



The CHIPS Prototype Water Cherenkov Detector

Simeon Bash

Doctoral Thesis

University College London

August 2022

Declaration

I, Simeon Bash, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Simeon Bash
August 30, 2022

Abstract

The CHIPS (CHerenkov detectors In mine PitS) detector was a large-scale water Cherenkov (WC) long baseline neutrino detector located in northern Minnesota. It was located 7 mrad off-axis, 712 km downstream from the NuMI neutrino source at Fermi National Accelerator Laboratory and with a fiducial detector mass of 5.9 kt. CHIPS was a research and development project aiming to demonstrate a reduction in the cost of construction of a WC neutrino detector to \$200k-\$300k per kt. CHIPS was constructed and deployed in 2019. The design, construction and deployment procedure of CHIPS are discussed.

CHIPS utilised photomultiplier tubes and pre-existing readout electronics for instrumentation. A new low-cost set of readout electronics and its associated DAQ system was developed as a successor to the existing CHIPS electronics. Its design is presented here.

A muon neutrino beam from an accelerator directed through the Earth produces a continuous flux of muons along the beamline due to the neutrino interactions with the rock. A technique is proposed to monitor the energy profile of such a neutrino beam by measuring the neutrino flux through the Earth using a small portable detector with a magnetic deflector. The study shows that as the off-axis angle changes, information about the kaon content of the parent hadron beam can be inferred using existing detector technologies.

Impact Statement

As high energy physics continues to advance and increase in scale, scope and complexity, the difficulties associated with subatomic particle physics - such as resources, time and financial backing - arise. Research also slows and becomes less accessible.

The CHIPS experiment has demonstrated that neutrino detectors could be constructed at a lower cost and in a shorter timescale than is the case with existing experiments.

In addition, the obvious impact of this work, on top of the potential to further our understanding of neutrino physics, is to increase the repeatability of experiments involving neutrinos. It would usually be feasible to construct only one or two experiments using cutting-edge detectors, but by using the CHIPS methodology, several experiments could be constructed and used together. This would lead to repeatability and more precise measurements and therefore make the scientific method more robust.

Outside of academia, the work presented is relevant as a contribution to the foundation of knowledge about neutrino physics. In the long term, this may contribute to advances in technology as well as advances in science that add to the benefit and well-being of humanity. The way in which the research has been conducted also demonstrates that research can be performed in a socially and financially responsible way. This frees resources for other projects and demonstrates that high-quality, cutting-edge physics can be conducted for relatively small sums.

Finally, the techniques proposed for monitoring neutrino beams and utilising neutrino interactions in the Earth increases the usefulness of existing neutrino beam-lines. The uncertainties in the composition of the neutrino beam could be reduced by using the rock muon flux to gather additional data about the beam. This improves the quality of existing research by reducing uncertainties. Also, every neutrino produced has an associated cost - not just in money but in environmental impact and human labour. Improving the value of each neutrino already produced and recording interactions that would otherwise be wasted improves the value of existing neutrino beams. This also applies to the CHIPS experiment as a whole.

Acknowledgements

Firstly, I would like to thank my supervisor Jenny for her help, support and wisdom during the course of my studies.

I would also like to thank my colleagues at CHIPS and in the wider physics community who made this work possible, in particular Jerry Meier and Richard Wade whose patience and willingness to teach me made me into a better, more practical physicist.

During my PhD I received support from my friends and fellow PhD students. Their support and motivation were invaluable, particularly during the final write up and I thank them for that support and motivation. I also thank the Anderson family whose warm hospitality and generosity made field work in Minnesota and beyond all the more pleasant.

Finally, I thank my family for their love and support.

Contents

Declaration	3
Abstract	5
Impact Statement	7
Acknowledgements	9
List of Figures	13
List of Tables	21
1 Introduction	22
2 Neutrino Oscillations	25
2.1 Neutrino Mixing	25
2.1.1 Two Flavour Vacuum Oscillations	27
2.1.2 The Matter Effect and Three Flavour Oscillations	27
3 Neutrino Interactions	30
3.1 The Weak Interaction	30
3.1.1 Neutral Current Elastic Scattering	31
3.1.2 Charged Current Quasi-Elastic Scattering	31
3.1.3 Resonant Pion Production	31
3.1.4 Deep Inelastic Scattering	31
3.1.5 Coherent Pion Scattering	31
3.1.6 Meson exchange current	31
3.2 Measuring Neutrinos	32
3.2.1 Cherenkov Light	32
3.2.2 Photomultiplier Tubes	35
3.3 Neutrino Production	35
3.3.1 Neutrino Beams	36
4 The CHIPS Detector	43
4.1 Experiment Overview	43

4.2	The PolyMet Mine	46
4.3	Detector Frame	47
4.4	Liner	52
4.5	Buoyancy and Floating Dock	54
4.6	Deployment Procedure	57
4.7	Structure Summary	58
4.8	Water Purity	59
4.9	Detector Planes	63
4.9.1	Nikhef Planes	64
4.9.2	PMT Potting	67
4.9.3	Madison Planes	70
4.10	Electronics Fanout	75
4.10.1	Nikhef Fanout	75
4.10.2	Madison Fanout	76
4.10.3	Main Control Box	78
4.10.4	Shore Electronics Hut	81
4.11	Timing System	81
4.12	The Conclusion of the CHIPS Experiment	83
5	Future Detector Electronics' Development	88
5.1	Detailed Description of the Badgerboard	89
5.2	Badgerboard Replacement	93
5.3	Software Delivery and Data Storage	99
5.4	Summary of Fanout Development	101
6	Neutrinos & Muons in Rock	102
6.1	Muons in Rock	104
6.2	Detector Concept	104
6.3	Simulation Overview	105
6.4	Generating & Propagating Neutrinos	107
6.5	Generating Muons	109
6.6	Detector Deflection	114
7	Summary & Conclusions	126
A	Off-axis Neutrino Beams	128
B	New Fanout Circuit Diagrams	132

List of Figures

2.1	Neutrinos scattering in matter. All neutrinos can undergo NC scattering. Only electron neutrinos can undergo coherent forward scattering in electronic matter.	28
3.1	Examples of NC and CC neutrino interactions. The NC interaction: the exchange of Z^0 during neutrino-electron scattering. The CC interaction: a ν_μ interacts with a neutron in a nucleus exchanging a W^- to give a μ and a proton.	30
3.2	An example of MEC where a muon neutrino interacts with a proton in a nucleus to produce a muon and neutron. The two protons are correlated through the exchange of a π^0 leading to both resultant nucleons being ejected from the nucleus.	32
3.3	A charged particle (denoted by the red line) travelling through a medium faster than the speed of light through the medium. The black circles show the expanding spheres of light that are emitted as the charged particle passes. The spheres expand and interfere to form the characteristic cone of Cherenkov light.	33
3.4	Construction showing the distances a charged particle and its Cherenkov light can travel; this gives the Cherenkov angle.	34
3.5	The internal structure of a PMT showing the photocathode, the chain of dynodes and the readout anode [36].	35
3.6	The internal dynode structure of a PMT. The first dynode is at the front covered by a copper mesh and the chain of dynodes can be seen behind in an alternating ladder arrangement.	36
3.7	A representative diagram of the NuMI beam facility. Protons from the main injector strike a graphite target. The resulting pions and kaons are focused and decay to give muons and muon neutrinos. The muons are absorbed and the neutrinos continue towards particle physics experiments [39].	38

3.8	The flux of neutrinos at the MINOS near detector [42]. The solid line corresponds to a low energy neutrino beam, the dashed line corresponds to a medium energy neutrino beam and the dotted line corresponds to a high energy neutrino beam.	39
3.9	A simulation of the NOVA neutrino energy spectrum due to kaons and pions. The off-axis effect causes the peak energy peak to be lower but narrowed.	40
3.10	A simulated spectrum of ν_μ CC events in the MINOS near detector which shows the large number of neutrinos produced by kaon parents. Figure from [58].	41
3.11	The ratios of pions and kaons across the range of longitudinal momenta with different transverse momenta cuts. The larger error bars for plots with kaon ratios can be seen. Figure from [58] [59]. To experimentally measure the pion-kaon production ratio, a secondary proton beam was fired at test targets and the hadrons produced were counted and identified using a number of instruments in the MIPP detector system (including time projection chambers and Cherenkov detectors) [57] [60].	42
4.1	Map showing the NuMI beam source at Fermilab and the mine pit where CHIPS was located in Hoyt Lakes, Minnesota [64].	44
4.2	A rendering of the CHIPS detector showing the overall structure: two steel end caps supporting detector planes with buoyancy in the top cap. The detector stretched out to be a long cylinder and the complete detector was wrapped in liner. Here detector planes on the walls of the cylinder are shown which was a future upgrade path for CHIPS and not installed for the initial deployment. Image courtesy of Thomas Dodwell.	45
4.3	A map of the NuMI beam exiting the Earth's surface showing the positions of CHIPS, NOVA and MINOS. The z-axis shows the expected number of neutrino events per year per kiloton of water assuming no neutrino oscillations. Contours of constant L/E are shown. Image taken from [67].	47
4.4	The PolyMet mine site. CHIPS was built beside the mine pit on a flat construction with additional assembly done in the PolyMet building. The detector was then deployed underwater in the mine pit [27].	48

4.5	The Wentworth 2W mine pit. The body of water in the foreground shows the mine pit itself and the CHIPS construction site which was the slipway leading into the water. The slipway was built on the road where the excavators would have driven in and out of the pit when the mine was still operational. The detector can be seen partially constructed along with the building site apparatus.	49
4.6	Pieces of the bottom end cap. The gusset plates used to join pieces are visible as are the tyres which the frame sits on. The hollow lattice style design saves weight as less steel is required than for a solid detector.	50
4.7	Rows of stringers bolted to the main support structure. Detector planes would be attached on top of the stringers.	51
4.8	A mounting bracket welded to the main frame to which a stringer was bolted.	51
4.9	The completed frame. The pillars supporting the top cap are clearly visible.	52
4.10	Two pieces of liner fused together (unfinished). Hot air is blown in the gap between the two pieces until the melting point is reached. The two pieces of liner are clamped together. The joints were not individually tested due to the circumstances of construction. However, prior testing showed they were expected to be light-tight if properly performed.	53
4.11	The detector completely wrapped in liner. Additional liner was folded around the base below the water line so that the detector could stretch to be 12 m tall.	54
4.12	The 12 inch PVC buoyancy pipes mounted in the top cap to provide flotation.	55
4.13	Two one-quarter segments of the floating dock. Three winches spaced between six floats were mounted on top.	56
4.14	One of the winches used in CHIPS mounted onto the floating dock. The 12 winches could lift a total of 19 tonnes of dry mass.	57
4.15	The detector being towed from the construction site to the deeper water.	58
4.16	The interior of the water hut showing the filters. Image courtesy of Thomas Dodwell.	60
4.17	The exterior of the water hut showing the passthrough between the black HDPE pipe and the inside of the hut.	60

4.18	The equipment used to butt weld two sections of HDPE pipe together. The ends of the two pipes were shaved down to be smooth and flush with each other. A heated plate was placed in between until the ends had melted. The two sections of pipes were pressed together and the melted sections were fused together. The process was precisely controlled for standardised joints.	61
4.19	The completed inflow and outflow water pipes for the detector. The short pipe on the surface is the water intake for filling the detector during deployment.	62
4.20	The plate where the flanges on the end of the HDPE pipe were bolted onto the detector. The liner was clamped between the plate and the flanges to maintain light and water proofing. This is a standard technique to make waterproof and light-tight seals.	63
4.21	A fully disassembled Nikhef PMT showing from left to right: the PMT insert, the PMT itself with its high-voltage base attached, the waterproof cover complete with its O-ring and then the plastic reflector for light.	65
4.22	The assembled PMT with its plastic reflector attached ready to be glued onto a plane.	65
4.23	The Nikhef electronics which sat inside the electronics container: the green PCB is the CLB and the protruding red board is a Calamari board. An SFP port connects the CLB to the White Rabbit network. The 8P8C connectors are for the PMTs. There is a total of 30 connectors on the front and rear Calamari boards.	66
4.24	A Nikhef plane installed inside the detector. The PMTs were angled forwards into the beam and the aluminium electronics container can be seen in the middle. The PVC hose exiting the electronics container was for connecting the plane to cable manifolds which lead to fanout containers.	67
4.25	A potted Nikhef PMT. The presence of the optical gel between the plastic and the glass reduced the reflectivity between the two layers. .	68
4.26	The jig for PMT potting. Alternating layers of PMTs and covers were stacked and the tubes were held straight during the curing process. .	69
4.27	A row of PMTs after the potting process showing the consistency of the transparent finish that can be achieved by potting the tubes. . .	70
4.28	A Madison PMT showing the Cockcroft-Walton base attached to the end of the tube and the Microdaq on top of that.	71
4.29	The photocathode of a Madison PMT exposed to the environment. The tube was glued into the white insert with black resin that created a waterproof barrier.	72

4.30 The Badgerboard. The Cat 5 input cable enters at the bottom left and the BeagleBone Green sits at the top left. The Cat 5 cables to the Microdaq leave on the right hand side of the board.	74
4.31 A rendering of a Madison plane. The PMTs are mounted onto rungs of PVC pipe. The large cylinder in the centre of the plane is the electronics container. 15 PMTs were mounted despite the electronics supporting 16. The 16 th channel was used as a spare if problems arose during the assembly of the plane.	75
4.32 The electronics container for a Madison plane. The entire container was made from glued PVC which could then be glued to the rest of the plane. The Cat 5 cables between the electronics and the Microdaqs can be seen leaving the container.	76
4.33 A fanout container, the same kind of aluminium pressure cylinder used for the electronics containers on the Nikhef planes.	76
4.34 The inside of five Nikhef fanout containers. The relays and power converters are visible. The White Rabbit switches sit under the relays attached to the aluminium block which acts as a heatsink.	77
4.35 A Danout board complete with its BeagleBone Green.	78
4.36 The interior of the main control box.	79
4.37 The main control box.	80
4.38 The network and power distribution systems of CHIPS. Here POM stands for <i>Planar Optical Modules</i> - the detector planes. Image from [27].	82
4.39 A Nova TDU [85].	83
4.40 The network topology of the CHIPS time distribution system [83]. The Nova TDU transmits timestamped signals to the CHIPS spill forwarder which then traverses the Fermilab Firewall and internet to the CHIPS site. The chipsshore03 computer decodes the timing signals.	84
4.41 A CHIPS bitmap showing what is likely to be a cosmic ray muon event. A cosmic ray would be expected to form a ring, which can be seen in this pattern of pixels. Each pixel corresponds to a photomultiplier tube in the detector and the boxes show the position of each plane. The damage sustained to the detector during deployment gave little time to assess the efficiency of the detector. This event is assumed to be a cosmic ray as it is unlikely to be noise because each plane had been tested before installation to ensure low noise. At the time of this event, the NuMI beam was not running. Image courtesy of Stefano Germani.	85

4.42 The exterior of the detector after completion. The liner can be seen supporting the detector and the floating dock forms a perimeter around the detector. The two boats tow and steer the detector during deployment. The scale of the detector is visible.	85
4.43 A panorama inside the detector. The detector planes and the manifolds which connect the planes can be seen.	86
4.44 The interior of the detector: the rows of Nikhef PMTs angled to point into the neutrino beam can be seen at the back with the Madison planes in front. The top cap buoyancy can be seen above.	86
4.45 The detector from the front during the final liner sealing process so the detector was and waterproof.	87
5.1 Top view of a BeagleBone Green showing the 96 I/O pins.	91
5.2 A BeagleBone Green mounted onto a Badgerboard using the two 48-pin connectors. The U-turn cable is shown in the top left. Here, a longer cable was used and coiled up.	91
5.3 WR-LEN OEM edition for embedded use.	94
5.4 The front and rear of the RPCM4. The two 100-pin connectors can be seen at the top and bottom of the rear side of the board to connect the module to a carrier board.	95
5.5 The prototype of the new fanout system. The motherboard is carrying 7 daughterboards with ports for four Microdaqs each.	98
5.7 A manufactured and assembled fanout.	99
6.1 The hadron flux uncertainties on neutrino production in the NOVA experiment. $pC \rightarrow \pi X$ represents proton-carbon interactions yielding pions and $pC \rightarrow KX$ represents proton-carbon interactions yielding kaons [92]. For muon neutrinos, the pion production uncertainty dominates but the kaon production uncertainty does increase as the energy increases. For electron neutrinos, initially the pion production uncertainty dominates. However, at approximately 3.5 GeV, the kaon production uncertainty overtakes the pion production uncertainty and continues to rise. This means that at higher energies, the uncertainty in the number of electron neutrinos produced becomes increasingly dominated by the kaon production uncertainty. For experiments that measure muon neutrino disappearance and electron neutrino appearance, the uncertainty in kaon production is important. Figure from [93].	103

6.2	A cartoon of neutrinos interacting through a solid medium (rock). The neutrinos interact to create muons. The high energy muons travel far in the rock whereas the low energy muons only travel short distances. The more energy a muon has, the more likely it is to travel far enough to intersect with the detector plane. There is therefore a high-energy bias when counting the muons in the detector compared with the energy distribution of the muons or the original neutrinos. . .	105
6.3	A diagram of the detector concept. The muons travel through the rock and then enter a detector plane which measures the direction of travel. The muon is influenced by a narrow region of magnetic field perpendicular to the detector plane which bends the path of the muon. As the muon leaves the magnetic field region, it continues travelling tangentially until it reaches the second detector plane which measures the direction of travel upon exiting the detector. Between the two detector fields, there is empty space for the muon to travel through.	106
6.4	The energy spectrum of the parent hadrons of the neutrino beam approximating the NuMI beam in low energy configuration. Pions are more abundant than kaons.	108
6.5	The energy spectrum of muon neutrinos intersecting a 1 m radius circle 1036 m downbeam on-axis.	109
6.6	The cumulative effect of different weightings on the neutrino flux. The DK2NU neutrinos are initially different due to the reverse ray-tracing of the DK2NU detector neutrinos. Adding in the hadron weighting <code>Nimpwt</code> affects both the DK2NU detector and the custom crosscheck detector. Adding the neutrino location weighting to both changes the DK2NU spectrum but not the crosscheck spectrum. This because the crosscheck detector already correctly randomly distributes the neutrinos so the dashed green and dashed blue lines directly overlap. The DK2NU detector with both hadron and location weightings agrees with the crosscheck detector with just the hadron weighting. The location weighting of the crosscheck detectors was fixed at 1 so each event recorded represented a single real event.	110
6.7	The probability of a neutrino interacting via a charged current interaction per metre of travel through rock and iron (per 10^{20} POTS). . .	111
6.8	The muon energy spectrum for both the DK2NU and crosscheck detector.	112
6.9	The effect of the muon's travel distance on the muon energy spectrum. The low energy muons are less likely to travel far enough to reach the detector so are weighted downwards in the spectrum.	113

6.10	The effective mass of the detector. High energy particles can travel further through the existing rock which gives the detector a larger effective mass without altering the experimental setup.	114
6.11	The fraction of muons where the angle between the muon's direction vector and the parent neutrino's direction vector is less than θ	115
6.12	The deflection of muons due to the magnetic field. The solid line at approximately 0.45 m is due to the muons being unable to exceed the speed of light. In total there are 7261 DK2NU entries but due to the weighting the entries are spread across the entire energy range. There are 373375 crosscheck entries. The method used to generate neutrinos for the crosscheck detectors changed the weighting system. The distribution of the hits is not spread out evenly across the entire energy range. This is visible in at the high-deflection region of the plot where statistical fluctuations become significant. This is an inherent flaw in the method used to generate the neutrinos for detectors that can be moved.	117
6.13	The form of the muon energy vs deflection curve. The horizontal line at approximately 0.45 m is due to the fastest muons approaching but being unable to exceed the speed of light. Muons that have a velocity that causes them to be deflected almost 90° would have a deflection approaching ∞ in the second region with no magnetic field. The detector has been restricted to 2 m.	118
6.14	Neutrino flux energy spectra through each detector.	120
6.15	Muon flux energy spectra through each detector. Here, the muons have lost no energy in the rock. The counting correction which applies the distance a muon of a given energy can actually travel through rock is included.	121
6.16	The effective mass of each detector.	122
6.17	The deflection of each muon through each detector.	123
6.18	The deflection of each muon through each detector, normalised by the total number of muons.	124
A.1	A pion decay from the pion rest frame. The muon and neutrino have equal and opposite momenta.	128
A.2	Diagram showing the component of the neutrino's velocity which receives a Lorentz boost.	129
A.3	Simulated energy spectra of charged current events at a far detector location 735 km from Fermilab at various off-axis angles for the NuMI beam at the medium-energy setting [104]. A similar set of spectra with separate pion and kaon contributions at different off-axis angles can be seen in figure 6.14.	131

List of Tables

- | | | |
|-----|---|-----|
| 4.1 | The buoyancy budget for CHIPS. The top and bottom halves were not the same. The top half had a steel frame, the steel stilts, the PVC flotation, half of the liner, a set of detector planes and the electronics boxes. The bottom cap had a steel frame, half of the liner and a different set of detector planes. The detector planes were assumed to contribute no flotation or dead weight to the final system. This was because the raw materials were approximately neutrally buoyant but the planes contained air so had the potential to provide buoyancy. However, it was unknown how many planes would leak so no buoyancy was assumed. It would be easier to add ballast to the final detector compared with adding additional flotation. The table shows that the total detector would sink to the bottom of the lake with the top cap floating above the bottom and that, with the addition of a floating dock, the detector would not sink. | 56 |
| 6.1 | The total number of muons, the number of muons from pion ancestors and the number of muons from kaon ancestors at each detector location. The number of muons is per 10^{20} POTS (equivalent to approximately 150 days worth of Numi beam operation). | 125 |

Chapter 1

Introduction

In 1956, Chien-Shiung Wu tested the parity conservation of the weak interaction [1]. A β -decaying source of ^{60}Co was used to create electrons and electron antineutrinos. As β decay is mediated *only* by the weak interaction, β -decay is suitable for testing parity violation of the weak interaction in isolation. The spins in the ^{60}Co atoms were aligned using a strong magnetic field and intense cooling to maintain the strongly polarised state. ^{60}Co decays into ^{60}Ni and emits an electron and an electron antineutrino (shown in Equation 1.1).



^{60}Co has a spin of 5 and ^{60}Ni has a spin of 4, and the electron and electron antineutrino have spins of 1/2; thus spin is conserved.

Wu observed that electrons were emitted preferentially in the direction opposite to their spin (*negative helicity* or *left-handed*) and, by conservation of momentum, the antineutrinos must have been emitted preferentially in the direction of their spin (*positive helicity* or *right-handed*).

The spin vectors and magnetic field vector are axial and do not change direction under a parity transformation. However, the momentum vector of the resultant electron does change direction under a parity transformation and therefore so does the overall helicity. If both left-handed and right-handed neutrinos existed, the neutrino would not prefer a particular direction with respect to the original spin vector and the electrons would have been emitted isotropically. If neutrinos had a single handedness, the neutrino would prefer a single direction with respect to the original spin vector and an anisotropic electron distribution would be observed. This would be in violation of parity symmetry. Wu's experiment provided the first evidence of parity violation in the weak interaction and that neutrinos are right-handed and antineutrinos are left-handed. This result was corroborated by Goldhaber in 1957 [2] by measuring only neutrinos with left-handed helicity.

In summary, it can be inferred from experiments that the weak interaction only produces left-handed neutrinos and right-handed antineutrinos.

In 1968, the Homestake experiment observed a deficit of electron neutrinos from the Sun [3]. This won Raymond Davis Jr. the 2002 Nobel Prize in Physics for the first direct observation of neutrinos produced by the sun [4]. This turned out to provide the first experimental evidence of neutrino flavour oscillation - a process that can only occur if neutrinos have mass. Follow up experiments investigated the phenomenon and in 2015 the Nobel Prize in Physics [5] was awarded to Arthur B. McDonald of the Sudbury Neutrino Observatory [6] and to Takaaki Kajita of the Super-Kamiokande Neutrino Detection Experiment [7] for the discovery of neutrino oscillations.

In the Standard Model, the Yukawa coupling and the Higgs mechanism may be used to generate the masses of fermions [8] [9] [10]. The mass term in the Dirac Lagrangian can be expressed as the sum of the chiral states [10], shown in Equation [11] 1.2.

$$\begin{aligned}
 -m\bar{\psi}\psi &= -m\bar{e} \left[\frac{1}{2} (1 - \gamma^5) + \frac{1}{2} (1 + \gamma^5) \right] e \\
 &= -m\bar{e} \left[\frac{1}{2} (1 - \gamma^5) e_L + \frac{1}{2} (1 + \gamma^5) e_R \right] \\
 &= -m (\bar{e}_R e_L + \bar{e}_L e_R)
 \end{aligned} \tag{1.2}$$

where L and R represent left-handed and right-handed components respectively. This requires that *either* neutrinos are left-handed and no mass is generated through this process *or* there are right-handed neutrinos and neutrinos may gain mass. This mechanism would be the ‘traditional’ method of neutrinos acquiring mass as it is the same mechanism used for the other fermions.

As no right-handed neutrinos have been observed and other methods of generating neutrino mass have not been experimentally verified (such as, amongst other theories, Majorana neutrinos and the Seesaw mechanism [12] [13]), research is ongoing with a heavy emphasis on measuring the neutrino mass differences and oscillation parameters using long baseline neutrino experiments such as the MINOS and NOvA experiments [14] [15].

An overview of this thesis and my specific contributions are as follows. CHIPS is a small collaboration and accordingly each collaborator does a larger, more generalised proportion of the work than is usual in particle physics.

- Chapters 2 and 3 motivate the theory and experimental techniques for long baseline neutrino experiments and the CHIPS experiment.
- Chapter 4 details the design, construction and deployment process of the CHIPS detector. I was personally involved with many aspects of the design and almost every facet of the construction and deployment process. This thesis details my contribution to the construction process. I also operated and monitored the Madison section of the detector before its decommissioning. I

then contributed to the effort to recover the detector.

- Chapter 5 details the development of electronics and DAQ systems for future detectors. I designed, manufactured and tested the electronics and played a leading role in the overall direction of the project. Also, I was involved in most aspects of the software architecture, testing process and calibration.
- Chapter 6 describes a technique for monitoring the profile of a neutrino beam by measuring the muon flux in rock. Significant portions of the simulation were written by me with CHIPS collaborators completing the rest. The sensitivity study conducted using the code is my original work.
- Chapter 7 summarises and reflects upon the research presented in this thesis.

Overall, the three distinct components of my contribution represent three different stages in the life cycle of experimental particle physics:

1. Studies to see if effects are visible using a specific detector technique – a vital prerequisite before experiments can be conceived and developed. A physics phenomenon may cause an effect which may or may not be detectable; studies need to ascertain if an effect can be observed and measured and if any conclusions could be drawn from such measurements.
2. Designing and testing instrumentation at a small scale and preparing DAQ systems. This has to be done before experiments can be performed at full-scale.
3. Designing, constructing and deploying (and decommissioning) large-scale detectors.

Chapter 2

Neutrino Oscillations

2.1 Neutrino Mixing

Neutrinos interact according to the weak interaction under their three flavour eigenstates $|\nu_\alpha\rangle$, $\alpha = e, \mu, \tau$. However, neutrinos can also be defined in terms of three mass eigenstates $|\nu_i\rangle$, $i = 1, 2, 3$. However, there is not a direct one-to-one mapping between flavour and mass eigenstates. Instead, the flavour and mass eigenstates are linear combinations of each other:

$$|\nu_\alpha\rangle = \sum_{i=1}^3 U_{\alpha i}^* |\nu_i\rangle \quad (2.1)$$

$$|\nu_i\rangle = \sum_{\alpha=e,\mu,\tau} U_{\alpha i} |\nu_\alpha\rangle \quad (2.2)$$

where the coefficients U are elements in the unitary mixing matrix.

The conditions:

$$\langle \nu_i | \nu_j \rangle = \delta_{ij} \quad (2.3)$$

$$\langle \nu_\alpha | \nu_\beta \rangle = \delta_{\alpha\beta} \quad (2.4)$$

are enforced. If one set of neutrino basis states is required to be orthonormal and the coefficients of the matrix U are required to be unitary, then as a consequence the second set of neutrino basis states is also orthonormal.

Neutrinos, despite being observed in terms of the flavour eigenstates, propagate as a combination of the mass eigenstates. This means that the specific combination of mass eigenstates evolves as a neutrino travels and therefore the probability of measuring a given flavour of neutrino also changes with time or distance.

The neutrino mass eigenstates evolve in time according to the Time Dependent Schrödinger Equation with energy eigenvalues. Hence the time evolution of the flavour eigenstates can be expressed in terms of evolving mass eigenstates.

$$|\nu_i(t)\rangle = e^{-iE_i t} |\nu_i\rangle \quad (2.5)$$

$$|\nu_\alpha(t)\rangle = \sum_{i=1}^3 U_{\alpha i}^* e^{-iE_i t} |\nu_i\rangle \quad (2.6)$$

Combining 2.6 and 2.2 gives the time evolution of flavour states in terms of the initial flavour eigenstate.

$$|\nu_\alpha(t)\rangle = \sum_{\alpha=e,\mu,\tau} \left(\sum_{i=1}^3 U_{\alpha i}^* e^{-iE_i t} U_{\beta i} \right) |\nu_\beta\rangle \quad (2.7)$$

In natural units, using the relativistic limit and the binomial approximation, the energy of a given mass eigenstate is given by

$$E_i = E_\nu + \frac{m_i^2}{2E_\nu} \quad (2.8)$$

and the difference between two mass eigenstates becomes

$$E_i - E_j \approx \frac{\Delta m_{ij}^2}{2E_\nu} \quad (2.9)$$

where

$$\Delta m_{ij}^2 = m_i^2 - m_j^2 \quad (2.10)$$

is the mass-squared difference between two mass eigenstates.

Finally, in natural units, as neutrinos are travelling at almost the speed of light, the time a neutrino is travelling t can be approximated to the distance travelled L (the baseline). Therefore a transition probability can be constructed:

$$P_{\nu_\alpha \rightarrow \nu_\beta} = |\langle \nu_\beta | \nu_\alpha(t) \rangle|^2 = \sum_{i,j=1}^3 U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* e^{-i \frac{\Delta m_{ij}^2 L}{2E_\nu}} \quad (2.11)$$

This presents the probability of a neutrino oscillating between measurable flavour states in terms of two variables: the neutrino energy and the distance travelled. Furthermore, the oscillation probability depends on a number of fundamental constants: the neutrino mass eigenstate values and the parameters of the mixing matrix. This means neutrino oscillation experiments can be used to measure the mixing parameters and probe the neutrino masses [16]. Note, however, that oscillation experiments are sensitive to the difference of the mass eigenstates squared and not the absolute values. To directly measure neutrino masses other experimental techniques are required [17]. Although this formalism is described in terms of three mass eigenstates and three flavour eigenstates, neutrino mixing can be generalised for an arbitrary number of neutrinos.

2.1.1 Two Flavour Vacuum Oscillations

It is simpler to express neutrino oscillations if there are only two flavours and there is no medium for the neutrinos to interact with as they propagate. The orthonormal unitary mixing matrix (U) is represented as a rotation with mixing parameter θ and the difference in mass between the two mass eigenstates is given by Δm^2 .

$$U = \begin{pmatrix} U_{e1} & U_{e2} \\ U_{\mu 1} & U_{\mu 2} \end{pmatrix} \quad (2.12)$$

$$\begin{pmatrix} |\nu_e\rangle \\ |\nu_\mu\rangle \end{pmatrix} = \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix} \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \end{pmatrix} \quad (2.13)$$

The transition probability between the two flavour states is:

$$P(\nu_\mu \rightarrow \nu_e) = \frac{1}{2} \sin^2(2\theta) \left[1 - \cos\left(\frac{\Delta m^2 L}{2E}\right) \right] = \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 L}{4E}\right) \quad (2.14)$$

in natural units. In units of km for the baseline L , units of eV for the mass-squared difference Δm^2 and units of GeV for the neutrino energy E , the transition probability [18] becomes

$$P(\nu_\mu \rightarrow \nu_e) = \sin^2(2\theta) \sin^2\left(1.27 \frac{\Delta m^2 L}{4E}\right) \quad (2.15)$$

which is suitable for estimating muon-electron neutrino oscillations in neutrino beams.

2.1.2 The Matter Effect and Three Flavour Oscillations

For the full treatment of three-flavour neutrino oscillations, the flavour and mass eigenstates are related by:

$$\begin{pmatrix} |\nu_e\rangle \\ |\nu_\mu\rangle \\ |\nu_\tau\rangle \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \\ |\nu_3\rangle \end{pmatrix} \quad (2.16)$$

where the 3×3 mixing matrix is the Pontecorvo-Maki-Nakagama-Sakata (PMNS) matrix [19] [20]. The PMNS matrix can be expressed in terms of three mixing angles θ_{12} , θ_{23} , θ_{13} and a complex charge-parity (CP) violating phase δ (Equation 2.17). This assumes neutrinos are Dirac particles. If Neutrinos are Majorana particles, two more phases are required [21].

$$U_{\text{PMNS}} = \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \quad (2.17)$$



Figure 2.1: Neutrinos scattering in matter. All neutrinos can undergo NC scattering. Only electron neutrinos can undergo coherent forward scattering in electronic matter.

The PMNS matrix can be represented as:

$$U_{\text{PMNS}} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \quad (2.18)$$

where $s_{ij} = \sin(\theta_{ij})$ and $c_{ij} = \cos(\theta_{ij})$.

So far, neutrino oscillations have been presented as neutrino flavour mixing as a neutrino travels without interacting with a medium. In reality, neutrinos can interact with a medium: for example, rock in a beamline experiment. Since matter contains electrons, electron neutrinos have an additional interaction channel compared with muon or tau neutrinos. Neutrinos of all flavours can undergo neutral current (NC) coherent forward scattering (shown in Figure 2.1a) whereas only electron neutrinos can undergo charge current (CC) coherent forward scattering (shown in Figure 2.1b).

The NC interaction affects all neutrinos in the same way whereas the CC ν_e interaction affects only electron neutrinos. Electron neutrinos and electron antineutrinos are also affected differently as the universe is biased towards containing matter. This can enhance an experiment's sensitivity to the CP δ_{CP} violating phase in the PMNS matrix. The matter effect can be treated as a perturbation in the Hamiltonian modifying the effective mass of the neutrino as it propagates. This is known as the matter effect or alternatively the Mikheyev-Smirnov-Wolfenstein (MSW) effect [22] [23].

The matter effect and the full three-flavour treatment of neutrino oscillations can be combined to give a robust set of oscillation probability equations. Particularly useful is the $P(\nu_\mu \rightarrow \nu_e)$ oscillation probability (Equation 2.19) for beamline neutrino experiments as the neutrino beams are muon neutrinos which are directed through the Earth. $P(\nu_\mu \rightarrow \nu_e)$ is sensitive to δ_{CP} as the electron neutrinos experi-

ence the matter effect.

$$\begin{aligned}
 P(\nu_\mu \rightarrow \nu_e) = & \sin^2(\theta_{23}) \sin^2(2\theta_{13}) \frac{\sin^2(\Delta(1-A))}{(1-A)^2} \\
 & + \alpha \tilde{J} \cos(\Delta \pm \delta_{CP}) \frac{\sin(\Delta A)}{A} \frac{\sin(\Delta(1-A))}{1-A} \\
 & + \alpha^2 \cos^2(\theta_{23}) \sin^2(2\theta_{12}) \frac{\sin^2(\Delta A)}{a^2}
 \end{aligned} \tag{2.19}$$

Where $A = \frac{2\sqrt{2}G_F N_e E}{\Delta m_{31}^2}$, $\tilde{J} = \cos(\theta_{13}) \sin(2\theta_{13}) \sin(2\theta_{12}) \sin(2\theta_{23})$, $\Delta = \frac{\Delta m_{31}^2 L}{4E}$ and $\alpha = \frac{\Delta m_{21}^2}{\Delta m_{32}^2}$. G_F is the Fermi coupling constant and N_e is the electron density in matter. The \pm symbol represents the neutrino (+) and antineutrino (-) [24].

This equation, in the absence of unexpected extra neutrino flavours, describes accurately neutrino oscillations through matter.

Long baseline neutrino experiments such as CHIPS specialise in measuring θ_{23} and δ_{CP} .

Chapter 3

Neutrino Interactions

3.1 The Weak Interaction

In the standard model, neutrinos only undergo the weak interaction (and gravitation). The weak interaction is mediated by the W^\pm bosons (80.379 ± 0.012 GeV) and the Z^0 boson mass (91.1876 ± 0.0021 GeV) [25]. The high masses of the W^\pm and Z^0 reduce the strength of the weak force due to a $1/M_{boson}$ dependency; hence the name ‘weak force’ [26] [11].

Z^0 mediates *neutral current* (NC) events - where an uncharged Z^0 is exchanged - for example, neutrino scattering. Therefore, no charged particles can be produced making the NC interactions difficult to observe. W^\pm bosons mediate *charged current* (CC) events. For example with beta decay, where a W^- boson is exchanged and an electron with its complementary electron antineutrino is produced, the electron can be easily observed in detectors. Examples of NC and CC interactions can be seen in Figure 3.1.

Neutrino interactions in matter vary depending on whether a CC or NC interaction is taking place and the energy of the incoming neutrino. As the neutrino energy increases, different categories of interaction can take place. The energy re-



Figure 3.1: Examples of NC and CC neutrino interactions. The NC interaction: the exchange of Z^0 during neutrino-electron scattering. The CC interaction: a ν_μ interacts with a neutron in a nucleus exchanging a W^- to give a μ and a proton.

gion of CHIPS is 1 GeV to 10 GeV where there are five main categories of interaction [27] which are given below.

3.1.1 Neutral Current Elastic Scattering

Beginning at the lowest energies, under 1 GeV, NC elastic scattering (ES) involves the scattering of a neutrino off a nucleon. Mediated by the Z boson, no charged lepton is produced so no information about the neutrino's flavour can be discerned [28].

3.1.2 Charged Current Quasi-Elastic Scattering

Also dominating at the 1 GeV level is CC quasi-elastic scattering (QE). A neutrino scatters off a neutron to give a lepton and proton. (For an antineutrino scattering off a proton, a neutron and antilepton are produced.) The need to create the mass of the lepton causes this type of interaction to be known as a quasielastic interaction [29] [30].

3.1.3 Resonant Pion Production

After the QES region, resonant pion production (RPP) dominates in the 1 GeV - 2 GeV region. In RPP, the neutrino interacts with a nucleon to form a resonant excited state such as a delta resonance which then de-excites and produces a pion [29].

3.1.4 Deep Inelastic Scattering

Deep inelastic scattering (DIS) dominates interactions for neutrinos over 3 GeV in energy. The neutrino has enough energy to probe inside a nucleon into the individual quarks. A quark is ejected from the nucleus which initiates a hadronic shower that can be detected.

3.1.5 Coherent Pion Scattering

Coherent pion scattering (CPS, also known as coherent forward scattering, CFS) dominates low-momentum transfer interactions. A neutrino scatters off a nucleus and produces a pion in the forward direction. Very little momentum is transferred to the nucleon [31] [32].

3.1.6 Meson exchange current

When investigating the interactions of neutrinos with a nucleus, it is incorrect to assume that all nucleons are moving independently from each other. Standard models

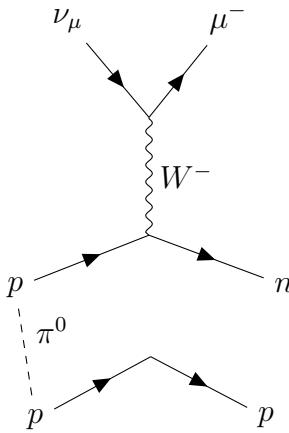


Figure 3.2: An example of MEC where a muon neutrino interacts with a proton in a nucleus to produce a muon and neutron. The two protons are correlated through the exchange of a π^0 leading to both resultant nucleons being ejected from the nucleus.

assume that all nucleons are independent but experimental data shows that this is not the case and nucleons may act in a *correlated* fashion - there may be a nucleus substructure. This advanced behaviour may be handled through more sophisticated modelling. A common model used is meson exchange current (MEC). In MEC, pairs of nucleons may be correlated - the two nucleons exchange a virtual meson (hence the name), a π^0 for example. A neutrino interacts with one of the nucleons through the exchange of a W boson, which gives rise to the complimentary lepton. In the process, *both* nucleons are ejected from the nucleus: two particles are ejected leaving two holes, which gives this process the alternative name 2p2h). An example of this process can be seen in Figure 3.2. [33] [34] [35].

3.2 Measuring Neutrinos

NC neutrino interactions do not produce a charged particle that leaves evidence of the neutrino's flavour. So, for neutrino oscillation experiments, CC interactions are used to detect neutrinos and infer the flavour. One technique used in neutrino oscillation experiments is measuring Cherenkov light.

3.2.1 Cherenkov Light

When a charged particle travels through a dielectric medium, the molecules become polarised - electrons are induced into an excited state. After the charged particle has passed, the medium de-excites which causes the emission of photons. The light is emitted isotropically. If the charged particle is travelling faster than the speed of light through the medium, the light constructively interferes to give a pattern in the shape of a cone expanding in the direction the particles travel (shown in Figure 3.3). This is analogous to the bow wave of a ship moving fast through water.

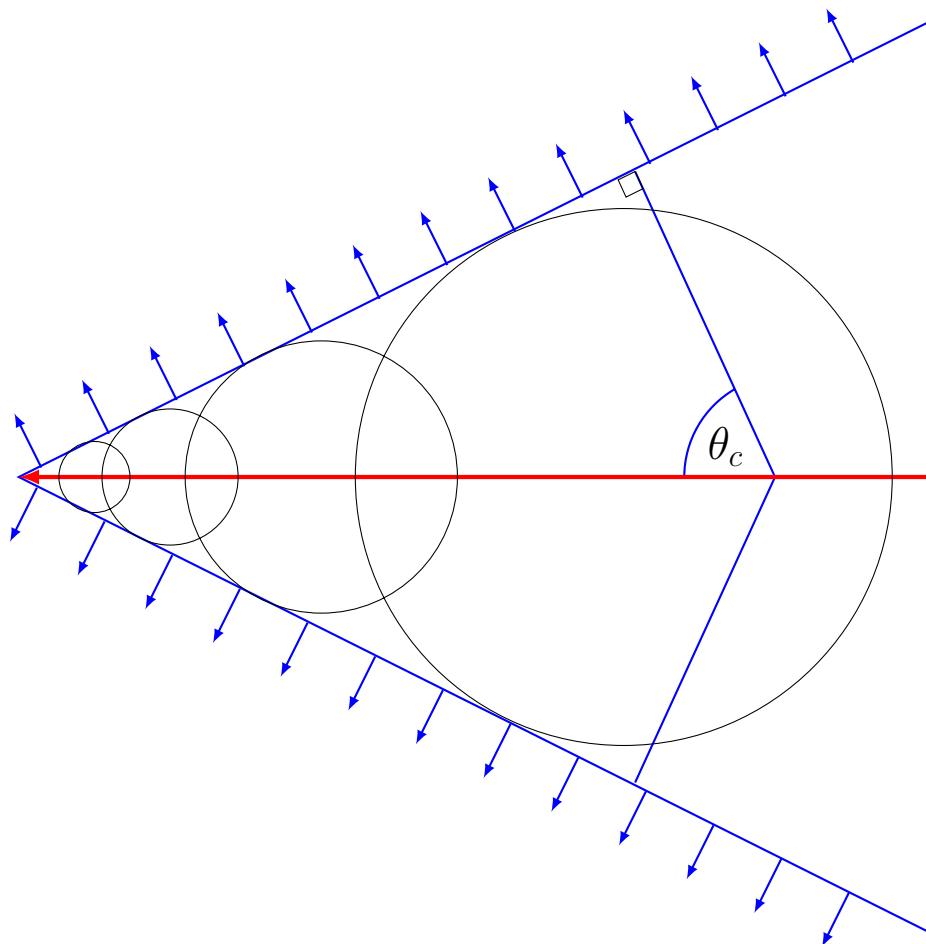


Figure 3.3: A charged particle (denoted by the red line) travelling through a medium faster than the speed of light through the medium. The black circles show the expanding spheres of light that are emitted as the charged particle passes. The spheres expand and interfere to form the characteristic cone of Cherenkov light.

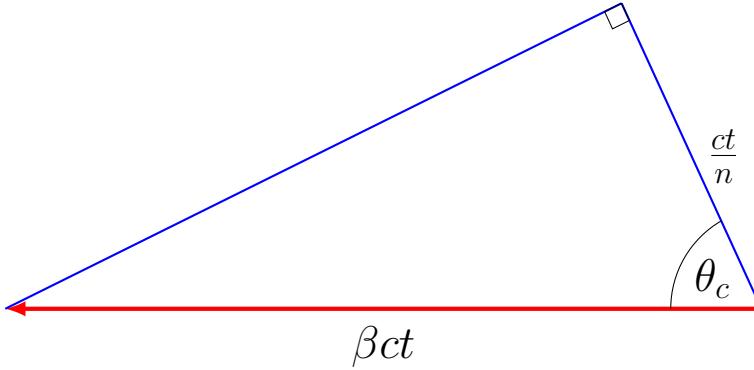


Figure 3.4: Construction showing the distances a charged particle and its Cherenkov light can travel; this gives the Cherenkov angle.

The shape of the cone is parameterised by the angle θ which is dependent on the refractive index of the medium n and the speed of the charged particle v .

Defining

$$\beta = \frac{v}{c} \quad (3.1)$$

the particle's velocity can be written as

$$v = \beta c \quad (3.2)$$

The speed of light through a medium is given by

$$c_{\text{medium}} = \frac{c_{\text{vacuum}}}{n} \quad (3.3)$$

Therefore, the following construction can be created using the distance the particle travels and the distance the emitted light travels as time progresses (shown in Figure 3.4).

$$\cos(\theta) = \frac{1}{n\beta} \quad (3.4)$$

This gives rise to a characteristic ‘ring’ shape that can be projected onto a flat surface along the length of the cone. Reconstructing the image of the ring allows the direction of the charged particle to be reconstructed. The total number of photons and their wavelengths are related to the velocity of the particle.

The number of photons emitted per unit distance travelled x and per unit wavelength λ is given by [25]

$$\frac{d^2N}{d\lambda dx} = \frac{2\pi\alpha z^2}{\lambda^2} \left(1 - \frac{1}{1 - \beta^2 n^2(\lambda)} \right), \quad (3.5)$$

which in terms of the Cherenkov angle is

$$\frac{d^2N}{d\lambda dx} = \frac{2\pi\alpha z^2}{\lambda^2} \sin^2(\theta_c), \quad (3.6)$$

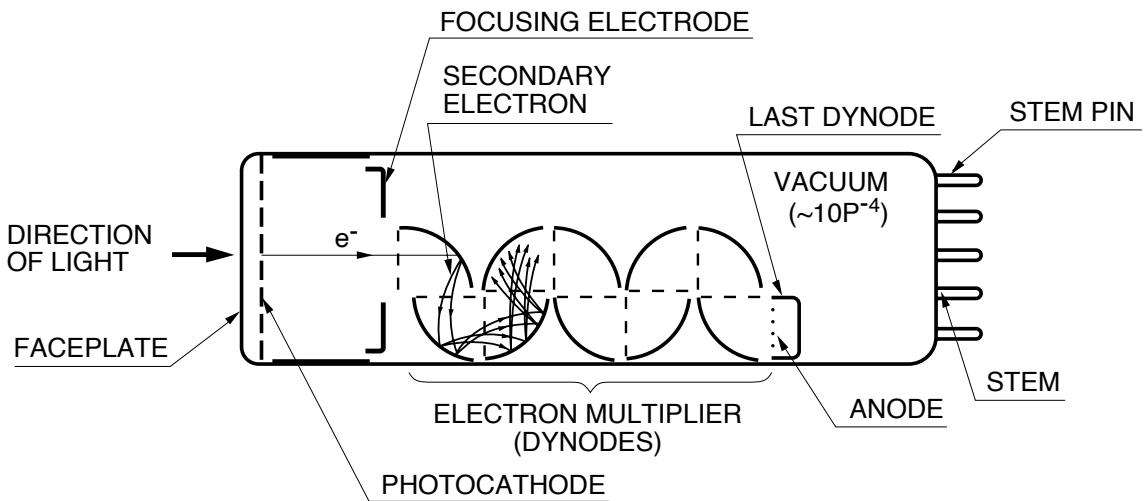


Figure 3.5: The internal structure of a PMT showing the photocathode, the chain of dynodes and the readout anode [36].

where α is the fine structure constant and ze is the charge of the particle. The interaction medium for CHIPS was water, a common medium for Cherenkov detectors. Water has a refractive index of 1.33 which gives the Cherenkov angle as 41.2° .

3.2.2 Photomultiplier Tubes

CHIPS uses PMTs to measure Cherenkov light produced by neutrino interactions.

PMTs are analogue instruments that produce an electrical signal when exposed to light. They are sensitive to single photons. PMTs consist of a photocathode - a negatively charged material which liberates an electron when struck by a photon - and a series of dynodes. When an electron is released, the electron is accelerated towards the first dynode and upon impact releases more electrons. These electrons are accelerated at the second dynode. This pattern continues causing an avalanche of electrons which eventually reaches the anode. Measuring the current of electrons at the anode indicates the arrival time and the number of photons that struck the photocathode. A diagram of the operation of a PMT can be seen in Figure 3.5.

An example of a PMT detailing the dynode staging can be seen in Figure 3.6.

Typically, PMTs are sensitive to light of approximately 400 nm to 700 nm in wavelength. This translates into approximately 210 photons/cm in a water Cherenkov detector, using Equation 3.6.

3.3 Neutrino Production

It has been shown that the neutrino oscillation probability is dependent on the neutrino baseline L/E . Furthermore, as neutrinos have a very small interaction cross section, an intense flux of neutrinos is required to observe a significant number of interactions. Therefore it is useful, when trying to measure the neutrino oscillation



Figure 3.6: The internal dynode structure of a PMT. The first dynode is at the front covered by a copper mesh and the chain of dynodes can be seen behind in an alternating ladder arrangement.

parameters, to have a very high neutrino flux and a baseline and energy which are tuned to maximise the expected neutrino oscillation probabilities. A neutrino beam provides a guaranteed source of neutrinos (with a high neutrino flux) that is potentially tuneable for oscillation experiments.

3.3.1 Neutrino Beams

All neutrino beams generate neutrinos using the same principles. Different neutrino beams have different optimisations and uncertainties but the basic concept remains the same [37] [38]. The Neutrinos at the Main Injector (NuMI) beam is the most powerful neutrino beam in the world [39]. The NuMI beam acts as the beam source for the CHIPS detector. Located at Fermilab, Illinois, the beam is aimed north-west through the Earth such that the beam emerges at the Soudan Mine in northern Minnesota.

The NuMI beam consists of a 120 GeV proton source which is fired at a graphite target to create a source of pions and kaons. The protons are delivered in ‘spills’ - bunches of protons $10\ \mu\text{s}$ wide delivered at $1.33\ \text{s}$ intervals. Each spill delivers approximately 10^{13} protons-on-target (POTs).

