250	3	80	80	80	1		I
1	250	250	3	70	80	80	1		I
2	250	250	3	80	80	80	1		I
3	250	250	3	170	90	70	1		II
4	250	250	3	90	90	180	1		II
8	250	250	1	220	170	170	1		III
9	250	250	3	80	70	110	1		II
10	250	250	3	80	80	120	1		II
11	250	250	1	220	200	170	1		III
16	250	250	3	80	60	80	1		I
17	250	250	3	150	100	150	1		I
18	250	250	3	150	80	150	1		I
19	250	250	3	80	130	160	1		I
20	250	250	3	150	150	80	1		I
24	250	250	3	150	140	140	1		I
25	250	250	3	140	110	170	1		I
26	250	250	3	90	80	150	1		I
27	250	250	3	90	80	90	1		I

Table B.1: The DAQ settings used for the CERN and LANL particle beam tests.

Appendix B

SuperFGD Prototype DAQ Settings

Table B.1 shows the DAQ settings used for the majority of the runs for the CERN and LANL beam tests for each FEB.

The units for each value are arbitrary and correspond with the numbers input into the DAQ settings console for the SuperFGD Prototype. As such, these numbers are more for reference to future users of the SuperFGD Prototype than to readers learning about the SuperFGD.

The DAC10b threshold is the threshold at which data acquisition starts if the signal amplitude rises above the set value.

The DAC10b_t threshold is the threshold at which the rising (falling) edge time is recorded if the signal rises above (falls below) the set value. As the DAC10b_t threshold was equal to the DAC10b threshold for all FEBs, there was no distinction between these threshold throughout the beam tests.

The HG shaper time constant affects the shape output by the HG slow-shaper.

The FPGA ADC thresholds determine the zero suppression levels applied by the FPGA for each CITIROC ASIC. The higher the value, the higher the minimum amplitude needed for a data hit to be retained. These values are relatively high for the Type III MPPCs which gave frequent dark counts.

Appendix C

GAMBIT T2K Likelihood Reconstruction Functions

Both Gaussian and log-normal distributions were used in the reconstruction functions for the T2K likelihood calculation in GAMBIT. Some interaction channels were also given different reconstruction function parameters depending on the true energy region.

Due to the integration over the reconstruction function with respect to the reconstructed neutrino energy, as shown in Eq. (6.1), the reconstruction function was implemented in the likelihood calculation as analytic solutions of the integration. For the Gaussian R function, R_G , this solution looked like

$$\begin{aligned} \int_{E_i}^{E_{i+1}} R_G(E_{\text{rec}}, E_{\nu}) dE_{\text{rec}} &= \int_{E_i}^{E_{i+1}} dE_{\text{rec}} \frac{1}{\sigma_G \sqrt{2\pi}} e^{-\frac{1}{2} \left(\frac{E_{\text{rec}} - (E_{\nu} + \Delta_G)}{\sigma_G} \right)^2} \\ &= \frac{1}{2} \left[\text{erf} \left(\frac{E_{i+1} - (E_{\nu} + \Delta_G)}{\sqrt{2}\sigma_G} \right) - \text{erf} \left(\frac{E_i - (E_{\nu} + \Delta_G)}{\sqrt{2}\sigma_G} \right) \right], \end{aligned} \quad (\text{C.1})$$

where σ_G is the corresponding width of the Gaussian and Δ_G is the offset of the Gaussian with respect to the true neutrino energy.

For the log-normal R function, R_{LN} , the analytical solution was

$$\begin{aligned} \int_{E_i}^{E_{i+1}} R_{\text{LN}}(E_{\text{rec}}, E_{\nu}) dE_{\text{rec}} &= \int_{E_i}^{E_{i+1}} dE_{\text{rec}} \frac{1}{E_{\text{rec}} \sigma_{\text{LN}} \sqrt{2\pi}} e^{-\frac{1}{2} \left(\frac{\ln E_{\text{rec}} - \mu_{\text{LN}}(E_{\nu}, \sigma_{\text{LN}})}{\sigma_{\text{LN}}} \right)^2} \\ &= \frac{1}{2} \left[\text{erf} \left(\frac{\ln E_{i+1} - \mu_{\text{LN}}(E_{\nu}, \sigma_{\text{LN}})}{\sqrt{2}\sigma_{\text{LN}}} \right) - \text{erf} \left(\frac{\ln E_i - \mu_{\text{LN}}(E_{\nu}, \sigma_{\text{LN}})}{\sqrt{2}\sigma_{\text{LN}}} \right) \right], \end{aligned} \quad (\text{C.2})$$

where σ_{LN} is the width of the log-normal and μ_{LN} is the mean of a normal distribution of $\ln(E_{\text{rec}})$ where

$$\mu_{\text{LN}}(E_{\nu}, \sigma_{\text{LN}}) = \ln(E_{\nu}) + \sigma_{\text{LN}}^2 + \Delta_{\text{LN}}, \quad (\text{C.3})$$

Sample	ϵ ($E_\nu < 0.5$ GeV)	σ_{LN} ($E_\nu < 0.5$ GeV)	Δ_{LN} ($E_\nu < 0.5$ GeV)	ϵ	σ_{LN}	Δ_{LN}
ν_μ FHC CCQE	0.885	0.0867	0.0957	0.999	0.157	-0.0189
ν_e FHC CC1 π	0.0392	0.202	0.797	0.0157	0.467	0.845
$\bar{\nu}_\mu$ RHC CCQE	0.995	0.0663	0.0895	0.997	0.0979	0.00619
$\bar{\nu}_e$ RHC CCQE	0.993	0.152	-0.117	0.998	0.112	-0.0499

Table C.1: CCQE nominal reconstruction function parameters found using a χ^2 fit to T2K MC. The R function consisted of a single log-normal which had unique parameters for high energy and low energy regions. The parameters for the ν_e FHC CCQE data sample are shown in Tab. C.2.

Energy Region	ϵ	LN/G Ratio	σ_{LN}	Δ_{LN}	σ_G	Δ_G
$E_\nu < 0.5$ GeV	0.879	0.430/0.570	0.470	-0.319	0.294	0.193
$0.5 < E_\nu < 0.9$ GeV	0.992	0.814/0.186	0.250	-0.130	0.0611	-0.00805
$E_\nu > 0.9$ GeV	0.999	0.845/0.155	0.297	-0.724	0.293	-0.393

Table C.2: CCQE nominal reconstruction function parameters for the ν_e FHC CCQE data sample. The parameters were found using a χ^2 fit to T2K MC. The R function consisted of a single log-normal and a single Gaussian, which had unique parameters for high energy, medium and low energy regions.

assuming E_ν corresponds to the mode of the log-normal. Δ_{LN} is the offset of the log-normal with respect to the true neutrino energy.

After a χ^2 fit to T2K's MC over several oscillation parameter points (see Sec. 6.4), the nominal parameters for the reconstruction function parameters for each interaction channel of each data sample were found. Tabs. C.1, C.3 and C.4 display the nominal parameters for the CCQE, 2p2h and CC1 π interaction channels respectively for all data samples, except the parameters for the ν_e FHC CCQE sample for the CCQE interaction channel are shown in Tab. C.2 as the parameters are unique for this specific sample.

Sample	ϵ	Peak1/Peak2 Ratio	Peak 1 σ_{LN}	Peak 1 Δ_{LN}	Peak 2 σ_{LN}	Peak 2 Δ_{LN}
ν_μ FHC CCQE	1.00	0.815/0.185	0.484	-0.562	0.0822	0.0140
ν_e FHC CCQE	0.996	0.602/0.398	0.424	-0.843	0.0751	-0.0985
ν_e FHC CC1 π	0.0308	0.290/0.710	0.250	-1.480	0.101	0.0560
$\bar{\nu}_\mu$ RHC CCQE	0.999	0.162/0.838	0.236	-0.984	0.107	-0.0706
$\bar{\nu}_e$ RHC CCQE	0.912	0.620/0.380	0.469	-0.683	0.0530	-0.0964

Table C.3: 2p2h nominal reconstruction function parameters found using a χ^2 fit to T2K MC. The R function consisted of a double log-normal.