Neutrinos cannot be directly focused but their parent hadrons can be focused to form a narrower cone of neutrinos. The charged pions and kaons are focused by a pair of electromagnetic horns. Changing the polarity of the focusing horns will either focus positive particles and de-focus negative particles (known as *forward horn current*) or focus negative particles and de-focus positive particles (known as *reverse horn current*). This allows a beam of neutrinos or antineutrinos to be toggled. Particles which are not parallel to the magnetic field axis receive a smaller Lorentz boost which leads to lower energies off the beam axis.

The pions and kaons travel through a 675 m decay pipe where they decay into

muon neutrinos (shown in Equations 3.7 and 3.8 for a neutrino beam).

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \quad (3.7)$$

$$K^+ \rightarrow \mu^+ + \nu_\mu \quad (3.8)$$

The branching ratio of a pion decaying into a muon and a muon neutrino is 99.99% and the branching ratio of a kaon decaying directly into a muon and a muon neutrino is 63.56%. The next two dominant kaon decays are

$$K^+ \rightarrow \pi^0 + e^+ + \nu_e \quad (3.9)$$

$$K^+ \rightarrow \pi^0 + \mu^+ + \nu_\mu \quad (3.10)$$

with branching ratios of 5.07% and 3.35% respectively [25]. These are sources of an intrinsic electron-neutrino component of the beam. The muons produced in the decay pipe may decay into electrons and electron neutrinos which contribute to an unavoidable electron neutrino contamination in the beam. The difference in probability between Equation 3.7 and Equation 3.8 is because kaons have additional decay channels available. The mass of a kaon is large enough that kaons can decay into pions (the lightest mesons) on top of the direct decay into a muon and neutrino with no other particles. The production of positrons is heavily suppressed in the cases shown in Equations 3.7 and 3.8 compared with the case of Equation 3.10 because the former are two-body decays whereas the latter is a three-body decay; there is more phase space available for the electrons to have a lower momentum to avoid suppression by helicity in the three-body decay as the extra pion carries momentum. The three-body decay also results in a wider variety of neutrino energies as the pion also carries momentum from the parent kaon [10].

At the end of the decay pipe there are hadron absorbers to catch any un-decayed particles and muons; the neutrinos pass through into the Earth's rock. The muons are absorbed and the neutrino beam is complete. A diagram of the NuMI beam can be seen in Figure 3.7.

The NuMI beam has had several different configurations so the precise specifications of the beam have evolved over time and different detectors receive different flux spectra.

When in ‘low energy’ forward horn current configuration at the MINOS near detector, the NuMI beam is 92.8% ν_μ , 5.8% $\bar{\nu}_\mu$ and 1.3% ν_e & $\bar{\nu}_e$. The ν_μ contribution was produced through 87% π^+ decay and 13% K^+ decay. K^0 and μ decay contribute negligible fractions to ν_μ but are significant in the production of ν_e) [40]. There is an approximate 10% uncertainty on the ratio of π^+/K^+ [41].

For the off-axis experiment, NovA, the NuMI beam energy was increased to ‘medium energy’ levels. The energy peak in the MINOS near detector increased from

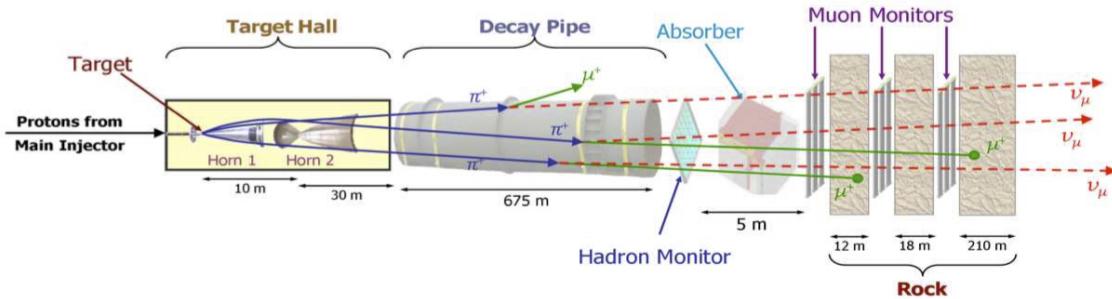


Figure 3.7: A representative diagram of the NuMI beam facility. Protons from the main injector strike a graphite target. The resulting pions and kaons are focused and decay to give muons and muon neutrinos. The muons are absorbed and the neutrinos continue towards particle physics experiments [39].

3 GeV to 7 GeV [42]. This energy was chosen because the Nova far detector, which is located 14.6 mrad off-axis and 810 km away from the proton target [43], receives neutrino energies at 2 GeV, close to 1.6 GeV - the oscillation maxima energy for the baseline [44] [45]. The energy spectra of the NuMI beam for different energy levels can be seen in Figure 3.8. In medium energy forward horn current configuration, the NuMI beam consists of 96.9% ν_μ , 1.9% $\bar{\nu}_\mu$ and 1.2% ν_e & $\bar{\nu}_e$ at the MINOS near detector [14].

Off-axis, the neutrino beam energy spectrum narrows as the angle from the axis increases. This allows a specific energy region to be selected close to the oscillation maximum and also reduces the high energy neutrino backgrounds, chiefly NC feed-down. Off-axis, the relative number of neutrinos at higher energies produced from kaons increases. An example of an off-axis neutrino experiment is Nova; the near detector neutrino energy spectrum can be seen in Figure 3.9. Like Nova, CHIPS was an experiment in the path of the NuMI beam but while Nova was 14.6 mrad off-axis, CHIPS was 7 mrad so the off-axis principle applied to both. The full off-axis effect can be seen in Appendix A [46].

A contributing source of neutrino flux uncertainty is the fraction of kaons in the parent hadron beam. Rather than attempting to distinguish between the type of hadron (pion or kaon) that is produced in the target and then recording them during beam operations (this is impossible to do), neutrino experiments use data from other external experiments to model the beam and target response. Proton beams are used to bombard targets and detect the hadrons produced and π^+ and K^+ so that the interaction cross sections can be measured [47] [48]. Both Nova and MINOS use neutrino flux data from their respective near detectors in conjunction with hadron measurements from dedicated experiments to tune Monte-Carlo simulations to model the NuMI beam. However, this technique introduces uncertainties due to the energy range of the hadron measurements not spanning the entire energy range of the NuMI beam and variations in the graphite target geometry. Furthermore, even less is known about the hadron yields in reverse horn current configuration.

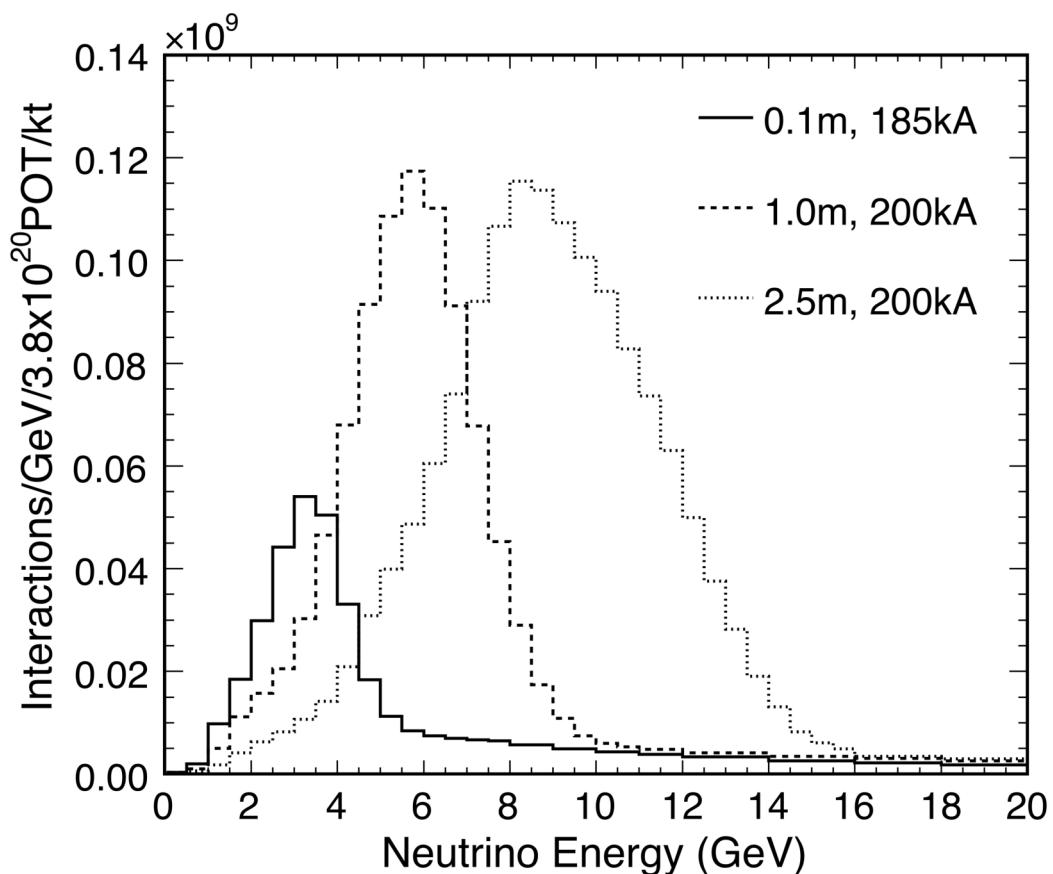


Figure 3.8: The flux of neutrinos at the MINOS near detector [42]. The solid line corresponds to a low energy neutrino beam, the dashed line corresponds to a medium energy neutrino beam and the dotted line corresponds to a high energy neutrino beam.

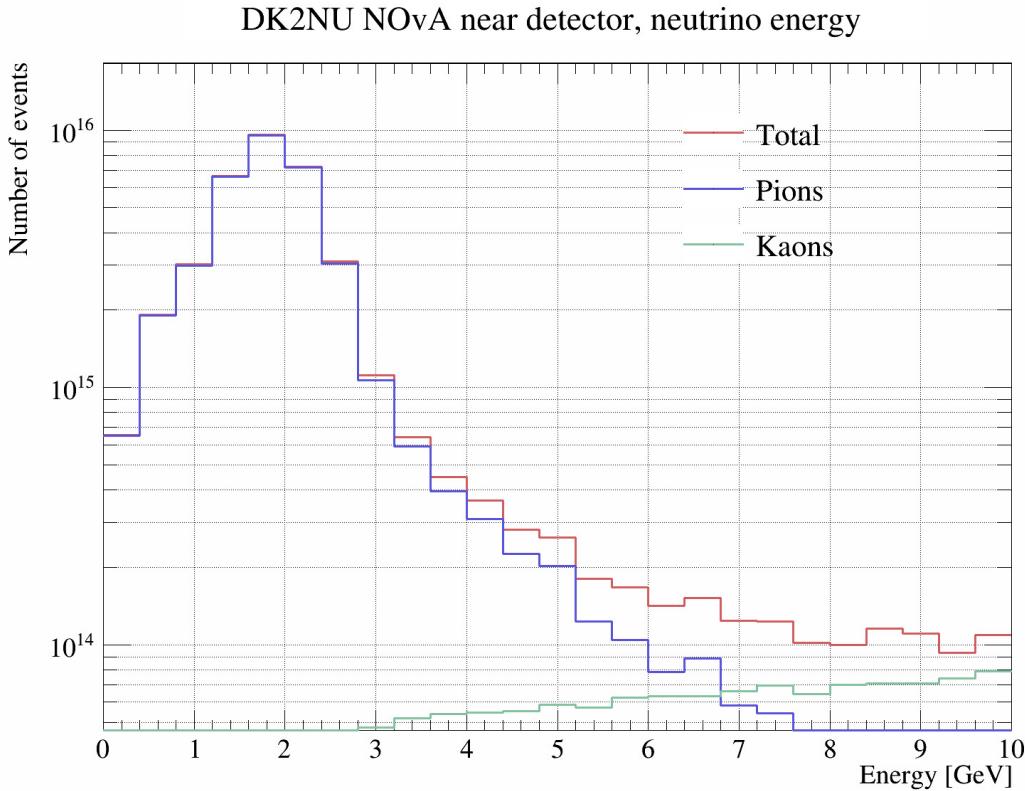


Figure 3.9: A simulation of the NOvA neutrino energy spectrum due to kaons and pions. The off-axis effect causes the peak energy peak to be lower but narrowed.

Because the fraction of neutrinos produced by kaons increases as the off-axis angle increases and energy increases, the uncertainty in the kaon content of the beam becomes increasingly important. In practical terms, there are two reasons why. Firstly, because the kaon component contributes ν_e events even at oscillation energy levels [49] which need to be modelled and accounted for. Secondly, high energy NC events will feed-down to lower energies and masquerade as signal events when they are in fact background events [50] [51] – NC scattering is difficult to distinguish from electron neutrino events causing electromagnetic showers [52]. Although this effect is reduced by a narrower neutrino energy spectrum, it is still important to understand.

Initially, in the NOvA experiment, the total hadron uncertainty was 21%, which contributed to a 6% uncertainty in the neutrino energy [53] - although this has been improved over time. In the MINOS experiment, the simulated spectrum of near detector CC ν_μ events can be seen in Figure 3.10 which shows that the number of kaons is not significant and becomes dominant at energies over 26 GeV. Figure 3.11 shows the ratio of pions and kaons in Fluka simulations (which was used in MINOS) [54] [55] and experimental observations from NA49 [56] and MIPP [57]. Figure 3.11 show that the uncertainties in the observations of the K^+/π^+ and K^-/π^- ratios are much larger than the uncertainties of just the pion only ratio π^-/π^+ . (The kaon only ratio, K^-/K^+ , also has large observational uncertainties.) These

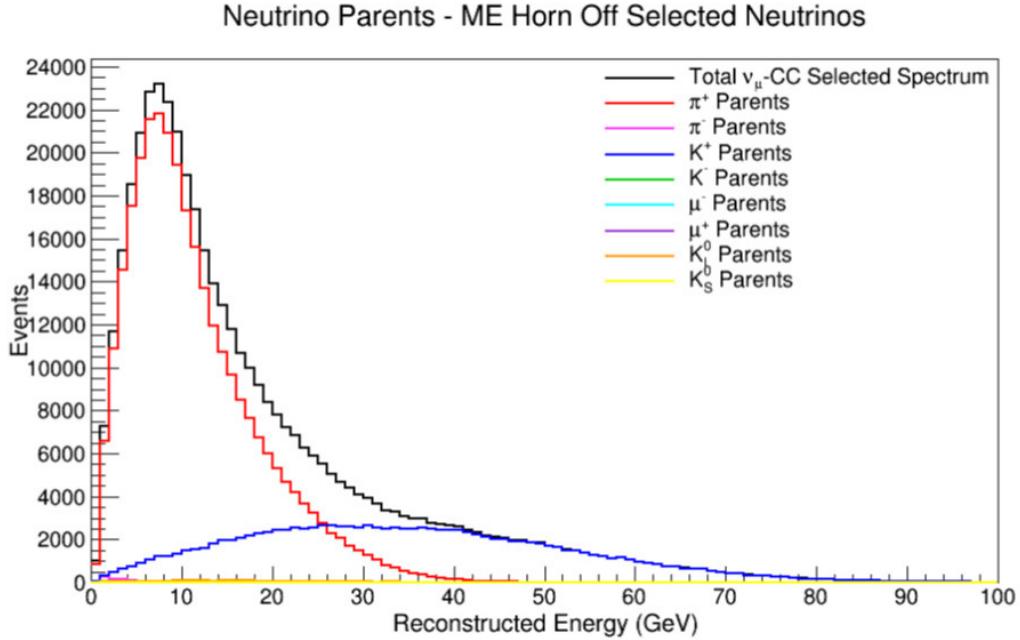


Figure 3.10: A simulated spectrum of ν_μ CC events in the MINOS near detector which shows the large number of neutrinos produced by kaon parents. Figure from [58].

measurements are used to inform and refine the tuning of the simulation. Therefore, large uncertainties in the measurements cascade into inaccuracies of the simulations. Overall these figures show that the number of kaons significantly contribute to the neutrinos produced and the uncertainty in the number of kaons is also significant. Furthermore, the branching ratio shown in Equation 3.9 combined with the fact that π^0 can decay into (unfocused) electron and muon neutrinos means that the kaon composition of the hadron beam impacts the muon neutrino purity of the neutrino beam.

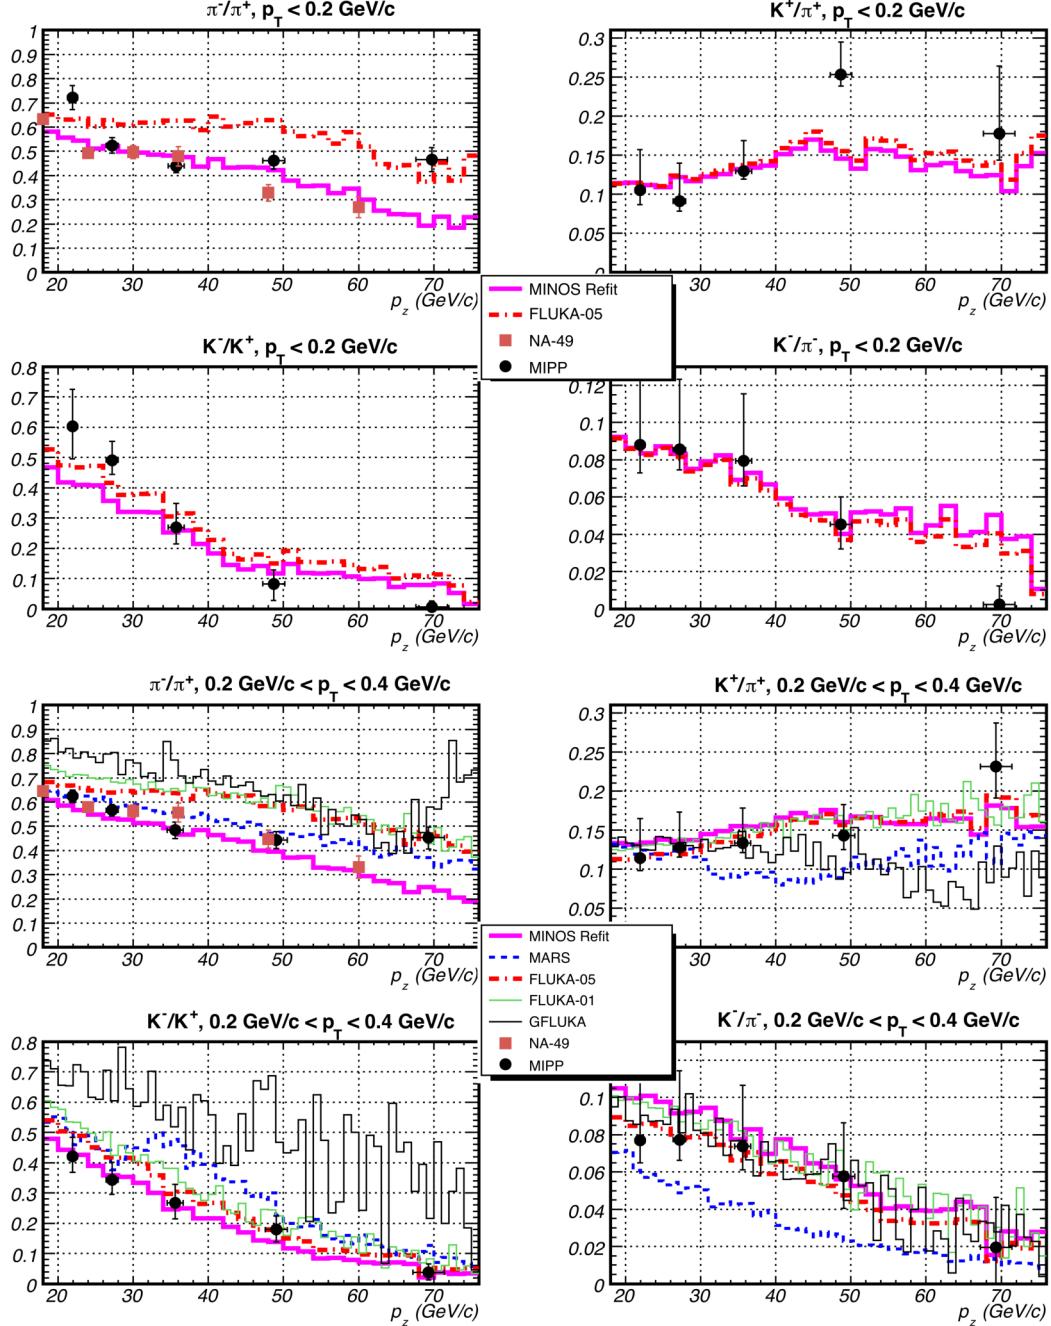


Figure 3.11: The ratios of pions and kaons across the range of longitudinal momenta with different transverse momenta cuts. The larger error bars for plots with kaon ratios can be seen. Figure from [58] [59]. To experimentally measure the pion-kaon production ratio, a secondary proton beam was fired at test targets and the hadrons produced were counted and identified using a number of instruments in the MIPP detector system (including time projection chambers and Cherenkov detectors) [57] [60].

Chapter 4

The CHIPS Detector

Due to the unlikeliness of neutrino interactions, neutrino detectors are traditionally very large and therefore expensive. Due to cosmic ray background, these detectors are often in locations very deep underground which are expensive to create and difficult to access. Hyper-Kamiokande is projected to cost over \$600m for a 258 kt detector ($\$2.3\text{m kt}^{-1}$) [61] [62] and the Deep Underground Neutrino Experiment (DUNE) is projected to cost over \$3bn (including the neutrino beam facilities) for a 70 kt module ($\$43\text{m kt}^{-1}$). Of this, over \$2bn will be used for excavating underground caverns and constructing buildings [63]. The CHIPS experiment was a prototype research and development project focusing on lowering the cost and barriers of entry in neutrino detector construction both through careful placement and design of the detector to ease construction and optimising the instrumentation for cost efficiency.

4.1 Experiment Overview

CHIPS was a prototype long baseline neutrino experiment consisting of a large water Cherenkov detector in a mine pit in Northern Minnesota, 712 km away from the NuMI beam at Fermilab near Chicago (shown in Figure 4.1). The goal of CHIPS was to measure ν_e appearance and ν_μ disappearance to demonstrate the observance of neutrino oscillations within a budget of about $\$200\text{k kt}^{-1}$.

The detector design comprised a volume of purified water wrapped in a waterproof and light-tight plastic liner with 1860 PMTs (in phase one) lining the walls facing inwards arranged in detector planes. Neutrinos interacting with the water within the detector would produce Cherenkov light in the darkened environment which could then be observed by the PMTs. The instrumentation was mounted to steel frames which formed a cylinder which was then wrapped in light-tight plastic. The detector was to be placed at the bottom of a flooded open-pit iron mine. The detector used as many standardised off-the-shelf components and readily available building materials as possible. It was designed in such a way that non-specialists

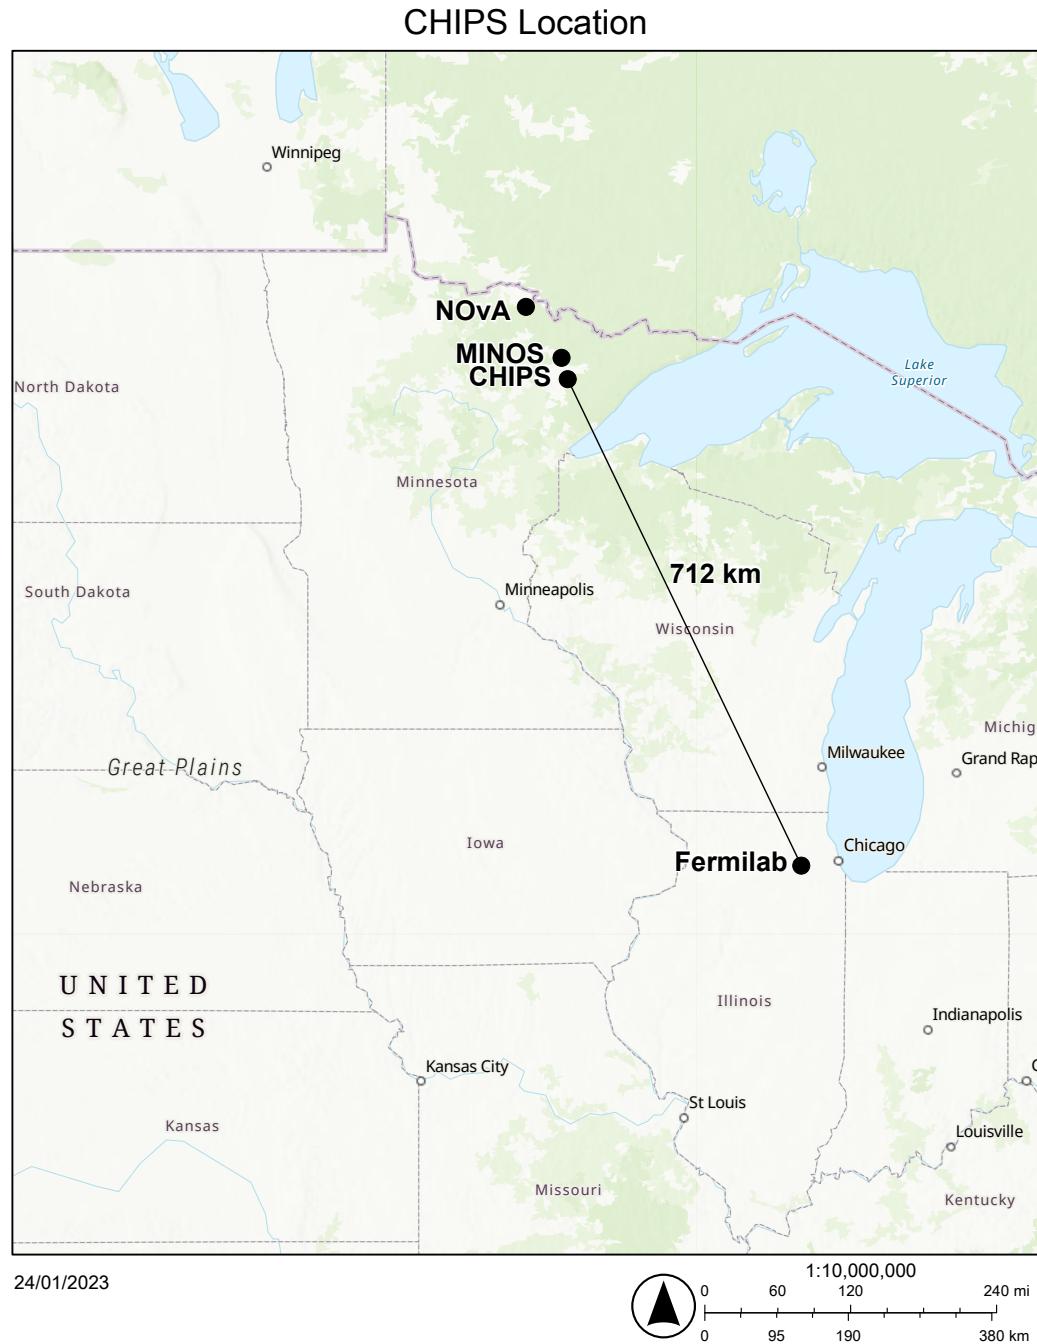


Figure 4.1: Map showing the NuMI beam source at Fermilab and the mine pit where CHIPS was located in Hoyt Lakes, Minnesota [64].

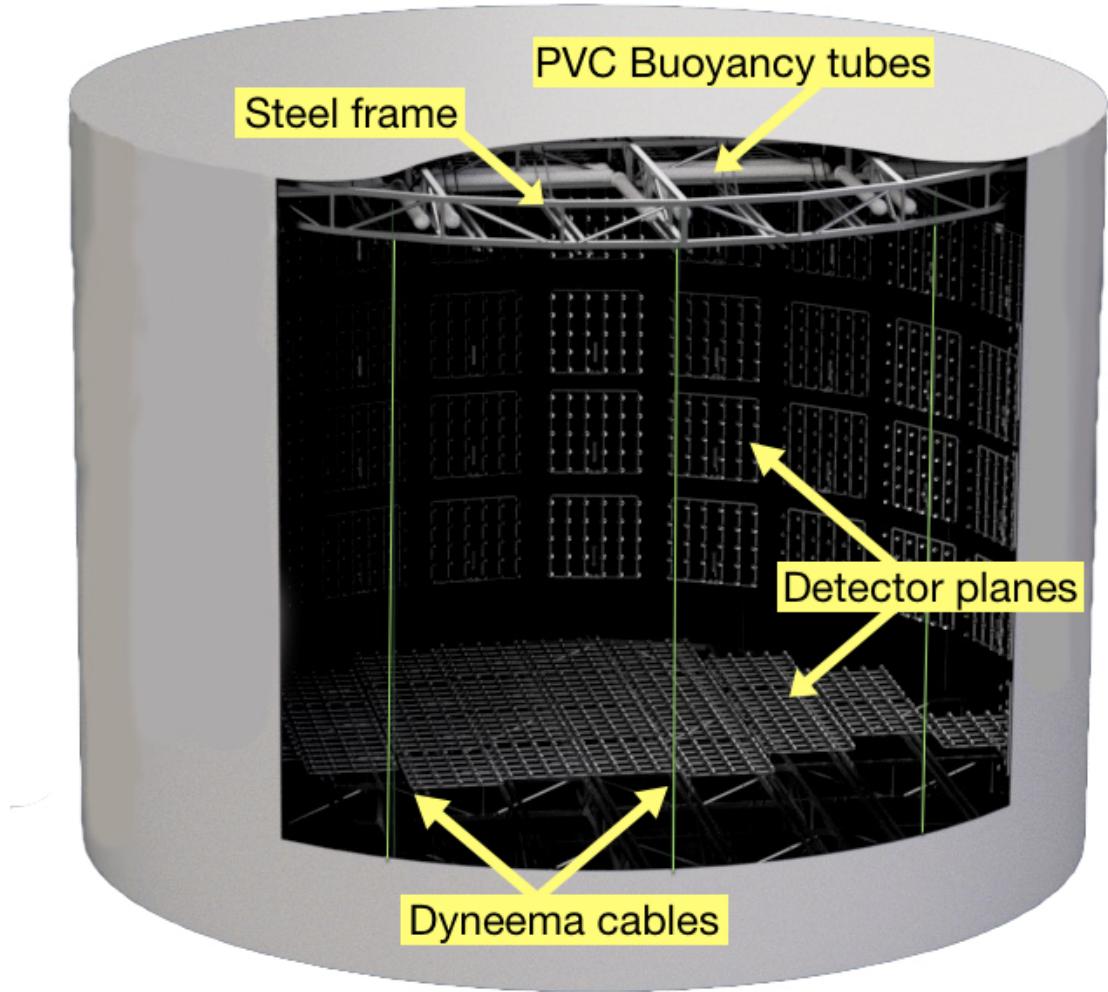


Figure 4.2: A rendering of the CHIPS detector showing the overall structure: two steel end caps supporting detector planes with buoyancy in the top cap. The detector stretched out to be a long cylinder and the complete detector was wrapped in liner. Here detector planes on the walls of the cylinder are shown which was a future upgrade path for CHIPS and not installed for the initial deployment. Image courtesy of Thomas Dodwell.

(including students) could perform the majority of the construction tasks safely. As a consequence, there was a significant saving in construction costs in comparison with those of other neutrino detectors. An overview of the detector can be seen in Figure 4.2

Overall, this concept provided three main advantages:

- Free overburden: the detector did not need to be buried underground so no expensive mining was needed to be performed or rocks placed on top. The detector was submerged in approximately 60 m of water - compared with 2700 m water-equivalent overburden for the Super-Kamiokande detector [65]. A detector 24 m in diameter and 24 m tall with no overburden results in a cosmic ray rate of approximately 80 kHz. The same detector resting at the bottom

of a 60 m lake gives a cosmic ray rate of approximately 20 kHz representing a factor of four decrease due to the overburden [66].

- Ease of construction: the detector could be built in its final form above ground and deployed to its underwater location in one piece (see Section 4.6). This saved design and construction time as all the work could be performed in facilities at ground level and the detector did not need to fit in a disassembled form into a mine shaft elevator.
- Commonly-available building and easily-found materials as well as off-the-shelf components were cheaper and easier to procure than bespoke electronics and unusual materials. This also made it faster to complete a detector from conception to deployment.

The detector was assembled from April to October 2019 when it was then deployed. The detector was decommissioned in July 2020 due to external factors relating to the COVID-19 pandemic. The instrumentation is in the process of being reused.

4.2 The PolyMet Mine

CHIPS was based within the PolyMet mine site, in a disused and flooded iron mine, now used as a buffer pit for managing water levels across the site. Previously the iron was extracted from the ground using open pit mining. This became unprofitable in the late 20th century and the mine fell into disuse in the early 2000s. The mine site holds the second largest copper reserve in the world, as well as substantial nickel and rare earth metal reserves that are more valuable in the present day than they were when the mine was decommissioned. Consequently, the mine site is being recommissioned and the infrastructure refurbished for fresh mining.

The PolyMet site was uniquely useful to CHIPS for several reasons. Firstly, the Wentworth 2W (w2w) abandoned flooded mine pit provided a suitably deep location (60 m) for the detector. Secondly, the existing infrastructure that was being maintained for future mining operations provided access by road and rail despite the remoteness of the location and also gave access to heavy machinery and industry for construction. Finally, the location of the mine site itself was important. The site was within accessible distance to the MINOS and Nova sites for accessing equipment and tools. The site was also in the path of the NuMI beam as it re-emerged from the Earth (shown in Figure 4.3). It was for these reasons that the w2w was selected as the location of the CHIPS detector.

During the lifecycle of CHIPS, the w2w mine pit was deliberately flooded by the mining company. The pH level and balance of chemicals in the water of the mine pits were tightly regulated for environmental reasons. As rain water ran into the

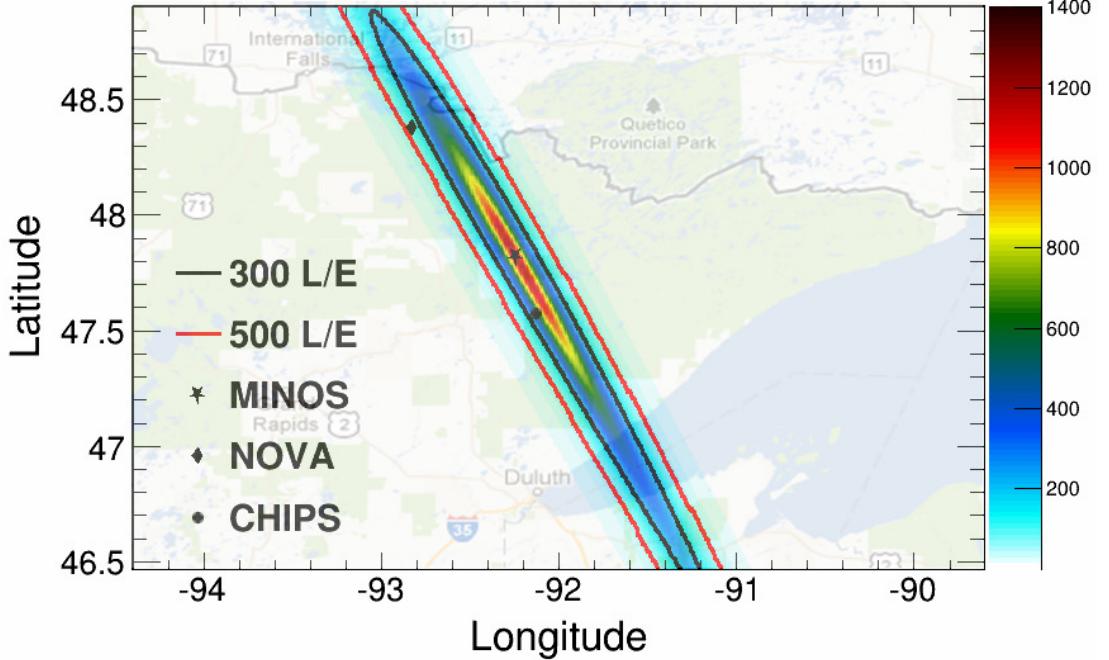


Figure 4.3: A map of the Numi beam exiting the Earth’s surface showing the positions of CHIPS, NOVA and MINOS. The z-axis shows the expected number of neutrino events per year per kiloton of water assuming no neutrino oscillations. Contours of constant L/E are shown. Image taken from [67].

mine pits through different sections of the mine site, the pH and chemical balance of the water changed. Water was pumped between the different mine pits to maintain safe levels and pH. In autumn, water was pumped into the CHIPS pit and then pumped out over winter into other pits. Therefore, the water level of w2w changed seasonally and it was possible to use this to advantage. A map of the PolyMet site including the w2w mine pit and the construction area can be seen in Figure 4.4. The Construction area at the w2w pit can be seen in Figure 4.5.

4.3 Detector Frame

The CHIPS detector was cylindrical in shape (25 m in diameter and extendable to 12 m tall). The two circular ends of the cylinder were made up of two lattice frames of stainless steel, which were known as the *end caps*. The lattice construction was for a high strength-to-weight ratio and the stainless steel was to resist corrosion in the water so that rust did not inhibit the water’s clarity. Pieces of the frame’s main structure can be seen in Figure 4.6. Rather than manufacturing perfect 25 m circles, the end caps were approximated as two icosikaioctagons (28-sided regular polygons) and consisted of a main frame which gave each end cap its rigidity, strength and physical shape. Steel beams, known as *stringers*, were mounted facing the interior of the cylinder in rows to which the instrumentation was attached (shown in Figure



Figure 4.4: The PolyMet mine site. CHIPS was built beside the mine pit on a flat construction with additional assembly done in the PolyMet building. The detector was then deployed underwater in the mine pit [27].

4.7). Overall each end cap weighed 14 t and used approximately 3600 bolts.

The frame was designed to be prefabricated in a factory off-site and shipped to the mine pit. Each piece was designed to be movable and assembled with a single telescopic forklift rather than an industrial crane or multiple pieces of machinery; the stringers were movable by either a single person or two people depending on the size of the stringers. This meant that, other than a forklift operator, the frame could be assembled by a team with as few as three people; it also meant that only two workers at any time needed specialist training (forklift driver and forklift spotter/director). The main structural pieces were of gusset plate construction with bolts and slotted bolt holes which are visible in Figure 4.6. The stringers were bolted to brackets attached to the frame, an example of which can be seen in Figure 4.8.

Since cylinders have two parallel circular caps, the CHIPS frame end caps were constructed on top of each other. The bottom cap was built on the ground with legs that had feet made from tyres (also visible in Figure 4.6). The bottom cap had stainless steel pegs pointing upwards; on top of the pegs, stilts made of steel tubing were mounted and the top cap was built on top of the tubing. The top cap sat on pins which were inserted into the stilts which could be removed during detector deployment. The stilts could then be removed from the detector once it was in the water. The legs were only necessary during construction because, once deployed, the bottom cap sank and the top cap floated and the caps no longer needed to be rigidly



Figure 4.5: The Wentworth 2W mine pit. The body of water in the foreground shows the mine pit itself and the CHIPS construction site which was the slipway leading into the water. The slipway was built on the road where the excavators would have driven in and out of the pit when the mine was still operational. The detector can be seen partially constructed along with the building site apparatus.



Figure 4.6: Pieces of the bottom end cap. The gusset plates used to join pieces are visible as are the tyres which the frame sits on. The hollow lattice style design saves weight as less steel is required than for a solid detector.



Figure 4.7: Rows of stringers bolted to the main support structure. Detector planes would be attached on top of the stringers.



Figure 4.8: A mounting bracket welded to the main frame to which a stringer was bolted.



Figure 4.9: The completed frame. The pillars supporting the top cap are clearly visible.

attached - once that happened the detector could be extended longer than the steel tubing length. The completed frame can be seen in Figure 4.9. 28 Cables made from Dyneema connected the top and bottom cap together. When the bottom cap hung from the top cap, the length of the cables determined the maximum difference in height between the two caps and therefore the length of the detector. Dyneema was selected due to its low creep (the extension of a rope when subject to long-term loads).

4.4 Liner

The CHIPS detector needed sealing for two reasons. Firstly, the detector needed to be light-tight so the only light detected was Cherenkov light. Secondly, pure water inside the detector had to be separated from the dirty water outside the detector for clarity. In addition, during the deployment procedure, the liner acted as a hull for the detector which displaced enough water for the heavy detector to float like a boat.

The liner was made from flexible fibreglass reinforced plastic. It was strong and light-tight, usually used as roofing material for flat roofs. Special welding equipment was used to fuse different pieces of liner using heat, the results of which can be seen in Figure 4.10. Although the plastic liner was strong, great care was taken to protect it during the construction phase of the experiment. Sharp metal on the detector was covered with round plastic to ensure that the liner was not pierced, and wooden boards were put over the floor liner on the ground to protect it from hazards such as falling tools, sharp ladders or personnel walking on it.

The liner was delivered in 3 m wide rolls but the detector was 25 m wide and

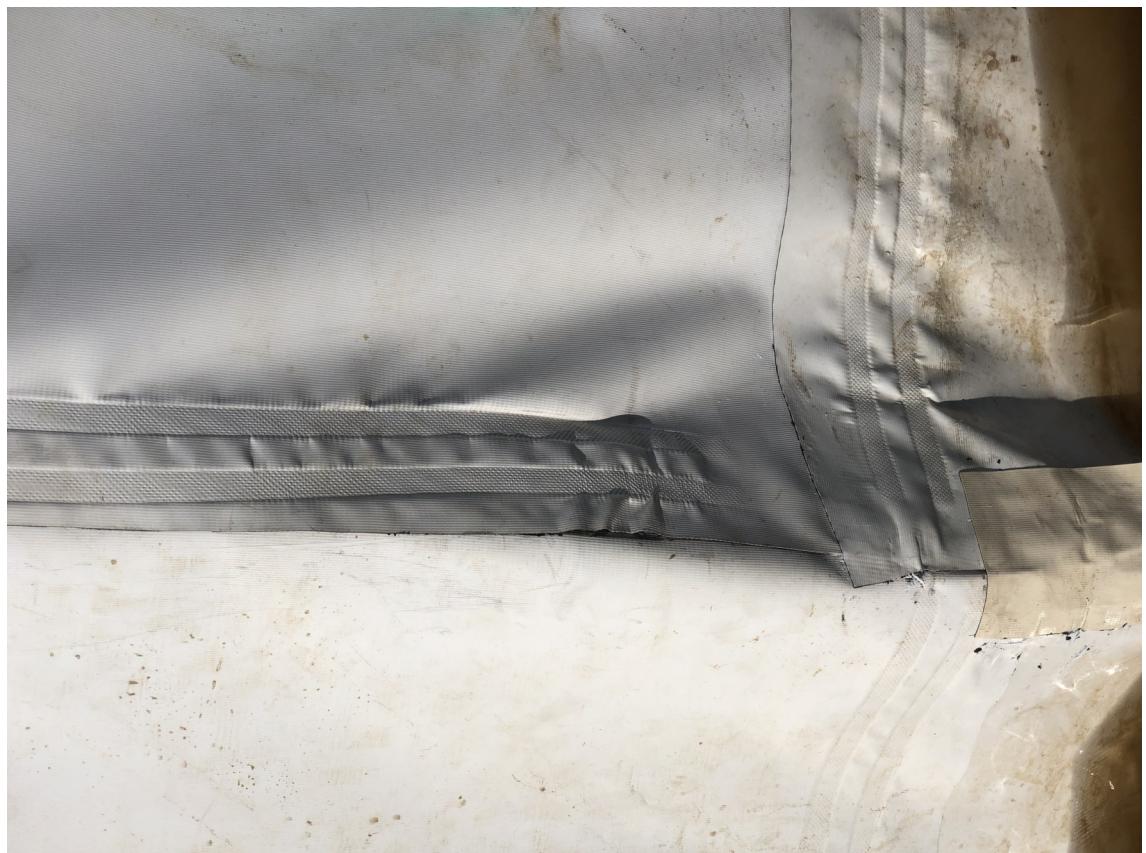


Figure 4.10: Two pieces of liner fused together (unfinished). Hot air is blown in the gap between the two pieces until the melting point is reached. The two pieces of liner are clamped together. The joints were not individually tested due to the circumstances of construction. However, prior testing showed they were expected to be light-tight if properly performed.



Figure 4.11: The detector completely wrapped in liner. Additional liner was folded around the base below the water line so that the detector could stretch to be 12 m tall.

12 m tall so the liner was assembled in the style of kirigami. Strips of liner were cut and welded together into two large circles. One was placed under the detector to form the bottom of the detector; the detector's rubber feet were on the outside of the liner and bolted through to steel plates on the inside of the detector. Expanding rubber washers were used to seal the bolt holes in the liner from light and water ingress. The second circle of liner was hoisted onto the top of the detector with the fork lift to form the top. 12.5 m strips of liner were welded vertically to form the walls and the vertical strips were welded to each of the end cap circles to seal the detector. As the detector was not 12.5 m tall until its final deployment the liner was stored (like a rolled down sock) around the base of the detector.

4.5 Buoyancy and Floating Dock

The design of the detector was guided by three axioms:

1. The bottom cap would sink in water.
2. The top cap would float in water.
3. The total detector would sink in water once it had been filled with water.

The bottom and top caps were made of steel and minimal buoyancy was introduced by the hollow detector planes so that both caps would sink. The bottom cap was left without additional buoyancy whereas the top cap required additional buoyancy so that it would float. 12 inch PVC piping was used for flotation in the



Figure 4.12: The 12 inch PVC buoyancy pipes mounted in the top cap to provide flotation.

top cap. The pipe was sealed at both ends with PVC end caps which were glued onto the ends. The top cap, in water, weighed 11.6 tonnes, which needed offsetting. PVC pipes equating to 14 tonnes of buoyancy were installed in the detector and can be seen in 4.12. The buoyancy and weight budget can be seen in Table 4.1.

The detector needed to sink but it needed to sink in a controlled manner. Since overall the detector was heavier than water, the sinking could not be controlled without some external intervention. A floating dock with winches was used to lower the detector to the lake bed.

The floating dock consisted of a ring of floats that were welded to a circular frame which surrounded the detector. One quarter of this can be seen in Figure 4.13. The floating dock provided 17.6 tonnes of buoyancy when the floats became fully submerged. This offset the entire weight of the detector when in the water. The floating dock had 12 winches attached to it which ran steel cables down to the detector. The cables were attached to pulleys bolted to the outside of the detector through the liner. A winch can be seen in Figure 4.14. They were designed to lower the detector to the bottom of the lake and raise it back up for upgrades and recovery.

The floating dock was custom-fabricated using steel tubing welded together and standard pontoon dock floats bolted to the bottom. Winches were added through the centre so the load was in the middle of the floats. The dock was designed to support the detector when it was static and in the centre of the floating dock.

Buoyancy Budget		
Object	Weight (air) [kg]	Weight (water) [kg]
Frame caps	26750	23272
Stilts	960	835
PVC flotation	4236	-14700
Liner	2000	0
Electronics boxes	200	200
Detector planes	3880	0
Floating dock	1000	-13500
Complete top cap	21591	-2229
Complete bottom cap	16435	11636
Combined caps	38026	9407
Total system	39026	-4093

Table 4.1: The buoyancy budget for CHIPS. The top and bottom halves were not the same. The top half had a steel frame, the steel stilts, the PVC flotation, half of the liner, a set of detector planes and the electronics boxes. The bottom cap had a steel frame, half of the liner and a different set of detector planes. The detector planes were assumed to contribute no flotation or dead weight to the final system. This was because the raw materials were approximately neutrally buoyant but the planes contained air so had the potential to provide buoyancy. However, it was unknown how many planes would leak so no buoyancy was assumed. It would be easier to add ballast to the final detector compared with adding additional flotation. The table shows that the total detector would sink to the bottom of the lake with the top cap floating above the bottom and that, with the addition of a floating dock, the detector would not sink.



Figure 4.13: Two one-quarter segments of the floating dock. Three winches spaced between six floats were mounted on top.



Figure 4.14: One of the winches used in CHIPS mounted onto the floating dock. The 12 winches could lift a total of 19 tonnes of dry mass.

4.6 Deployment Procedure

To recap, the CHIPS detector was designed to be built on dry land and transferred to the water. This required a special procedure due to the size and weight of the detector. The mine pit was flooded with water. The water level level changed seasonally because the mining company transferred water between mine pits in a predictable manner. The detector was built at the bottom of a slipway that was above the water level in the summer months but below the water level during the autumn months. The detector was completely encapsulated by the waterproof liner. Despite the weight, if the detector was sealed, the liner displaced enough water with the result that the detector floated. If the liner was filled with water, the detector would sink. The bottom cap of the detector acted as dead weight while the top cap floated and the liner was required to maintain net buoyancy.

The sequence of the procedure was as follows

1. The detector would be built on dry land and sealed; when the water level rose the detector would be flooded so that it would remain in place until deployment.
2. The water would be drained from the detector using electric pumps so that it would float at the appropriate moment.
3. The detector would be towed by a boat to its deployment location (shown in Figure 4.15).
4. The umbilical pipe containing the power and data connections would be connected to the passthrough plate. The two high-density polyethylene (HDPE)



Figure 4.15: The detector being towed from the construction site to the deeper water.

pipes for circulating the water would be connected so that water could be added and removed from the detector using electric pumps.

5. The detector would be filled slowly with water. This would cause the bottom cap to sink and keep sinking as more water was added. This would stretch out the detector to its full length as the liner expanded as more water would be added. The top cap would keep floating on the surface until the detector had reached its maximum height.
6. When the detector was full of water, the detectors would be held by the floating dock which would begin to lower down the detector slowly.
7. The detector would be lowered to the bottom and disconnected from the floating dock. The weight of the bottom cap would keep the detector in place; the buoyancy of the top cap would keep the top cap floating above the bottom cap to give the detector its height. Overall the detector would remain sunken. At this point, the detector would be fully deployed.

4.7 Structure Summary

The steel frame, the liner, water and the flotation all worked together to give the detector its final structural form. The end caps gave the cylinder its rigid round cross-sectional shape and acted as a mounting platform for detector planes. The length of the detector came from a combination of the liner, water and Dyneema ropes. The liner was manufactured as a cylinder and when it was filled with water it expanded and this stretched the detector to its final length. The Dyneema ropes

stopped the bottom and top caps expanding too far apart in order to avoid damaging the liner. Since the detector was filled with water and the detector itself was under water, the water on the outside of the detector supported the water on the inside of the detector. This meant the frame only needed to bear its own weight and the weight of the instrumentation. The water surrounding the detector bore the weight of the water within. Therefore, the detector needed no more than to be strong enough to hold itself together rather than being a load-bearing structure to support the weight of the water contained within.