Sample	ϵ	σ_{LN}	Δ_{LN}
ν_μ FHC CCQE	0.381	0.678	-0.765
ν_e FHC CCQE	0.324	0.391	-1.00
ν_e FHC CC1 π	0.325	0.0597	0.0369
$\bar{\nu}_\mu$ RHC CCQE	0.429	0.488	-0.506
$\bar{\nu}_e$ RHC CCQE	0.319	0.405	-1.04

Table C.4: CC1Pi nominal reconstruction function parameters found using a χ^2 fit to T2K MC. The R function consisted of a single log-normal.

C.1 T2K Resources Used in GAMBIT

Several T2K resources were used in the T2K likelihood calculation in GAMBIT. Some of these resources are included directly in the GAMBIT code, some were used passively to inform GAMBIT developers of the accuracy of the GAMBIT predicted energy spectra. All resources are either publicly available or were explicitly granted permission for use from the T2K collaboration. Tab. C.5 shows each resource used, where it was sourced from and what it was used for in the GAMBIT T2K likelihood calculation.

Resource	Source	Use in GAMBIT
MaCh3 MC energy spectrum plots for 5 different parameter space points, broken down by interaction type and neutrino type	Member of Oscillation Analysis working group	Not encoded directly in GAMBIT, but used to inform us of energy reconstruction parameters, efficiency parameters and accuracy of predicted energy spectra in GAMBIT
Flux predictions at SK	https://t2k-experiment.org/result_category/flux/ (2016 fluxes)	Used directly in code to predict energy spectra
NEUT cross-sections	NEUT 5.4	Extracted cross-sections from NEUT and used directly in GAMBIT
BANFF parameter central values and uncertainties	TN395 [105]	Used to re-weight energy bins to correct flux values
Neutrino 2020 analysis contours	https://t2k-experiment.org/result_category/prelim/	Used to compare to GAMBIT contours
Run1-10 data points	https://t2k-experiment.org/result_category/prelim/	Used to calculate likelihood values
Plot showing reconstructed energy distribution for single true energy	http://dx.doi.org/10.5281/zenodo.1286752	Advised the shape used for reconstruction functions
BANFF fit OA2020 covariance matrix	https://nextcloud.nms.kcl.ac.uk/remote.php/dav/files/ASGReader/ASG/asg_backup/asg/asg2019oa/ND280/BANFF_fits/1278BANFF_postfit_OA2020_fixed2p2h_fixedIHalfLowPPi_20200409.root	Used in PCA decomposition of flux BANFF parameters to reduce number of nuisance parameters in GAMBIT scans
Prob3++	https://github.com/rogerwendell/Prob3plusplus	Used in oscillation probability calculations

Table C.5: An exhaustive list of T2K resources used in the GAMBIT T2K likelihood calculation. Not all resources are directly used in the GAMBIT software, but were used to inform the calculation as to its accuracy compared to T2K.

Bibliography

- [1] K. Abe et al. (Hyper-Kamiokande), *Hyper-Kamiokande Design Report* (2018), [arXiv:1805.04163](https://arxiv.org/abs/1805.04163).
- [2] R. Acciarri et al. (DUNE), *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 1: The LBNF and DUNE Projects* (2016), [arXiv:1601.05471](https://arxiv.org/abs/1601.05471).
- [3] Z. Maki, M. Nakagawa, and S. Sakata, *Remarks on the unified model of elementary particles*, *Prog. Theor. Phys.*, **28**, 870–880 (1962).
- [4] B. Pontecorvo, *Inverse beta processes and nonconservation of lepton charge*, *Sov. Phys. JETP*, **7**, 172–173 (1958).
- [5] B. Pontecorvo, *Neutrino Experiments and the Problem of Conservation of Leptonic Charge*, *Sov. Phys. JETP*, **26**, 984–988 (1968).
- [6] Mark Scott for the T2K Collaboration, *Latest oscillation results from T2K*, *Journal of Physics: Conference Series*, **1342**, 012043 (2020).
- [7] K. Abe et al. *Proposal for an Extended Run of T2K to 20×10^{21} POT* (2016). [arXiv:1609.04111](https://arxiv.org/abs/1609.04111).
- [8] A. D. Dolgov, *Cosmological matter antimatter asymmetry and antimatter in the universe*, In *14th Rencontres de Blois on Matter - Anti-matter Asymmetry* (2002), [arXiv:hep-ph/0211260](https://arxiv.org/abs/hep-ph/0211260).
- [9] J. Christenson, J. Cronin, V. Fitch, and R. Turlay, *Evidence for the 2π Decay of the K_2^0 Meson*, *Phys. Rev. Lett.*, **13**, 138–140 (1964).
- [10] A. Alavi-Harati et al. (KTeV Collaboration), *Observation of Direct CP Violation in $K_{S,L} \rightarrow \pi\pi$ Decays*, *Phys. Rev. Lett.*, **83**, 22–27 (1999).
- [11] V. Fanti et al., *A new measurement of direct CP violation in two pion decays of the neutral kaon*, *Physics Letters B*, **465** (1), 335–348 (1999).

- [12] K. Abe et al. (Belle Collaboration), *Observation of Large CP Violation in the Neutral B Meson System*, Phys. Rev. Lett., **87**, 091802 (2001).
- [13] B. Aubert et al. (BABAR Collaboration), *Measurement of CP-Violating Asymmetries in B^0 Decays to CP Eigenstates*, Phys. Rev. Lett., **86**, 2515–2522 (2001).
- [14] R. Aaij et al. (LHCb Collaboration), *First Observation of CP Violation in the Decays of B_s^0 Mesons*, Phys. Rev. Lett., **110**, 221601 (2013).
- [15] R. Aaij et al. (LHCb Collaboration), *Observation of CP Violation in Charm Decays*, Phys. Rev. Lett., **122**, 211803 (2019).
- [16] M. Fukugita and T. Yanagida, *Baryogenesis Without Grand Unification*, Phys. Lett. B, **174**, 45–47 (1986).
- [17] S. Davidson, E. Nardi, and Y. Nir, *Leptogenesis*, Physics Reports, **466** (4), 105–177 (2008).
- [18] A. D. Sakharov, *Baryon asymmetry of the universe*, Soviet Physics Uspekhi, **34** (5), 417–421 (1991).
- [19] E. A. Paschos, *Leptogenesis with Majorana neutrinos*, Nucl. Phys. B Proc. Suppl., **112**, 36–41 (2002), [arXiv:hep-ph/0204137](https://arxiv.org/abs/hep-ph/0204137).
- [20] W. Pauli, *Rapports du Septième Conseil de Physique Solvay*, Brussels (Gauthier Villars, Paris, 1934) (1933).
- [21] J. Chadwick, *The existence of a neutron*, Proc. R. Soc. Lond. A, **136**, 692–708 (1932).
- [22] F. L. Wilson, *Fermi's Theory of Beta Decay*, American Journal of Physics, **36** (12), 1150–1160 (1968).
- [23] P. A. M. Dirac and R. H. Fowler, *The quantum theory of the electron*, Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character, **117** (778), 610–624 (1928), <https://royalsocietypublishing.org/doi/10.1098/rspa.1928.0023>.
- [24] W. Heisenberg, *Über den Bau der Atomkerne. I*, Z. Physik, **77** (1–2), 1–11 (1932).