4.8 Water Purity

To maintain water clarity inside the detector, the water would have to be filtered. A water filtration plant was housed in a hut on the shore of the pit (known as the *water hut*). The hut housed ten pairs of $10\text{ }\mu\text{m}$ and $0.5\text{ }\mu\text{m}$ filters in parallel (shown in Figure 4.16) and the water was to be pumped into the hut from the detector, through the filters, then back out to the detector. The input and output pipes into the hut can be seen in Figure 4.17. During nominal operation, the water system would have formed a closed loop so the water could be reused. To fill the detector for the first time during deployment, the water would be pumped into the detector from the lake. The system was designed to pump up to half a million litres per day; this was achieved during testing. Additionally, attenuation lengths of 133 m were achieved [68], an order of magnitude larger than the 25 m diameter of the detector.

To get the clean water to and from the detector, HDPE pipe was used. Two lengths of HDPE pipe were prepared by butt welding segments together until each was over a kilometer. The butt welding equipment can be seen in Figure 4.18 and the two completed pipes can be seen in 4.19. HDPE was chosen as a compromise between its price, its strength and its availability and had the advantage that, unlike a flexible hose, it could resist the partial vacuum produced by the pump unlike a flexible hose.

To get the water into the detector, a metal bulkhead plate was installed onto the edge of bottom cap (shown in Figure 4.20) which could be used to attach pipes and cables while maintaining the integrity of the liner. The two HDPE pipes were attached to the plate using flanges. The plate was also used for the pipe containing the main power supply and optical fibre for the detector's main control box.

In winter, when the temperature dropped below freezing, the water at the bottom of the lake remained warmer than water higher up in the lake. As long as the pump remained turned on then the warm water from the bottom would circulate through the system and stop ice forming. If the pumps had stopped then the water would have stopped moving and then frozen which would render the system irrecoverable. This technique was tested over previous winters during prototyping.



Figure 4.16: The interior of the water hut showing the filters. Image courtesy of Thomas Dodwell.



Figure 4.17: The exterior of the water hut showing the passthrough between the black HDPE pipe and the inside of the hut.



Figure 4.18: The equipment used to butt weld two sections of HDPE pipe together. The ends of the two pipes were shaved down to be smooth and flush with each other. A heated plate was placed in between until the ends had melted. The two sections of pipes were pressed together and the melted sections were fused together. The process was precisely controlled for standardised joints.



Figure 4.19: The completed inflow and outflow water pipes for the detector. The short pipe on the surface is the water intake for filling the detector during deployment.



Figure 4.20: The plate where the flanges on the end of the HDPE pipe were bolted onto the detector. The liner was clamped between the plate and the flanges to maintain light and water proofing. This is a standard technique to make waterproof and light-tight seals.

4.9 Detector Planes

The detector was filled with PMTs pointing inwards. The PMTs required a high voltage power supply; they also required a system for digitising and logging their outputs as well as a trigger system. This meant that electronics were needed in the detector. The PMTs were arranged into flat planes with an electronics container at the centre. The requirement of a detector plane was to provide a mounting system for the PMTs and a safe waterproof area for the electronics.

There were two kinds of detector planes in CHIPS, named after the locations at which they were created:

1. Nikhef planes: which were based on KM3NET electronics [69] modified for the CHIPS geometry.
2. Madison planes: which were designed by CHIPS at the University of Wisconsin-Madison in conjunction with the IceCube experiment [70].

Both kinds of plane placed value for money and ease of manufacture as key design priorities. Also, both had two aspects of their design in common: PVC pipe construction and Cat 5 cabling.

The planes were made using standard off-the-shelf schedule 40 2.5" PVC piping, the same kind used in home plumbing. PVC pipe is inexpensive and easy to work

with because it conforms to existing fixtures and fittings. Standard *tee* fittings were used to attach PMT housings to the plane. The PVC was assembled by using standard PVC primer (acetone) and PVC cement. When done properly this made perfectly fused waterproof PVC joints. Each plane had an electronics container which was linked to the PMTs via Cat 5 cables. Cat 5 cables are standardised networking cables so they are cheap, easy to acquire and easy to work with. The combination of the PVC piping and the Cat 5 cables made the planes very simple and cost-effective.

During assembly, before any PMTs or electronics were installed, the planes were pressure-tested by sealing all openings in the plane and filling the plane with air at a positive pressure. The planes were then submerged in water and inspected for leaks at the PVC joints by looking for bubbles. Any leaking joints were repaired with PVC primer and cement.

The Madison planes had the Cat 5 cables pre-threaded before any joints were glued making cabling fast and easy but the gluing more difficult. The Nikhef planes were glued and then cabled which made the gluing fast but the cabling much more challenging. All the plane assembly, cabling and pressure testing was done by members of a non-specialist workforce with no technical training.

In total 62 planes were placed into the detector, 56 Nikhef planes spread out between the top and bottom caps and 6 Madison planes in just the bottom cap. Approximately 60 more Nikhef planes and 24 more Madison planes were prepared for a detector upgrade which would have increased the PMT coverage in the detector.

4.9.1 Nikhef Planes

The Nikhef planes consisted of 30 XP82B20FNB PMTs mounted inside plastic housings. The PMTs were angled to point into the neutrino beam. (The angle increases the apparent surface area of the light collectors in the direction that the Cherenkov cones form). The PMT model was selected due to its optimal price per unit area of photocathode.

Each PMT had an electronics base attached. This provided the high voltage power source and amplified the small PMT signal. The high voltage was generated using a Cockcroft-Walton voltage multiplier [71] which is driven by the CoCo custom application-specific integrated circuit ASIC [72]. A second custom ASIC, PROMis, amplified the readout signal, compared the signal with a threshold voltage and converted it into a digital signal whose amplitude was proportional to the time over threshold. Finally, PROMis converted the digital signal into a low-voltage differential signal LVDS signal for long-range transmission and sent it to the electronics container for data acquisition and storage. The high voltage value and threshold voltage could be changed by the upstream controller.

The Cockcroft-Walton generated a high voltage that was negative with respect to the Earth. This meant the PMT had an anode at earth and the photocathode was at



Figure 4.21: A fully disassembled Nikhef PMT showing from left to right: the PMT insert, the PMT itself with its high-voltage base attached, the waterproof cover complete with its O-ring and then the plastic reflector for light.



Figure 4.22: The assembled PMT with its plastic reflector attached ready to be glued onto a plane.

a very large potential difference with respect to ground and potentially to the water around the PMT face. This would have introduced noise into the PMT measurements. Therefore, a cover was needed over the face of the PMT to insulate it from the water. Each PMT was covered with a waterproof cover made from transparent acrylic to protect it from the water. The acrylic was bonded to the PMT's glass with optical gel to maximise the light transmission between the different mediums: the glue matched the refractive indices of glass and acrylic and removed the layer of air in between. A disassembled and fully assembled PMT can be seen in Figures 4.21 and 4.22 respectively.

At the centre of the plane, the PVC was attached to an aluminium electronics container which housed the Km3NeT electronics. The electronics consisted of two Calamari boards, descendants of the Km3NeT Octopus board [73], which aggregated

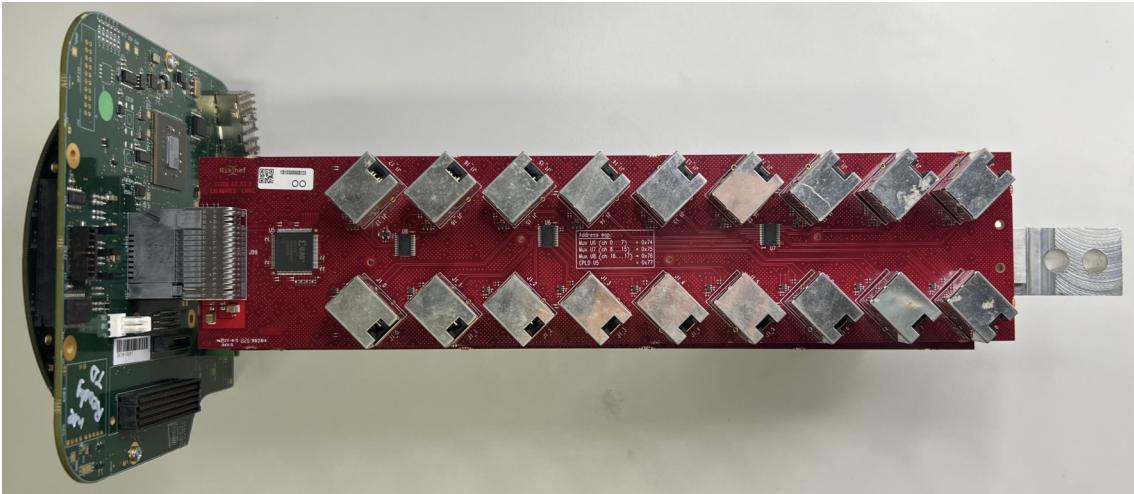


Figure 4.23: The Nikhef electronics which sat inside the electronics container: the green PCB is the CLB and the protruding red board is a Calamari board. An SFP port connects the CLB to the White Rabbit network. The 8P8C connectors are for the PMTs. There is a total of 30 connectors on the front and rear Calamari boards.

the PMT output signals. The Calamari boards were attached to a Central Logic Board (CLB) [73] which measured the time-over-threshold (ToT) of the PMT signals from the PROMis ASIC as well as providing a control signal to the Cockcroft-Walton driver. The CLB also handled communication with the upstream (DAQ) computers over a White Rabbit network [74] [75] [76]. White Rabbit Networks synchronise clocks precisely using an Ethernet-based connection for data transfer. The White Rabbit network allowed the CLB to receive very precise clock and time signals which were used to timestamp the ToT hits accurately (under 2 ns accuracy corresponding to approximately 0.6 m of light travel distance [72]). Finally, an LED flasher was attached to the plane which gave both a variable brightness and a variable frequency light pulse, synchronised with the precise White Rabbit clock to calibrate the detector. The Nikhef plane electronics can be seen in Figure 4.23.

The CLB was powered by a field-programmable gate array (FPGA) which ran the White Rabbit Precision Time Protocol Core [77] and two LatticeMico32 (LM32) soft cores [78]. One LM32 core controlled the White Rabbit and the other controlled the functionality of the plane. The CLB was connected to a 12 V power supply and to an upstream network using a small form-factor pluggable transceiver (SFP) with a fibre optic cable. A small 240 V to 12 V power supply was glued to the CLB and power was provided over 240 V. These power and data connectors left the electronics container through a waterproof bulkhead known as a *water block*.

To waterproof the planes, every joint of the PVC was sealed with PVC cement and every PMT cover had two waterproof gaskets secured with bolts. The power and data links leaving the electronics container went through a water block. The water block consisted of a piece of PVC glued to the output of the electronics container for power and data passthrough; a small fibre and brass rod passthrough was installed



Figure 4.24: A Nikhef plane installed inside the detector. The PMTs were angled forwards into the beam and the aluminium electronics container can be seen in the middle. The PVC hose exiting the electronics container was for connecting the plane to cable manifolds which lead to fanout containers.

into the PVC; and then potting compound was poured in to make a complete seal. The plane was then complete and ready to be installed into the detector.

A completed Nikhef plane in the detector can be seen in Figure 4.24.

4.9.2 PMT Potting

The Nikhef PMTs required a cover between the photocathode and the water surrounding the tube. Ideally, as much light as possible would be transmitted through the cover and onto the face of the PMT for maximum collection efficiency. If there was a gap between the cover and the photocathode, light could be reflected or refracted and lost, thus reducing the efficiency. To mitigate this, a potting compound was placed between the face of the PMT and the acrylic cover that bridged the gap of the refractive indices to improve light transmission. A potted Nikhef PMT can be seen in Figure 4.25

A special procedure was required to pot the PMTs in their covers with the compound. The potting compound came in two parts: a resin and a hardener. The two compounds were mixed together to make a gel and the ratio of the two parts determined the hardness of the gel. Great care had to be taken to completely mix the two halves. Imperfections greatly modified the final properties of the material, such as it never hardening or becoming runny at cold temperatures. The potting compound had to be degassed so bubbles did not form in front of the face of the PMT. To do this a vacuum was drawn to remove any air before potting. A special

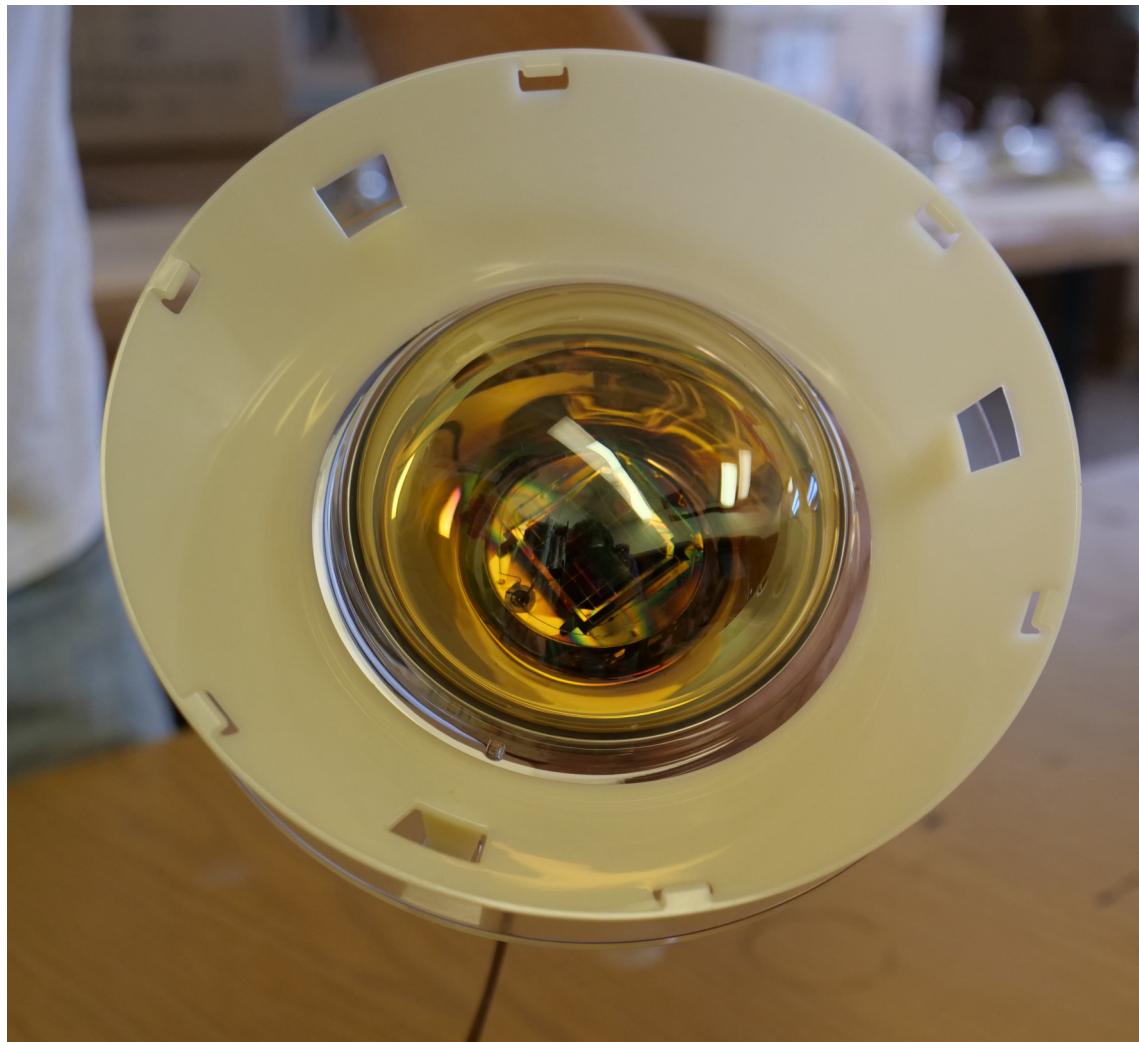


Figure 4.25: A potted Nikhef PMT. The presence of the optical gel between the plastic and the glass reduced the reflectivity between the two layers.



Figure 4.26: The jig for PMT potting. Alternating layers of PMTs and covers were stacked and the tubes were held straight during the curing process.

jig, shown in Figure 4.26, was used to pot the PMTs in a repeatable fashion to the desired standard using 15 mL of potting fluid per PMT.

The primary purpose of the potting jig was to hold the PMTs at a minimum separation from the cover without actually touching the cover. This was to minimise the amount of potting compound required while PMTs were being glued. The jig's secondary purpose was to hold the PMTs straight while the potting compound dried.

The potting jig held a rack of the waterproof covers pointing downwards; 30ml of potting compound was put inside. A second rack was then filled with the PMTs and placed on top, also facing downwards. The PMTs were pushed all the way to the bottom of the covers, so the glass face was touching the plastic; the natural buoyancy of the PMTs raised the tubes by a small amount (~ 0.5 mm) before the potting compound hardened. The role of the jig was then to hold the PMTs both gently enough to allow this process to happen and also strongly enough so that they were straight and not disturbed by activity in the surroundings. Completed PMTs can be seen in Figure 4.27.

In total, approximately 3500 PMTs were potted in batches of 155. Each batch took approximately 2.5 hours to complete and then was left overnight to cure.



Figure 4.27: A row of PMTs after the potting process showing the consistency of the transparent finish that can be achieved by potting the tubes.

4.9.3 Madison Planes

Each Madison plane consisted of 16 PMTs also mounted in plastic housings. Custom electronics were designed in conjunction with the Wisconsin IceCube Particle Astrophysics Center (WIPAC) [70]. Hamamatsu R6091 PMTs with a custom Cockcroft-Walton high voltage generating base were used. The R6091 PMTs were surplus from the NEMO-3 experiment decommissioning. On top of the high voltage base there was another custom printed circuit board (PCB) containing a microcontroller which controlled the PMT. This was known as the *Microdaq*. This setup is shown in Figure 4.28. The purpose of the Madison planes was to lower the cost and complexity of the planes compared with the Nikhef design. The Madison planes were simpler and cheaper both in terms of the techniques used for their construction and also in terms of the electronics used in them.

The Cockcroft-Walton used a pulse-width-modulated (PWM [79]) signal to modulate a 3.3 V supply to generate a low voltage alternating current. The alternating current was used in a resonance circuit to step up the voltage before going into the Cockcroft-Walton to generate the full high voltage.

Unlike the Nikhef planes, the Madison planes used a positive high voltage with respect to the Earth. This meant that the anode was at a large positive voltage and the photocathode was at earth and therefore there was no potential difference between the face of the tube and the water in the detector. As long as there was no water leaking into the plane, the face of the PMT could be directly exposed

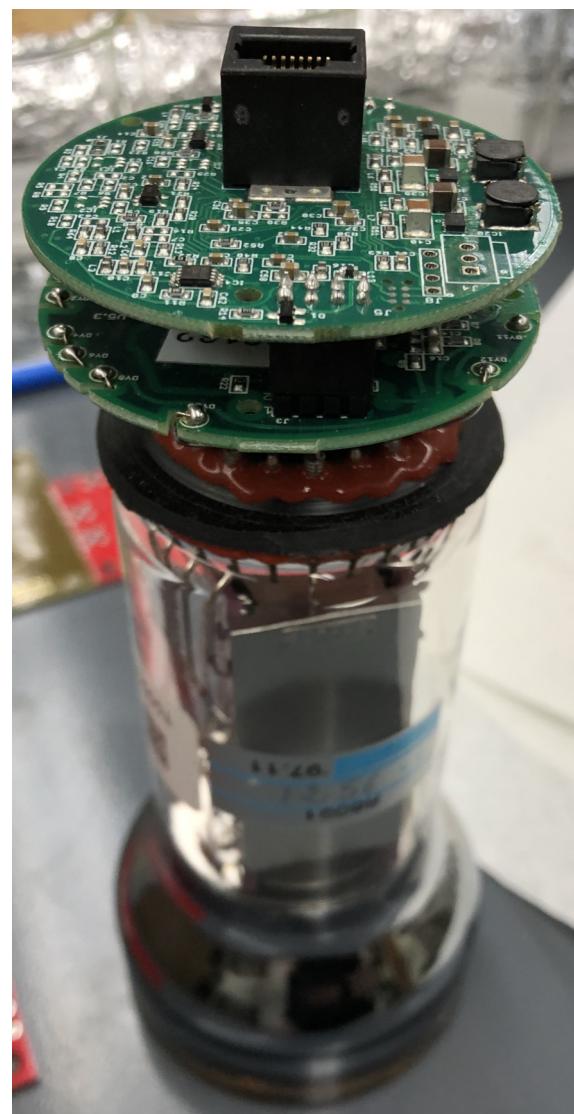


Figure 4.28: A Madison PMT showing the Cockcroft-Walton base attached to the end of the tube and the Microdaq on top of that.



Figure 4.29: The photocathode of a Madison PMT exposed to the environment. The tube was glued into the white insert with black resin that created a waterproof barrier.

to the water and there was no need for an insulating cover over the face of the tube. This represented a major design simplification as fewer custom-made parts were required to construct a Madison plane. The tubes could simply be glued into their inserts with the photocathode pointing outwards rather than a custom fitted transparent acrylic cover with gaskets. This can be seen in Figure 4.29. Fewer parts and assembly steps made the final preparation and assembly faster.

The Cockcroft-Walton base also had a preamplifier on board. The base was soldered directly to the pins of the photomultiplier tube so the readout signal passed to the Cockcroft-Walton base first. The preamplifier amplified and shaped the signal and passed it onwards to the Microdaq.

The Microdaq was a small data logger and controller based around the STM32 microcontroller platform [80]. The Microdaq handled the powering, triggering and data acquisition of the PMT. It acted as a semi-autonomous smart device: an upstream DAQ computer issued commands to the Microdaq describing the PMT itself (the gain and voltage levels), triggering and high voltage levels. Upon command, the Microdaq then gathered data and sent them back to the DAQ computer. A serial link over LVDS was used to communicate between the two.

The Microdaq had an analogue-to-digital-converter (ADC) which integrated the charge of the PMT hit. The Microdaq also had a discriminator circuit which measured the time of the hit over a threshold voltage (the TOT). The Microdaq had a 180 MHz core speed, which gave a clock cycle time period of 5.5 ns. A time res-

olution of 5.5 ns would be extremely poor for an experiment that measured light propagation (corresponding to approximately 1.65 m of light travel distance) using expensive high-speed components or FPGAs and would not have been in keeping with the CHIPS ethos. One of the major improvements the Microdaq provided was its low-cost method of improving the time resolution. Each clock cycle, the Microdaq produced a pulse which was fed through 8 time capturing registers that had a propagation delay of approximately 0.5 ns. When a hit was recorded, the registers were sampled and the number of registers that the leading edge had passed through was counted. This gave nanosecond timing precision using a standard low-speed microcontroller and cheap shift register and no high speed electronics were required. This system was used to timestamp ADC hits and measure the time over threshold. The discriminator threshold voltage could be adjusted by the Microdaq to include or disallow PMT hits.

Additionally, as well as recording data, the Microdaq had a PWM output which was used to drive an H-Bridge [79] to switch an on-board power supply to generate a 3.3 VAC current. This current was then used to drive the resonator circuit on the Cockcroft-Walton base to generate the high voltage for a PMT. The Microdaq then had an auxiliary ADC which measured an attenuated form of the high voltage so the tube voltage could be measured. Changing the frequency of the PWM signal changed the voltage on the tube.

The Microdaqs required a reference clock so that timestamps were comparable to each other and reflected the true time. A 10 MHz reference clock was provided using LVDS and the Microdaq used this to synchronise its core clock speed. An inter-range instrumentation group code B (IRIG-B) signal was also provided over LVDS. The IRIG-B signal marked the moment a second ticked over and encoded the data and time of the current second. These two signals were used to synchronise the clocks on the Microdaqs.

The Microdaq used a standard 8P8C connector (commonly referred to as an RJ45) with 8 conductors. Three differential pairs were used for the 10 MHz IRIG-B and serial data links and the final conductors were used for power and ground. The use of a standard connector reduced the cost of manufacturing the Microdaq.

The Microdaq firmware was written in the C programming language, and used the Hardware Abstraction Layer (HAL) standard library provided for the STM32 to access the features of the microcontroller.

The STM32 microcontroller is a low speed, low power, high reliability and high value-for-money device. The supporting components on the Microdaq were all standard off-the-shelf low-end electronics. No dense ball-grid-array packages were used so the printed circuit board manufacturing process did not need to be optimised for fine-pitched components and the number of layers could be kept to four layers. Overall this meant the cost of components and manufacturing was low due to the

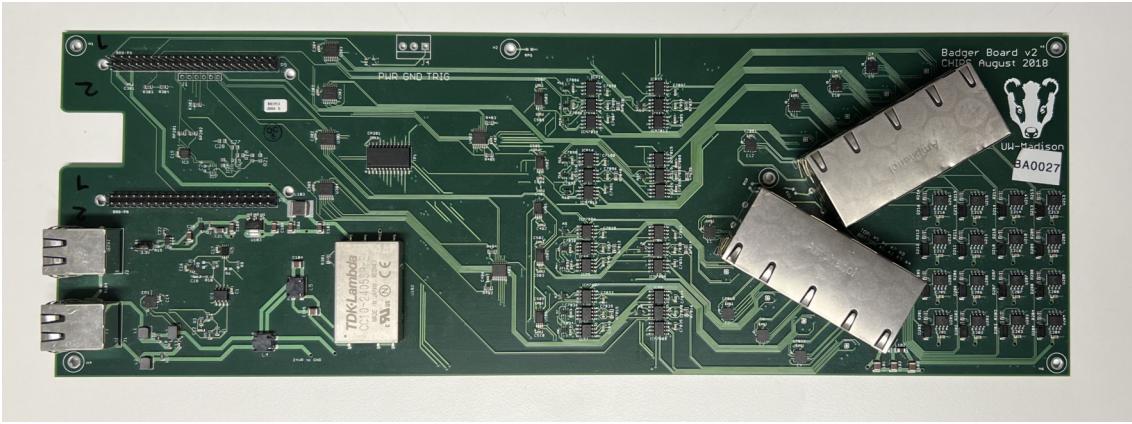


Figure 4.30: The Badgerboard. The Cat 5 input cable enters at the bottom left and the BeagleBone Green sits at the top left. The Cat 5 cables to the Microdaq leave on the right hand side of the board.

simplicity and standardised requirements.

A fanout board, known as the *Badgerboard*, (shown in Figure 4.30) sat in the centre of the plane to manage the 16 PMTs in the place of a CLB. The Badgerboard received power (24 V) and ground, a 10 MHz clock signal, an IRIG-B time-of-day signal and an Ethernet connection. All of these arrived in a single cat 5 cable through an 8P8C connector from the upstream system. The 10 MHz and IRIG-B signals were multiplexed 16 times and forwarded to the Microdaqs. The Ethernet was connected to a BeagleBone Green [81] single-board-computer. The BeagleBone had four serial universal asynchronous receiver-transmitter (UART) ports which were multiplexed four times each. The 16 serial links were distributed to the 16 Microdaqs. The BeagleBone used the Debian operating system so could run standard software applications rather than a hardwired FPGA data acquisition system. The BeagleBone hosted two sets of DAQsoftware: a set of Python scripts that were for hardware tests and simple single-tube control, and then a software suite called *Fieldhub* which was intended to efficiently control an entire plane and was written in C.

The Badgerboard had sensors for temperature, humidity, orientation and acceleration which were queried by the BeagleBone and saved for logging. It also had a port for plugging in a nanobeacon (an LED flasher for calibration). The nanobeacon had variable brightness by changing the voltage and a variable flash frequency. However, the flashing frequency was derived from a PWM signal generated by BeagleBone which had a floating clock rather than a clock synchronised to the CHIPS-wide White Rabbit clock. The result of this was that the times at which light pulses occurred could not be considered a true time reference.

Like the Nikhef planes, the Madison planes were made of rungs of PVC with PVC PMT inserts glued directly in. An electronics container was attached to the centre of the plane containing the Badgerboard board. Unlike in the Nikhef planes, the container was made from a PVC cylinder. This meant that the electronics container could be glued on quickly and reliably rather than being fixed by a system of gaskets

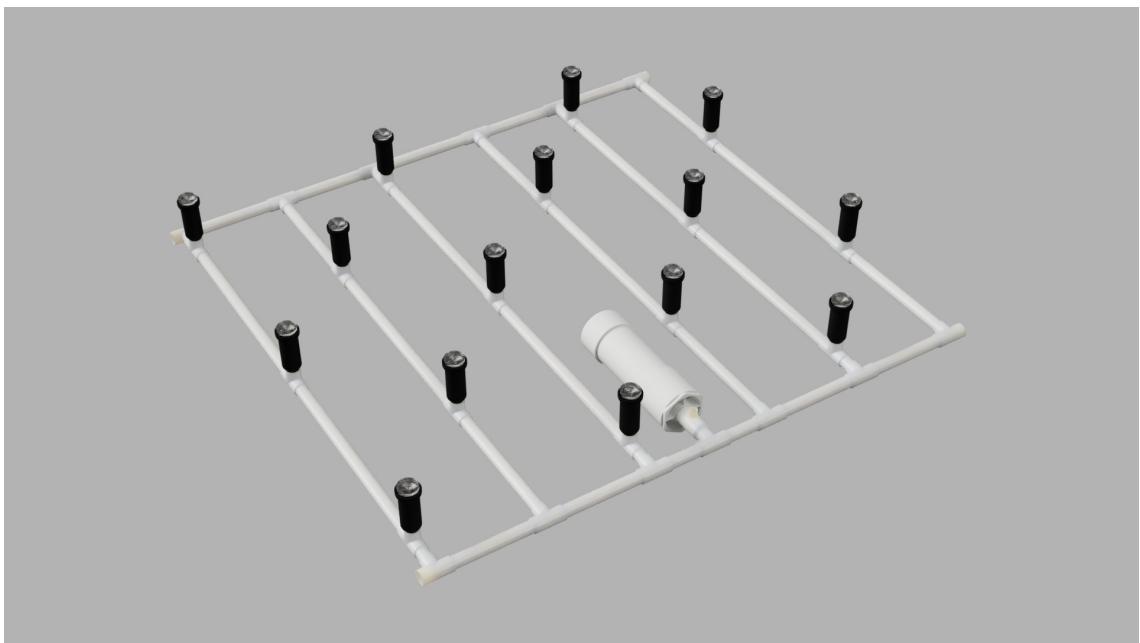


Figure 4.31: A rendering of a Madison plane. The PMTs are mounted onto rungs of PVC pipe. The large cylinder in the centre of the plane is the electronics container. 15 PMTs were mounted despite the electronics supporting 16. The 16th channel was used as a spare if problems arose during the assembly of the plane.

and bolts. This method also reduced the total number of parts. A completed Madison plane can be seen in Figure 4.31 and the electronics container in Figure 4.32.

4.10 Electronics Fanout

The detector planes required a set of supporting electronics: the Madison and Nikhef planes were connected to fanouts which linked many planes to a single device. The fanouts were then connected to a central control box that was in turn connected to a hut on the shore of the mine pit. The shore hut housed the central DAQ computers, the Global Positioning System (GPS), antennas and the main internet connection.

4.10.1 Nikhef Fanout

The Nikhef planes required two pieces of infrastructure to function: a power supply and a White Rabbit network to provide the data connection and timing information. To manage the 64 Nikhef planes in the detector, five fanout devices were used. These were waterproof canisters (shown in Figure 4.33) which contained a White Rabbit network switch and a set of power relays. A single power and network connection went into the canister and up to 18 power and network connections left via a manifold. The manifolds were made from PVC hose and used standard plumbing fixtures. The power cables and fibre optics were threaded through the hose. The Relays were networked so they could be controlled by the up-stream DAQ

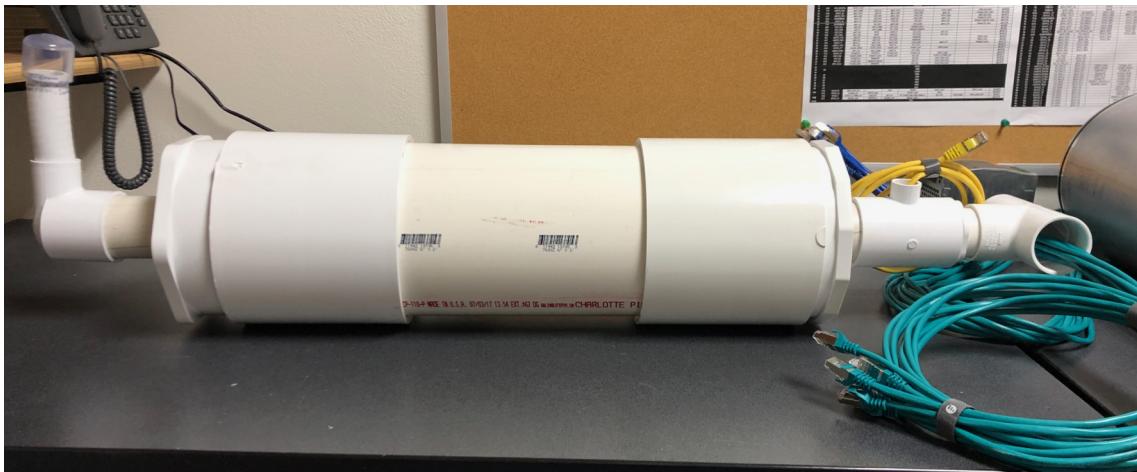


Figure 4.32: The electronics container for a Madison plane. The entire container was made from glued PVC which could then be glued to the rest of the plane. The Cat 5 cables between the electronics and the Microdaqs can be seen leaving the container.



Figure 4.33: A fanout container, the same kind of aluminium pressure cylinder used for the electronics containers on the Nikhef planes.

computer via the main detector control box: the reason for this was to be able to turn planes on or off in the event of a hardware problem or leak. The interior of the fanout devices can be seen in shown in Figure 4.34.

4.10.2 Madison Fanout

The Madison planes also required a fanout device inside the detector which needed extra functionality compared with the Nikhef fanout system. The Madison planes did not have the White Rabbit network built directly into them; they used standard Ethernet, but they did require a 10 MHz and IRIG-B signal which could be generated by a White Rabbit enabled device. The fanout needed to provide the 10 MHz and IRIG-B timing signals (an extra feature not required by the Nikhef system). The fanout also needed to provide power, power switching as well as Ethernet to the

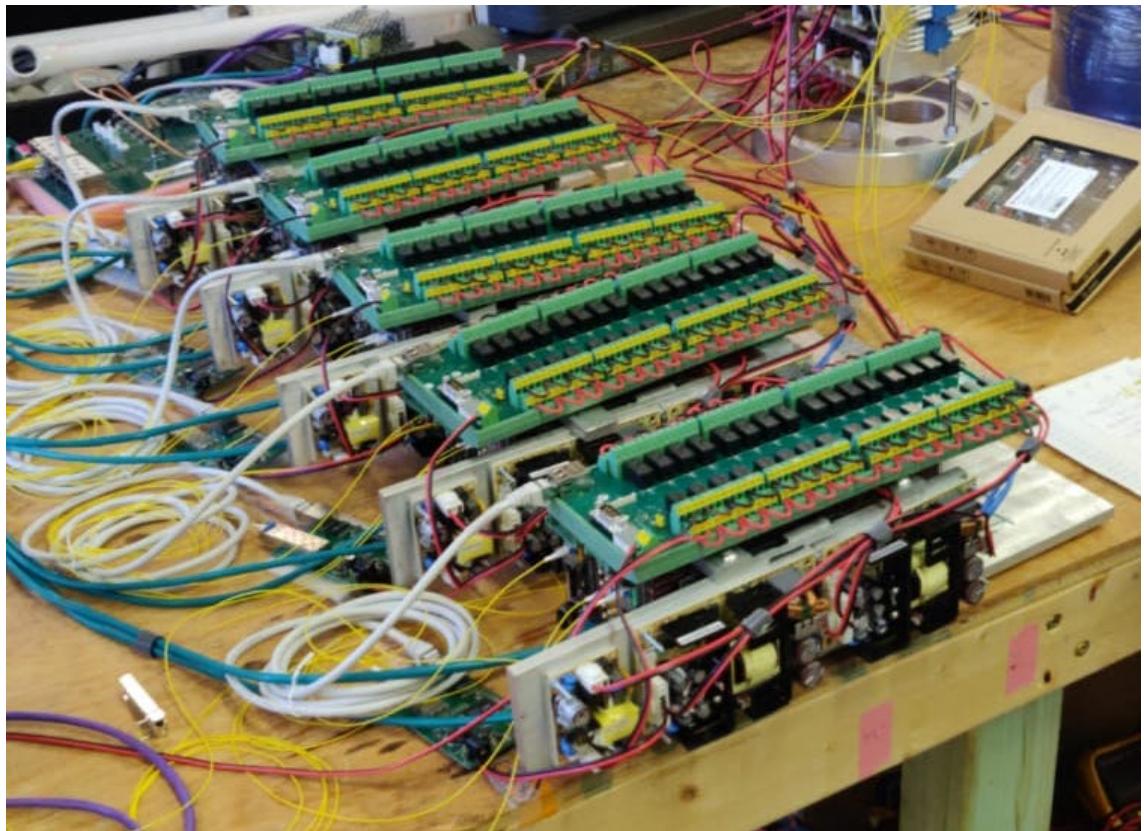


Figure 4.34: The inside of five Nikhef fanout containers. The relays and power converters are visible. The White Rabbit switches sit under the relays attached to the aluminium block which acts as a heatsink.

planes.

Power was supplied to the fanout container at 240 V as it had to travel a long distance. Power converters converted 240 V to 24 V for powering the devices inside the fanout container as well as sending it to the planes (which were a short distance away).

The network topology of the Madison fanout was as follows:

- A White Rabbit network (from the main CHIPS network) arrived into the fanout container.
- A fibre connected the network to a White Rabbit Light End Node (LEN) [82].
- The White Rabbit LEN generated 10 MHz and IRIG-B signals.
- The White Rabbit LEN spawned an Ethernet connection which was fed into a small router and network switch which gave each plane an Ethernet connection.
- The 10 MHz, IRIG-B and Ethernet connections were fed into a custom fanout board.

A custom fanout board, known as the *Danout* board (shown in Figure 4.35), was developed to accommodate the needs of the Madison planes. The Danout hosted a



Figure 4.35: A Danout board complete with its BeagleBone Green.

BeagleBone Green which operated a set of Inter-Integrated Circuit (I^2C) controlled power switches; these enabled power to selectively flow to the planes. The Danout board also had inputs for the 16 Ethernet connections and the 10 MHz and IRIG-B signals. The Danout duplicated the 10 MHz and IRIG-B signals 16 times, converted them to LVDS signals for long-range transmission and then superimposed them onto the 24 V and ground rails which were used for sending power to the planes. This provided the timing signals to 16 planes while only using one White Rabbit enabled device. The 16 Ethernet connections (for the planes) were fed into the Danout board; there, the power and timing signals were arranged so they could be fed back out into 8P8C connectors with the Ethernet to go to the detector planes over a single Cat 5 cable.

The White Rabbit LEN was a commercially available device and cheaper than an 18-port White Rabbit switch. A networking router and switch were used in the fanout container and high value-for-money standard cables and connectors were used between the planes and fanout rather than separate fibres and power cables. The result was that the design was simple and the components used were inexpensive.

Like the Nikhef fanout, the Madison fanout was housed in an aluminium waterproof canister inside the detector and was connected to the main detector control box.

4.10.3 Main Control Box

Inside the detector, there was a central control box which managed the main network connection into the detector and the power distribution.

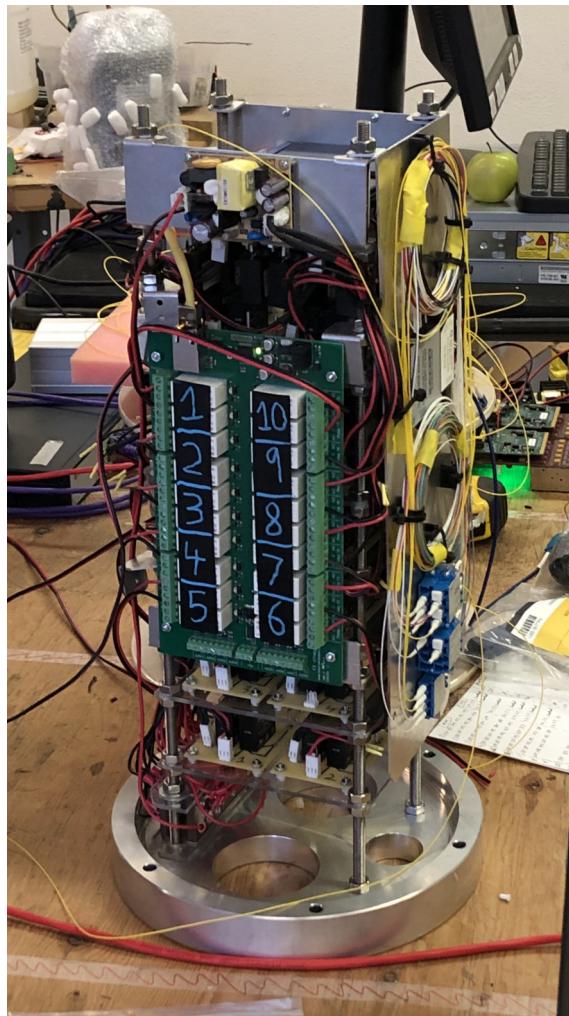


Figure 4.36: The interior of the main control box.

A large armoured multi-mode fibre was fed from the shore, through the water, into the detector and into the control box. Inside the control box a demultiplexer split the different wavelengths into different fibres and passed them out of the control box to the fanout boxes. An additional component channel of the main fibre was fed into a media converter inside the control box; this gave a copper Ethernet connection which was used to operate a set of relays. The relays were controlled over the network and switched the power to the fanout boxes. The main fibre and control box could have been the largest single point of failure in the CHIPS detector so great care was taken to make it as water resistant as possible. Every input and output was filled with a potting compound to prevent water ingress and the container itself was pressure-tested after it was sealed. The contents of the control box can be seen in Figure 4.36 and the control box installed inside the detector can be seen in Figure 4.37.



(a) A rendering of the main control box. The silver cylinder is the waterproof pressure vessel and the white ring of PVC tubing contains the waterproof passthroughs for cables and fibres. The fixtures at the top attach the control box to the detector.



(b) The exterior of the main control box installed inside the detector. The PVC fitting attached to the aluminium container contains resin for preventing any leaks in the manifolds entering the container.

Figure 4.37: The main control box.

4.10.4 Shore Electronics Hut

The final piece of physical infrastructure needed to run the detector was the electronics hut on the shore of the lake. The hut contained the DAQ computers which controlled the detector and received the data to process and store. The hut also contained a GPS antenna which was plugged into a White Rabbit Grandmaster switch which acted as the definitive detector reference clock. The White Rabbit Grandmaster switch linked all the detector White Rabbit switches onto the same network as the main DAQ computer. This connected everything in the CHIPS system into one network with a single reference clock source.

The hut also housed the connection to the wider internet. This linked the detector to Fermilab for access to Fermilab computing and data storage as well as the beam spill server for triggering the detector.

Finally, the hut was connected directly to the power grid and had a 240 V power supply. This was passed through to the detector.

These data and power connections were inside an umbilical made from PVC hose which ran from the inside of the hut to the detector.

The complete topology of the power and data connections from the hut into the detector can be seen in Figure 4.38. Additionally, the data connection between the shore hut and Fermilab can also be seen.

4.11 Timing System

The CHIPS timing and trigger system is beyond the scope of this work. For further reading, see reference [83]. The description detailed below is a brief summary for the sake of a complete description of the CHIPS detector.

The NuMI beam releases neutrinos in periodic beam spills. If a detector is continuously taking data then it wastes computing time and storage space. Also, if data are only recorded in the brief time window when beam neutrinos are anticipated to be present, some background events are automatically rejected without the need to identify them in analysis. It is therefore advantageous to implement a trigger system to warn the detector when neutrinos will arrive. A trigger system would have to warn the detector of the arrival time of a beam spill *before* the spill arrives. By extension of this concept, the detector and accelerator clocks would need to be very precisely synchronised otherwise the neutrino arrival window would be missed.

The NuMI beam has a system of signals which are linked to different stages of the NuMI beam cycle [84]. The idea of the CHIPS trigger system is to use a precursor signal to the neutrinos being emitted and transmit its timestamp over the internet to the CHIPS detector. The time the signal can then be used to activate the detector moments before the neutrinos arrive. Note: using the internet means that the time between transmitting and receiving the signal is non-deterministic as the pathway

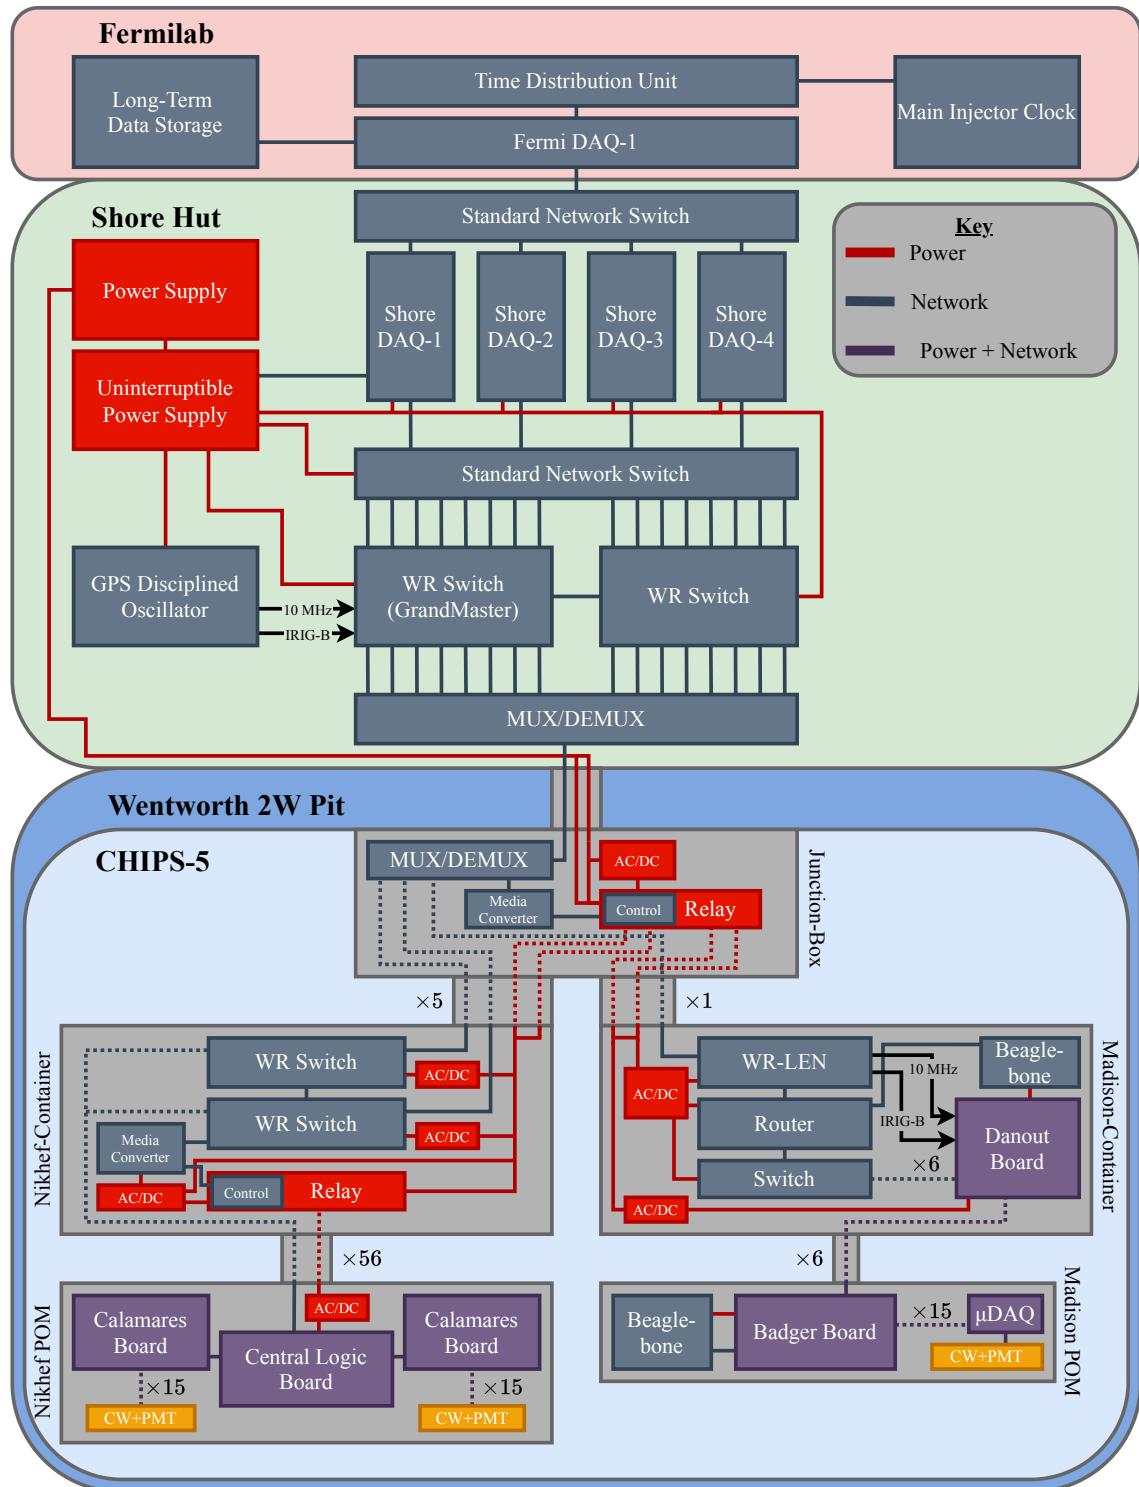


Figure 4.38: The network and power distribution systems of CHIPS. Here POM stands for *Planar Optical Modules* - the detector planes. Image from [27].

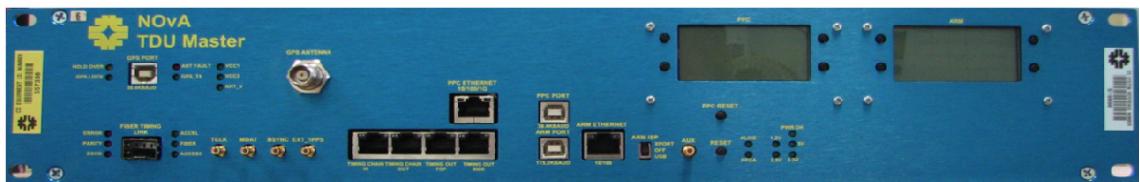


Figure 4.39: A Nova TDU [85].

through the internet is not fixed. This is why it is necessary to transmit a timestamp rather than a pulse.

Fermilab and Nova have a pre-existing time distribution system [85]. A number of accelerator signals are generated representing different states of the NuMI beam. The signals are decoded and timestamped for the Nova experiment by the master Time Distribution Unit (TDU) (shown in Figure 4.39) which is synchronised to Co-ordinated Universal Time by GPS. The TDU then distributes the coordinated Nova time and accelerator events to slave TDUS.

CHIPS has a prototype Nova TDU to act as a grandmaster clock at the accelerator. The accelerator signals are timestamped, decoded and sent through the internal Fermilab Ethernet network to CHIPS. The computer has software (the spill forwarder) which forwards the beam spills through the internet to a DAQ computer at the detector which has a spill receiver linked to the detector control software. The network topology is shown in Figure 4.40. The spill forwarding and receiving software is built using standardised networking components rather than proprietary software and the signals are then distributed over the internet rather than over purpose built infrastructure. The detector itself and detector-computer clocks were time synchronised using the White Rabbit network which was locked to a GPS disciplined oscillator.

The trigger signals were transmitted to the detector and decoded. The timing information was then used to schedule a trigger window for the detector. The detector lay dormant when no neutrinos were anticipated to arrive and was reactivated at the start the trigger window.

4.12 The Conclusion of the CHIPS Experiment

The detector was deployed in October 2019. The interior of the completed detector can be seen in Figures 4.42 - 4.45. During the deployment procedure, the liner sustained damage on sharp rocks which caused a hole to form and water to ingress. This meant that as the detector was being towed, the floating dock was taking the load of the detector whilst still in motion despite being designed for a central load. Consequently, the floating dock was unable to hold the load without deforming and was bent inwards towards the detector. The detector was still successfully deployed. The light leak caused was not repaired as winter was imminent and the intention

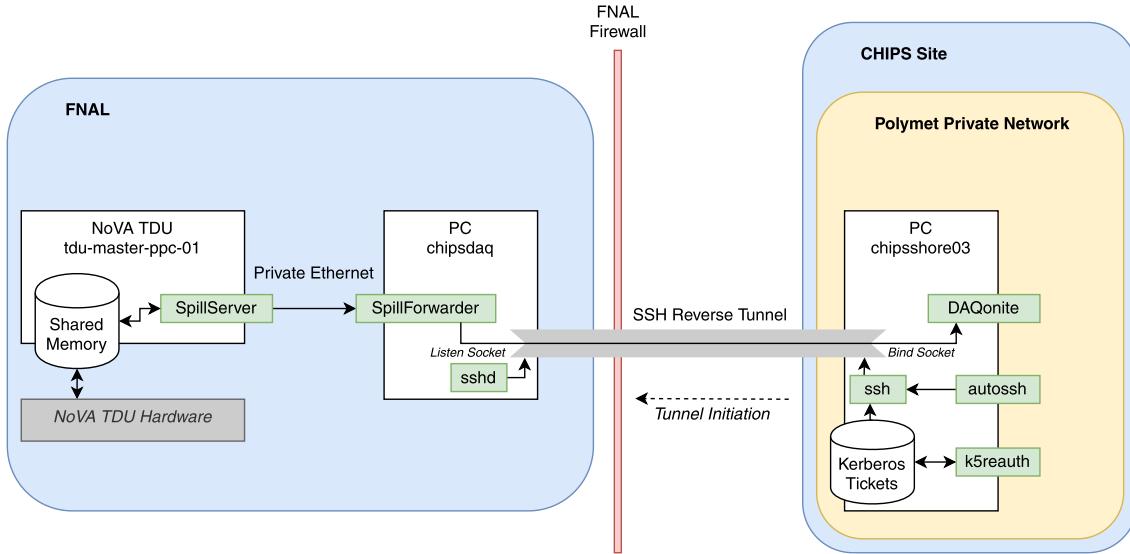


Figure 4.40: The network topology of the CHIPS time distribution system [83]. The NOVA TDU transmits timestamped signals to the CHIPS spill forwarder which then traverses the Fermilab Firewall and internet to the CHIPS site. The chipsshore03 computer decodes the timing signals.

was to return in spring the following year when the lake thawed.

The COVID-19 pandemic halted repair efforts as no personnel could access the site and the Numi beam shutdown was extended [86]. In September 2020, the detector was removed from the water as there was no timeline to resume activities and the pandemic had not resolved. An example of a hitmap for a cosmic ray event can be seen in Figure 4.41 which demonstrated that the detector was operational after deployment.

The detector was salvaged in such a way that the instrumentation could be reused for further projects and the mine site was returned to its original state.

The detector was constructed, built and deployed using commonly-available building materials, affordable technologies and with a workforce that, on the whole, was non-specialist. The total cost was under €4m and, despite not observing neutrinos, the detector demonstrated light could be observed in a large-scale water Cherenkov detector. Overall, despite not fulfilling its full original purpose of recording neutrinos from the Numi beam, the detector still demonstrated most of its original goals of designing, constructing and deploying a low cost neutrino detector. A detector was built and was operational: this alone was a major success of the CHIPS project.

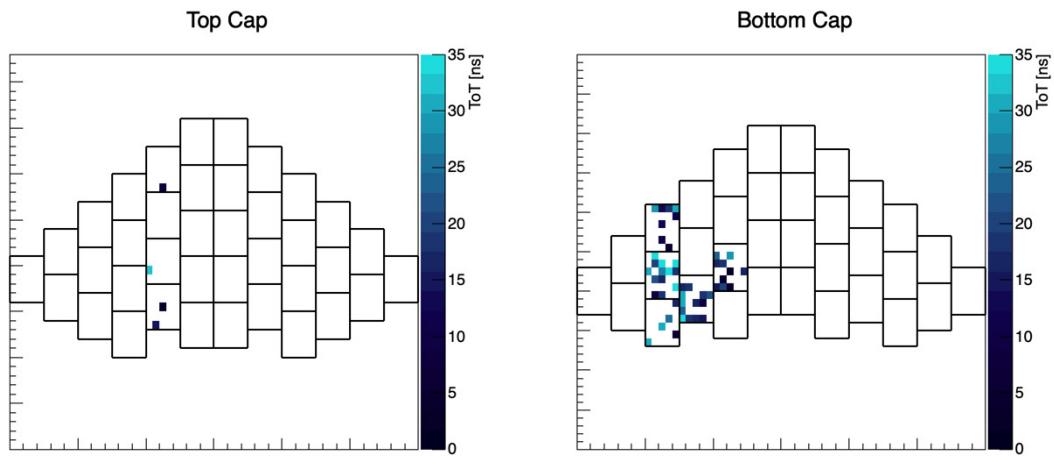


Figure 4.41: A CHIPS hitmap showing what is likely to be a cosmic ray muon event. A cosmic ray would be expected to form a ring, which can be seen in this pattern of pixels. Each pixel corresponds to a photomultiplier tube in the detector and the boxes show the position of each plane. The damage sustained to the detector during deployment gave little time to assess the efficiency of the detector. This event is assumed to be a cosmic ray as it is unlikely to be noise because each plane had been tested before installation to ensure low noise. At the time of this event, the NuMI beam was not running. Image courtesy of Stefano Germani.



Figure 4.42: The exterior of the detector after completion. The liner can be seen supporting the detector and the floating dock forms a perimeter around the detector. The two boats tow and steer the detector during deployment. The scale of the detector is visible.



Figure 4.43: A panorama inside the detector. The detector planes and the manifolds which connect the planes can be seen.

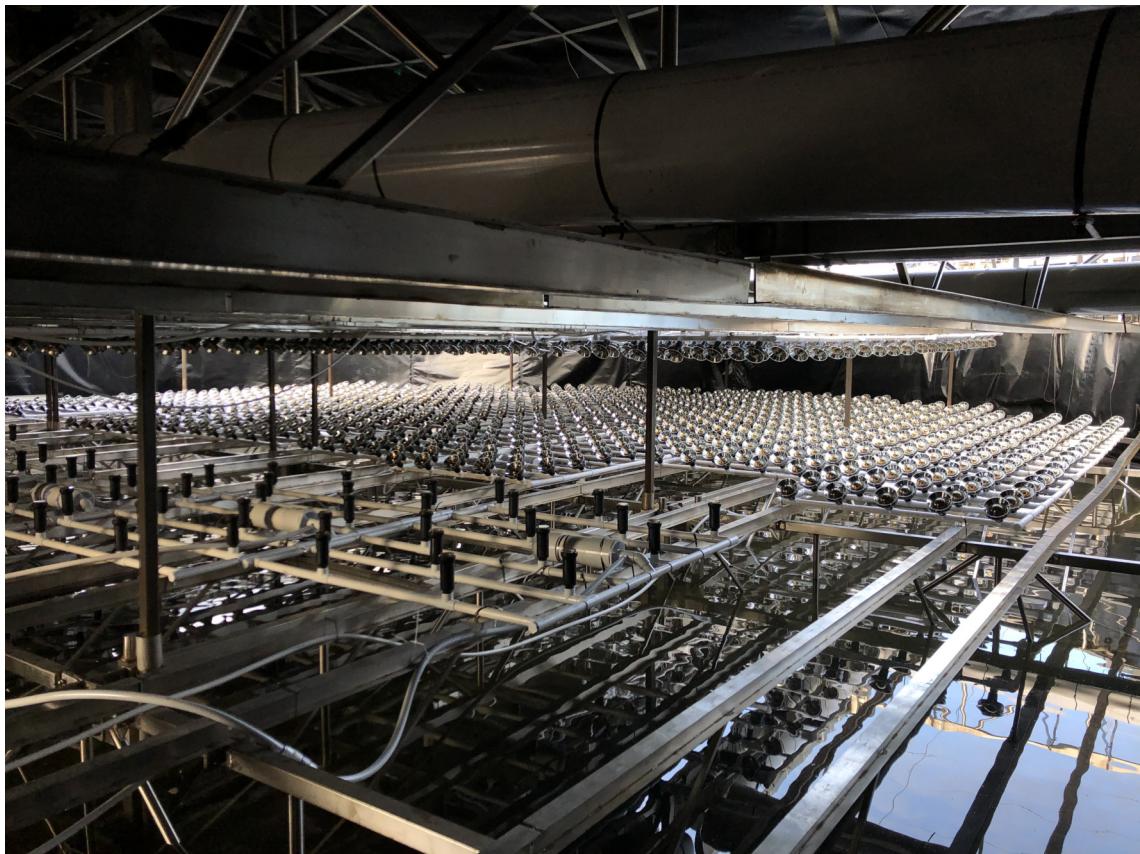


Figure 4.44: The interior of the detector: the rows of Nikhef PMTs angled to point into the neutrino beam can be seen at the back with the Madison planes in front. The top cap buoyancy can be seen above.



Figure 4.45: The detector from the front during the final liner sealing process so the detector was and waterproof.

Chapter 5

Future Detector Electronics' Development

As described in Section 4.9, there were two PMT readout systems used in CHIPS: the Nikhef and the Madison systems. Although both systems fulfilled their purpose, they both have limitations.

The Nikhef system utilised high speed, high precision electronics which meant it was expensive and harder to manufacture. For example, the CLB used a 12-layer PCB with tightly controlled impedance that had a high cost of manufacture and long lead times. The cost per board was over an order of magnitude more expensive than the competing Madison electronics. Furthermore, although the FPGA performed excellently, the firmware development was slow and required unique skills rather than common programming talent. The CLB and Calamari boards were however standalone, providing they had a White Rabbit network. In short, the Nikhef system was highly performant but also expensive and exotic.

The Madison system utilised low cost electronic components and standard PCB manufacturing technology making it cheaper to manufacture but suffered a different set of problems. The Madison system required a custom fanout board as well as an external White Rabbit enabled device to operate, so each Microdaq required an electronics container that had a custom board, a custom fanout board upstream, and an additional external White Rabbit enabled device. This meant that the system was cumbersome and difficult to set up and test. To set up a single Microdaq three sets of electronics were required. Once the fanout and White Rabbit Network were set up this scaled well but the initial setup cost was high and time-consuming.

The Madison electronics utilised on board BeagleBone computers which meant that each plane could act as a standalone computer if required which was convenient. However, the BeagleBones were slow, single cored, low on memory and had a difficult connection. The BeagleBone was also difficult to support with software as more modern systems are now more popular. The Nikhef planes had a soft microprocessor (LM32) implemented on-board but this just controlled the plane functionality

and did not act as a fully fledged data acquisition computer. Therefore an upstream computer was required. Ideally a future system would be able to operate as a standalone device with an on-board computer and be able to fall back to remote operation for more flexibility. Also, the computer would be a well supported platform to utilise existing codebases and designed for embedded use (as well as being sufficiently powerful for its intended use).

A new electronics system would have been useful for the CHIPS detector to balance development costs with performance and ease of use, rather than having either only the one or only the other. An upgrade to the Madison electronics was proposed to improve functionality and reduce the infrastructure requirements – bringing the functionality closer to the Nikhef electronics but with a simplified back end – but for a lower cost than the Nikhef electronics.

5.1 Detailed Description of the Badgerboard

All signals into the Badgerboard are in the form of LVDS for reliable long range communication. There are two differential pairs. Two pairs for Ethernet (10/100) and a pair each for IRIG-B and 10 MHz which completes the 8 conductors in a Cat 5 cable. However, two conductors are needed for the power supply: power and ground. Consequently, the IRIG-B is superimposed upon the power and ground signals so just two conductors carry the reference voltages and the reference time. The signals were split and filtered to make smooth 24 V and ground references and the differential IRIG-B signal was separated from the higher voltage with capacitors and then re-offset to correct LVDS levels before being fed into a receiver. The corrected IRIG-B LVDS signal and the 10 MHz LVDS signal were fed into an LVDS receiver to give the final form of the timing signals.

The Microdaqs used a 24 V power supply as the high voltage is suitable for providing power over long distances. The BeagleBone required a 5 V power supply and the electronic components on the Badgerboard used 3.3 V. The power supply into Badgerboard was at 24 V so that it could be split two ways: once to remain at 24 V for the Microdaqs and once to a 5 V switch mode power supply to generate the 5 V supply for the BeagleBone, which required the most current in the entire system (up to 1 A). The 5 V supply was then tapped and fed into a linear low-dropout regulator to generate a smooth 3.3 V to power the components on the Badgerboard itself. However, many of the logical devices on the Badgerboard were powered by the external 3.3 V supply and had a virtually instantaneous boot-up time and could pull lines in the Badgerboard `high` or `low`. The BeagleBone had a number of pins connected to the Badgerboard for usage after startup. However, the BeagleBone has a ‘slow’ startup time (on the order of a few hundred milliseconds) and if pins on the BeagleBone were not what the bootloader expected then the boot process

could be interrupted and the BeagleBone would freeze. Therefore, a delay circuit was added to the 3.3 V power supply so that it would only be enabled after a delay long enough for the BeagleBone to finish its initial boot process.

To power each Microdaq, a power switch was used to toggle the 24 V supply to each Microdaq on and off. Each switch used a signal line `high` or `low` to determine if the power was on or off respectively. The signal was provided by a I²C general-purpose input/output GPIO extender which was a logical device controlled by the BeagleBone's I²C. The I²C GPIO extender allowed 16 outputs to be toggled on or off (using just two pins on the BeagleBone for I²C rather than 16 pins, one for each channel). Once the power was enabled on a channel, it was considered 'active' and the signal used to toggle the switch was used as an indicator that the channel had been activated. The signal was known as the `power-enable` signal.

The IRIG-B and 10 MHz signals were fanned out 16 times each and fed into LVDS transmitters. The transmitters were enabled by the `power-enable` signal if the Microdaq had power. If a Microdaq was active then the LVDS transmitter was also activated so the IRIG-B and 10 MHz were transmitted to the Microdaq over the cat 5 cable.

The Badgerboard received Ethernet through the 8P8C connector. The BeagleBone needed to receive the Ethernet signal that was already within the Badgerboard. However, the BeagleBone could only receive Ethernet over another 8P8C connector. To solve this, once the Ethernet had been split from the timing and power signals, it was immediately fed into an 8P8C connector as a standard Ethernet output and then a small cable in the shape of the letter U (U-turn) was used to immediately plug into the BeagleBone (shown in Figure 5.2). Although this worked, it was seen as an area for future improvement.

The BeagleBone used on the Badgerboard is shown in Figure 5.1. It had 96 general purpose input output (GPIO) pins (female), some of which had special functions on top of the GPIO functionality including:

- I²C — used to control power switches and environment sensors.
- PWM — used to control the nanobeacon.
- UART — used to communicate with Microdaqs.

The Badgerboard had 96 pins (male) mounted to it so the BeagleBone could be directly attached (also shown in Figure 5.2). The BeagleBone was powered over the pins and the only cable required was the U-turn Ethernet cable. The BeagleBone required a micro-SD card for storage and a copy of Debian Linux (7.1 Wheezy) was stored on it. Requiring an SD card was considered a limitation as SD cards are not designed for continuous access and in this use case acted as a single point of failure. Henceforth, all functionality of the BeagleBone came through these pin headers.

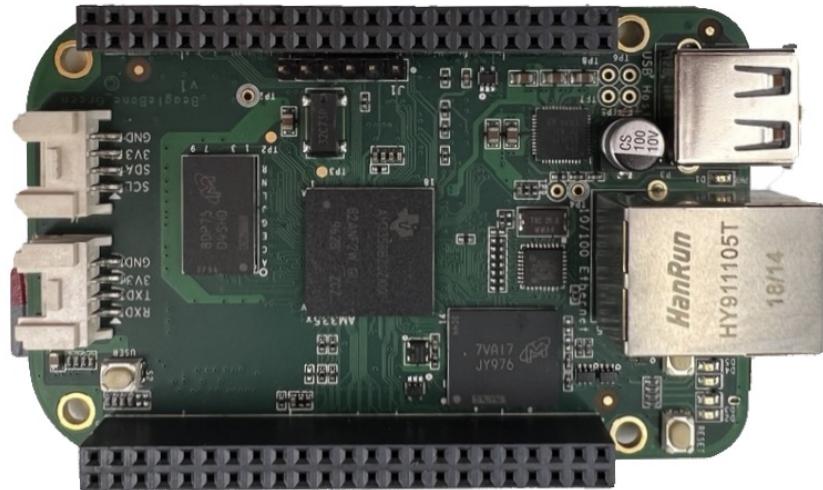


Figure 5.1: Top view of a BeagleBone Green showing the 96 I/O pins.



Figure 5.2: A BeagleBone Green mounted onto a Badgerboard using the two 48-pin connectors. The U-turn cable is shown in the top left. Here, a longer cable was used and coiled up.

The Microdaqs used half-duplex communication: the Microdaq and BeagleBone took turns transmitting and receiving over the same differential pair rather than doing both simultaneously. This was done over UART. The BeagleBone only had four UARTs and usually UART communication is in full-duplex mode so two problems had to be solved: communicating with 16 Microdaqs over four communication channels and doing it in a way that half-duplex communication could be coordinated. The four UARTs from the BeagleBone were shared between four Microdaqs each. To support each UART bank, five signals are used:

1. A data transmit (**tx**) line.
2. A data receive (**rx**) line for the UART communication.
3. A **transmit-enable** line which determined the mode for half-duplex communication.
4. Two selection signals which encoded a two bit number (0,1,2,3) to determine which one of the four Microdaqs in a UART bank was being used.

The **tx** and **rx** lines went into a four-way switch where the two selection lines toggled which of the four outputs were used. Each output then went to an LVDS

transceiver. These could not transmit and receive at the same time. The default mode was to receive data, but toggling the `transmit-enable` to the activated position changed the transceiver to the transmit mode. The Microdaq did not proactively communicate with the BeagleBone, but only replied to messages when probed. Therefore, this method of half-duplex communication worked as the Badgerboard only needed to listen to one specific channel immediately after transmitting a message to that same channel. The receive and transmit modes could be rapidly toggled as could the channel in question. This setup was replicated four times on the Badgerboard to give the 16 channels. To control the `transmit-enable` pins while simultaneously using the UART, custom Linux Kernel drivers (provided by the WIPAC) for RS485 had to be installed which were cumbersome to support and complicated the software development for the BeagleBone. While effective, this was also seen as an opportunity for future improvement

The STM32 microcontroller on the Microdaq used firmware that was flashed onto the memory remotely. Once the firmware was stored on the Microdaq and the Microdaq was powered, a special command to begin firmware execution was transmitted to the Microdaq which indicated to it that it should load the firmware and start running. The process to flash the memory and start the application followed a unique process required by the manufacturer which required full duplex communication. Full duplex UART communication over LVDS would require a minimum of four conductors and only two were allocated for communication; so temporarily, during Microdaq initialisation, two more were used. The Microdaq used the 10 MHz signal as its reference clock so the 10 MHz signal could not be interrupted, and the power and ground signals could not be used. As a result, only the IRIG-B conductors could be used as the temporary communication lines. If the Microdaq had no IRIG-B signal, the Microdaq did not know the true time of day, but this was irrelevant during the boot up process.

During the startup process the IRIG-B differential pair was commandeered by the UART. A switch was inserted before the input of the LVDS transceiver which received the IRIG-B signal and the UART `tx` signal. A final signal called `program-enable` (to indicate when the Microdaq was being reprogrammed) was used to toggle between the two modes. Under normal operation, the Microdaq was program-disabled and IRIG-B was transmitted on one differential pair and half duplex communication through another differential pair. When the Microdaq was being reprogrammed and the application being initialised after powerup, the Microdaq was put into program-enable mode. The regular communication differential pair was forced into receive mode continuously (with the `transmit-enable` pin held `low`) and the IRIG-B differential pair was forced into the UART `tx` mode continuously with the `program-enable` pin. The Microdaq was then reprogrammed and the full duplex handshake was performed. Once the Microdaq indicated that the process had been successful and the

on-board application had started, the half duplex communication was restored along with the IRIG-B timing signal. This was an unavoidable process due to the number of conductors physically available in a Cat-5 cable. The Microdaq had been developed with this procedure integral to its functionality so any future devices designed to control Microdaqs needed to be capable of this procedure.

Since the Badgerboard was designed to be deployed remotely without human supervision, environmental sensors were used to monitor the surroundings. The sensors measured pressure, temperature, humidity, orientation and acceleration. These sensors were controlled using the Beaglebone's I²C bus and the BeagleBone recorded and processed the data.

Finally, the Badgerboard was designed to support the operation of a nanobeacon which was external to the Badgerboard. The nanobeacon required power (which determined its brightness), a trigger signal (which determined its frequency) and ground. A three pin connector was used to connect the nanobeacon to the Badgerboard on the end of a wire. The trigger signal was provided by the BeagleBone's built-in PWM generator. The power signal was provided by a step-up boost converter which used the Badgerboard's 3.3 V and generated a voltage up to 36 V. The boost converter was controlled using I²C and allowed arbitrary voltages to be selected. A power switch between the boost converter's output and the board's connector was used to activate the nanobeacon after the PWM and voltage had been set. The power switch was controlled using a GPIO pin on the BeagleBone. The PWM signal was referenced using the BeagleBone's internal clock rather than the 10 MHz synchronised clock. Therefore nanobeacons could not be synchronised with each other.

5.2 Badgerboard Replacement

The first part of this chapter described the Madison electronics. The following sections of this chapter describe a new set of electronics which I designed for a new detector readout system which is being deployed in several remote locations around the world.

A successor to the Badgerboard was designed which merged features of Badgerboard with the CLB & Calamari boards known simply as the Fanout board. The Fanout board used two commercial system-on-modules: a Raspberry Pi Compute Module 4 (RPCM4) [87] and a White Rabbit LEN OEM edition.

The White Rabbit LEN OEM (WR-LEN OEM) (shown in Figure 5.3) was a White Rabbit enabled device designed to be embedded in products produced by third parties rather than the manufacturers of the White Rabbit WR-LEN OEM themselves. It had two SFP cages which were White Rabbit enabled. These provided gigabit networking over fibre. The network could both enter and exit the WR-LEN OEM in which case the WR-LEN OEM acted as a transparent network device and simply

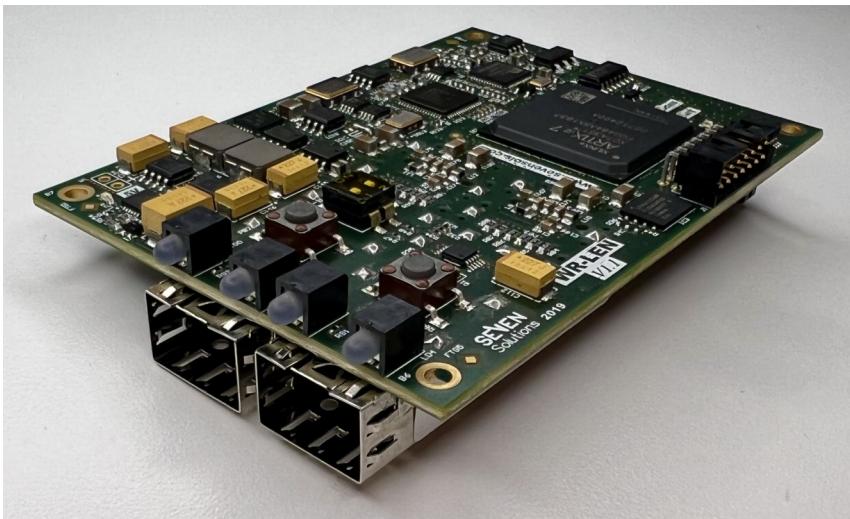


Figure 5.3: WR-LEN OEM edition for embedded use.

extracted the timing information while cascading the fibre optic network. The WR-LEN OEM's on board clock synchronised with the White Rabbit Grandmaster clock and generated a 10 MHz and IRIG-B signal. The WR-LEN OEM also had an additional feature: it could spawn an auxiliary Ethernet connection which was forwarded to the upstream network.

Other than the SFPs, the WR-LEN OEM had a single connector which carried power, the timing signals, the auxiliary Ethernet and an auxiliary UART for controlling the WR-LEN OEM itself. This meant that the WR-LEN OEM could be mounted onto a carrier board and a White Rabbit network connected to the LEN. Then the precise 10 MHz, IRIG-B and Ethernet signals could be used for functionality on the carrier board itself. A final advantage of the White Rabbit LEN was that it supported several network topologies. The LEN had a small routing table built in and two SFP ports rather than one. Up to five LENS could be daisy-chained together which removed the need for a full-sized 18-port White Rabbit Switch in small networks and could save cost and complexity. Alternatively, White Rabbit LENS could be daisy-chained with a White Rabbit Switch to increase the number of available ports. The daisy-chaining could even be disregarded entirely as another option. Overall this gave greater flexibility by supporting a variety of network typologies and reduced the startup cost as no White Rabbit Switch acting as grandmaster was needed in small networks.

The Raspberry Pi is a family of single-board computers optimised for a high performance to value ratio. For embedded use, the Raspberry Pi company produce a compute module (version 4) with more streamlined functionality. The RPCM4 used is shown in Figure 5.4. The RPCM4 had a low power quad-core 64-bit 1.5 GHz ARM based processor and the model used had 2 GB of on-board random-access memory (RAM) (up to 8 GB was supported). The RPCM4 had many built-in interfaces including (but not limited to): Ethernet, PCIE, Universal Serial Bus (USB) 2.0,



Figure 5.4: The front and rear of the RPCM4. The two 100-pin connectors can be seen at the top and bottom of the rear side of the board to connect the module to a carrier board.

UART, I²C and GPIO pins. Two 100-pin connectors (shown in Figure 5.4b) were used to connect the RPCM4 onto a carrier board. The RPCM4 supported a number of operating systems including Raspberry Pi OS (a Debian based Linux distribution).

The RPCM4 was selected as it had a number of key advantages over the BeagleBone:

1. As the RPCM4 was designed for embedded applications, the only physical connectors were the two 100-pin board to board connectors which were ideal for this application. The BeagleBone had pin-headers which were not designed to act as a BeagleBone carrier and the Ethernet used a separate physical connector necessitating the small U-turn Ethernet cable.
2. The Raspberry Pi was more performant for the cost. The BeagleBone had a 1 GHz 32-bit single-core processor and 0.5 GB of RAM whilst RPCM4 had a 1.5 GHz 64-bit quad-core processor and 0.5 GB despite costing approximately the same price (approximately \$35).
3. The Raspberry Pi is a young platform with an active community of developers for both hardware and software so is more readily available. This makes development faster and cheaper as existing resources can be more readily used.

The combination of the WR-LEN OEM and the RPCM4 meant that the White Rabbit delivered a precise time and clock onto the carrier board. The Ethernet spawned by the WR-LEN OEM, which was still connected to the main White Rabbit network despite not delivering the precision timing, could then be bridged to the RPCM4. The fact that both devices used board-to-board mezzanine connectors meant this could be done using just the circuit traces on the carrier board rather than small U-turn cables. This solution delivered both the timing and networking elegantly.

A drawback of the White Rabbit LEN - Raspberry Pi system was that due to the network being delivered over fibre, a separate copper power cable was required. This was considered a worthwhile trade off for the faster networking speeds while maintaining White Rabbit timing. If the White Rabbit was integrated onto the fanout rather than in an upstream system, any timing delays to the fanout would already be accounted for automatically by the White Rabbit network.

To improve upon the UART system used by the BeagleBone to control the Microdaqs, the system was redesigned for the new Fanout. Rather than using four UARTs, the new system used the built-in USB 2.0 interface which was comparatively high speed (480 Mbit s^{-1} [88] compared with the BeagleBone's UARTs which were up to 3.6 Mbit s^{-1} [89]). The USB bus was connected to a 7-port USB hub and each of the 7 USB channels was then attached to a USB to quad-UART bridge. This gave a total of 28 UART channels. Because the USB bus was so much faster than UART and the USB-to-UART bridges had in-built I/O buffers, from the Raspberry Pi's point of view, there were 28 discrete UARTs that could be accessed simultaneously rather than by sequential sampling. In reality sampling was occurring but at speeds so fast that the buffers bridged any communication gaps. Therefore 28 full speed UARTs were achieved through a simple $1 \times 7 \times 4$ multiplication rather than a complicated system of multiplexing. Furthermore, the number of channels went up from 16 to 28.

A further advantage of the USB-to-UART bridges was that they were designed with the expectation of advanced configuration, including transceiver control and support for both half and full-duplex communication. The bridges could be configured to control the `transmit-enable` pin built into the bridge itself. The bridge could control an LVDS transceiver while it was in use. This allowed half duplex communication without custom Linux Kernel drivers to control the transceiver itself; instead controlling the transceiver was managed by changing the mode of the UART bridge and the data flow control was managed in hardware. From the point of view of the Raspberry Pi, it was just a standard UART with no unusual functionality; this greatly simplified software development for Microdaq control. Like the Badgerboard, each UART was routed into an LVDS transceiver, but unlike the Badgerboard, the transmit and receive modes were toggled by the UART bridge. The differential UART signals were then routed to the output connectors.

Since the UART bridges came in quartets, all other aspects of controlling Microdaq channels were also done in quartets so the supporting electronics could be grouped into fours. Like the Badgerboard, 8P8C connectors (for Cat 5 cables) were used for the outputs to the Microdaqs, and blocks of four connectors were used. The 10 MHz and IRIG-B signals were each fanned out 7 times and routed to each block of connectors where they were fanned out another four times each towards LVDS transceivers.

The implementation of full-duplex communication for reprogramming and initialising Microdaqs was recreated using the same method as the Badgerboard. The USB-to-UART bridge was used for both the **tx** and **rx** communication lines; however the **tx** lines were redirected through the IRIG-B differential pair which was perpetually in transmit mode. The **rx** lines were routed through the standard communication channel.

Therefore, during the startup process the IRIG-B differential pair was commandeered by the UART. A switch was inserted before the input of the LVDS transceiver which received the IRIG-B signal and the UART **tx** signal. A final signal called **program-enable** (to indicate when the Microdaq was being reprogrammed) was used to toggle between the two modes. Under normal operation, the Microdaq was program-disabled and IRIG-B was being transmitted on one differential pair along with half duplex communication through another. When the Microdaq was being reprogrammed and the application was being initialised after powerup, the Microdaq was put into program-enable mode where the IRIG-B differential pair was forced into the UART **tx** mode continuously. Unlike the Badgerboard, the regular communication differential pair was still controlled by the USB-to-UART bridge rather than forced into a fixed mode with a GPIO pin. The Microdaq was then reprogrammed and the full duplex handshake was performed. Once the Microdaq indicated that the process had been successful and the on-board application had started, the half duplex communication was restored along with the IRIG-B timing signal.

As with the Badgerboard, a power switch was used to switch the 24 V supply to each Microdaq. The signal used to turn on and off each channel's power switch was used as an indicator to show if a channel was turned to the on or off state (the **POWER-ENABLE** signal).

Since every channel required a **program-enable** signal and a **power-enable** signal and there were 28 channels, 56 signals were needed in total. The Raspberry Pi only had 28 GPIO pins. The Microdaq channels were controlled in blocks of four and therefore seven sets of 8 GPIO signals would be required. Rather than using the Raspberry Pi's GPIO signal pins, seven I²C GPIO extenders were used which had eight I²C controlled GPIO pins. Therefore each block of four Microdaqs had four **program-enable** pins and four **POWER-ENABLE** pins. Every single pin was controlled by the I²C bus.

The fanout system broadly consisted of a set of central electronics including the Raspberry Pi, the White Rabbit LEN and the power system, and then a set of 7 groups of the circuits for managing four Microdaq channels. To ease prototyping, the fanout was split into two kinds of circuit board that were connected together with board-to-board mezzanine connectors. This was to enable the electronics to be advanced to the prototyping and testing phase faster and the designs could be merged into a single larger PCB at a later date. The two kinds of board were known

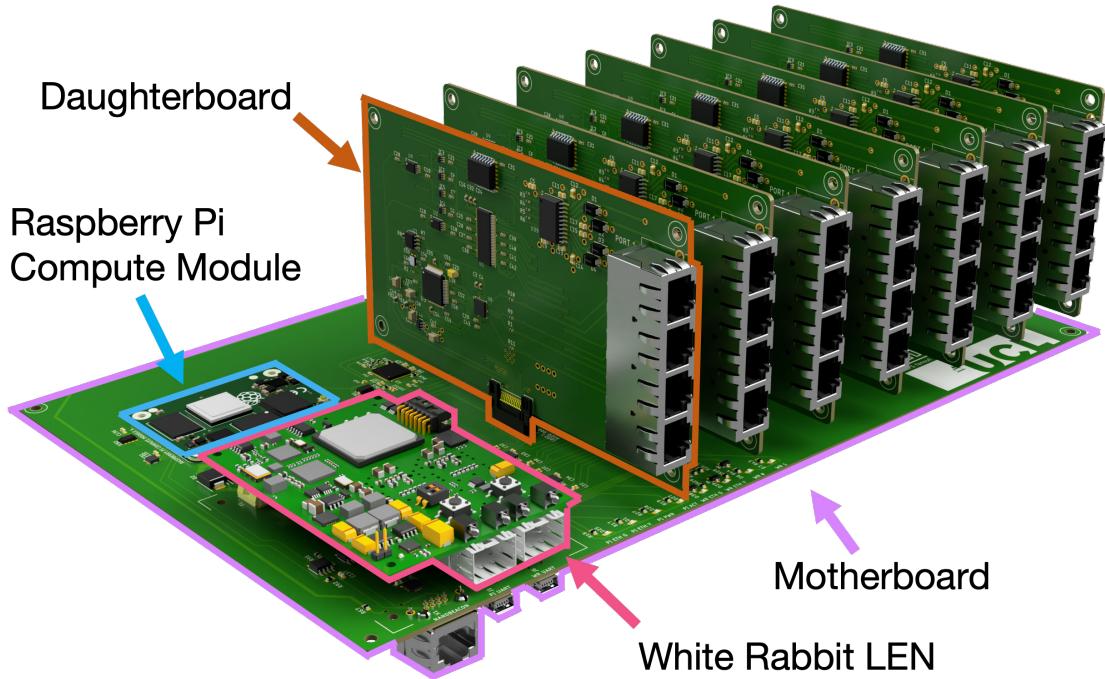
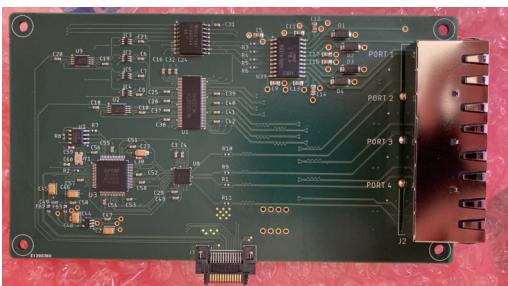


Figure 5.5: The prototype of the new fanout system. The motherboard is carrying 7 daughterboards with ports for four Microdaqs each.



(a) Prototype daughterboard.



(b) Prototype motherboard.

as the *motherboard* and the *daughterboard*.

The motherboard had the power supply, the Raspberry Pi, the White Rabbit LEN, the USB hub and the initial 1×7 timing fanout. The daughterboards had the USB quad-UART bridge, the I²C GPIO expanders, the power switches, the 1×4 timing fanout and the LVDS transceivers. The motherboard supported seven daughterboards. A rendering of the completed system can be seen in Figure 5.5. A prototype motherboard and prototype daughterboards can be seen in Figures 5.6a and 5.6b respectively.

The motherboard had sensors that measured the same environmental parameters as the Badgerboard. The sensors were the MS8607-02BA01 to measure pressure, temperature and humidity (chosen for its high resolution over the entire room temperature range and a humidity range exceeding 60% relative humidity.) and the MXC4005XC to measure orientation and acceleration. The motherboard added auxiliary USB UART ports for both the Raspberry Pi and the White Rabbit LEN. The ports exposed the serial consoles of both of the devices so that in the event

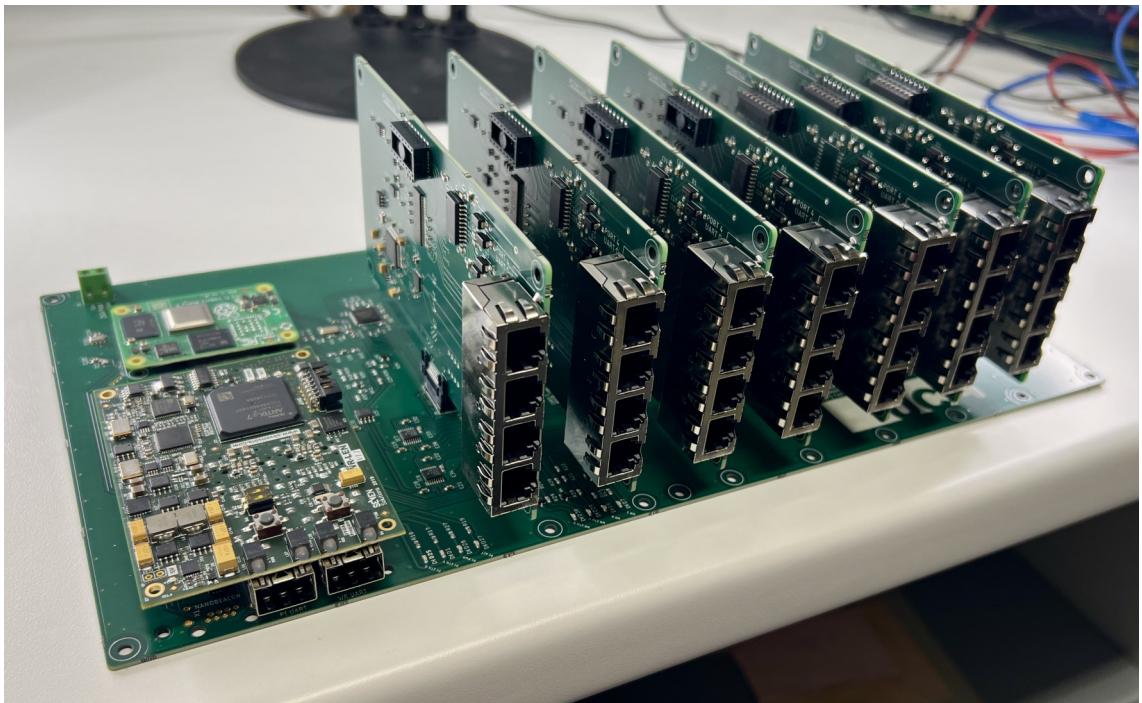


Figure 5.7: A manufactured and assembled fanout.

of a network outage, local debugging could be performed which was not possible using the Badgerboard. Furthermore, a number of LED indicator lights indicating the statuses of the devices and the network were added for status information.

The new fanout had no nanobeacon circuitry built in but to allow for a nanobeacon (with precise timing, unlike the Badgerboard), a connector was added that distributed the 10 MHz and IRIG-B signals, 24 V power and ground reference, and the Raspberry Pi's I²C bus using a long range I²C extender so it could be transmitted over a long cable. This allowed an expansion board to be plugged in to power and trigger a nanobeacon using the standardised CHIPS time.

The system was manufactured and tested using Microdaqs. An assembled fanout can be seen in Figure 5.7. The circuit diagrams can be seen in Appendix B.

5.3 Software Delivery and Data Storage

While using the Madison system of electronics in CHIPS, the challenges of vital software maintenance arose on the BeagleBones as there was no system for automatic software updates and synchronisation. Furthermore, the data gathered were forwarded to the DAQ storage but if the connection was lost then the data were also lost because the storage available locally was insufficient for the data gathered. The hardware updates presented an opportunity to improve the convenience of synchronising the software and storing the data.

The RPCM4 supported booting from local eMMC storage as well as booting over the network (with an adjustable boot order). The RPCM4 was configured to boot

from the network with the local storage as a secondary fallback option. If the network was available and a boot server was visible then the operating system was automatically retrieved and booted. If no network was available then an emergency local operating system was booted. All RPCM4s used the same boot server so they all received the same operating system and only one system needed to be maintained.

To prepare the operating system used on the RPCM4, a system of scripts was used to provision an operating system for the boot server to distribute. The scripts simulated a RPCM4 using QEMU [90], virtually installed the latest version of the Raspberry Pi OS, automatically compiled the latest version of the DAQ software for the RPCM4 and then created an image of this new operating system from the simulation. The boot server was then updated to use the latest operating system and the RPCM4s deployed automatically retrieved the latest version. The data was not tagged with the DAQ software version used during operations; however, the option exists to tag data in the future and will be used as the software matures and development slows. In summary, scripts built a fresh version of the operating system using the latest DAQ software and the RPCM4s automatically received an up-to-date version without the need to manually update any individual device.

The filesystem on each RPCM4 was upgraded in such a way that a storage server was used rather than a local filesystem. When the operating system booted, the remote filesystem was fetched and mounted locally (using NFS - Network File System [91]). The fileserver had access to all of the filesystems of all of the RPCM4s so they were automatically backed up and by design the latest files from the RPCM4s were sent to the central storage system. In the event of network loss, the RPCM4s could fall back to temporary local storage.

To access the remote operating system server and the remote filesystem server, the devices were added to a virtual private network of devices (VPN). As the hardware was all attached to a White Rabbit network, the router managing the White Rabbit network could act as the VPN gateway and therefore anything added to the White Rabbit network was automatically already part of the virtual CHIPS network (but this was not a requirement, just a convenience). Fanout Devices could be added to the virtual CHIPS network individually. The virtual network had the result that all of the devices on the network acted as one collective unit even if they were not physically connected or even in the same physical location. The management servers could be in a remote central location while the data-acquisition nodes could be at a number of remote locations (for example underground). Adding a new fanout board to the system was as simple as adding it to the VPN, powering it on and assigning it a name in the configuration file. The software and filesystem were then automatically loaded onto the RCM4.

Nodes on the VPN did not have to be physically connected using CHIPS infrastructure; they could be connected over the internet. As long as disjointed portions

of the network both receive a reference White Rabbit time that is synchronised, then the system acted as if it was a contiguous local network rather than a set of disjointed nodes.

From the point of view of any central control software, the `RPCM4s` were all on the same local network and in the same physical location as the `DAQ`. A disjointed set of fanouts appeared to be one homogeneous detector regardless of location or physical connectivity.

A test system has been constructed with boot and filesystem servers at University College London (UCL). An assembled prototype fanout board is located at UCL. A second fanout had been configured at Czech Technical University (CTU) in Prague, Czech Republic. A third fanout is in the process of being commissioned in Qui Nhon, Vietnam. Each node use Microdaqs and PMTs reused from `CHIPS` and the performance of the hardware and software system is being assessed over time.

5.4 Summary of Fanout Development

Using the lessons learned from `CHIPS`, a new fanout board has been developed. The fanout can act as a standalone device with no supporting infrastructure or part of an ensemble with coordinated timing, filesystems and software. White Rabbit timing is provided directly to the fanout and the on-board computing has been upgraded from a BeagleBone Green to a Raspberry Pi Compute Module 4. The number of channels has been upgraded from 16 to 28. The hardware is being evaluated at University College London and at Czech Technical University and further development and time synchronisation between the two locations is currently being carried out.

Chapter 6

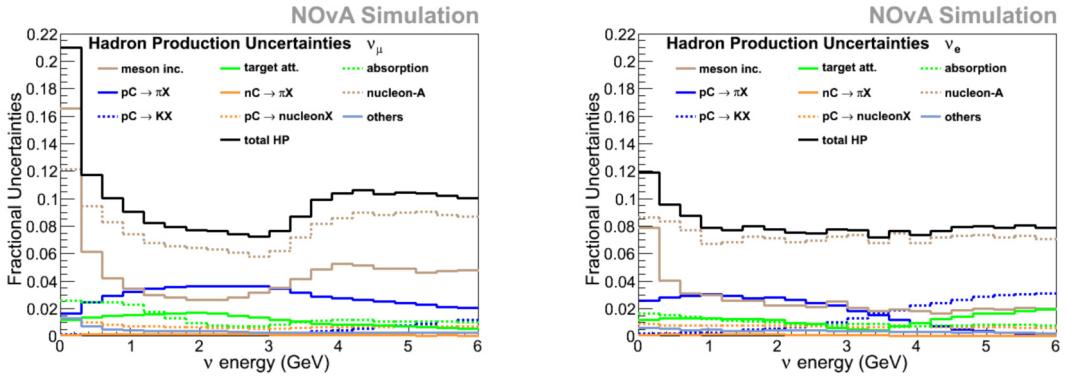
Neutrinos & Muons in Rock

Neutrino beams produce a very large flux of neutrinos which pass through the Earth. As the neutrinos travel through the rock under the Earth's surface, they may interact - muon neutrino CC interactions produce muons. The probability of an interaction is dependent on the density of rock. The muons themselves then interact with the rock, the probability of which is also proportional to the density of the rock. Along the path of the neutrino beam there therefore also exists a muon flux.

In existing detectors, muons produced in the Earth may be recorded at the beam-facing end of the detector and have to be filtered out as an unwanted background. However, as the rock muons are an intrinsic and unavoidable product of a neutrino beam travelling through the Earth they may be harnessed and the rock used as 'free' extra detector mass.

Only instrumentation is needed to record and measure the rock muons - the Earth is already in place. This means that wherever there is rock there is a potential detector. As previously discussed in Section 3.3.1, there is a significant uncertainty in the fraction of neutrinos produced by kaons (larger than the uncertainty due to pions) in the NuMI beam. This is important because at off-axis angles, the energy of the neutrinos becomes increasingly directly dependent on the mass of the parent hadron (as explained in Appendix A). At off-axis angles, the high energy tail of the neutrino energy distribution becomes dominated by neutrinos produced by kaons rather than pions. Consequently, it is important to investigate techniques for estimating or measuring the kaon content of neutrino beams - this is shown in Figure 6.1.

A possible use for muons in rock would be to monitor a neutrino beam across a wide range of off-axis angles - multiple detectors at different off-axis locations or a locally mobile detector could be used. As discussed in Section 3.3.1 (and Appendix A), the energy spectrum of the neutrinos varies as the off-axis angle changes, the form of which is related to the hadron composition of the beam used to create the neutrinos. Understanding the composition of the hadron beam is important for reproducing the neutrino beam in simulations (see Figure 3.11). A potential



(a) Hadron flux uncertainties on ν_μ production.

(b) Hadron flux uncertainties on ν_e production.

Figure 6.1: The hadron flux uncertainties on neutrino production in the NOvA experiment. $pC \rightarrow \pi X$ represents proton-carbon interactions yielding pions and $pC \rightarrow KX$ represents proton-carbon interactions yielding kaons [92]. For muon neutrinos, the pion production uncertainty dominates but the kaon production uncertainty does increase as the energy increases. For electron neutrinos, initially the pion production uncertainty dominates. However, at approximately 3.5 GeV, the kaon production uncertainty overtakes the pion production uncertainty and continues to rise. This means that at higher energies, the uncertainty in the number of electron neutrinos produced becomes increasingly dominated by the kaon production uncertainty. For experiments that measure muon neutrino disappearance and electron neutrino appearance, the uncertainty in kaon production is important. Figure from [93].

solution would be to count neutrinos as the off-axis angle changes and use the flux to infer information about the number of pions and kaons in the hadron beam. The advantage of using rock muons would be that only the instrumentation would be needed at each off-axis angle rather than an entire detector medium. Or, the same instrumentation could be moved along a rock face.

Another possible use for rock muons is as a receiver for neutrino communication. Neutrinos have been considered for communications that could travel over long distances [94] and through objects that would block traditional communication methods, but this is not feasible with current methods of producing neutrinos [95]. Neutrino communication is limited by the low probability of a neutrino interaction occurring and the infrastructure needed to produce a high power neutrino beam. Furthermore, the large detectors used for detecting neutrinos cannot be moved [96]. Rock muons are a potential candidate for a neutrino communications' detector as the rock (or water) acts as a large detector medium, is available in many land (and sea) locations as an existing resource and the flux of rock muons follows any muon-neutrino beam. An example potential use case is stealthy communication between a submarine and its home base separated by a large deep ocean. The issue of the neutrino source, however, remains.

6.1 Muons in Rock

Once a muon neutrino interacts via a CC interaction, the charged muon will travel through the rock for a limited distance. The distance the muon travels depends on the density of the rock and the energy of the muon. The *stopping power* S is the energy the muon loses per unit distance travelled and is given by

$$S = \frac{-dE}{dx} \quad (6.1)$$

where E is the energy of the muon and x is the travel distance. The stopping distance of the muon R (the *range*) is given by

$$R = \frac{E}{\frac{-dE}{dx}} = \frac{E}{S} \quad (6.2)$$

Conventionally, the quantity *mass stopping power* is used to characterise a material rather than simply the stopping power. The mass stopping power P is defined as

$$P = \frac{S}{\rho} \quad (6.3)$$

where ρ is the density of the material. Therefore, overall, the stopping distance d of a muon in a material is given by

$$d = \rho \frac{E_\mu}{P(E_\mu)} \quad (6.4)$$

It is therefore clear that the distance a muon travels once it has been created is dependent on its energy and the density of the rock. High energy muons travel further and are therefore more likely to travel far enough to intersect with any detectors. Low energy muons, although they may be more likely to occur, do not travel as far and so their abundance is offset by the reduced probability that they will travel far in rock. This is illustrated in Figure 6.2. Furthermore, the probability of a neutrino interacting and creating a muon is dependent on the neutrino cross section which is dependent on the neutrino's energy. So, there is a double dependency on energy for the number of muons at a given energy. Therefore it is important to weight any muons created for the energy of the neutrino because of the travel distance and the neutrino cross section.

6.2 Detector Concept

A simple detector concept is proposed as a tool for studying rock muon fluxes in simulations. The imaginary detector design is shown in Figures 6.2 and 6.3.