- [25] W. Heisenberg, *Über den Bau der Atomkerne. II*, Z. Physik, **78** (3–4), 156–165 (1932).
- [26] W. Heisenberg, *Über den Bau der Atomkerne. III*, Z. Physik, **80** (9–10), 587–596 (1932).
- [27] F. Reines and C. L. Cowan, *Detection of the Free Neutrino*, Phys. Rev., **92**, 830–831 (1953).
- [28] C. L. Cowan et al., *Detection of the Free Neutrino: a Confirmation*, Science, **124** (3212), 103–104 (1956).
- [29] G. Danby et al., *Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos*, Phys. Rev. Lett., **9**, 36–44 (1962).
- [30] M. L. Perl, *The Discovery of The Tau Lepton*, NATO Sci. Ser. B, **352**, 277–302 (1996).
- [31] K. Kodama et al., *Observation of tau neutrino interactions*, Physics Letters B, **504** (3), 218–224 (2001).
- [32] A. A. Aguilar-Arevalo et al. (MiniBooNE Collaboration), *First measurement of the muon neutrino charged current quasielastic double differential cross section*, Phys. Rev. D, **81**, 092005 (2010).
- [33] B. Kayser, *Neutrino physics*, eConf, **C040802**, L004 (2004), [arXiv:hep-ph/0506165](https://arxiv.org/abs/hep-ph/0506165).
- [34] B. Kayser, *B-Meson and Neutrino Oscillation: A Unified Treatment*, AIP Conf. Proc., **1441**, 464 (2011), [arXiv:1110.3047](https://arxiv.org/abs/1110.3047).
- [35] Y. Fukuda et al. (Super-Kamiokande), *Evidence for oscillation of atmospheric neutrinos*, Phys. Rev. Lett., **81**, 1562–1567 (1998), [arXiv:hep-ex/9807003](https://arxiv.org/abs/hep-ex/9807003).
- [36] A. W. P. Poon et al. (SNO), *Neutrino observations from the Sudbury Neutrino Observatory*, AIP Conf. Proc., **610** (1), 218–230 (2002), [arXiv:nucl-ex/0110005](https://arxiv.org/abs/nucl-ex/0110005).
- [37] J. K. Ahn et al. (RENO), *Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment*, Phys. Rev. Lett., **108**, 191802 (2012), [arXiv:1204.0626](https://arxiv.org/abs/1204.0626).

- [38] H. de Kerret et al. (The Double Chooz Collaboration), *Double Chooz θ_{13} measurement via Total Neutron Capture Detection*, Nat. Phys., **16**, 558–564 (2020).
- [39] M. Decowski, *KamLAND’s precision neutrino oscillation measurements*, Nuclear Physics B, **908**, 52–61 (2016), neutrino Oscillations: Celebrating the Nobel Prize in Physics 2015.
- [40] F. P. An et al. (Daya Bay), *Observation of electron-antineutrino disappearance at Daya Bay*, Phys. Rev. Lett., **108**, 171803 (2012), [arXiv:1203.1669](https://arxiv.org/abs/1203.1669).
- [41] M. H. Ahn et al. (K2K Collaboration), *Measurement of neutrino oscillation by the K2K experiment*, Phys. Rev. D, **74**, 072003 (2006).
- [42] D. G. Michael et al. (MINOS), *Observation of muon neutrino disappearance with the MINOS detectors and the NuMI neutrino beam*, Phys. Rev. Lett., **97**, 191801 (2006), [arXiv:hep-ex/0607088](https://arxiv.org/abs/hep-ex/0607088).
- [43] M. A. Acero et al. (NOvA), *First Measurement of Neutrino Oscillation Parameters using Neutrinos and Antineutrinos by NOvA*, Phys. Rev. Lett., **123** (15), 151803 (2019), [arXiv:1906.04907](https://arxiv.org/abs/1906.04907).
- [44] P. A. Zyla et al. (Particle Data Group), *Review of Particle Physics*, Progress of Theoretical and Experimental Physics, **2020** (8) (2020), 083C01.
- [45] S. Mertens, *Direct Neutrino Mass Experiments*, J. Phys.: Conf. Ser., **718**, 022013 (2016).
- [46] P. F. de Salas et al., *Neutrino Mass Ordering from Oscillations and Beyond: 2018 Status and Future Prospects*, Front. Astron. Space Sci., **5**, 36 (2018).
- [47] K. Abe et al., *Constraint on the Matter–Antimatter Symmetry-Violating Phase in Neutrino Oscillations*, Nature, **580**, 339–344 (2020).
- [48] K. Abe et al. (T2K), *The T2K Experiment*, Nucl. Instrum. Meth., **A659**, 106–135 (2011), [arXiv:1106.1238](https://arxiv.org/abs/1106.1238).
- [49] K. Abe et al. (The T2K Collaboration), *T2K measurements of muon neutrino and antineutrino disappearance using 3.13×10^{21} protons on target*, Phys. Rev. D, **103**, L011101 (2021).