The detector consists of two thin tracking plates (for example, amongst other

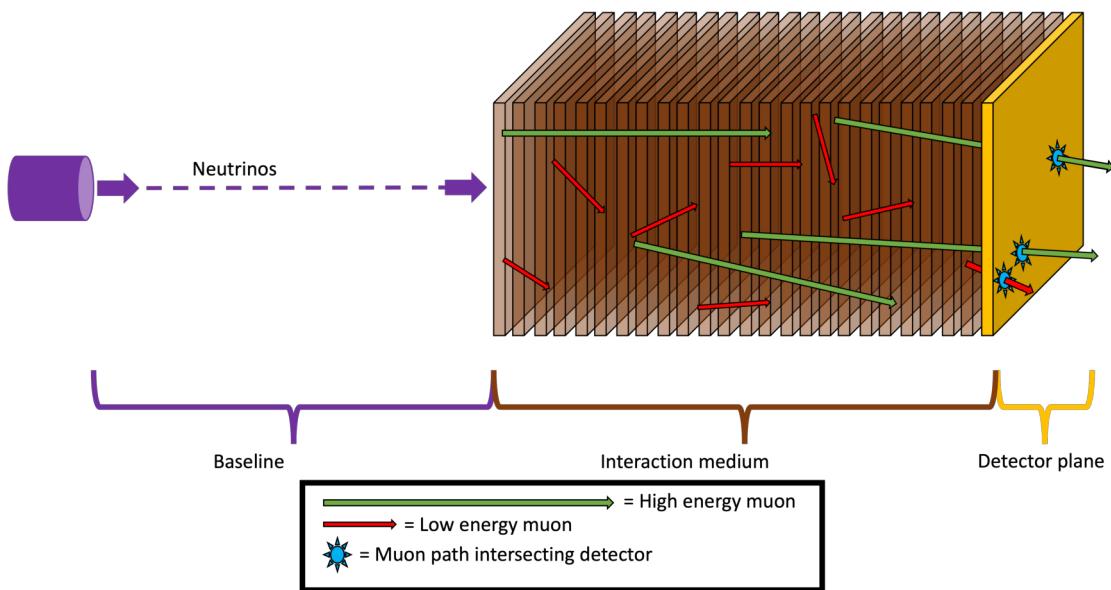


Figure 6.2: A cartoon of neutrinos interacting through a solid medium (rock). The neutrinos interact to create muons. The high energy muons travel far in the rock whereas the low energy muons only travel short distances. The more energy a muon has, the more likely it is to travel far enough to intersect with the detector plane. There is therefore a high-energy bias when counting the muons in the detector compared with the energy distribution of the muons or the original neutrinos.

technologies, scintillator strips and resistive plate chambers) that detect when a muon passes through and the direction the muon is travelling. The first detector plate is exposed to the rock and there is empty space between the first and second plate. Between the two plates there is a narrow region of magnetic field perpendicular to the plates (provided by, for example, a Helmholtz coil).

The first detector plate measures the direction the muon enters the detector; the magnetic field deflects the muon as it travels; and the second plate measures the position-after-deflection and new direction vector of the muon. The sign of the muon's charge can be inferred from the deflection direction and the muon's energy can be inferred from the deflection magnitude.

6.3 Simulation Overview

A simulation was written to investigate utilising the rock muons and the concept detector to assist in profiling neutrino beams; in particular, to investigate the kaon-parent contribution in the muon flux off axis and at higher energies.

To simulate this, a source of parent hadrons and their corresponding neutrinos was needed. The neutrinos were required to be propagated through the Earth and muons created. The muons then had to be propagated into the first detector plane and deflected by the magnetic field. The position of the muon at the second detector plane then needed to be calculated. The final muon deflections and directions could then be investigated with respect to the original neutrino beam and original parent

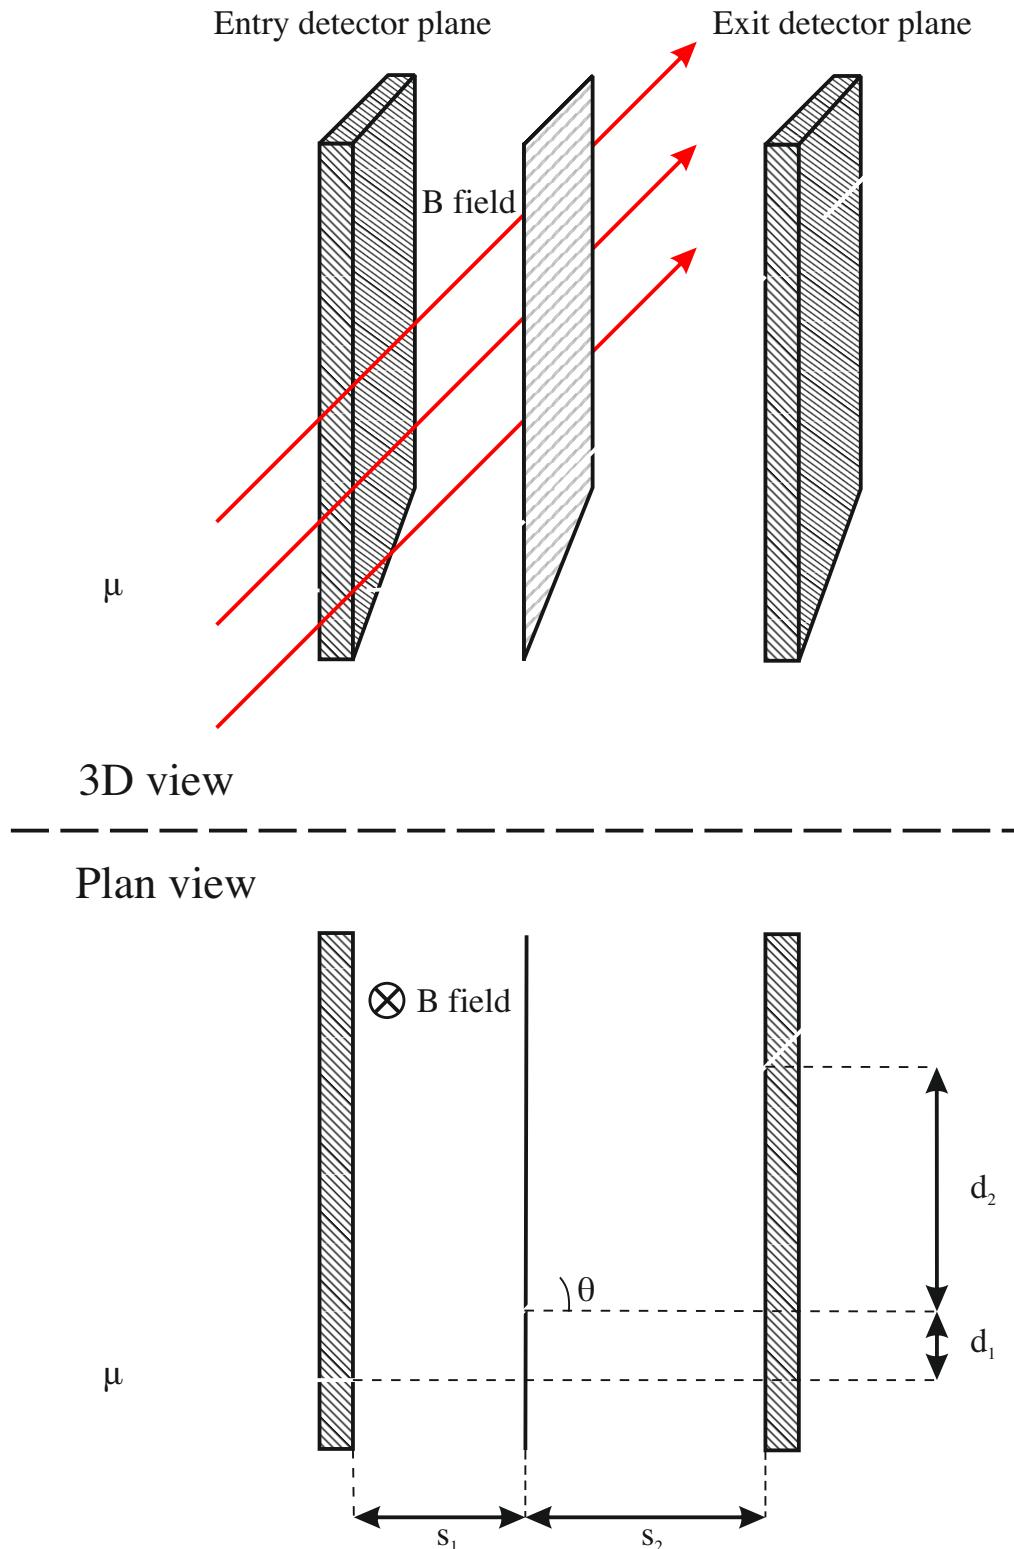


Figure 6.3: A diagram of the detector concept. The muons travel through the rock and then enter a detector plane which measures the direction of travel. The muon is influenced by a narrow region of magnetic field perpendicular to the detector plane which bends the path of the muon. As the muon leaves the magnetic field region, it continues travelling tangentially until it reaches the second detector plane which measures the direction of travel upon exiting the detector. Between the two detector fields, there is empty space for the muon to travel through.

hadrons.

The simulation was conducted for detectors at multiple locations in a single neutrino beam and only considered muon neutrinos and muons.

6.4 Generating & Propagating Neutrinos

To generate the neutrinos, G4NuMI, a Geant4 [97] [98] [99] based simulation for the NuMI beam, was used to generate neutrino flux spectra in the form of Dk2NU files [100]. Dk2NU are ROOT [101] files containing information on a neutrino beamline simulation in a predetermined format. The files contained:

- A parent hadron, specifying the type, a momentum vector and a decay vertex. Also, an importance weighting for the hadron.
- A neutrino produced by the parent hadron which specified the neutrino's energy and momentum vector and a weighting for a given detector location.

In order to generate a Dk2NU file, the entire production chain of grandparent protons, parent hadrons of pions and kaons. Then, the production of the neutrino and complementary lepton, needed to be simulated (by G4NuMI). A 120 GeV Gaussian provided the source for the protons fired at the target [102]. The parent hadrons were produced by interacting the grandparent protons in a graphite target to create parent hadrons. To represent as much of the spectrum of parent hadrons as possible, the least likely decays were over-simulated and the most likely were under-simulated. Each decay was assigned an importance weighting known as `Nimpwt`. `Nimpwt` weights each decay so that they are correctly distributed but allows the least likely decays to still be included. This would be an alternative to having to generate impossibly large numbers of decays so the least likely relevant processes occur in statistically significant quantities. The hadrons were then decayed and used to generate a list of neutrinos, neutrino properties and detector location weightings. The detector location weighting corresponds to the likelihood of a neutrino being forced through the detector at a given location [103].

For detectors that move to arbitrary locations, the problem arose that the entire NuMI beam simulation would have to be repeated for every single possible location and the entire beam regenerated to obtain each location's neutrino weightings. This was not feasible. Instead, one set of Dk2NU geometry was created which contained the correct (weighted) parent hadron distributions and the hadrons were then manually decayed. The neutrinos that were produced were then considered the definitive product the neutrino beam and the detector was moved around. A check was performed to determine if a neutrino intersected the detector. This allowed the detector to be mobile without regenerating a new neutrino beam but also introduced the random possibility that neutrinos could miss the detector entirely.

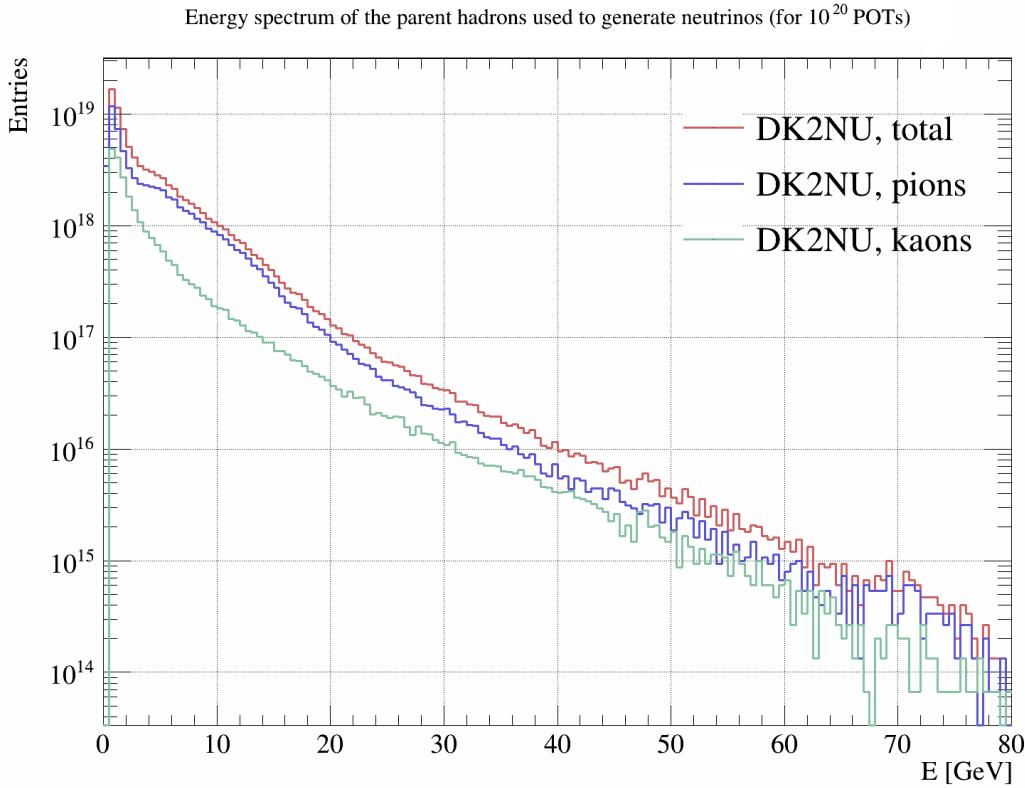


Figure 6.4: The energy spectrum of the parent hadrons of the neutrino beam approximating the NuMI beam in low energy configuration. Pions are more abundant than kaons.

Because neutrinos could miss the detectors, to increase the number of neutrinos actually arriving at a detector, the hadrons were reused with fresh random neutrinos generated. This again reduced wastage while maintaining the correct shape of the neutrino distribution.

A crosscheck was performed to confirm that manually creating neutrinos through parent hadron decay yielded the same results as the pre-produced neutrinos in the `DK2NU` file. Detectors were created to match the sizes (1 m radius circles) and locations (1036 m downbeam on axis, the position of the MINOS near detector) provided in the `DK2NU` file, and the neutrinos through each detector region were recorded. The hadron spectrum due to the proton beam is shown in Figure 6.4. The neutrino spectrum generated from the hadron decay is shown in Figure 6.5. The effect of the hadron weighting `Nimpwt` and detector location weighting `wgt` on the pre-provided `DK2NU` and the custom-made crosscheck detector are shown in Figure 6.6. The hadron weighting affected both the `DK2NU` detector and the crosscheck detector. The detector location weighting was only required for the `DK2NU` detector as the random generation of neutrinos correctly matched the weighted `DK2NU` distribution in the custom-generated detector.

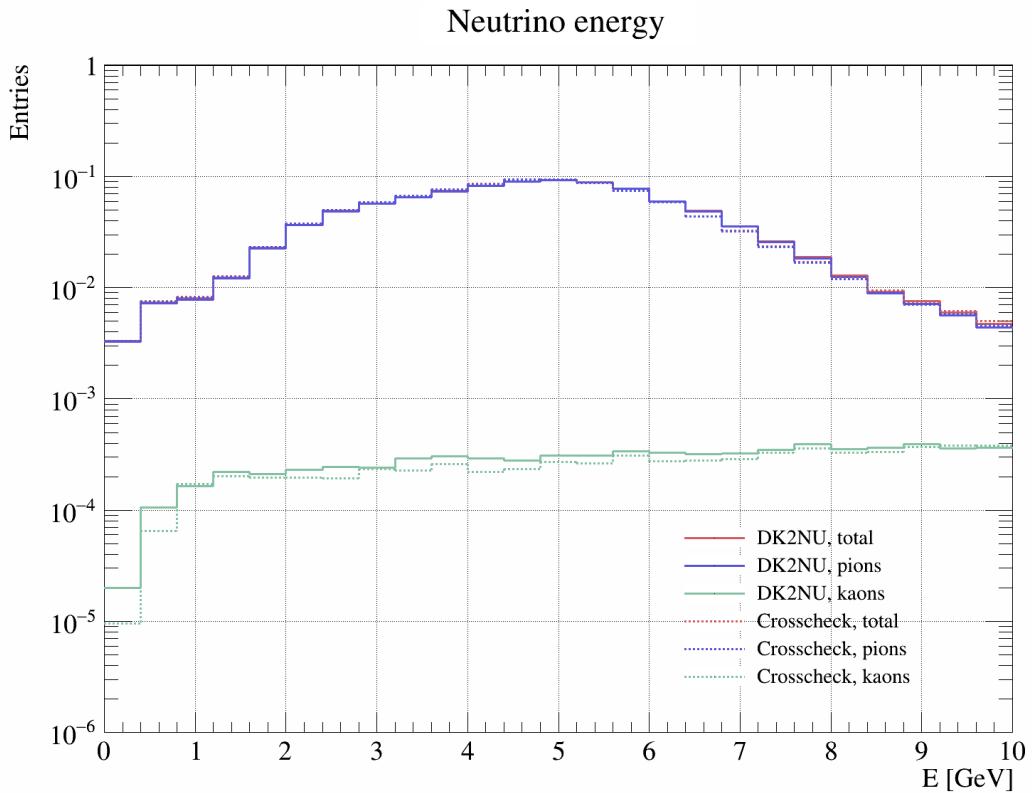


Figure 6.5: The energy spectrum of muon neutrinos intersecting a 1 m radius circle 1036 m downbeam on-axis.

6.5 Generating Muons

After being generated, the neutrinos' paths were traced downbeam. The neutrinos may have interacted at any point in their path of travel. After interacting, the muon produced will have a definite momentum vector - the direction of travel. The simulation needed to take into account the position at which the neutrino interaction occurred and the direction and energy of the muon after the interaction took place - i.e. did the muon reach a detector or not? Ideally, each neutrino would have been generated and then individually propagated, checking at every point along the path if an interaction occurred; and, once the muon had been created, propagating the muon and checking if it reached a detector. This would have been computationally expensive. To simplify this problem, an important approximation was made. Any detectors that were simulated must be close to the neutrino beam source such that virtually no neutrinos had already interacted and been lost. The neutrino flux through the region containing detectors is the same as the neutrino flux at the beam source.

The rock was assumed to be of uniform density. If this assumption was not made, a detailed finite element analysis would have to be performed to check the probability of a neutrino interacting at every point along its path. Assuming constant density means a simple neutrino survival probability can be calculated and muon fluxes can

Effect of weightings on neutrinos

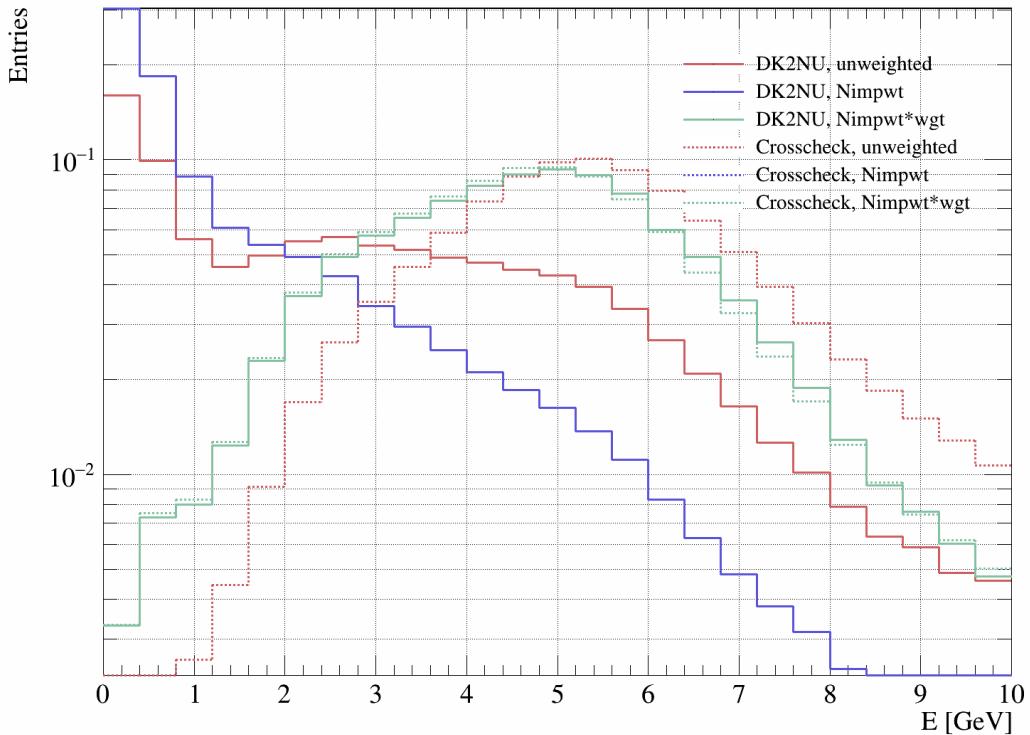


Figure 6.6: The cumulative effect of different weightings on the neutrino flux. The DK2NU neutrinos are initially different due to the reverse ray-tracing of the DK2NU detector neutrinos. Adding in the hadron weighting `Nimpwt` affects both the DK2NU detector and the custom crosscheck detector. Adding the neutrino location weighting to both changes the DK2NU spectrum but not the crosscheck spectrum. This because the crosscheck detector already correctly randomly distributes the neutrinos so the dashed green and dashed blue lines directly overlap. The DK2NU detector with both hadron and location weightings agrees with the crosscheck detector with just the hadron weighting. The location weighting of the crosscheck detectors was fixed at 1 so each event recorded represented a single real event.

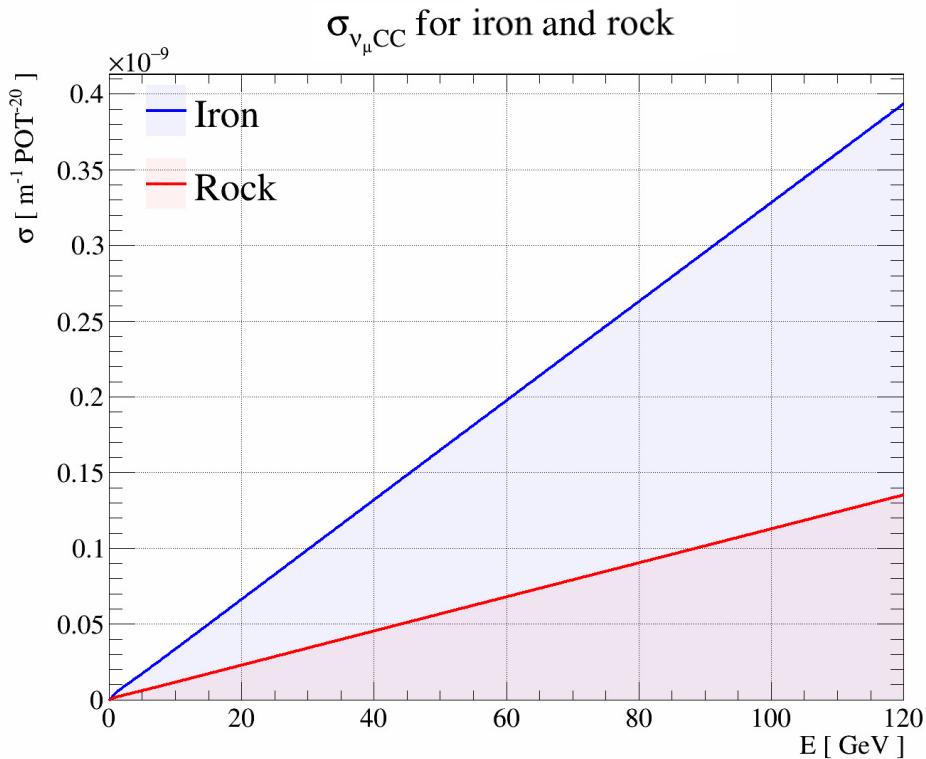


Figure 6.7: The probability of a neutrino interacting via a charged current interaction per metre of travel through rock and iron (per 10^{20} POTs).

then be inferred. Although limiting, the uniform rock density allowed the creation of a simple and performant simulation and a more realistic representation of the Earth's rock is an area for future improvements.

Once the source of neutrinos had been generated, muons needed to be produced. The energy of every neutrino was randomly thrown to give a muon's energy. In the energy range of 1 GeV to 120 GeV covered by the NuMI beam, the energy a muon loses per unit of travel is approximately linear [25]. The mass stopping power, the density of the rock and the energy of the muon could then be used to create a stopping distance for a muon of a particular energy. This was the effective length of the detector for the muon. The cross section of the neutrino was calculated for the rock medium and converted into units of m^{-1} (see Figure 6.7), the probability of a neutrino interacting per metre of travel in rock. The effective length of the detector and the neutrino cross section then gave an interaction probability for a muon of a specific energy in rock.

However, the flux of muons through the detector plane was biased. A muon may have been created within the effective detector volume but it may not necessarily have travelled far enough to intercept the instrumentation. Low energy muons have a shorter effective path so although they may have been created (more often) inside the detector volume, they will not have travelled as far and therefore will not have reached the instrumentation as often.

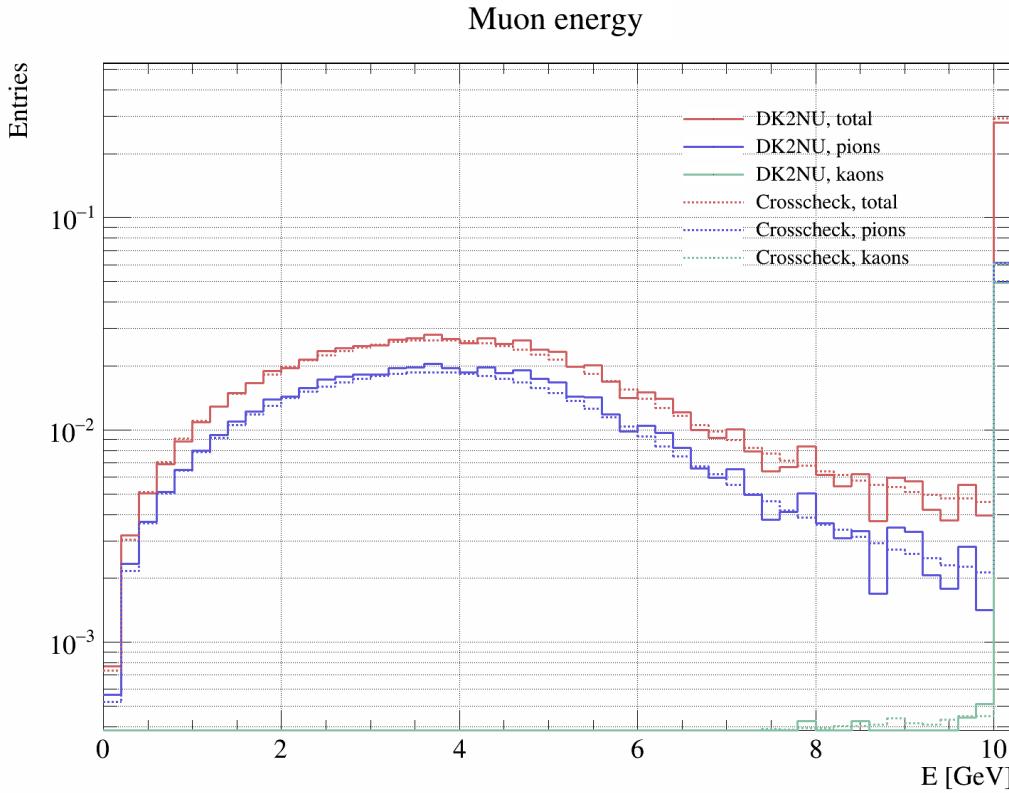


Figure 6.8: The muon energy spectrum for both the DK2NU and crosscheck detector.

The muons needed to be weighted by their path length through the rock (as a function of energy) so that the short path of low energy muons and the longer path of higher energy muons are correctly represented in the detector muon energy distributions. The path length of the muons was the detector's effective length so that was used as the weighting factor.

Fluctuations in the energy loss of the muons due to random energy losses in the rock before the muons arrive at the detector were ignored and presumed to be small compared with the total energy of the muons. This could be improved in future simulations and would make the low energy tail of the muon distribution more accurate.

The weighted energy spectrum of the muons passing through the detector is shown in Figure 6.8 and a comparison with and without the muon path weighting is shown in Figure 6.9. The muon path length weighting shows that, when the travel distance of the muon is taken into account, low energy muons are less frequent due to the reduced probability of the muons intersecting the detector.

The effective length of the detector is derived from the density of the medium (rock) and combined with the cross sectional area of the detector face to give a detector with an effective mass. This figure represents how large the detector is despite the fact that the medium is provided by the Earth at no cost to expand. The effective mass is shown in Figure 6.10 and shows that the effective mass of the detector expands arbitrarily as the energy of the muons increases because of the

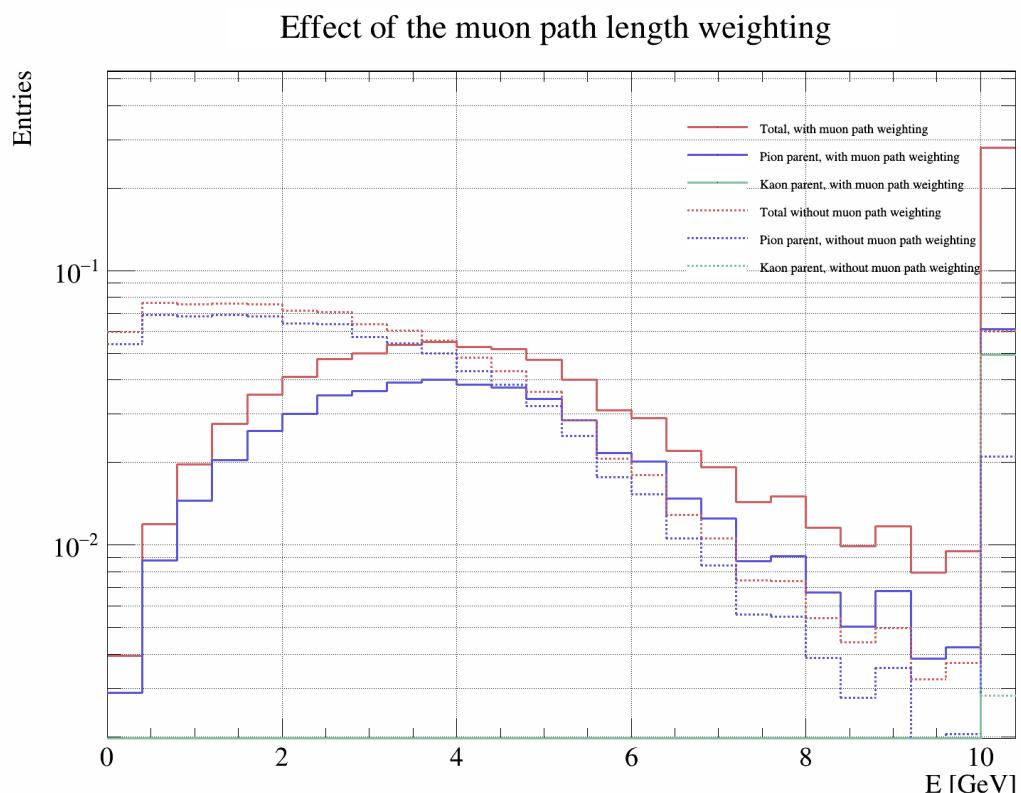


Figure 6.9: The effect of the muon's travel distance on the muon energy spectrum. The low energy muons are less likely to travel far enough to reach the detector so are weighted downwards in the spectrum.

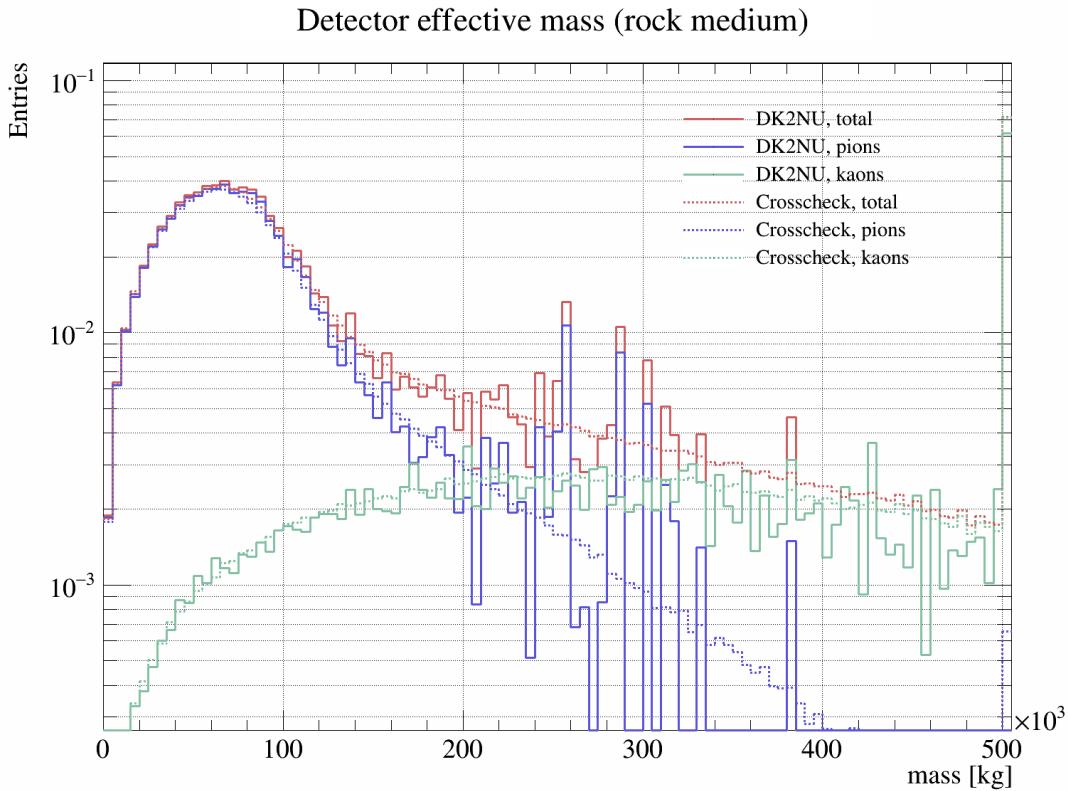


Figure 6.10: The effective mass of the detector. High energy particles can travel further through the existing rock which gives the detector a larger effective mass without altering the experimental setup.

mass of the Earth in front of the first detector plane.

Once the muons were generated with specific energies, their directions were randomly chosen using a Gaussian with a mean of zero and a width of 0.28 rad. This value was motivated by Figure 6.11 which shows that 1σ of muons fall within 0.28 radians of each neutrino's direction. All of the neutrinos were assumed to be going directly forwards at this stage.

6.6 Detector Deflection

The muons were created in the rock in front of the detector. Rather than fully modelling the path into the detector instrumentation and then the path between the two detector planes, an approximation was made. An angular cut was performed such that the muon's momentum had to be at least 99% in the forward direction and at most 1% in the transverse direction - approximately 13% of muons survived this cut. This means that for the muons produced in the rock in front of the detector and that survived the cut, the majority of the muons' momentum vector is perpendicular to the detector plane itself. For a physical application, as an example, this cut could be performed by alternating scintillator strips at each detector plane checking coarsely the direction of the muon. The second approximation made was

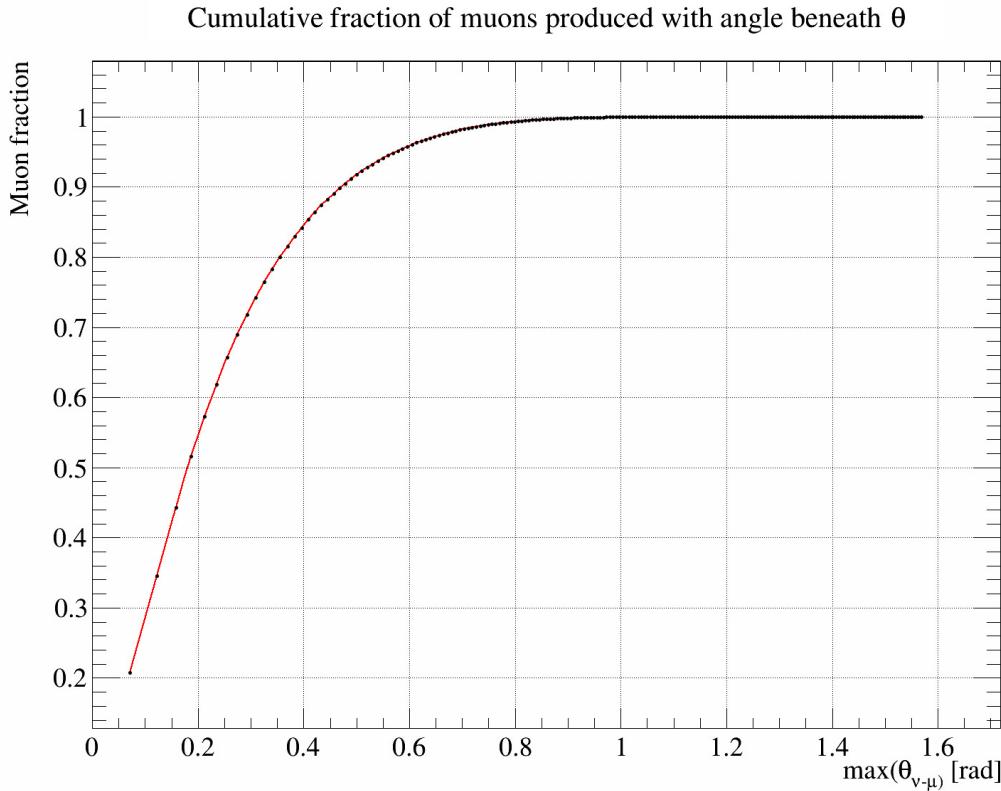


Figure 6.11: The fraction of muons where the angle between the muon's direction vector and the parent neutrino's direction vector is less than θ .

that the muons that entered the detector would survive long enough to leave the detector.

To deflect the muons, a magnetic field was used. The magnetic field was perpendicular to the detector plane and the incoming muons. The magnetic field formed a uniform region that was distance s_1 wide; the remaining space between the two plates not taken up by the magnetic field was s_2 wide. (The total distance $s = s_1 + s_2$ is the total separation between the two detector planes.) The experimental setup can be seen in Figure 6.3. Since the magnetic field region was narrow and the muons were travelling perpendicular to the field lines, the magnetic field provided a rapid 'kick' to deflect the muons rather than a long continuous region to curve the muons which would be hard to maintain at a uniform magnetic field strength. As the muons left the region of the magnetic field, the muons travelled tangentially through the detector until they reached the second detector plane. The deflection can also be seen in Figure 6.3. The deflection in the region of width s_1 due to the magnetic field was d_1 and the deflection due to the tangential trajectory after the magnetic field in the region of width s_2 was d_2 .

The velocity v of the muon is given by

$$v = \frac{c}{\sqrt{\left(\frac{mc}{p}\right)^2 + 1}} \quad (6.5)$$

where m is the mass of the muon, p is the magnitude of the muon's momentum and c is the speed of light. The radius of curvature r of the muon's path as it travels is given by

$$r = \frac{mv}{qB} \quad (6.6)$$

where q is the charge of a muon and B is the magnetic field strength. The equation for the total deflection d is shown in Equations 6.9, 6.7 and 6.8.

$$d_1 = r \pm \sqrt{r^2 - s_1^2} \quad (6.7)$$

$$d_2 = s_2 \tan \left[\arcsin \left(\frac{s_1}{r} \right) \right] \quad (6.8)$$

The \pm sign in Equation 6.7 refers to the two possible intersection point of the muon crossing s_1 . A muon could travel forwards while being deflected and intersect s_1 to yield a displacement. But, if the region of the magnetic field continued, the muon's path could continue to be bent and then the muon could circle backwards and intersect s_1 again. These are the positive and negative solutions of the quadratic equation. The negative solution was used as the muon was not returning. Therefore, the total muon deflection is given by:

$$d = d_1 + d_2 = r - \sqrt{r^2 - s_1^2} + s_2 \tan \left[\arcsin \left(\frac{s_1}{r} \right) \right] \quad (6.9)$$

The detectors had a magnetic field region 0.2 m wide with a field strength of 0.8 T was used. The detector planes were 1 m apart. The deflection due to the magnetic field in the DK2NU and crosscheck detectors is shown in Figure 6.12. These parameters were chosen as a compromise with bending the path of the muons with the most probable momenta across as much of the detector as possible and trying to contain as much of the low energy portion of muons within the detector volume. The fastest a muon can travel is the speed of light c and by Equation 6.6 there therefore exists a maximum radius by which a muon can be deflected. The largest radius corresponds to the straightest path a muon can travel and therefore a minimum deflection. This effect is visible as the solid cutoff where no muon is deflected below a certain distance. (Most muons are travelling very close to the speed of light and a direction cut has already been performed so the muons are biased towards having a velocity close to c .) The deflection as a function of the muon's energy for just the DK2NU detector is shown in Figure 6.13 to illustrate the form of the curve. This figure also shows that there is a minimum possible deflection due to the muons being unable to exceed the speed of light.

Overall the results of the DK2NU detector and crosscheck detector show that the neutrinos manually produced followed the correct distribution and the implementation was successful (without needing a detector location weighting). This allows the detectors to be mobile and the simulation of a number of detectors around a

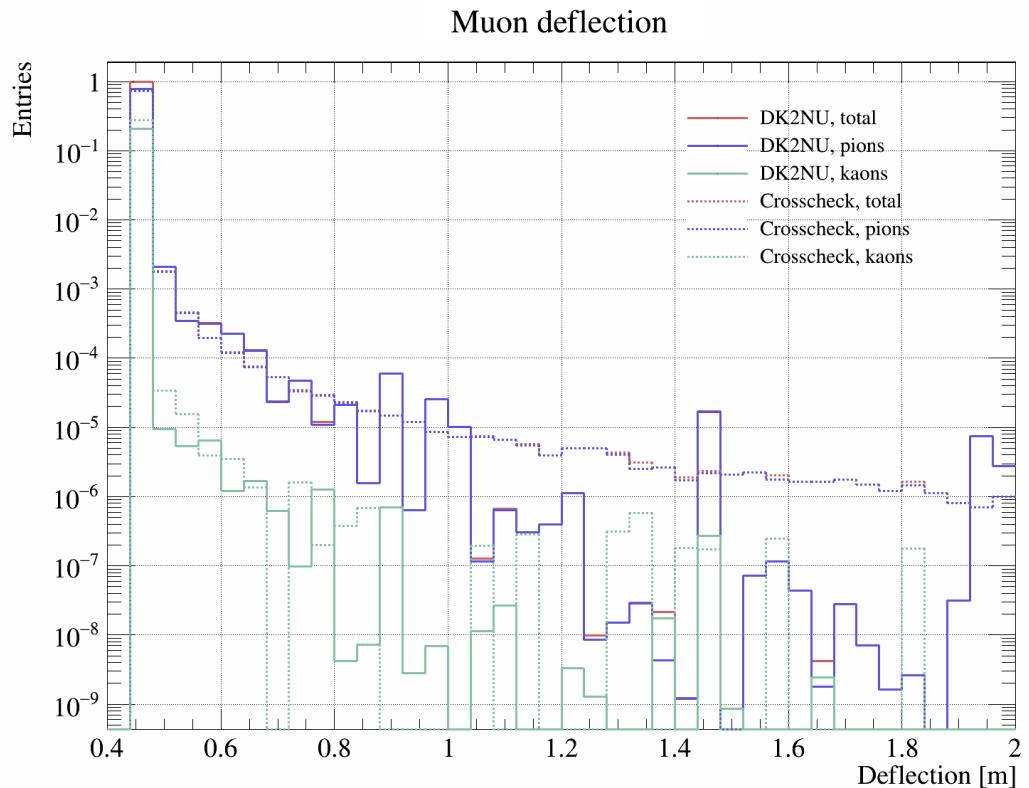


Figure 6.12: The deflection of muons due to the magnetic field. The solid line at approximately 0.45 m is due to the muons being unable to exceed the speed of light. In total there are 7261 DK2NU entries but due to the weighting the entries are spread across the entire energy range. There are 373375 crosscheck entries. The method used to generate neutrinos for the crosscheck detectors changed the weighting system. The distribution of the hits is not spread out evenly across the entire energy range. This is visible in at the high-deflection region of the plot where statistical fluctuations become significant. This is an inherent flaw in the method used to generate the neutrinos for detectors that can be moved.

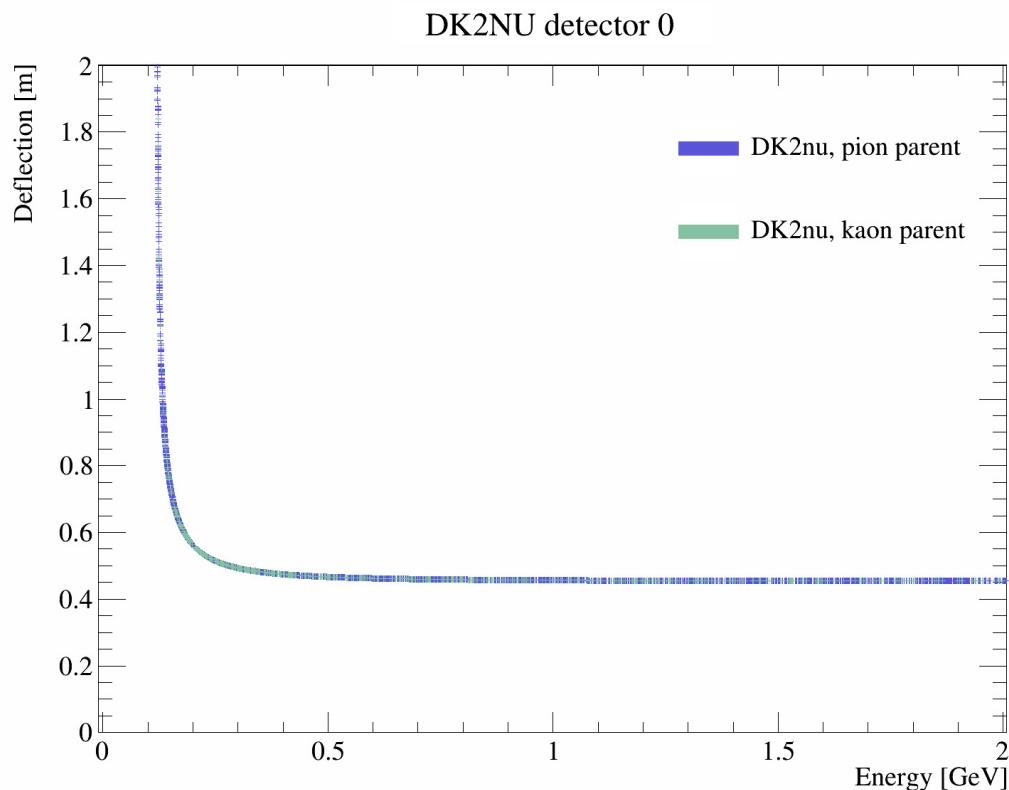


Figure 6.13: The form of the muon energy vs deflection curve. The horizontal line at approximately 0.45 m is due to the fastest muons approaching but being unable to exceed the speed of light. Muons that have a velocity that causes them to be deflected almost 90° would have a deflection approaching ∞ in the second region with no magnetic field. The detector has been restricted to 2 m.

neutrino beam to continue.

The question then remained, how large should the two detector planes be to intersect the muons? The larger the area of the detector planes, the more muons would intersect the detector and the less time required to gather a statistically significant sample of muons. To make a detector with as large a surface area as possible would be an engineering feasibility challenge, not a problem of physics. The detector plane separation however is related to the physics. The muons are not necessarily created going precisely forwards; so muons that strike the first detector plane could then leave the detector volume before arriving at the second. Figure 6.11 shows the cumulative fraction of muons produced by neutrinos with an angle between the neutrino and muon of less than θ . This shows that over 99% of muons are produced with an angle less than 0.8 rad. If a detector was 4 m wide, muons produced at the centre would have 2 m in each axis to travel before leaving the detector volume. If the detector planes were 1 m apart this would give an angle of 2 rad which would contain over 99% of the muons. Therefore this condition was used to constrain the detector dimensions.

Eight detectors of the form shown in Figure 6.3 were simulated which were 1 km downbeam and square in shape with 4 m edges. Each detector consisted of two parallel planes with 1 m of space between them. A 1.5 T magnetic field occupied the first 0.2 m of space between the two planes with the field direction perpendicular to the direction of the incoming muons. The detectors were arranged from 15 mrad off-axis to 50 mrad off-axis in 5 mrad increments (15 m to 50 m off-axis in 5 m steps).

The neutrino flux spectra through each detector are shown in Figure 6.14. The spectra show that as the off-axis angle increases, the neutrino peak energy reduces but also narrows (see Appendix A for the off-axis effect). The spectra also show that as the off-axis angle increases, the high energy tail becomes more and more dominated by the kaon component of the beam.

The muon energy spectra can be seen in Figure 6.15 which shows that, following on from the neutrino off-axis effect, the muon energy peak lowers and narrows as the off-axis angle increases.

The effective mass of each detector can be seen in Figure 6.16. The result shows the contribution to the largest effective detector sizes arises due to muons from kaons and also becomes more significant at larger off-axis angles. This effect is due to the kaons dominating the high energies at the off-axis positions.

The deflection for the muons in each detector can be seen in Figure 6.17 and the deflection normalised by the total number of events can be seen in Figure 6.18.