- [50] D. Beavis, A. Carroll, and I. Chiang (E889 Collaboration), *Long baseline neutrino oscillation experiment at the AGS. Physics design report* (1995).
- [51] K. Abe et al. (T2K Collaboration), *Measurements of neutrino oscillation in appearance and disappearance channels by the T2K experiment with 6.6×10^{20} protons on target*, Phys. Rev. D, **91**, 072010 (2015).
- [52] N. Abgrall et al. (NA61/SHINE Collaboration.), *Measurements of π^\pm , κ^\pm , κ_S^0 , Λ and proton production in proton–carbon interactions at 31 GeV/c with the NA61/SHINE spectrometer at the CERN SPS*, Eur. Phys. J. C, **76**, 84 (2016).
- [53] P. Dunne (T2K Collaboration), *Latest Neutrino Oscillation Results from T2K* (2020). <https://doi.org/10.5281/zenodo.3959558>.
- [54] *T2K Experiment Photograph Repository* (Last Accessed: 2021). <https://t2k-experiment.org/photo/super-kamiokande/>.
- [55] S. Fukuda et al., *The Super-Kamiokande Detector*, Nucl. Instrum. Methods Phys. Res., Sect. A, **501** (2), 418–462 (2003).
- [56] H. Nishino et al., *High-speed charge-to-time converter ASIC for the Super-Kamiokande detector*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, **610** (3), 710–717 (2009).
- [57] K. Abe et al., *Calibration of the Super-Kamiokande detector*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, **737**, 253–272 (2014).
- [58] A. Ferrero, *The ND280 Near Detector of the T2K Experiment*, AIP Conference Proceedings, **1189** (1), 77–82 (2009).
- [59] K. Abe et al. (T2K). *T2K ND280 Upgrade - Technical Design Report* (2019). [arXiv:1901.03750](https://arxiv.org/abs/1901.03750).
- [60] Y. Uchida, *ND280 Picture Repository* (Last Accessed: 2021). <https://www.t2kuk.org/wiki/YoshiUchida/ND280Pictures>.
- [61] S. Assylbekov et al., *The T2K ND280 Off-Axis Pi-Zero Detector*, Nucl. Instrum. Meth., **A686**, 48–63 (2012), [arXiv:1111.5030](https://arxiv.org/abs/1111.5030).

[62] K. Ieki, *MPPC for T2K Fine-Grained Detector*, PoS, **PD09**, 023 (2010).

[63] P. A. Amaudruz et al. (T2K ND280 FGD), *The T2K Fine-Grained Detectors*, Nucl. Instrum. Meth., **A696**, 1–31 (2012), [arXiv:1204.3666](https://arxiv.org/abs/1204.3666).

[64] N. Abgrall et al. (T2K ND280 TPC), *Time Projection Chambers for the T2K Near Detectors*, Nucl. Instrum. Meth., **A637**, 25–46 (2011), [arXiv:1012.0865](https://arxiv.org/abs/1012.0865).

[65] C. Bronner and K. Sakashita (T2K), *Joint Oscillation Analysis using Lepton Momentum and Angle for Neutrino Electron-like Events*, Tech. Rep., **TN-194** (2014).

[66] P. Hamacher-Baumann et al., *The T2K-ND280 upgrade proposal*, Technical Report CERN-SPSC-2018-001. SPSC-P-357, CERN, Geneva (2018).

[67] C. Andreopoulos et al., *The GENIE neutrino Monte Carlo generator*, Nucl. Instrum. Methods Phys. Res. Sect. A, **614** (1), 87–104 (2010).

[68] S. Agostinelli et al., *Geant4—a simulation toolkit*, Nucl. Instrum. Meth., **506** (3), 250–303 (2003).

[69] K. Abe et al. (T2K), *Measurement of neutrino and antineutrino oscillations by the T2K experiment including a new additional sample of ν_e interactions at the far detector*, Phys. Rev. D, **96** (9), 092006 (2017), [Erratum: Phys.Rev.D 98, 019902 (2018)], [arXiv:1707.01048](https://arxiv.org/abs/1707.01048).

[70] Yu. G. Kudenko et al., *Extruded plastic counters with WLS fiber readout*, Nucl. Instrum. Meth., **A469**, 340–346 (2001).

[71] *Kuraray Wavelength Shifting Fibers* (Last Accessed: 2021). <http://kuraraypsf.jp/psf/ws.html>.

[72] O. Mineev et al., *Scintillator detectors with long WLS fibers and multi-pixel photodiodes*, JINST, **6**, P12004 (2011), [arXiv:1110.2651](https://arxiv.org/abs/1110.2651).

[73] O. Basille et al., *Baby MIND Readout Electronics Architecture for Accelerator Neutrino Particle Physics Detectors Employing Silicon Photomultipliers*, JPS Conf. Proc., **27**, 011011 (2019).

- [74] E. Noah et al., *Readout Scheme for the Baby-MIND Detector*, PoS, **PhotoDet2015**, 031. 12 p (2016).
- [75] *Hamamatsu MPPCs* (Last Accessed: 2021). https://www.hamamatsu.com/eu/en/product/optical-sensors/mppc/mppc_mppc-array/index.html.
- [76] A. Blondel et al., *The SuperFGD Prototype Charged Particle Beam Tests*, JINST, **15** (12), P12003 (2020), [arXiv:2008.08861](https://arxiv.org/abs/2008.08861).
- [77] G. M. Mitev, L. T. Tsankov, M. G. Mitev, and E. N. Messomo, *Light pulse generator for multi-element scintillation detectors testing*, In *2016 XXV International Scientific Conference Electronics (ET)*, pages 1–4 (2016).
- [78] J. Adams et al., *The Cern Proton Synchrotron*, Nature, **185**, 569–572 (1960).
- [79] *Raspberry Pi 3 Sense HAT* (Last Accessed: 2021). <https://www.raspberrypi.org/products/sense-hat/>.
- [80] T. G. Mayerhöfer, H. Mutschke, and J. Popp, *Employing Theories Far beyond Their Limits—The Case of the (Boguer-) Beer–Lambert Law*, ChemPhysChem, **17** (13), 1948–1955 (2016).
- [81] P. F. de Salas et al., *2020 global reassessment of the neutrino oscillation picture*, JHEP, **02**, 071 (2021), [arXiv:2006.11237](https://arxiv.org/abs/2006.11237).
- [82] I. Esteban et al., *The fate of hints: updated global analysis of three-flavor neutrino oscillations*, JHEP, **09**, 178 (2020), [arXiv:2007.14792](https://arxiv.org/abs/2007.14792).
- [83] P. Athron et al. (GAMBIT), *GAMBIT: The Global and Modular Beyond-the-Standard-Model Inference Tool*, Eur. Phys. J. C, **77** (11), 784 (2017), [Addendum: Eur.Phys.J.C 78, 98 (2018)], [arXiv:1705.07908](https://arxiv.org/abs/1705.07908).
- [84] A. Kvellestad, P. Scott, and M. White, *GAMBIT and its Application in the Search for Physics Beyond the Standard Model* (2019), [arXiv:1912.04079](https://arxiv.org/abs/1912.04079).
- [85] P. Athron et al. (GAMBIT), *A global fit of the MSSM with GAMBIT*, Eur. Phys. J. C, **77** (12), 879 (2017), [arXiv:1705.07917](https://arxiv.org/abs/1705.07917).
- [86] P. Athron et al., *Global fits of axion-like particles to XENON1T and astrophysical data* (2020), [arXiv:2007.05517](https://arxiv.org/abs/2007.05517).