Initially, at smaller off-axis angles (starting at 15 mrad), as the deflection increases the number of muons derived from kaons declines faster than the number of muons derived from pions. As the off-axis angle increases, the decline rate of the kaon-produced-muon curve decreases (up to 50 mrad) (the number of muons

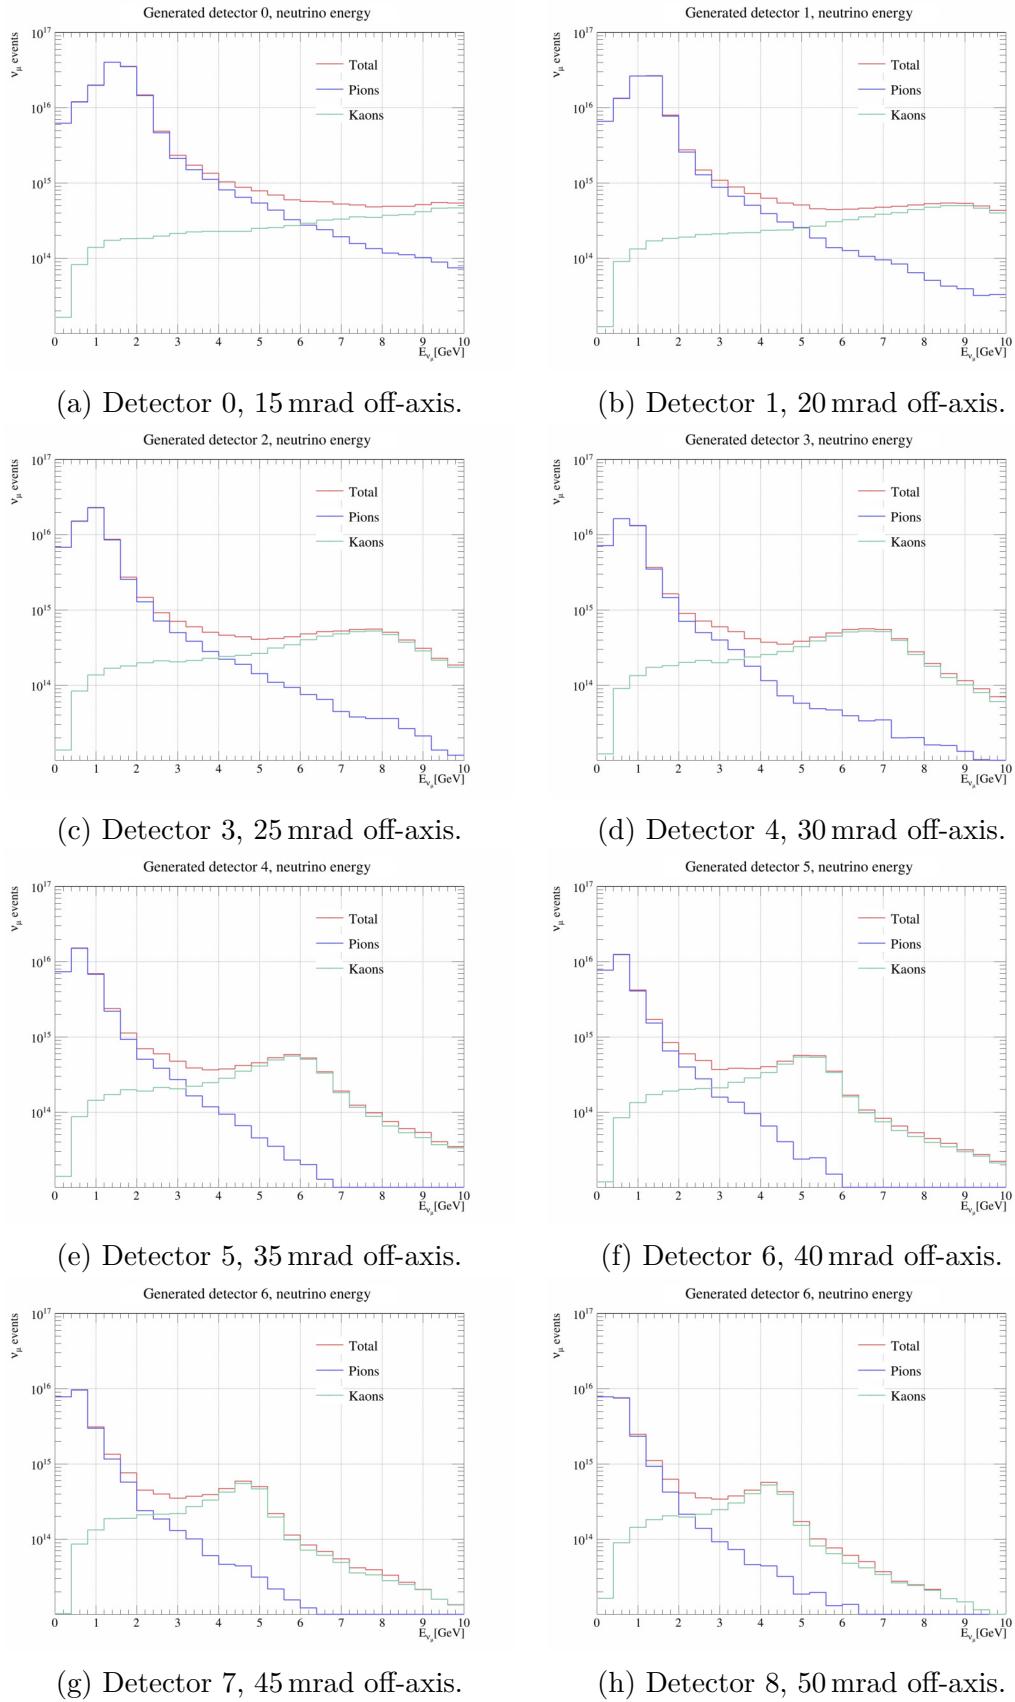


Figure 6.14: Neutrino flux energy spectra through each detector.

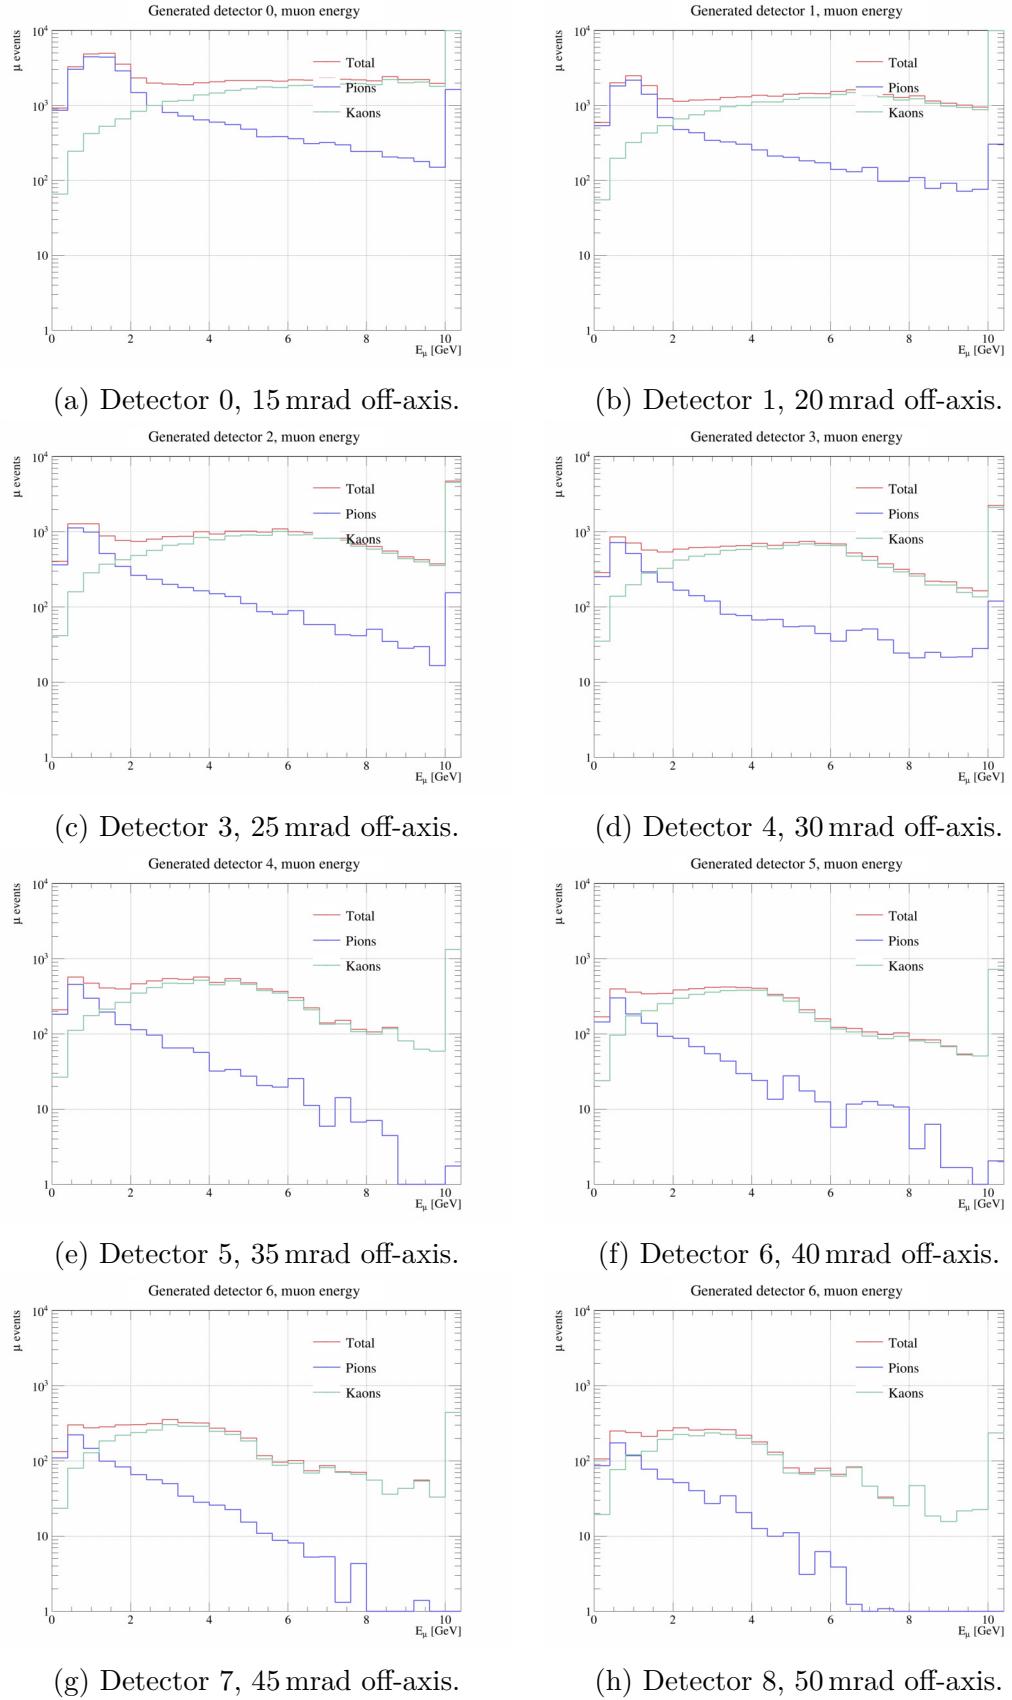


Figure 6.15: Muon flux energy spectra through each detector. Here, the muons have lost no energy in the rock. The counting correction which applies the distance a muon of a given energy can actually travel through rock is included.

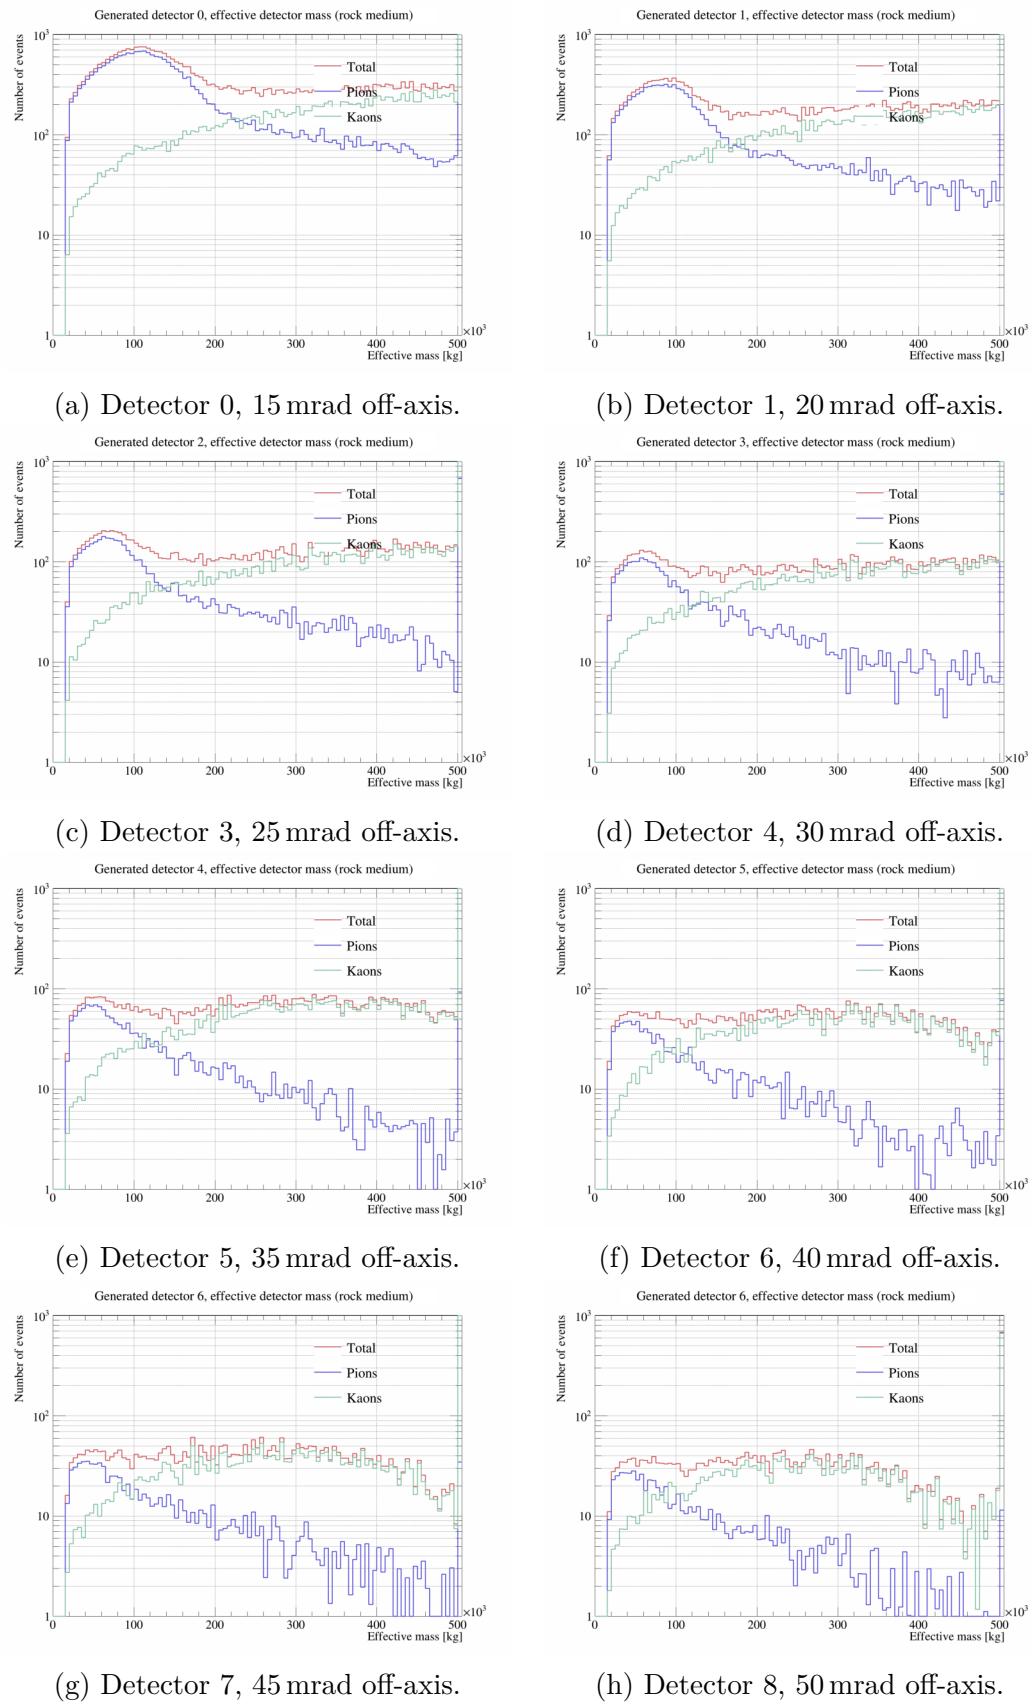


Figure 6.16: The effective mass of each detector.

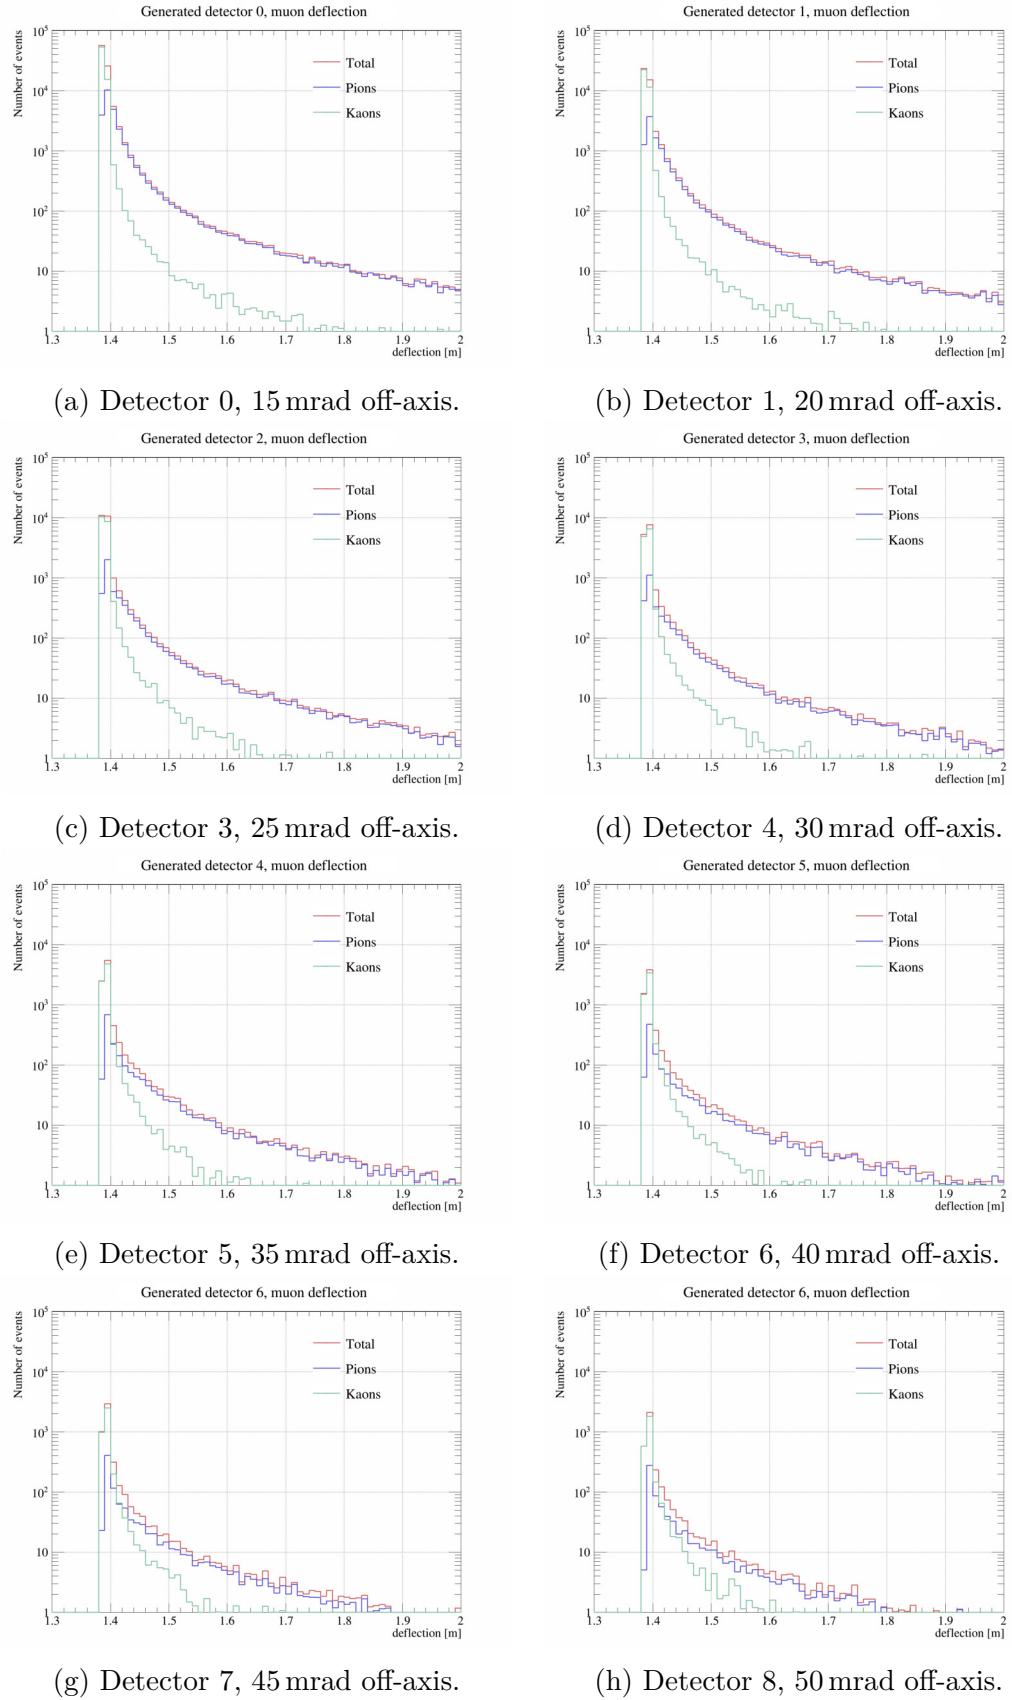


Figure 6.17: The deflection of each muon through each detector.

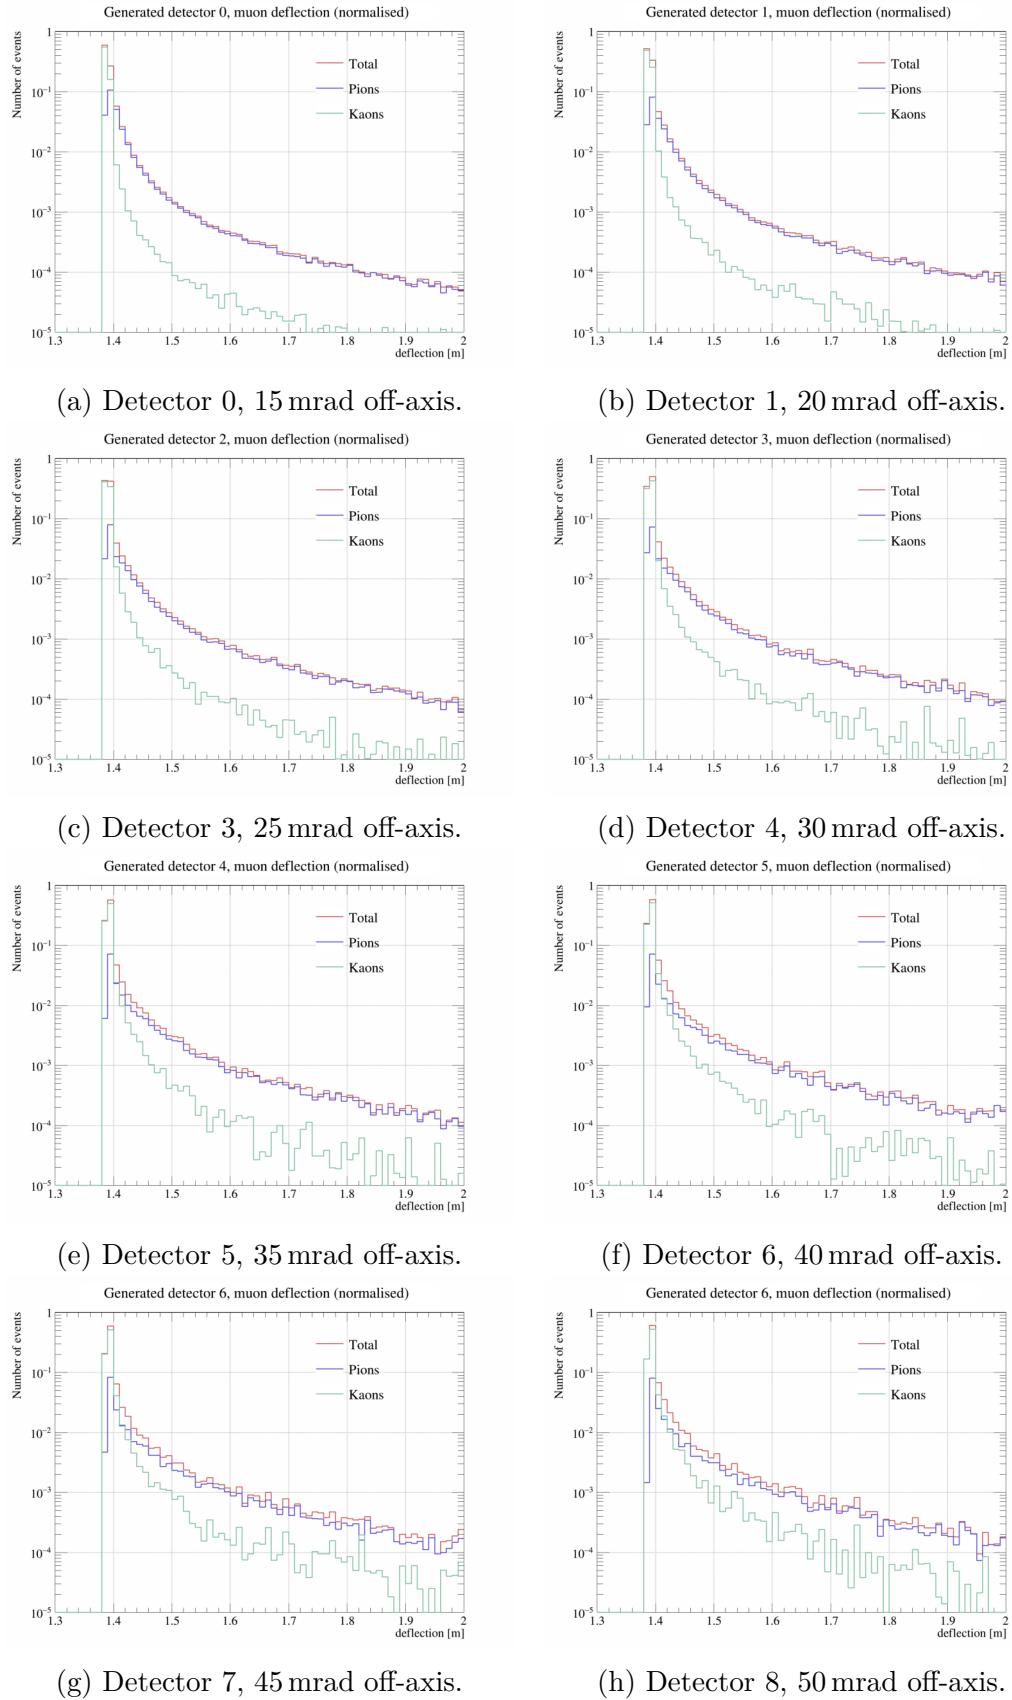


Figure 6.18: The deflection of each muon through each detector, normalised by the total number of muons.

Detector	Angle [mrad]	Muon Count			
		#Muons (total)	#Muons (π)	Muons (K)	Muon ratio (K/π)
0	15	96217	26573	69643	2.6
1	20	45486	10741	34744	3.2
2	25	25357	5444	19913	3.7
3	30	15218	3218	11999	3.7
4	35	9597	1819	7777	4.3
5	40	6631	1263	5368	4.3
6	45	4879	969	3909	4.0
7	50	3461	708	2753	3.9

Table 6.1: The total number of muons, the number of muons from pion ancestors and the number of muons from kaon ancestors at each detector location. The number of muons is per 10^{20} POTS (equivalent to approximately 150 days worth of Numi beam operation).

is more sustained, the magnitude of the gradients smaller) and the decline rate of the pion-produced-muon curve increases (the number of muons decreases faster, the magnitude of the gradient is larger). The effect of this is that the number of muons derived from kaons is more sustained as the deflection increases compared with muons derived from pions (and the effect is more pronounced as the off-axis angle increases). Additionally, the point at which the fraction of kaon-produced-muons and pion-produced-muons is equal increases in deflection position, although this effect is very small, a 2 bin (4 cm) change over the entire 30 mrad angle range.

The number of pions and kaons could be inferred by fitting two superimposed curves to the total muon deflection for each detector location. The form of the curves would be dependent on the original fractions of pions and kaons and the off-axis angle. By fitting for a range of angles, the kaon content could be estimated and used to tune Monte-Carlo simulations to more accurately model neutrino beams.

The number of muons that were deflected through the detector at each off-axis location for 10^{20} POTS can be seen in table 6.1. The table shows that between the smallest and largest off-axis angles (15 mrad to 50 mrad), the number of muons decreases by approximately a factor of 30 from tens-of-thousands to just thousands which is a change enough to be measurable. Also, the ratio of pions to kaons changes by a factor of 1.5 over such angle changes.

Overall, the simulation shows that resolving the differences in the pion and kaon content of the hadron ancestors of a neutrino beam at different off-axis angles could be achieved using magnetic field strengths that exist at current technology levels.

Chapter 7

Summary & Conclusions

This thesis presented the design and construction of the CHIPS detector. The goal of the project was to demonstrate that a neutrino detector could be built by a team of scientists themselves rather than by a dedicated team of engineers and construction workers. Also, the goal was to demonstrate that the cost of the detector could be substantially cheaper than the cost of existing detectors. Despite engineering challenges, the detector was successfully constructed by a team of physicists and students. The detector was deployed, although time and circumstances did not allow a full campaign of light-sealing and commissioning. Unfortunately, the detector had to be extracted and salvaged before full commissioning could take place due to the complications arising during the deployment process exacerbated by the COVID-19 pandemic. However, the initial goals of the project (stated above) were achieved and much was learned for improvements in the design of low-cost neutrino detectors. The extended goals of the project – to actually perform a measurement using the detector (for example, θ_{23}) – were not met and cannot be readdressed in the future as the project has concluded.

The CHIPS experiment was not designed to surpass the current level of precision measurements. However, the CHIPS experiment demonstrated the ability to deploy large detectors cheaply in a beamline that could have made a measurement of θ_{23} and δ_{CP} . The obvious impact of this work is to show that experiments in the future could be very large or have multiple detectors without being prohibitively costly. Future experiments could then be used for precision measurements. Overall, the field of particle physics could be advanced faster by utilising the example set by CHIPS, even though CHIPS itself did not directly produce any frontier-pushing measurements.

The construction of the CHIPS detector revealed that although the readout and DAQ electronics functioned as intended, there were cumbersome infrastructure requirements. A new set of electronics was designed to connect to the Microdaqs mounted onto the PMTs to manage the power, data and clock distribution and replace the cumbersome infrastructure of the detector. The new PMT readout and DAQ

systems presented here were designed and assembled using lessons learned and were successfully implemented. A modular decentralised readout system is being tested. Further development is ongoing to investigate using the DAQ and readout electronics to detect cosmic ray air showers. The electronics themselves continue to be tested in a lab setting at three disconnected locations around the world: the UK, the Czech Republic and Vietnam.

A technique for the analysis of muon fluxes in rocks to measure the kaon content of a neutrino beam has been investigated. The technique proposed using the off-axis effect and magnetic deflection to infer the number of parent pions and kaons in a neutrino beam. The study verified that the high-energy tail of the muon and of the neutrino distributions becomes more closely related to the mass of parent hadron as the off-axis angle increased rather than the energy of the hadron itself. Furthermore, the study showed that the deflection could be induced and observed at a scale that current detector technologies already exist at and in numbers large enough to count. The pre-existing rock of the Earth could be used as a detector medium making the monitoring of existing and future neutrino beams more attainable. The simulation assumed that the density of rock in the Earth was uniform and that the muons do not lose energy in the rock between the point they are produced and the point they enter the detector. These simplifications were made to extract the core physics of the problem but are areas that need to be improved and accounted for in future work. Also, the study did not perform a fit to estimate the kaon content of the neutrino beam used in the simulation – this was because of time constraints. Further work is required to ascertain if the kaon content could actually be inferred and at what precision level. If the kaon content could be inferred by measuring rock muon fluxes at different off-axis angles, then this technique could be used to monitor future neutrino beams.

Appendix A

Off-axis Neutrino Beams

For this section the speed of light, c , is taken to be 1 and the mass of the neutrino is approximated to 0.

Neutrino beams produce a broad spectrum of neutrinos. Due to kinematic effects of hadron decay, then moving off axis, the neutrino beam energy spectrum narrows and the peak drops in energy. This is known as the *off-axis effect*, the derivation of which is below.

A π^+ decays into a μ^+ and a ν_μ (shown in Figure A.1). In the rest frame of the pion, the 4-momenta are:

$$p_{\pi^+} = (E_{\pi^+}, \vec{0}) \quad (\text{A.1})$$

$$p_{\mu^+} = (E_{\mu^+}, \vec{P}_{\mu^+}) \quad (\text{A.2})$$

$$p_{\nu_\mu} = (E_{\nu_\mu}, \vec{P}_{\nu_\mu}) \quad (\text{A.3})$$

and therefore:

$$p_{\pi^+}^2 = m_{\pi^+}^2 \quad (\text{A.4})$$

$$p_{\mu^+}^2 = m_{\mu^+}^2 \quad (\text{A.5})$$

$$p_{\nu_\mu}^2 = 0 \quad (\text{A.6})$$

$$p_{\pi^+} p_{\nu_\mu} = E_{\pi^+} E_{\nu_\mu} \quad (\text{A.7})$$

Using the identity

$$p_{\mu^+} = p_{\pi^+} - p_{\nu_\mu} \quad (\text{A.8})$$

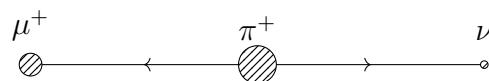


Figure A.1: A pion decay from the pion rest frame. The muon and neutrino have equal and opposite momenta.

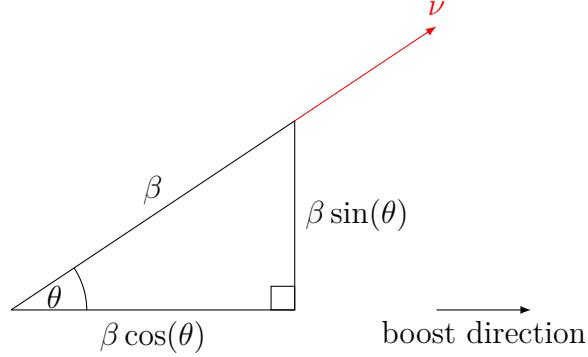


Figure A.2: Diagram showing the component of the neutrino's velocity which receives a Lorentz boost.

gives the momentum of the muon as

$$p_{\mu^+}^2 = p_{\pi^+}^2 + p_{\nu_\mu}^2 - 2p_{\pi^+}p_{\nu_\mu}^2 \quad (\text{A.9})$$

which is equivalent to

$$m_{\mu^+}^2 = m_{\pi^+}^2 - 2E_{\pi^+}E_{\nu_{\mu^+}} \quad (\text{A.10})$$

This gives, in the rest frame of the pion, the energy of the neutrino

$$E_{\nu_{\mu^+}}^{\text{Rest}} = \frac{m_{\pi^+}^2 - m_{\mu^+}^2}{2E_{\pi^+}} = \frac{m_{\pi^+}^2 - m_{\mu^+}^2}{2m_{\pi^+}} = 29.6 \text{ MeV} \quad (\text{A.11})$$

To boost the energy of the neutrino into the lab frame, the following equation is used

$$E_{\nu_{\mu^+}}^{\text{Rest}} = \gamma E^{\text{Lab}}[1 - \beta] \quad (\text{A.12})$$

which gives the energy as

$$E_{\nu_{\mu^+}}^{\text{Lab}} = \frac{E_{\nu_{\mu^+}}^{\text{Rest}}}{\gamma(1 - \beta)} = \frac{m_{\pi^+}^2 - m_{\mu^+}^2}{2m_{\pi^+}\gamma(1 - \beta)} \quad (\text{A.13})$$

$$(\text{A.14})$$

This assumes the neutrino is emitted in the same direction as the Lorentz boost. In reality, the neutrino can be emitted at an angle. The component of the neutrino's velocity in the direction of the Lorentz boost is given by the cosine of the emission angle θ . This is shown in Figure A.2. This has the effect of mapping β to $\beta \cos(\theta)$.

$$E_{\nu_{\mu^+}}^{\text{Lab}} = \frac{m_{\pi^+}^2 - m_{\mu^+}^2}{2m_{\pi^+}\gamma[1 - \beta \cos(\theta)]} \quad (\text{A.15})$$

To convert the energy of the neutrino in the lab frame into a convenient form

the following identities are used

$$E = \gamma m \quad (\text{A.16})$$

$$(\text{A.17})$$

$$P = \gamma \beta m \quad (\text{A.18})$$

which combine to give

$$\beta = \frac{P}{E} \quad (\text{A.19})$$

The angular Lorentz boost $\gamma[1 - \beta \cos(\theta)]$ becomes $\frac{1}{m_{\pi^+}} [E_{\pi^+} - P \cos(\theta)]$ which gives the neutrino energy in the lab frame as

$$E_{\nu_{\mu^+}}^{\text{Lab}} = \frac{m_{\pi^+}^2 - m_{\mu^+}^2}{2[E_{\pi^+} - P_{\pi^+} \cos(\theta)]} \quad (\text{A.20})$$

and using the mass-energy-momentum equivalence this becomes

$$E_{\nu_{\mu^+}}^{\text{Lab}} = \frac{m_{\pi^+}^2 - m_{\mu^+}^2}{2[E_{\pi^+} - \sqrt{E_{\pi^+}^2 - m_{\pi^+}^2} \cos(\theta)]} \quad (\text{A.21})$$

This equation shows the relationship between the energy of the neutrino, the energy of the pion and the off-axis angle. As θ increases, the energy decreases as will the neutrino flux. However, the width of the neutrino spectrum also decreases as it is less strongly dependent on the energy of the pion. This can be seen in Figure A.3. By reducing the widths of the tails of the neutrino distribution, the number of background events can be reduced. By using an off-axis detector, the beam can be tuned so that a very narrow peak at the oscillation maxima exists at the detector's position.

Furthermore, since the energy of the neutrino depends directly on the mass of the parent hadron and then more weakly on the energy of the parent hadron, changing the parent hadron type and therefore its mass more directly changes the neutrino energy when significantly off axis. This means that off-axis, kaons rather than pions contribute more to the long high energy tail of the neutrino energy distribution.

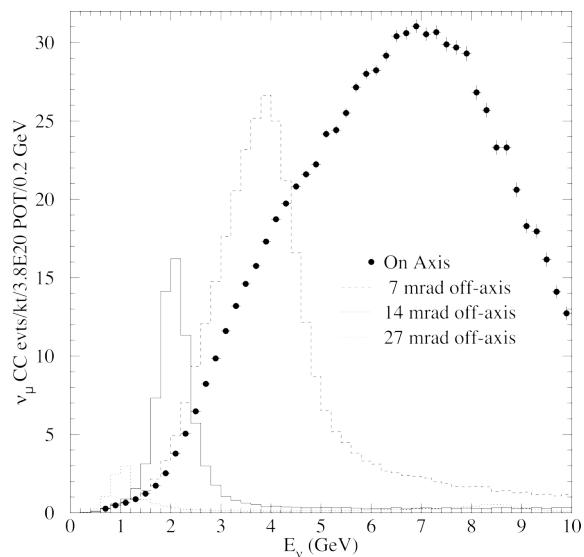
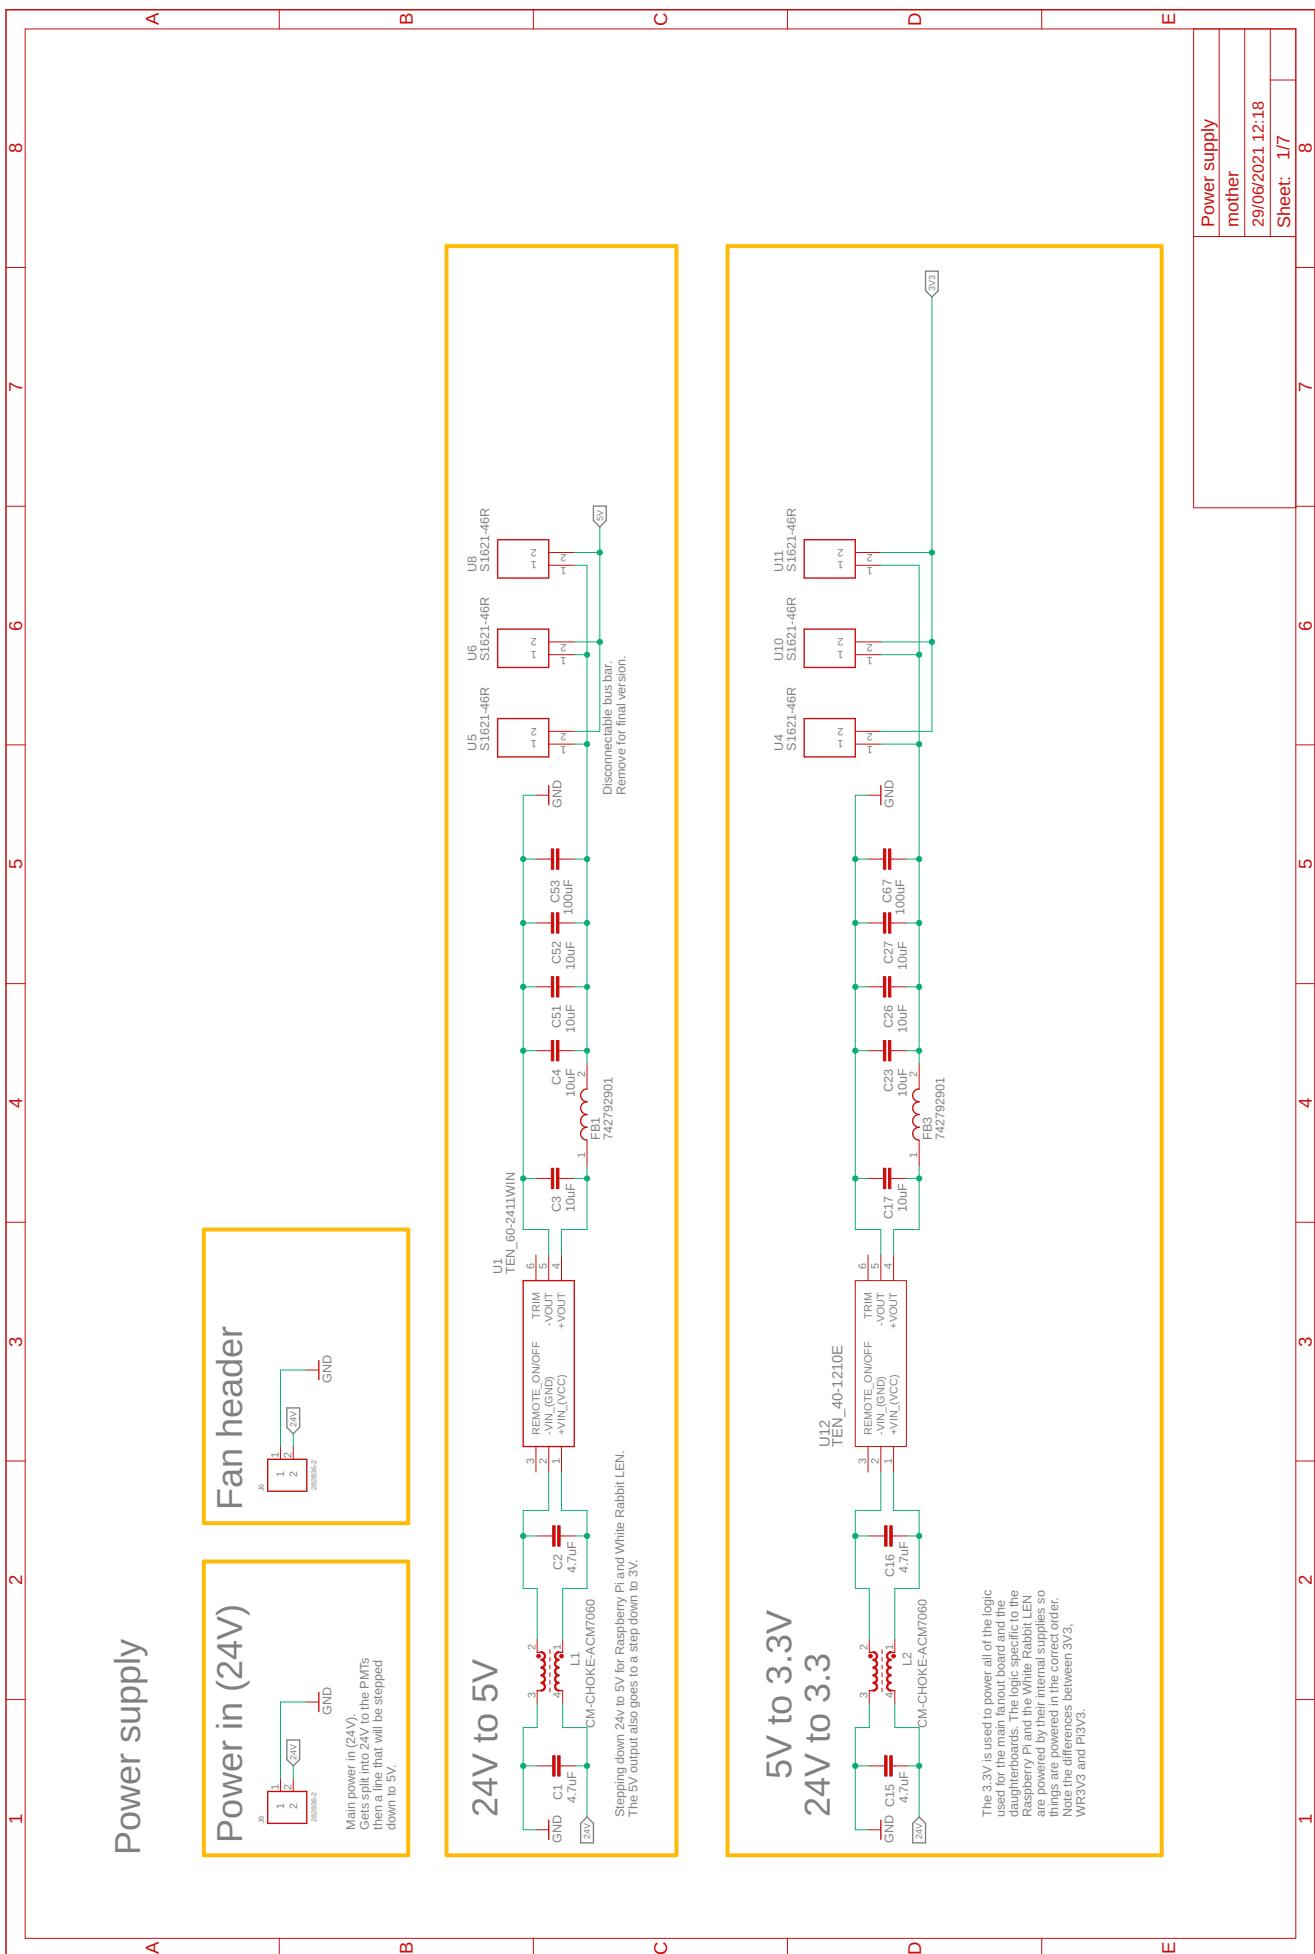


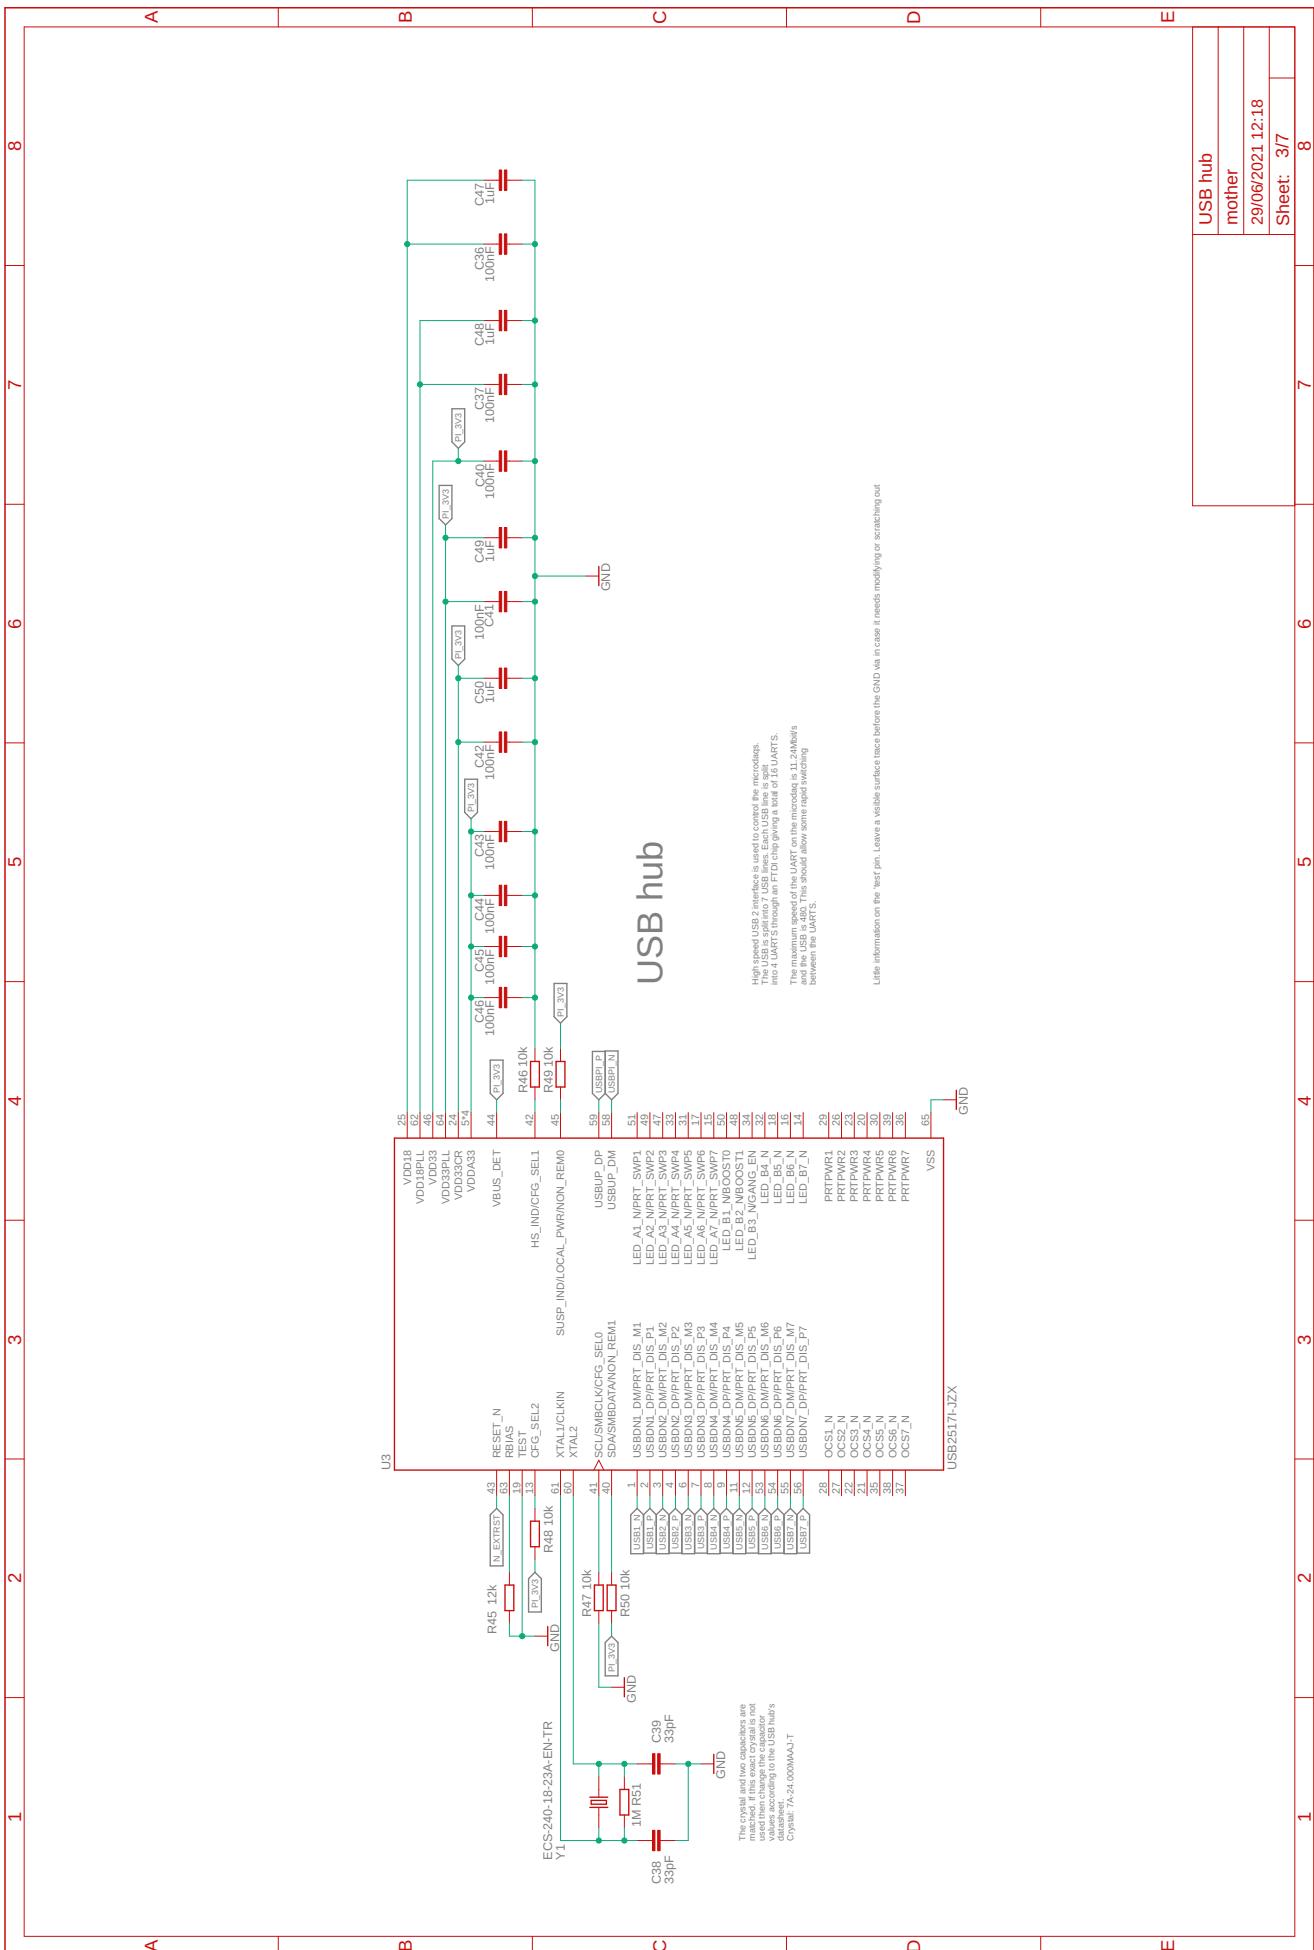
Figure A.3: Simulated energy spectra of charged current events at a far detector location 735 km from Fermilab at various off-axis angles for the NuMI beam at the medium-energy setting [104]. A similar set of spectra with separate pion and kaon contributions at different off-axis angles can be seen in figure 6.14.

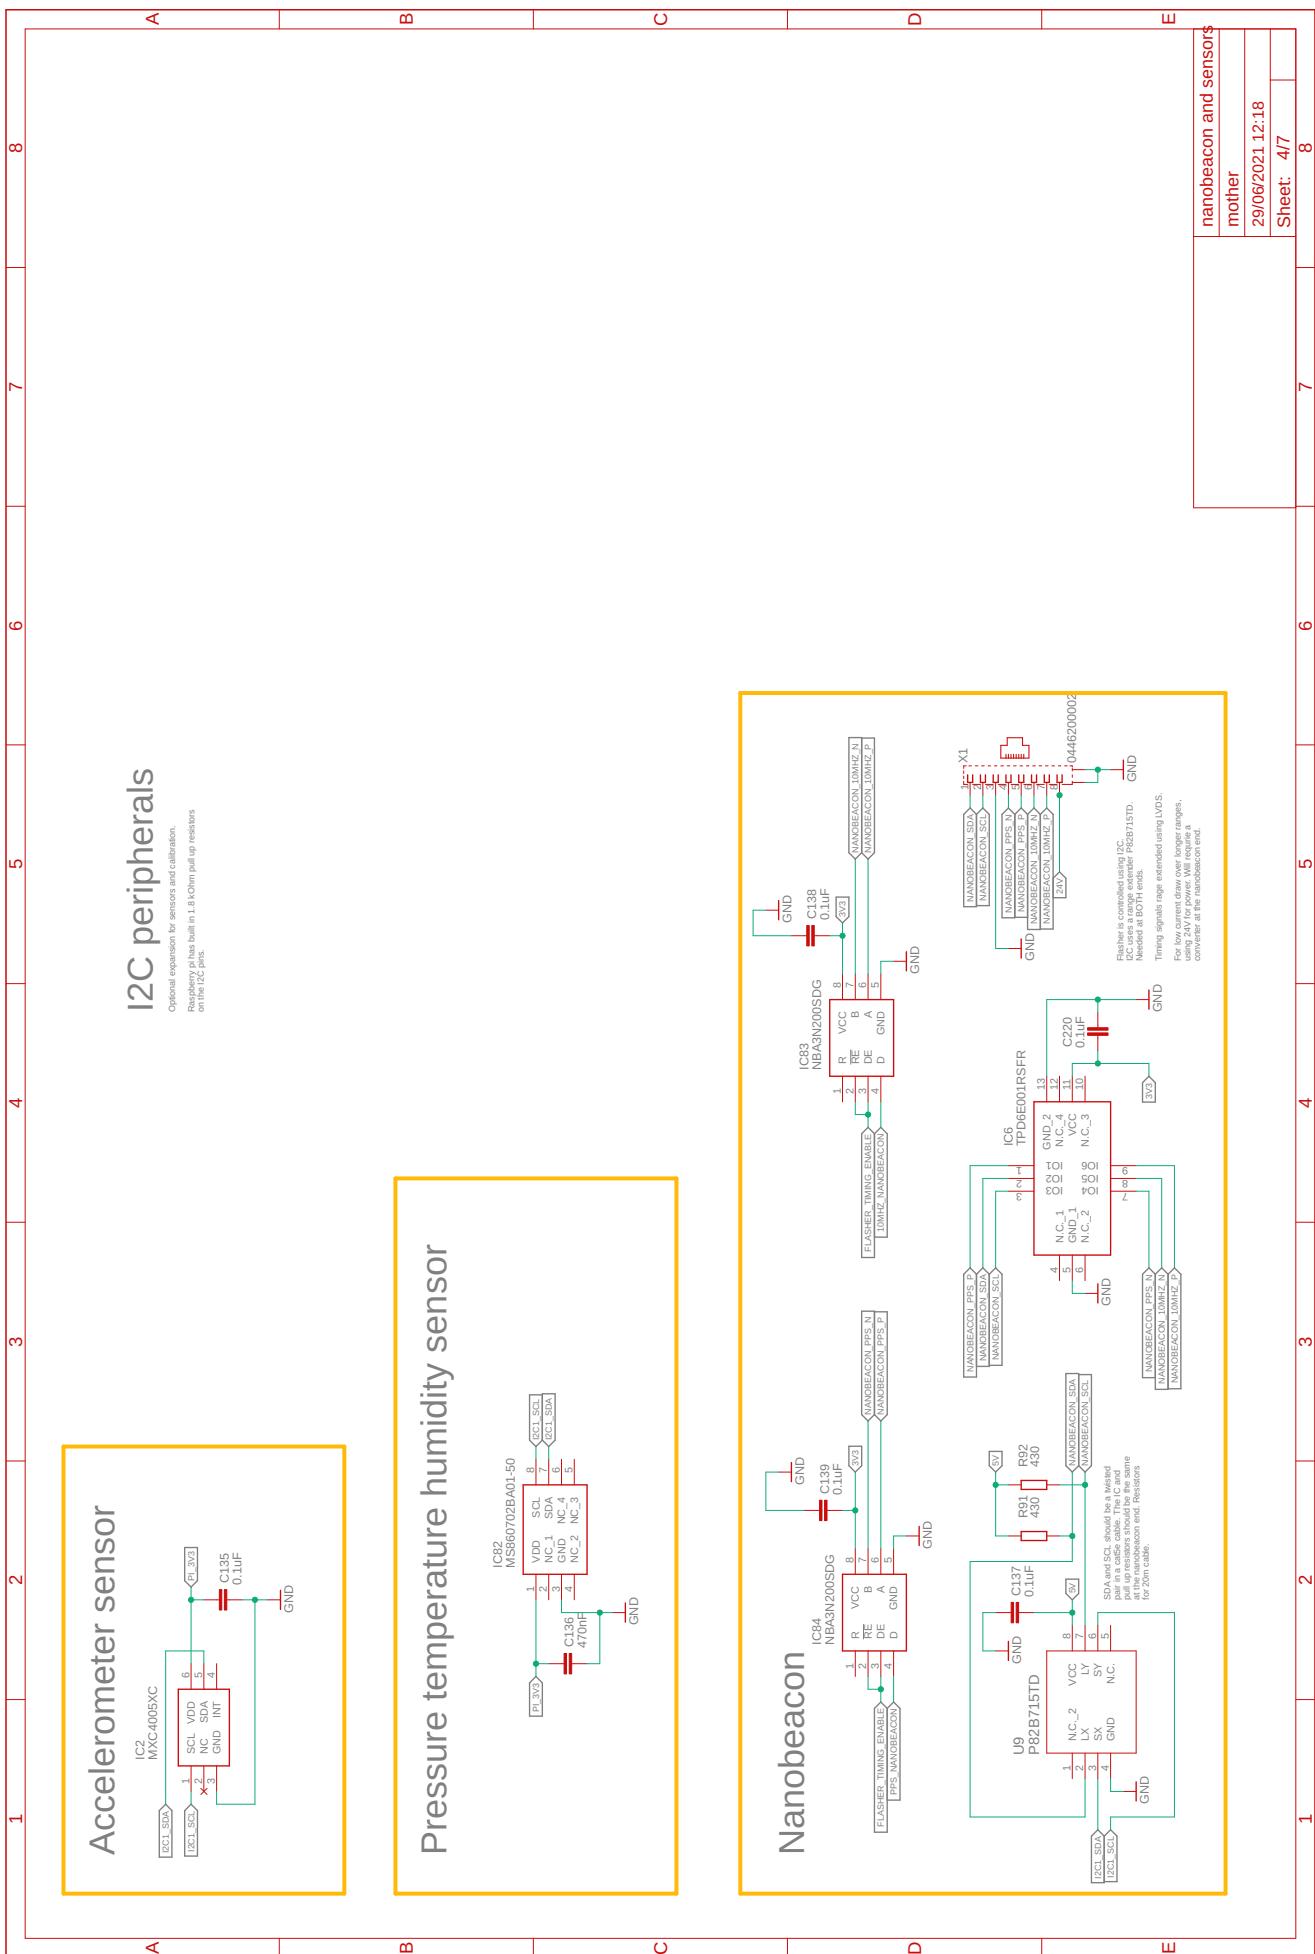
Appendix B

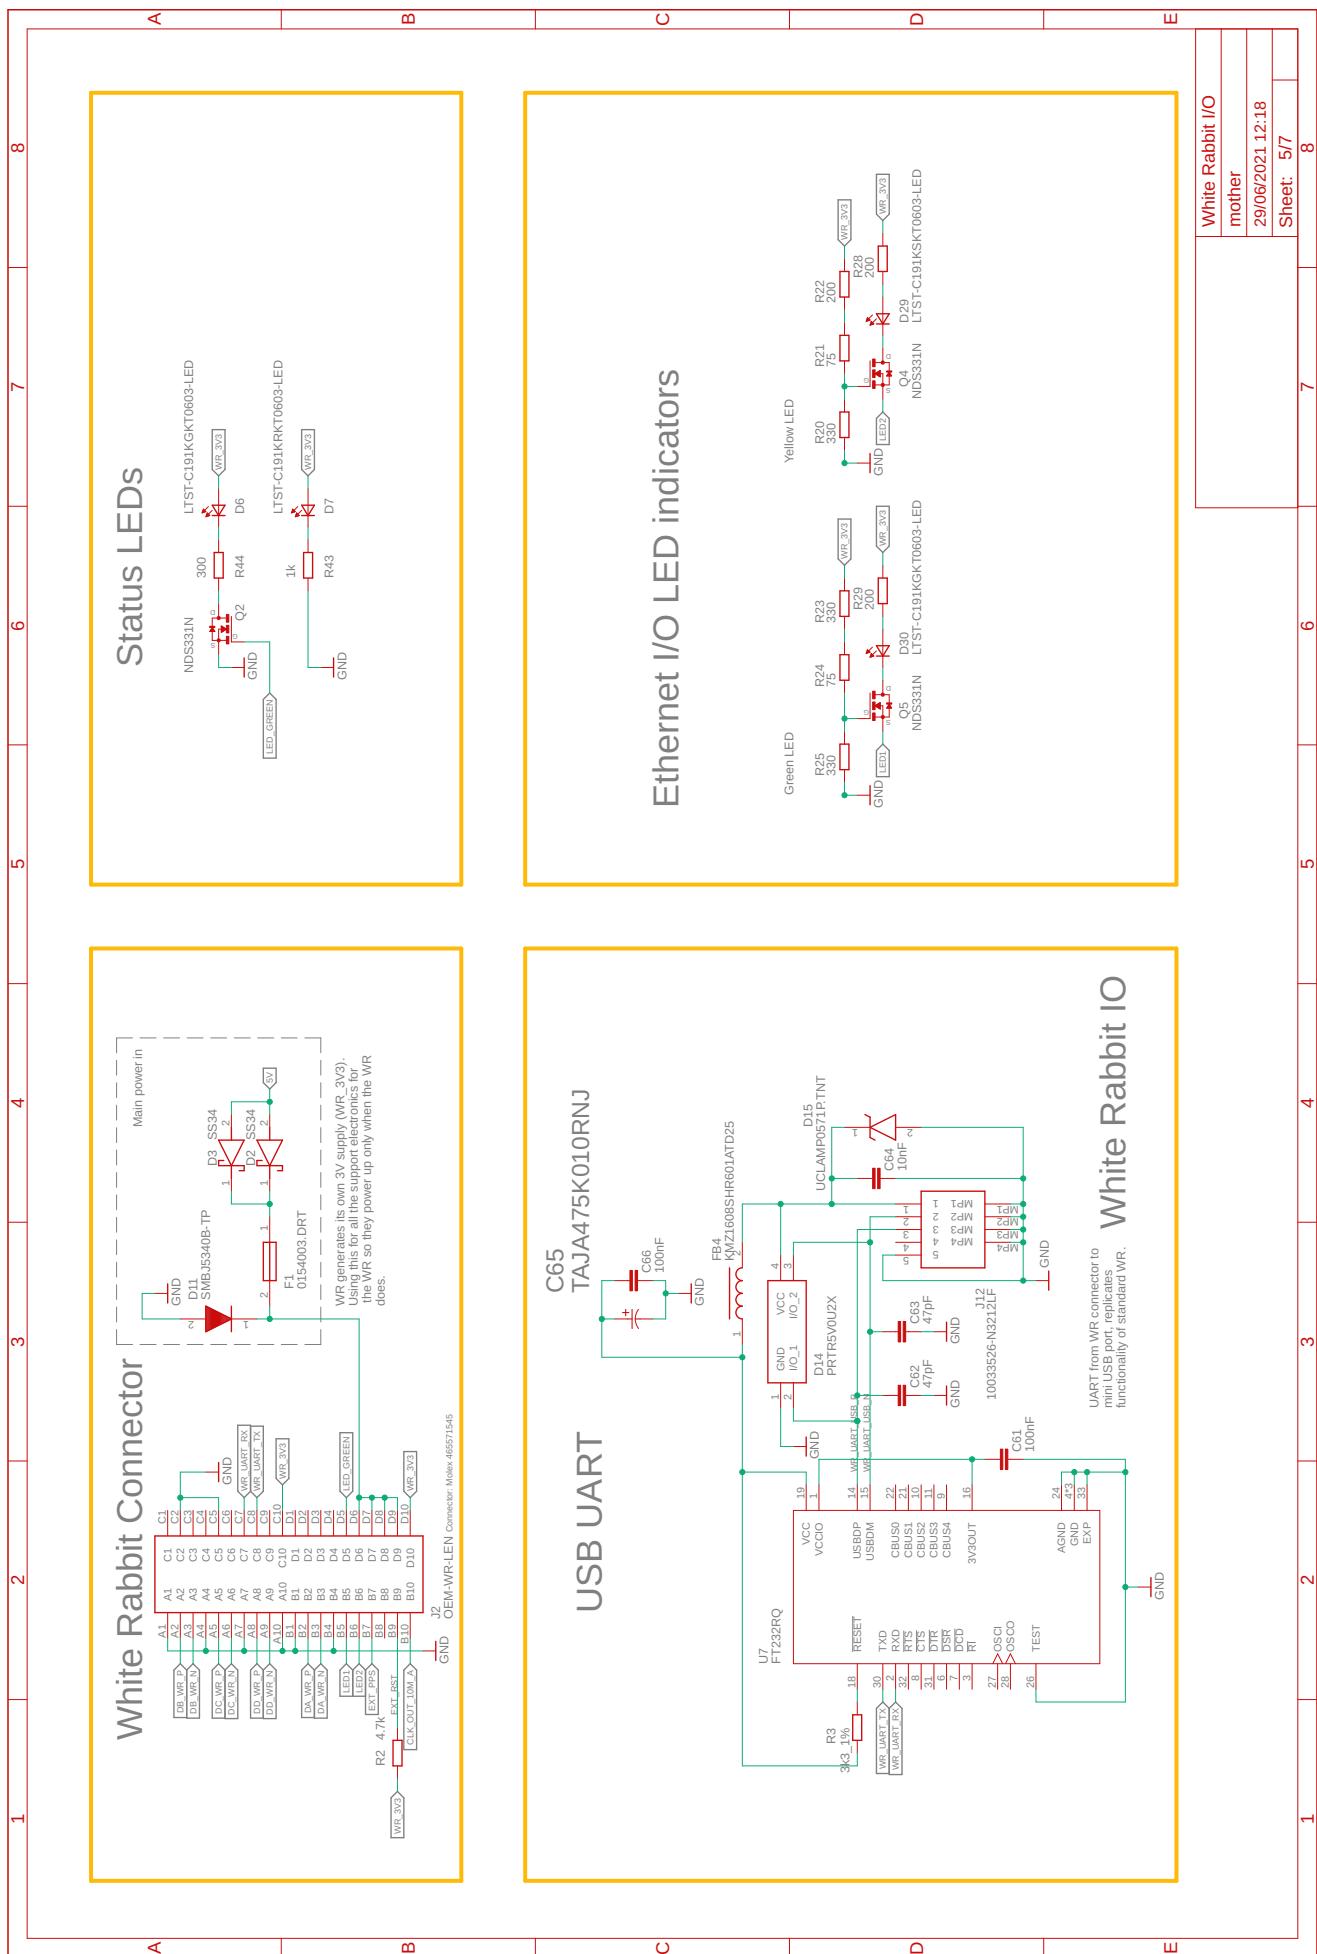
New Fanout Circuit Diagrams

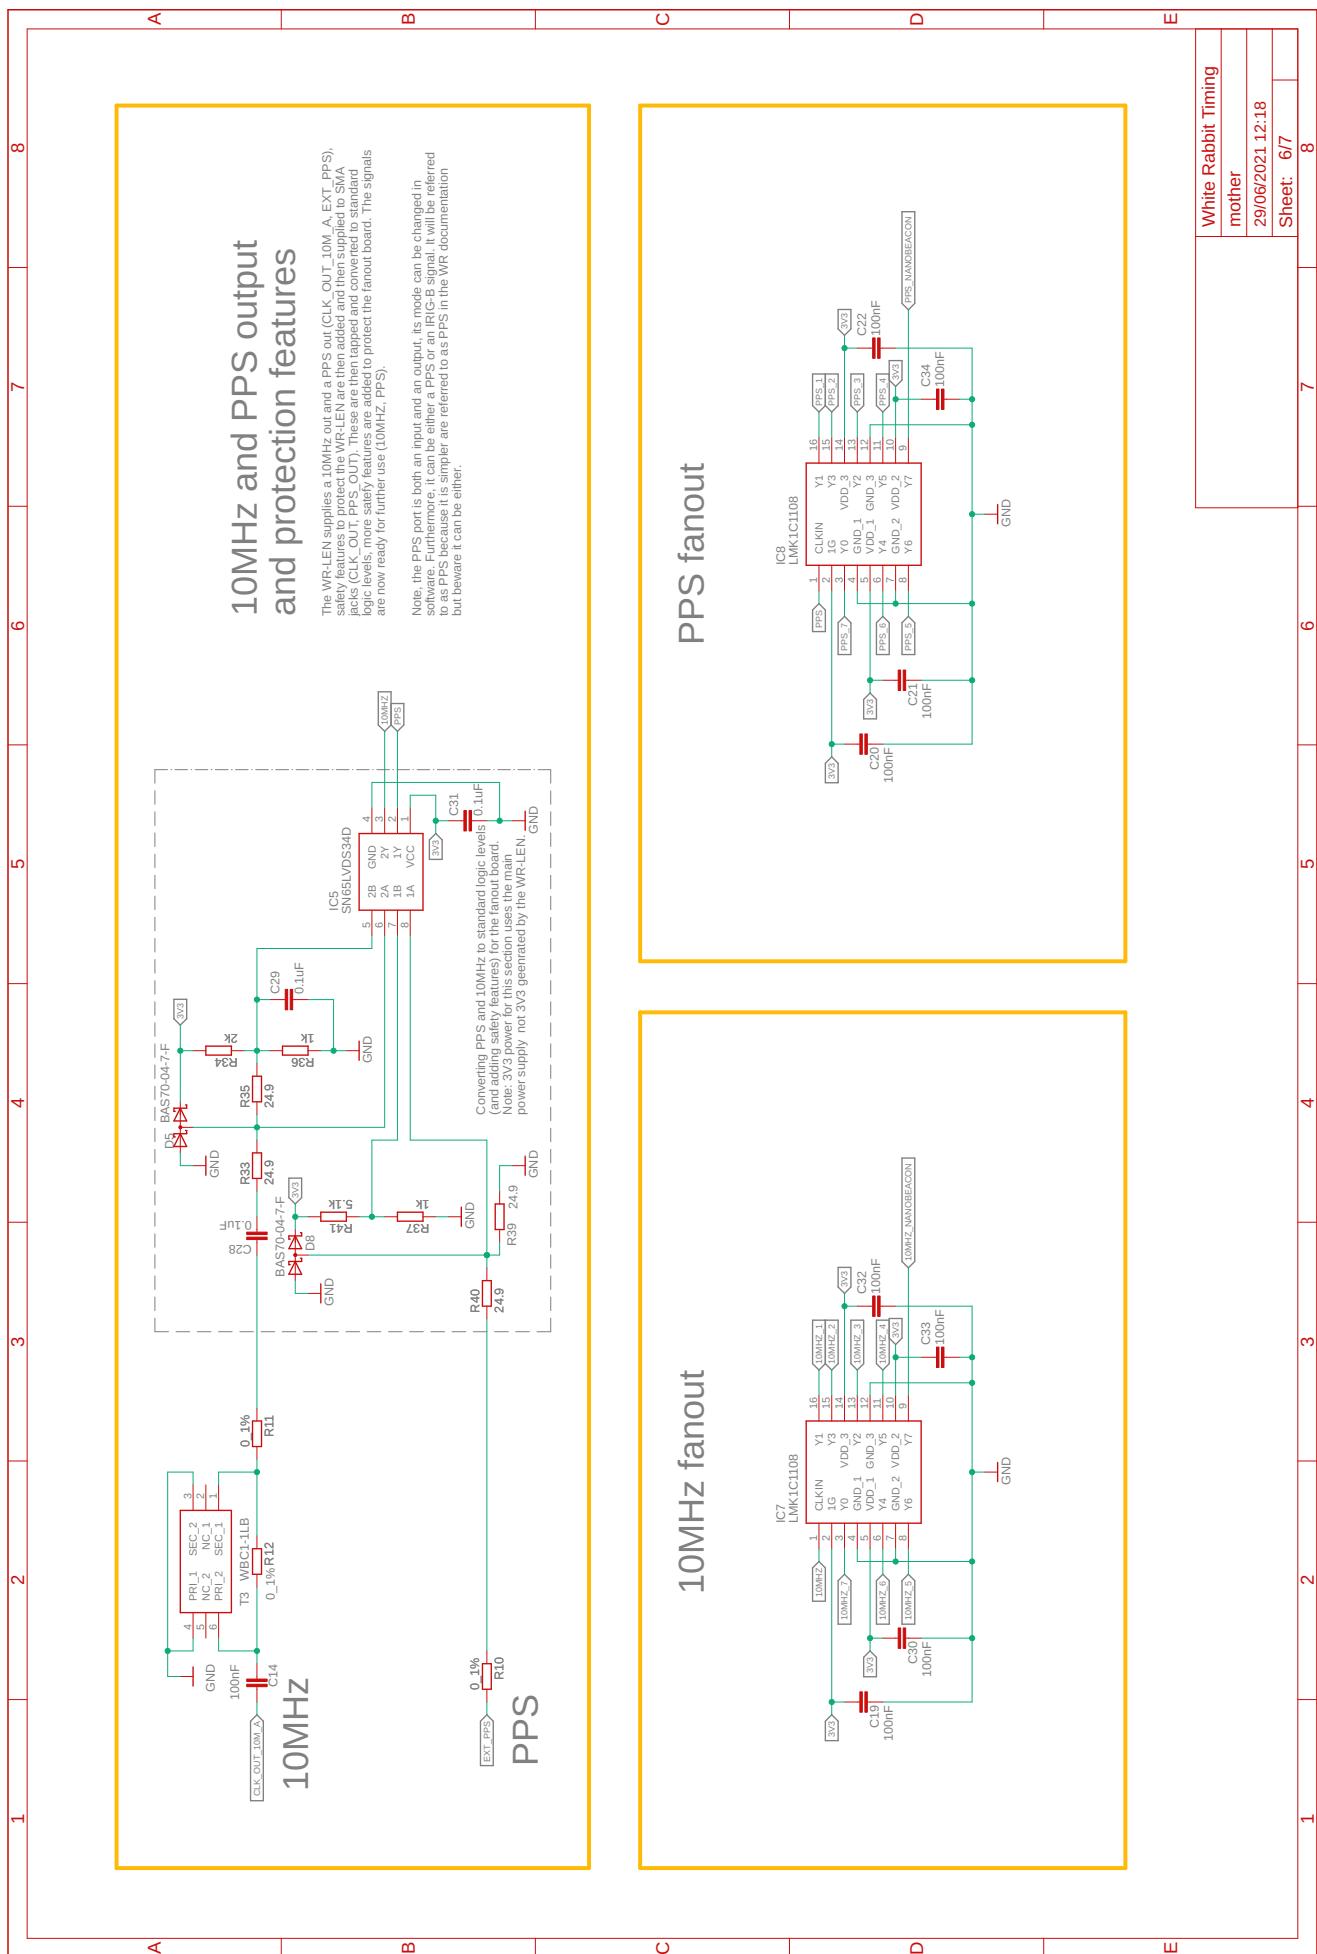


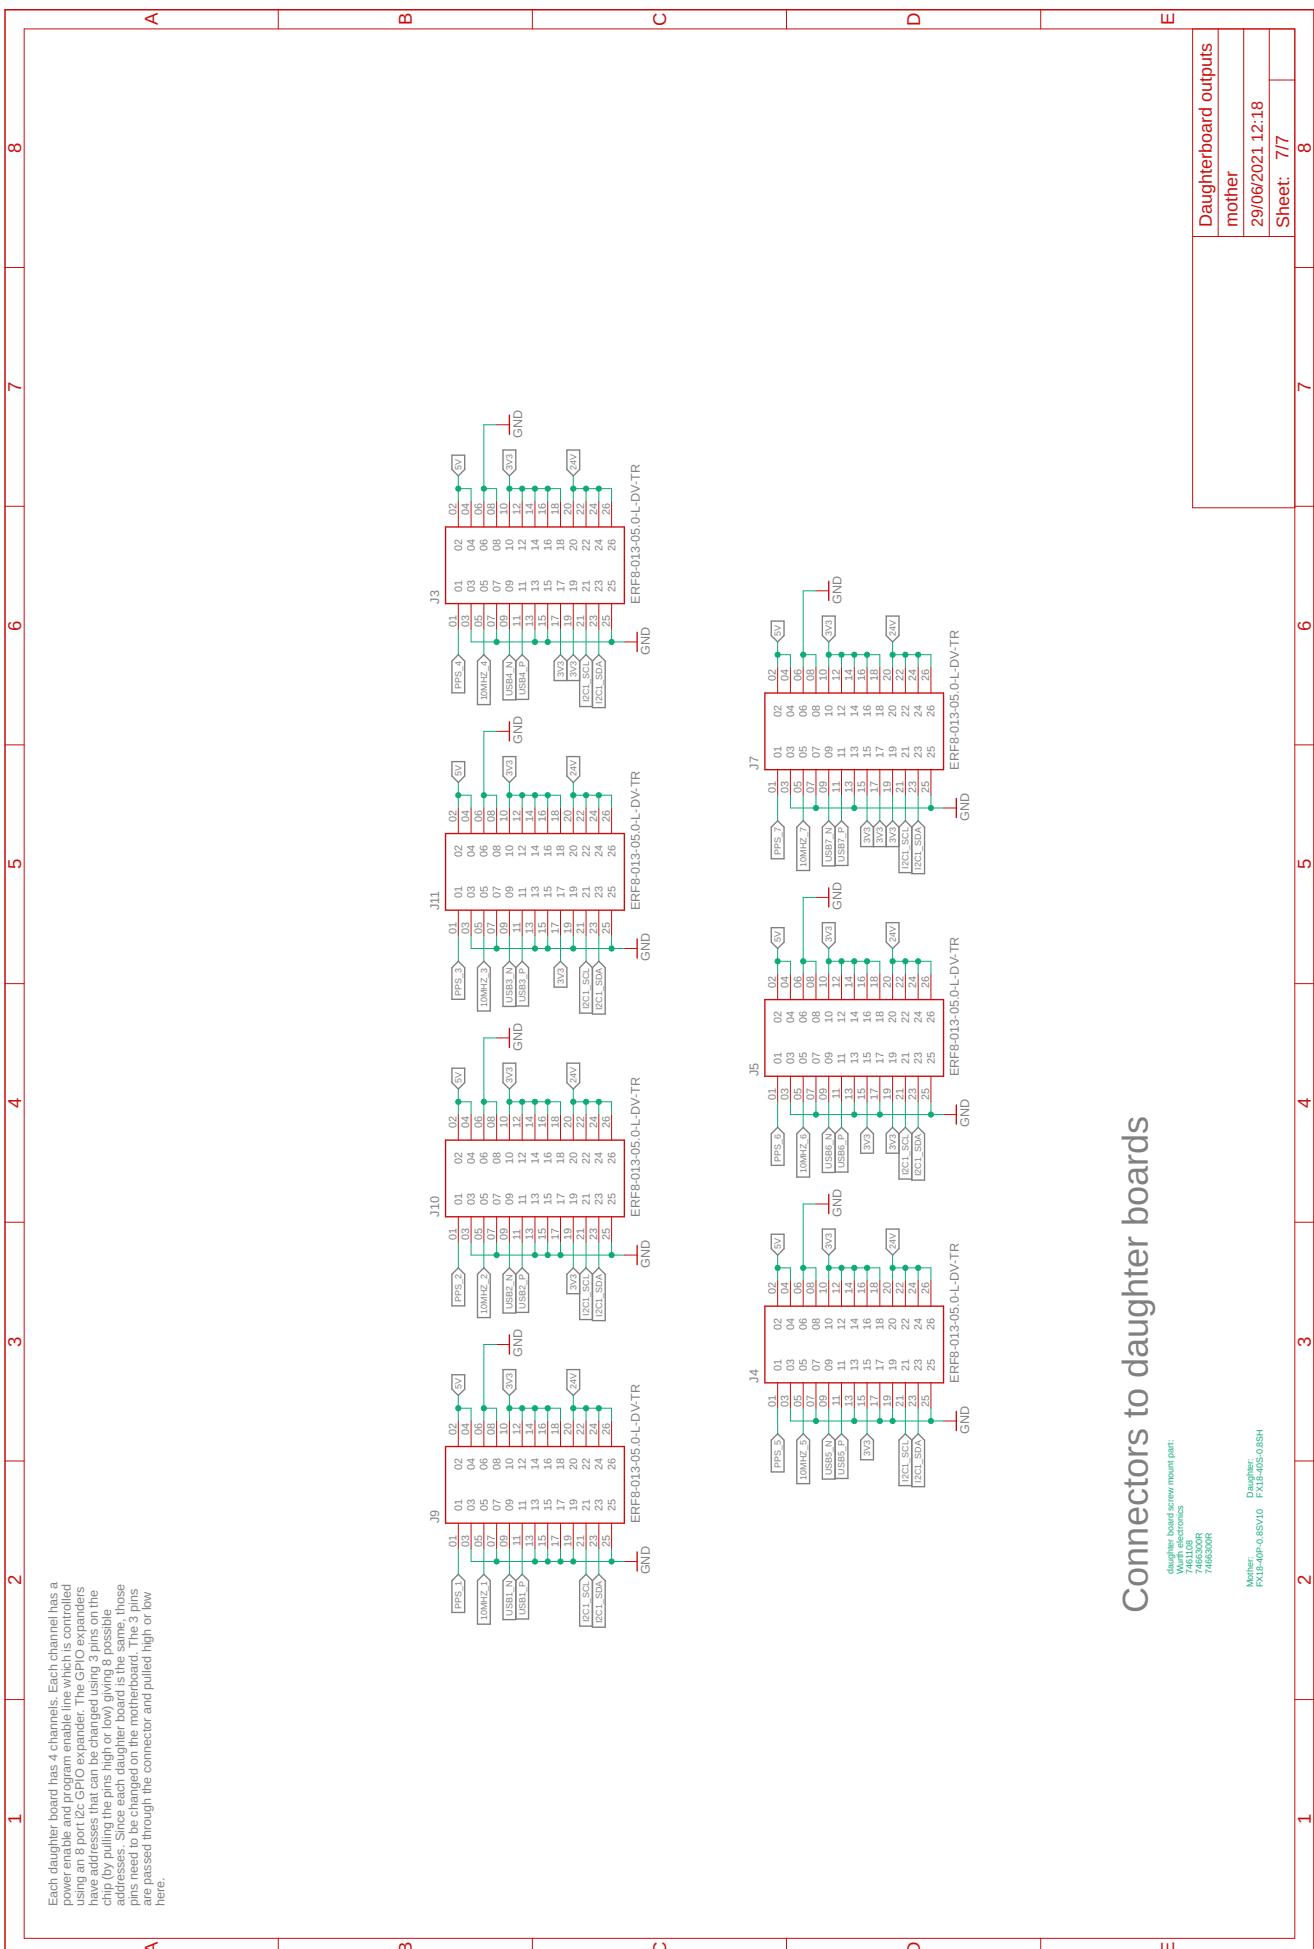


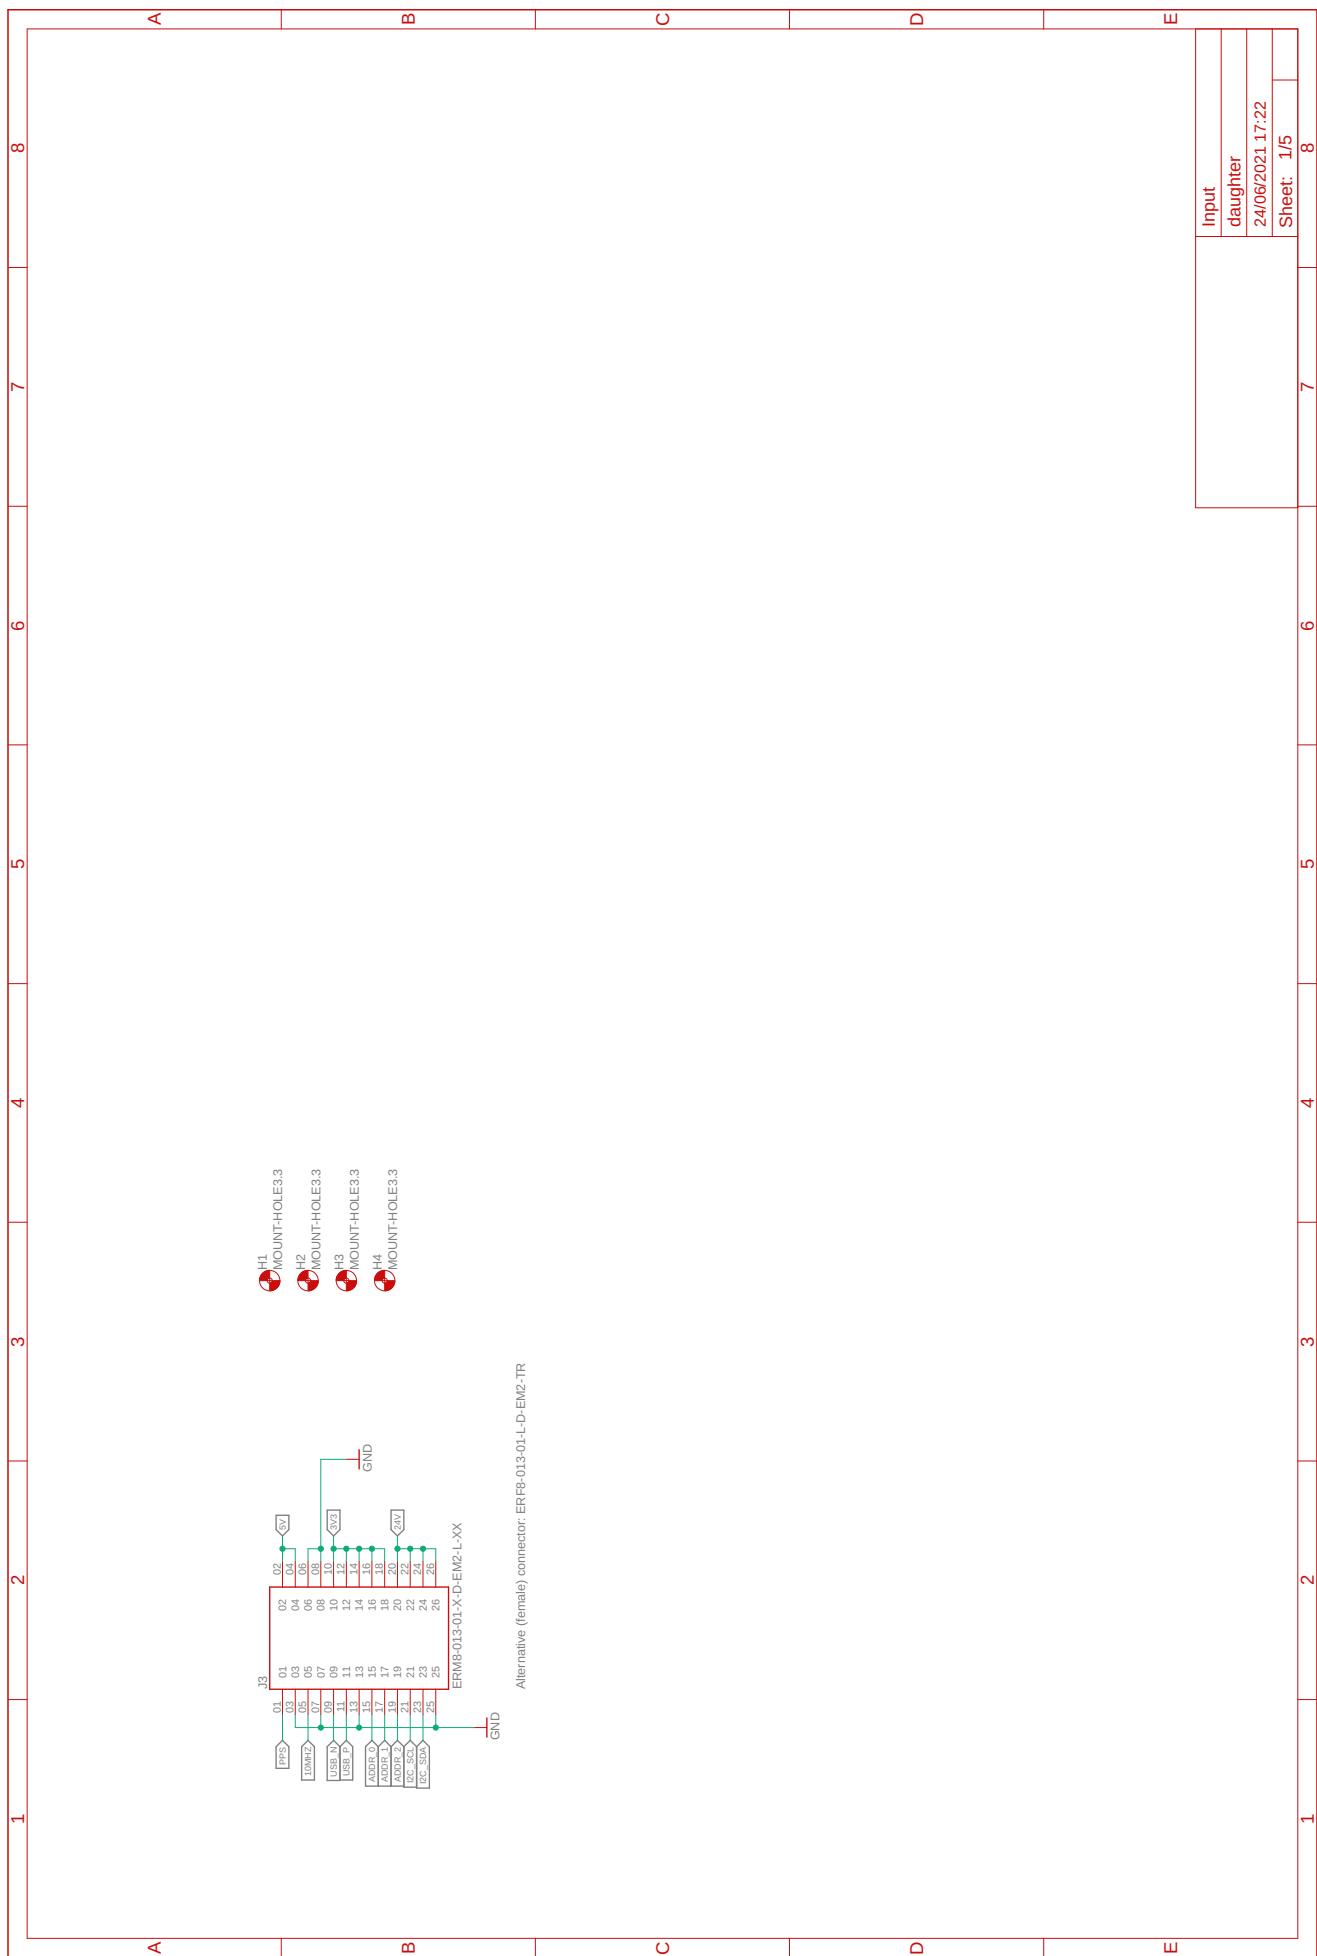


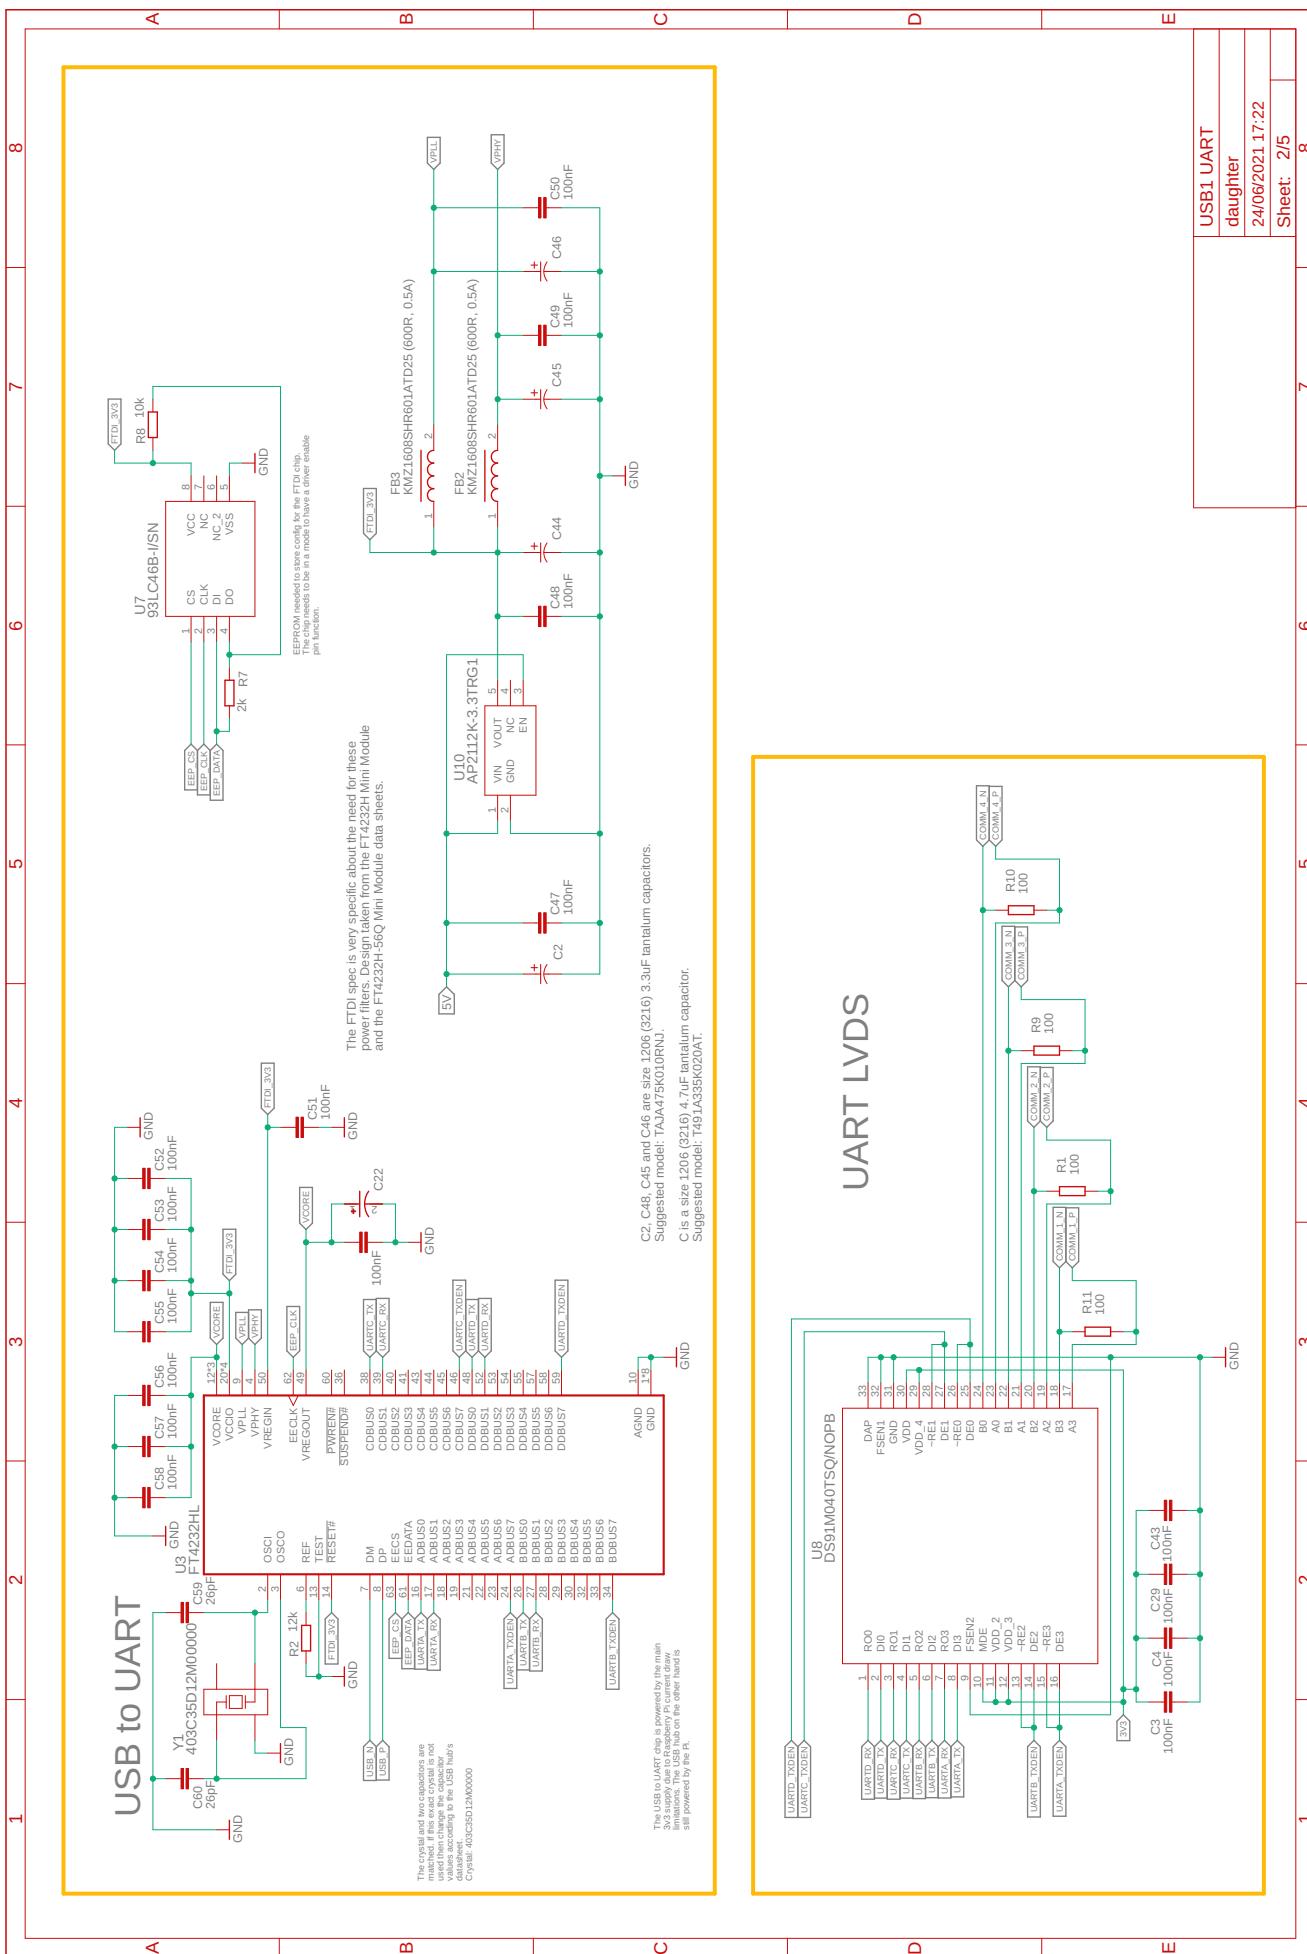


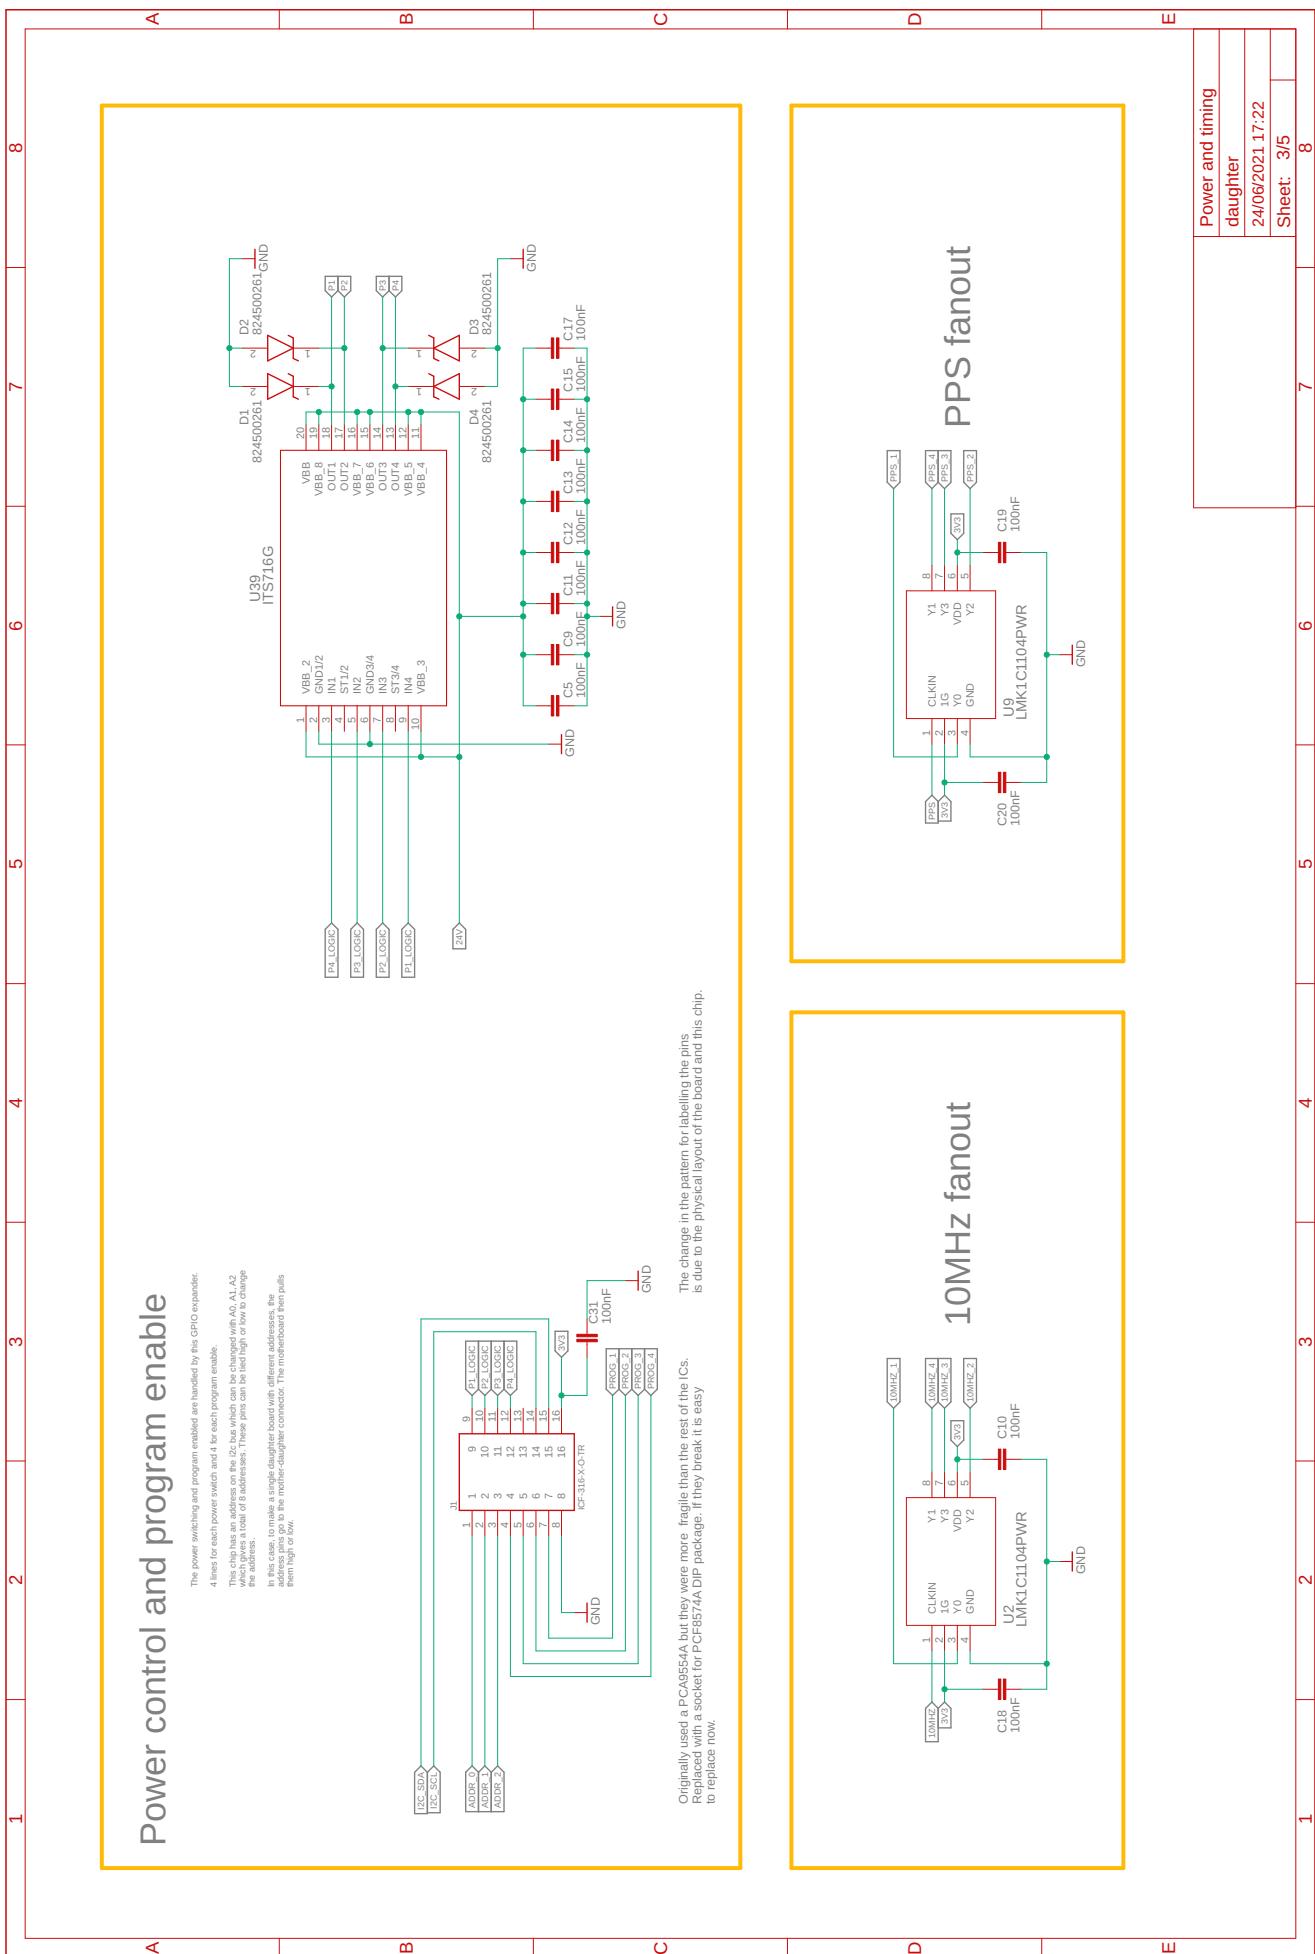


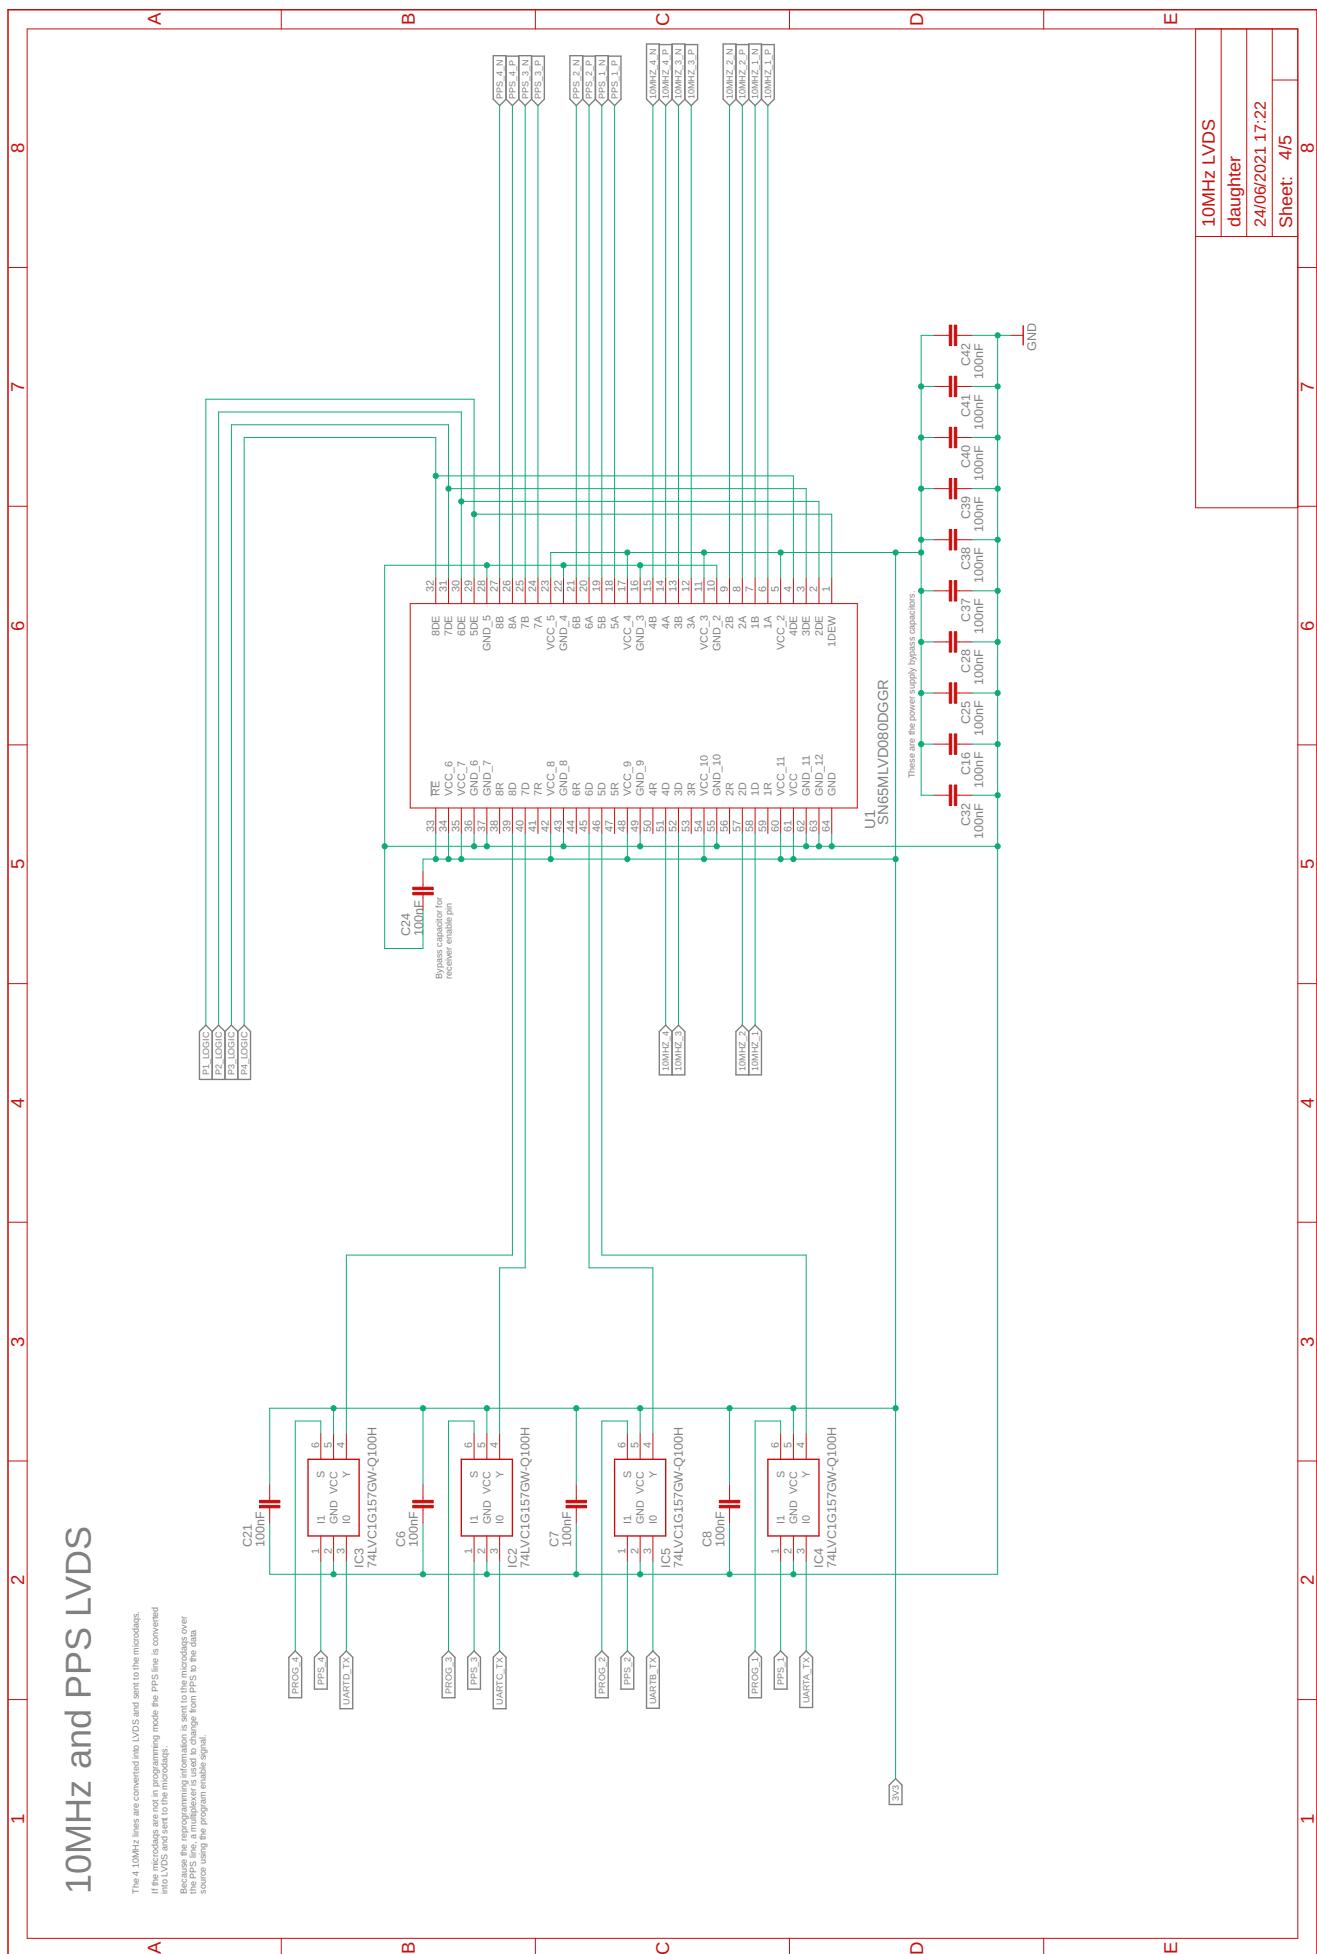


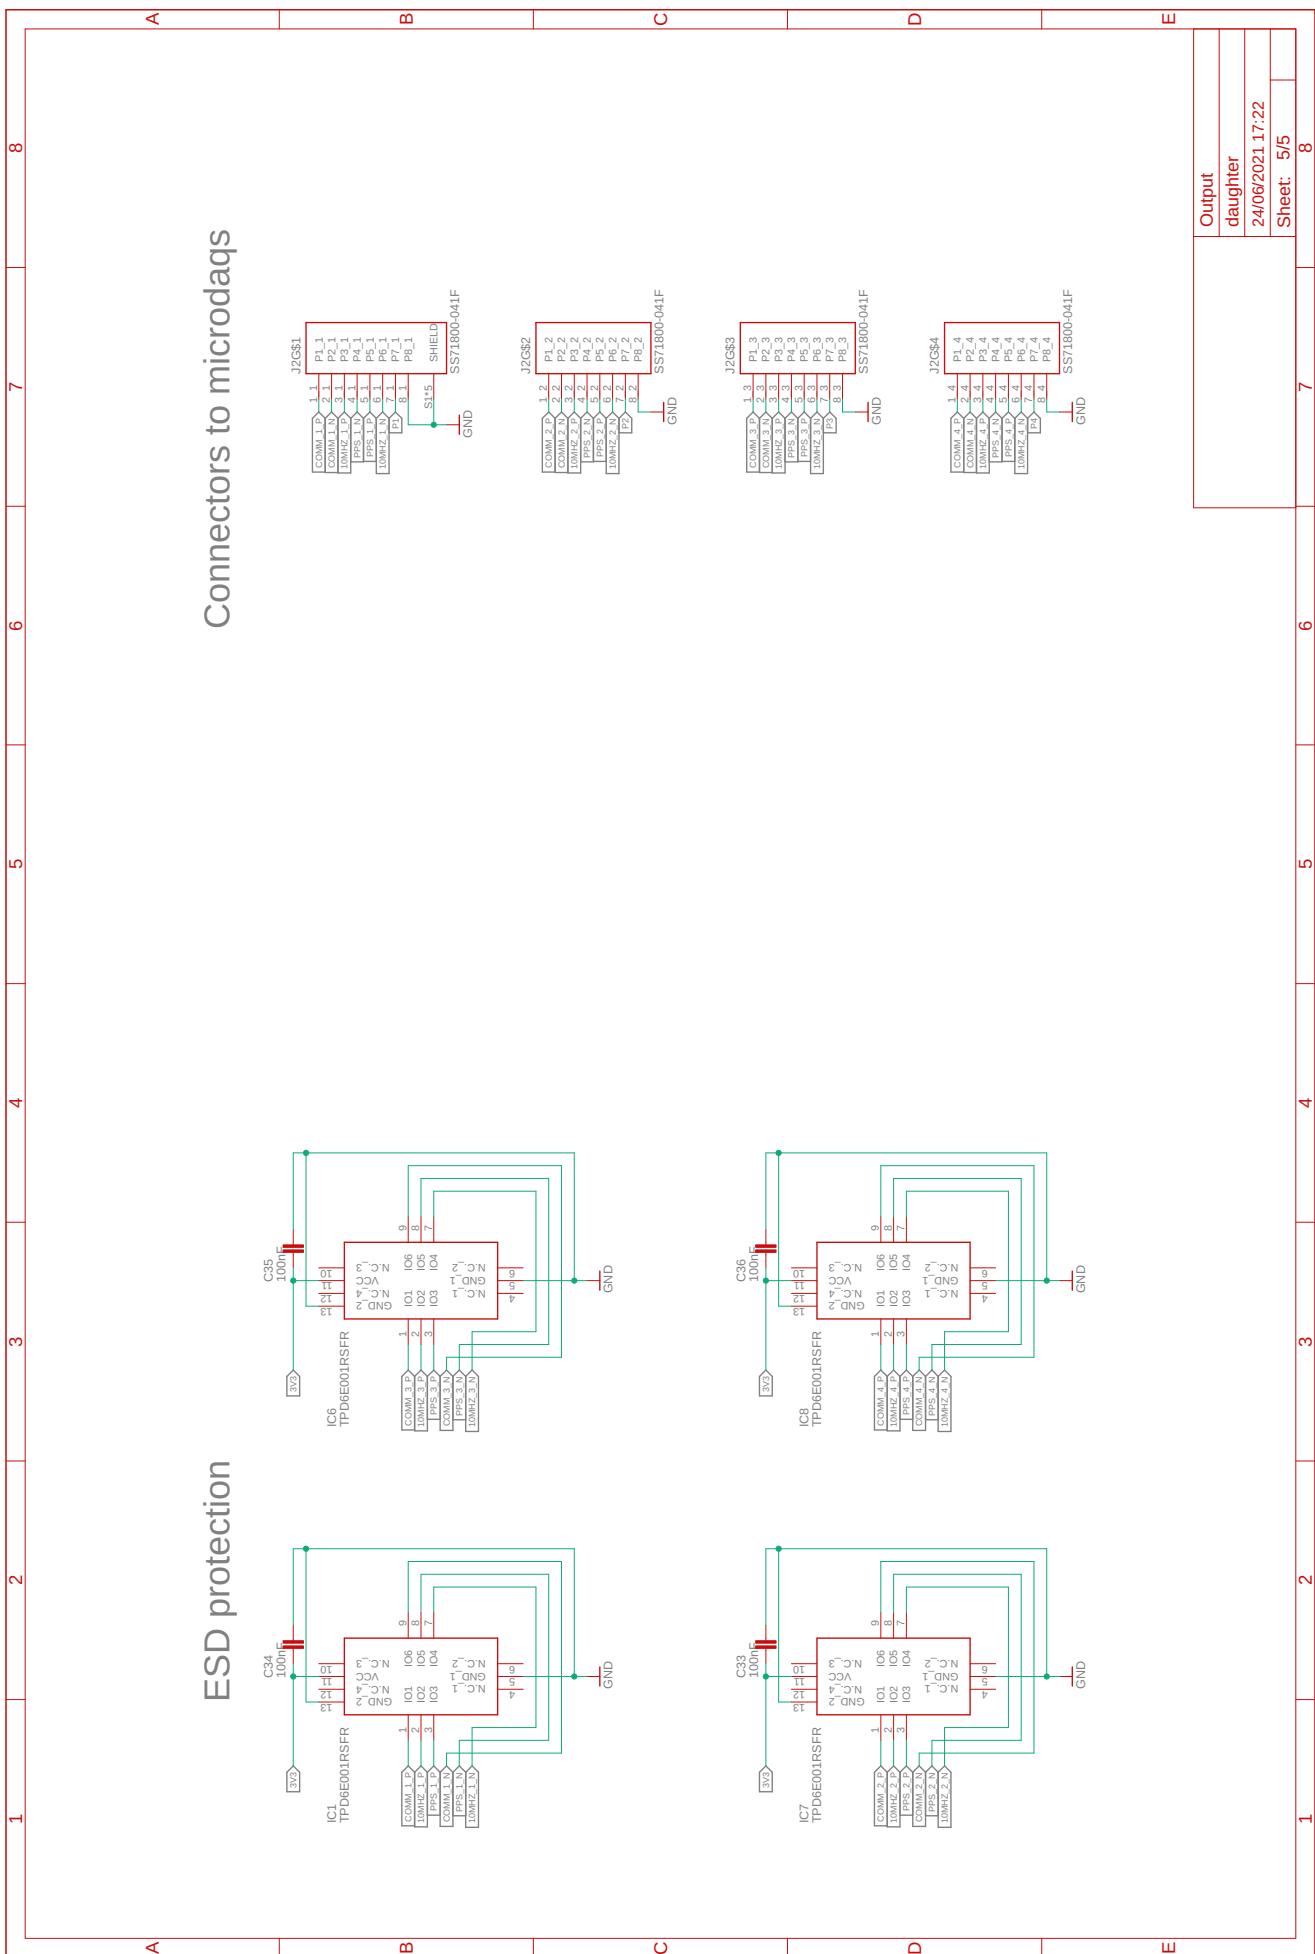












Bibliography

- [1] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, “Experimental Test of Parity Conservation in Beta Decay”, *Physical Review*, vol. 105, no. 4, pp. 1413–1415, Feb. 1957, ISSN: 0031-899X. DOI: [10.1103/PhysRev.105.1413](https://doi.org/10.1103/PhysRev.105.1413).
- [2] M. Goldhaber, L. Grodzins, and A. W. Sunyar, “Helicity of Neutrinos”, *Physical Review*, vol. 109, no. 3, pp. 1015–1017, Feb. 1958, ISSN: 0031-899X. DOI: [10.1103/PhysRev.109.1015](https://doi.org/10.1103/PhysRev.109.1015).
- [3] R. Davis, D. S. Harmer, and K. C. Hoffman, “Search for Neutrinos from the Sun”, *Physical Review Letters*, vol. 20, no. 21, pp. 1205–1209, May 1968, ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.20.1205](https://doi.org/10.1103/PhysRevLett.20.1205).
- [4] P. Ball, “Cosmic sleuths net physics Nobel”, *Nature*, no. 09, Oct. 2002. DOI: [10.1038/news021007-4](https://doi.org/10.1038/news021007-4).
- [5] A. Taroni, “Nobel Prize 2015: Kajita and McDonald”, *Nature Physics*, vol. 11, no. 11, pp. 891–891, Nov. 2015, ISSN: 1745-2473. DOI: [10.1038/nphys3543](https://doi.org/10.1038/nphys3543).
- [6] Q. R. Ahmad, R. C. Allen, T. C. Andersen, *et al.*, “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory”, *Physical Review Letters*, vol. 89, no. 1, p. 011301, Jun. 2002, ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.89.011301](https://doi.org/10.1103/PhysRevLett.89.011301).
- [7] Y. Fukuda, T. Hayakawa, E. Ichihara, *et al.*, “Evidence for Oscillation of Atmospheric Neutrinos”, *Physical Review Letters*, vol. 81, no. 8, pp. 1562–1567, Aug. 1998, ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562).
- [8] P. W. Higgs, “Broken Symmetries and the Masses of Gauge Bosons”, *Physical Review Letters*, vol. 13, no. 16, pp. 508–509, Oct. 1964, ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.13.508](https://doi.org/10.1103/PhysRevLett.13.508).
- [9] H. Yukawa, “On the Interaction of Elementary Particles. I”, *Progress of Theoretical Physics Supplement*, vol. 1, pp. 1–10, 195, ISSN: 0375-9687. DOI: [10.1143/PTPS.1.1](https://doi.org/10.1143/PTPS.1.1).
- [10] Mark Thompson, *Modern Particle Physics*. Cambridge University Press, Sep. 2013.

- [11] S. Weinberg, “A Model of Leptons”, *Physical Review Letters*, vol. 19, no. 21, pp. 1264–1266, Nov. 1967, ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.19.1264](https://doi.org/10.1103/PhysRevLett.19.1264).
- [12] C. Giunti and C. W. Kim, *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 2007.
- [13] J. Argyriades, R. Arnold, C. Augier, *et al.*, “Measurement of the Double- β Decay Half-Life of ^{150}Nd and Search for Neutrinoless Decay Modes with the NEMO-3 Detector”, *Physical Review C*, vol. 80, no. 3, p. 032501, Sep. 2009, ISSN: 0556-2813. DOI: [10.1103/PhysRevC.80.032501](https://doi.org/10.1103/PhysRevC.80.032501).
- [14] P. Adamson, I. Anghel, A. Aurisano, *et al.*, “Precision Constraints for Three-Flavor Neutrino Oscillations from the Full MINOS+ and MINOS Dataset”, *Physical Review Letters*, vol. 125, no. 13, p. 131802, Sep. 2020, ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.125.131802](https://doi.org/10.1103/PhysRevLett.125.131802).
- [15] M. A. Acero, P. Adamson, L. Aliaga, *et al.*, “First Measurement of Neutrino Oscillation Parameters using Neutrinos and Antineutrinos by NOvA”, *Physical Review Letters*, vol. 123, no. 15, p. 151803, Oct. 2019, ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.123.151803](https://doi.org/10.1103/PhysRevLett.123.151803).
- [16] M. H. Ahn, E. Aliu, S. Andringa, *et al.*, “Measurement of Neutrino Oscillation by the K2K Experiment”, *Physical Review D*, vol. 74, no. 7, p. 072003, Oct. 2006, ISSN: 1550-7998. DOI: [10.1103/PhysRevD.74.072003](https://doi.org/10.1103/PhysRevD.74.072003).
- [17] M. Aker, A. Beglarian, J. Behrens, *et al.*, “Direct Neutrino-Mass Measurement with Sub-Electronvolt Sensitivity”, *Nature Physics*, vol. 18, no. 2, pp. 160–166, Feb. 2022, ISSN: 1745-2473. DOI: [10.1038/s41567-021-01463-1](https://doi.org/10.1038/s41567-021-01463-1).
- [18] D. H. Perkins, *Introduction to High Energy Physics*, 4th. Cambridge University Press, 2000.
- [19] B. Pontecorvo, “Neutrino Experiments and the Problem of Conservation of Leptonic Charge”, *Zh. Eksp. Teor. Fiz.*, vol. 53, pp. 1717–1725, 1967.
- [20] Z. Maki, M. Nakagawa, and S. Sakata, “Remarks on the Unified Model of Elementary Particles”, *Progress of Theoretical Physics*, vol. 28, no. 5, pp. 870–880, Nov. 1962, ISSN: 0033-068X. DOI: [10.1143/PTP.28.870](https://doi.org/10.1143/PTP.28.870).
- [21] I. J. R. Aitchison and A. J. Hey, *Gauge Theories in Particle Physics: A Practical Introduction, Volume 2: Non-Abelian Gauge Theories*, Fourth. CRC Press, 2013, vol. 1&2, p. 352, ISBN: 9780429185397. DOI: [10.1201/9781466513105](https://doi.org/10.1201/9781466513105).
- [22] L. Wolfenstein, “Neutrino Oscillations in Matter”, *Physical Review D*, vol. 17, no. 9, pp. 2369–2374, May 1978, ISSN: 0556-2821. DOI: [10.1103/PhysRevD.17.2369](https://doi.org/10.1103/PhysRevD.17.2369).

- [23] S. P. Mikheyev and A. Y. Smirnov, “Resonant Amplification of ν Oscillations in Matter and Solar-Neutrino Spectroscopy”, *Il Nuovo Cimento C*, vol. 9, no. 1, pp. 17–26, Jan. 1986, ISSN: 0390-5551. DOI: 10.1007/BF02508049.
- [24] L. Whitehead, “Neutrino Oscillations with MINOS and MINOS+”, *Nuclear Physics B*, vol. 908, pp. 130–150, Jul. 2016, ISSN: 05503213. DOI: 10.1016/j.nuclphysb.2016.03.004.
- [25] P. A. Zyla, R. M. Barnett, J. Beringer, *et al.*, “Review of Particle Physics”, *Progress of Theoretical and Experimental Physics*, vol. 2020, no. 8, Aug. 2020, ISSN: 2050-3911. DOI: 10.1093/ptep/ptaa104.
- [26] D. Griffiths, *Introduction to Elementary Particles*. John Wiley & Sons, Inc., 1987.
- [27] J. C. Tingey, “Convolutional Neural Networks for the CHIPS Neutrino Detector R&D Project”, Ph.D. dissertation, University College London, 2021.
- [28] A. J. Perch, “Three-Flavour Neutrino Oscillations with MINOS and CHIPS”, Ph.D. dissertation, University College London, Jan. 2017.
- [29] D. I. Scully, “Neutrino Induced Coherent Pion Production”, Ph.D. dissertation, University of Warwick Department of Physics, Warwick, Sep. 2013.
- [30] D. Bhattacharya, “Neutrino and Antineutrino Inclusive Charged-Current Cross Section Measurement with the MINOS Near Detector”, Ph.D. dissertation, University of Pittsburgh, Batavia, IL (United States), Feb. 2009. DOI: 10.2172/952653.
- [31] E. A. Paschos, D. Schalla, F. Sanchez, M. Sorel, and L. Alvarez-Ruso, “Coherent Pion Production by Neutrinos”, in *AIP Conference Proceedings*, 2009, pp. 213–217. DOI: 10.1063/1.3274158.
- [32] K. Abe, C. Andreopoulos, M. Antonova, *et al.*, “Measurement of Coherent π^+ Production in Low Energy Neutrino-Carbon Scattering”, *Physical Review Letters*, vol. 117, no. 19, p. 192501, Nov. 2016, ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.117.192501.
- [33] J. Van Orden and T. Donnelly, “Mesonic Processes in Deep-Inelastic Electron Scattering from Nuclei”, *Annals of Physics*, vol. 131, no. 2, pp. 451–493, Feb. 1981, ISSN: 00034916. DOI: 10.1016/0003-4916(81)90038-5.
- [34] T. Katori, “Meson Exchange Current (MEC) Models in Neutrino Interaction Generators”, 2015, p. 030 001. DOI: 10.1063/1.4919465.
- [35] O. Benhar, D. Day, and I. Sick, “Inclusive Quasielastic Electron-Nucleus Scattering”, *Reviews of Modern Physics*, vol. 80, no. 1, pp. 189–224, Jan. 2008, ISSN: 0034-6861. DOI: 10.1103/RevModPhys.80.189.
- [36] *Photomultiplier Tubes Basics and Applications*, Third. Hamamatsu Photonics K.K., 2007.

- [37] R. J. Smith, “Measurement of Kaon and Pion Contributions to the T2K Neutrino Beam”, Ph.D. dissertation, University of Oxford Website Directions Save, Oxford, 2013.
- [38] A. Bueno, “The CERN Neutrino Beam to Gran Sasso and Future Long Baseline Experiments in Europe”, *Nuclear Physics B - Proceedings Supplements*, vol. 76, no. 1-3, pp. 463–469, Apr. 1999, ISSN: 09205632. DOI: 10.1016/S0920-5632(99)00508-3.
- [39] P. Adamson, K. Anderson, M. Andrews, *et al.*, “The NuMI Neutrino Beam”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 806, pp. 279–306, Jan. 2016, ISSN: 01689002. DOI: 10.1016/j.nima.2015.08.063.
- [40] M. Kordosky and D. Petyt, “A Study of Muon Neutrino Disappearance using the Fermilab Main Injector Neutrino Beam”, *Phys. Rev. D*, Nov. 2007. DOI: 10.1103/PhysRevD.77.072002.
- [41] L. J. Loiacono, “Measurement of the Muon Neutrino Inclusive Charged Current Cross Section on Iron using the MINOS Detector”, Fermi National Accelerator Laboratory (FNAL), Batavia, IL (United States), Tech. Rep., May 2010. DOI: 10.2172/1151630.
- [42] A. Timmons, “The Results of MINOS and the Future with MINOS+”, *Advances in High Energy Physics*, vol. 2016, pp. 1–25, 2016, ISSN: 1687-7357. DOI: 10.1155/2016/7064960.
- [43] P. Adamson, C. Ader, M. Andrews, *et al.*, “First Measurement of Muon-Neutrino Disappearance in NOvA”, *Physical Review D*, vol. 93, no. 5, p. 051104, Mar. 2016, ISSN: 2470-0010. DOI: 10.1103/PhysRevD.93.051104.
- [44] A. Holin, “Results from the MINOS Experiment and new MINOS+ Data”, in *16th International Workshop on Neutrino Factories and Future Neutrino Beam Facilities - NUFAC2014*, 2014.
- [45] D. Ayres, G. Drake, M. Goodman, *et al.*, “The NOvA Technical Design Report”, Fermi National Accelerator Laboratory (FNAL), Batavia, IL (United States), Tech. Rep., Oct. 2007. DOI: 10.2172/935497.
- [46] D. Beavis, A. Carroll, and I. Chiang, “Long Baseline Neutrino Oscillation Experiment at the AGS. Physics Design Report”, Brookhaven National Laboratory (BNL), Upton, NY, Tech. Rep., Apr. 1995. DOI: 10.2172/52878.
- [47] C. Alt, B. Baatar, D. Barna, *et al.*, “Inclusive Production of Charged Pions in p+C Collisions at 158GeV/c Beam Momentum”, *The European Physical Journal C*, vol. 49, no. 4, pp. 897–917, Mar. 2007, ISSN: 1434-6044. DOI: 10.1140/epjc/s10052-006-0165-7.

- [48] J. M. Paley, M. D. Messier, R. Raja, *et al.*, “Measurement of Charged Pion Production Yields off the NuMI Target”, *Physical Review D*, vol. 90, no. 3, p. 032001, Aug. 2014, ISSN: 1550-7998. DOI: [10.1103/PhysRevD.90.032001](https://doi.org/10.1103/PhysRevD.90.032001).