- [87] J. Bhom et al., *A model-independent analysis of $b \rightarrow s\mu^+\mu^-$ transitions with GAMBIT's FlavBit* (2020), [arXiv:2006.03489](https://arxiv.org/abs/2006.03489).
- [88] M. Chrzaszcz et al., *A frequentist analysis of three right-handed neutrinos with GAMBIT*, Eur. Phys. J. C, **80** (6), 569 (2020), [arXiv:1908.02302](https://arxiv.org/abs/1908.02302).
- [89] P. Stöcker et al. (GAMBIT Cosmology Workgroup), *Strengthening the bound on the mass of the lightest neutrino with terrestrial and cosmological experiments* (2020), [arXiv:2009.03287](https://arxiv.org/abs/2009.03287).
- [90] I. Esteban et al., *Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23} , δ_{CP} , and the mass ordering*, JHEP, **01**, 106 (2019), [arXiv:1811.05487](https://arxiv.org/abs/1811.05487).
- [91] *Neutrino Beam Flux Predictions from T2K* (Last Accessed: 2021). https://t2k-experiment.org/result_category/flux/.
- [92] P. de Perioet et al. (T2K), *Constraining the Flux and Cross Section with Data from the ND280 Detector for the 2012a Oscillation Analysis*, Tech. Rep., **TN-106** (2012).
- [93] H. R. Gallagher, *Neutrino Event Generators: A Review*, AIP Conference Proceedings, **1189** (1), 35–42 (2009).
- [94] P. Stowell et al., *Neutrino Monte-Carlo Event Generators and Cross-section Data*, In *18th International Workshop on Neutrino Factories and Future Neutrino Facilities Search* (2016), [arXiv:1611.03275](https://arxiv.org/abs/1611.03275).
- [95] R. Wendell, *Prob3++ Software Repository* (Last Accessed: 2021). <https://github.com/rogerwendell/Prob3plusplus>.
- [96] M. Wascko, *T2K Status, Results, and Plans*, In *XXVIII International Conference on Neutrino Physics and Astrophysics (Neutrino 2018)*, Zenodo (2018).
- [97] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, Eur. Phys. J. C, **71**, 1554 (2011), [Erratum: Eur.Phys.J.C 73, 2501 (2013)], [arXiv:1007.1727](https://arxiv.org/abs/1007.1727).
- [98] L. Berns et al. (T2K), *Five Sample Joint Oscillation Analysis with T2K Run1–10 Data*, Tech. Rep., **TN397** (2020).

- [99] G. D. Martinez et al. (GAMBIT), *Comparison of statistical sampling methods with ScannerBit, the GAMBIT scanning module*, Eur. Phys. J. C, **77** (11), 761 (2017), [arXiv:1705.07959](https://arxiv.org/abs/1705.07959).
- [100] R. Storn and K. Price, *Differential evolution: A simple and efficient heuristic for global optimization over continuous spaces*, Journal of Global Optimization, **11**, 341–359 (1997).
- [101] K. Price, R. M. Storn, and J. A. Lampinen. *Differential evolution: a practical approach to global optimization*. Springer (2005).
- [102] S. Das and P. N. Suganthan, *Differential Evolution: A Survey of the State-of-the-Art*, IEEE Transactions on Evolutionary Computation, **15** (1), 4–31 (2011).
- [103] K. Price, *Differential evolution*, In I. Zelinka, V. Snášel, and A. Abraham, editors, *Handbook of Optimization*, volume 38, pages 187–214, Springer Berlin Heidelberg (2013).
- [104] J. Brest et al., *Self-Adapting Control Parameters in Differential Evolution: A Comparative Study on Numerical Benchmark Problems*, IEEE Transactions on Evolutionary Computation, **10** (6), 646–657 (2006).
- [105] L. Munteanu et al. (T2K), *Constraining the Flux and Cross Section Models with Data from ND280 using FGD1 and FGD2 for the 2020 Oscillation Analysis*, Tech. Rep., **TN395** (2020).