- [49] N. Nayak, “A Joint Measurement of ν_μ -Disappearance and ν_e -Appearance in the NuMI Beam using the NOvA Experiment”, Ph.D. dissertation, Fermi National Accelerator Lab. (FNAL), 2021.
- [50] L. Vinton, “Measurement of Muon Neutrino Disappearance with the NOvA Experiment”, Fermi National Accelerator Laboratory (FNAL), Batavia, IL (United States), Tech. Rep., Jan. 2018. DOI: [10.2172/1423216](https://doi.org/10.2172/1423216).
- [51] J. G. Morfín, J. Nieves, and J. T. Sobczyk, “Recent Developments in Neutrino/Antineutrino-Nucleus Interactions”, *Advances in High Energy Physics*, vol. 2012, pp. 1–35, 2012, ISSN: 1687-7357. DOI: [10.1155/2012/934597](https://doi.org/10.1155/2012/934597).
- [52] T. Blackburn, “Measurement of Δm_{32}^2 and $\sin^2 \theta_{23}$ using Muon Neutrino and Antineutrino Beams in the NOvA Experiment”, Fermi National Accelerator Laboratory (FNAL), Batavia, IL (United States), Tech. Rep., Jan. 2019. DOI: [10.2172/1502822](https://doi.org/10.2172/1502822).
- [53] M. Muether, “NOvA Systematics”, in *Proceedings of the 10th International Workshop on Neutrino-Nucleus Interactions in Few-GeV Region (NuInt15)*, Journal of the Physical Society of Japan, Dec. 2016, ISBN: 4-89027-116-3. DOI: [10.7566/JPSCP.12.010011](https://doi.org/10.7566/JPSCP.12.010011).
- [54] A. Ferrari, P. Sala, A. Fasso, and J. Ranft, “FLUKA: A Multi-Particle Transport Code”, Stanford Linear Accelerator Center (SLAC), Menlo Park, CA, Tech. Rep., Dec. 2005. DOI: [10.2172/877507](https://doi.org/10.2172/877507).
- [55] T. Böhlen, F. Cerutti, M. Chin, *et al.*, “The FLUKA Code: Developments and Challenges for High Energy and Medical Applications”, *Nuclear Data Sheets*, vol. 120, pp. 211–214, Jun. 2014, ISSN: 00903752. DOI: [10.1016/j.nds.2014.07.049](https://doi.org/10.1016/j.nds.2014.07.049).
- [56] S. Afanasiev, T. Alber, H. Appelshäuser, *et al.*, “The NA49 Large Acceptance Hadron Detector”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 430, no. 2-3, pp. 210–244, Jul. 1999, ISSN: 01689002. DOI: [10.1016/S0168-9002\(99\)00239-9](https://doi.org/10.1016/S0168-9002(99)00239-9).
- [57] R. Raja, “The Main Injector Particle Physics Experiment (MIPP FNAL E-907) at Fermilab status and plans”, *Nuclear Physics B - Proceedings Supplements*, vol. 175-176, pp. 17–24, Jan. 2008, ISSN: 09205632. DOI: [10.1016/j.nuclphysbps.2007.10.003](https://doi.org/10.1016/j.nuclphysbps.2007.10.003).

- [58] P. Adamson, I. Anghel, A. Aurisano, *et al.*, “Search for Sterile Neutrinos in MINOS and MINOS+ using a Two-Detector Fit: Supplemental Discussion”, ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.122.091803.
- [59] A. V. Lebedev, “Ratio of Pion Kaon Production in Proton Carbon Interactions”, Ph.D. dissertation, Harvard University, May 2007. DOI: 10.2172/948174.
- [60] H. Meyer, “The main injector particle production experiment (MIPP) at fermilab”, *Journal of Physics: Conference Series*, vol. 69, p. 012025, May 2007. DOI: 10.1088/1742-6596/69/1/012025. [Online]. Available: <https://doi.org/10.1088/1742-6596/69/1/012025>.
- [61] D. Castelvecchi, “Japan Will Build the Worlds Largest Neutrino Detector”, *Nature*, Dec. 2019, ISSN: 0028-0836. DOI: 10.1038/d41586-019-03874-w.
- [62] F. Di Lodovico, “The Hyper-Kamiokande Experiment”, *Journal of Physics: Conference Series*, vol. 888, p. 012020, Sep. 2017, ISSN: 1742-6588. DOI: 10.1088/1742-6596/888/1/012020.
- [63] A. Cho, “Costs Balloon for U.S. Particle Physics Megaproject”, *Science*, vol. 373, no. 6562, pp. 1424–1425, Sep. 2021.
- [64] Esri. “"topographic" [basemap]. Scale Not Given. "world topographic map". january 12, 2023”. (Jan. 2023), [Online]. Available: <https://uclondon.maps.arcgis.com/home/item.html?id=7dc6cea0b1764a1f9af2e679f642f0f5>.
- [65] S. Fukuda, Y. Fukuda, T. Hayakawa, *et al.*, “The super-kamiokande detector”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 501, pp. 418–462, 2-3 Apr. 2003, ISSN: 01689002. DOI: 10.1016/S0168-9002(03)00425-X.
- [66] S. V. Cao, J. Huang, K. Lang, and F. Nova, “Cosmic ray muon rates in water cherenkov detectors at shallow depths”, The University of Texas at Austin, May 2013.
- [67] P. Adamson, J. Austin, S. V. Cao, *et al.*, “CHerenkov Detectors In Mine PitS (CHIPS) Letter of Intent to FNAL”, Fermi National Accelerator Laboratory (FNAL), Batavia, IL (United States), Tech. Rep., Dec. 2013. DOI: 10.2172/1342789.
- [68] M. Campbell, “Measuring neutrino oscillations in the nova and chips detectors”, University College London, 2020.
- [69] A. Margiotta, A. Marinelli, C. Markou, *et al.*, “KM3NeT Front-End and Readout Electronics System: Hardware, Firmware, and Software”, *Journal of Astronomical Telescopes, Instruments, and Systems*, vol. 5, no. 04, p. 1, Dec. 2019, ISSN: 2329-4124. DOI: 10.1117/1.JATIS.5.4.046001.

- [70] T. Huber, J. Kelley, S. Kunwar, and D. Tosi, “The IceTop Scintillator Upgrade”, in *Proceedings of 35th International Cosmic Ray Conference PoS(ICRC2017)*, Trieste, Italy: Sissa Medialab, Aug. 2017, p. 401. DOI: 10.22323/1.301.0401.
- [71] J. D. Cockcroft, Walton, and E. T. S., “Experiments with High Velocity Positive Ions.(I) Further Developments in the Method of Obtaining High Velocity Positive Ions”, *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, vol. 136, no. 830, pp. 619–630, Jun. 1932, ISSN: 0950-1207. DOI: 10.1098/rspa.1932.0107.
- [72] D. Gajana, V. Gromov, and P. Timmer, “ASIC Design in the KM3NeT Detector”, *Journal of Instrumentation*, vol. 8, no. 02, pp. C02030–C02030, Feb. 2013, ISSN: 1748-0221. DOI: 10.1088/1748-0221/8/02/C02030.
- [73] D. Real and D. Calvo, “Digital Optical Module Electronics of KM3NeT”, *Physics of Particles and Nuclei*, vol. 47, no. 6, pp. 918–925, Nov. 2016, ISSN: 1063-7796. DOI: 10.1134/S1063779616060216.
- [74] J. Serrano, P. Alvarez, Cattin M., E. G. Cota, P. M. J. H. Lewis, and T. Wostowski, “The White Rabbit Project”, in *ICALEPCS TUC004*, Kobe, 2009.
- [75] E. G. Cota, M. Lipiski, T. Wostowski, E. van der Bij, and J. Serrano, “White rabbit specification: Draft for comments”, version 2.0, Jul. 2011.
- [76] “IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems”, Standard IEEE 1588-2019, 2020.
- [77] G. Daniluk, “White Rabbit PTP Core the Sub-Nanosecond Time Synchronization Over Ethernet”, Ph.D. dissertation, Warsaw University of Technology, Warsaw, 2012.
- [78] Lattice Semiconductor Corporation, “LatticeMico32 Processor Reference Manual”, Dec. 2012. [Online]. Available: <https://www.latticesemi.com/Products/DesignSoftwareAndIP/EmbeddedDesignSoftware/LatticeMicoSystem>.
- [79] P. Horowitz and W. Hill, *The Art of Electronics*, 3rd. Cambridge University Press, 2015, ISBN: 0521809266.
- [80] *STM32F446xC/E*, 10th ed. STMicroelectronics, Jan. 2021.
- [81] G. Coley, *BeagleBone Green System Reference Manual*. BeagleBoard.org.
- [82] *User Guide WR-LEN*, 5th ed. Seven Solutions S.L, Jul. 2017.

- [83] S. Bash, J. Cesar, G. Deuerling, *et al.*, “Low-Latency NuMI Trigger for the Chips-5 Neutrino Detector”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1030, p. 166513, May 2022, ISSN: 01689002. DOI: 10.1016/j.nima.2022.166513.
- [84] D. G. Beechy and R. J. Ducar, “Time and data distribution systems at the fermilab accelerator”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 247, pp. 231–238, 1 Jun. 1986, ISSN: 01689002. DOI: 10.1016/0168-9002(86)90569-3.
- [85] A. Norman, R. Kwarciany, G. Deuerling, and N. Wilcer, “The NOvA Timing System: A System for Synchronizing a Long Baseline Neutrino Experiment”, *Journal of Physics: Conference Series*, vol. 396, no. 1, p. 012034, Dec. 2012, ISSN: 1742-6588. DOI: 10.1088/1742-6596/396/1/012034.
- [86] Fermi National Accelerator Laboratory, “Fermilab 2021 Site Sustainability Plan”, Tech. Rep.
- [87] Raspberry Pi Ltd, “Raspberry Pi Compute Module 4 Product Brief”, Apr. 2022. [Online]. Available: <https://datasheets.raspberrypi.com/cm4/cm4-product-brief.pdf>.
- [88] USB Implementers Forum, “Universal Serial Bus Specification Revision 2.0”, Apr. 2000.
- [89] Texas Instruments, “AM335x and AMIC110 Sitara Processors Technical Reference Manual”, Oct. 2011.
- [90] QEMU team, *QEMU A generic and Open Source Machine Emulator and Virtualizer*, Jul. 2021. [Online]. Available: <https://www.qemu.org>.
- [91] T. Haynes, “Network file system (nfs) version 4 minor version 2 protocol”, Nov. 2016. DOI: 10.17487/RFC7862.
- [92] L. Aliaga Soplin, “Neutrino Flux Prediction for the NuMI Beamline”, Fermi National Accelerator Laboratory (FNAL), Batavia, IL (United States), Tech. Rep., Jan. 2016. DOI: 10.2172/1250884.
- [93] L. Aliaga (Fermilab) and the Beam group, “NuMI Beam Prediction for the NOvA 2017 Analyses”, Tech. Rep., Sep. 2017.
- [94] M. Subotowicz, “Interstellar Communication by Neutrino Beams”, *Acta Astronautica*, vol. 6, no. 1-2, pp. 213–220, Jan. 1979, ISSN: 00945765. DOI: 10.1016/0094-5765(79)90157-7.
- [95] P. Huber, “Submarine Neutrino Communication”, *Physics Letters B*, vol. 692, no. 4, pp. 268–271, Sep. 2010, ISSN: 03702693. DOI: 10.1016/j.physletb.2010.08.003.

- [96] D. D. Stancil, P. Adamson, M. Alania, *et al.*, “Demonstration of Communication using Neutrinos”, English (US), *Modern Physics Letters A*, vol. 27, no. 12, Apr. 2012, ISSN: 0217-7323. DOI: 10.1142/S0217732312500770.
- [97] S. Agostinelli, J. Allison, K. Amako, *et al.*, “Geant4A Simulation Toolkit”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, no. 3, pp. 250–303, Jul. 2003, ISSN: 01689002. DOI: 10.1016/S0168-9002(03)01368-8.
- [98] J. Allison, K. Amako, J. Apostolakis, *et al.*, “Geant4 Developments and Applications”, *IEEE Transactions on Nuclear Science*, vol. 53, no. 1, pp. 270–278, Feb. 2006, ISSN: 0018-9499. DOI: 10.1109/TNS.2006.869826.
- [99] J. Allison, K. Amako, J. Apostolakis, *et al.*, “Recent Developments in Geant4”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 835, pp. 186–225, Nov. 2016, ISSN: 01689002. DOI: 10.1016/j.nima.2016.06.125.
- [100] R. Hatcher, “Proposal for a Unified Flux N-tuple Format”, Apr. 2012.
- [101] R. Brun and F. Rademakers, “Root - an object-oriented data analysis framework”, in *Proceedings AIHENP’96 Workshop, Lausanne, Sep. 1996*, Apr. 1997, pp. 81–86. DOI: 10.1016/S0168-9002(97)00048-X.
- [102] A. Booth, “Electron Neutrino Appearance at the NOvA Experiment”, Ph.D. dissertation, University of Sussex, Sussex, 2021.
- [103] M. Del Tutto, “Neutrino Beam Simulations and Data Checks for the NOvA Experiment”, Laurea thesis, Rome U., 2015. DOI: 10.2172/1230042.
- [104] S. E. Kopp, “The NuMI Neutrino Beam and Potential for an Off-Axis Experiment”, *Journal of Physics G: Nuclear and Particle Physics*, vol. 29, no. 8, pp. 1757–1761, Aug. 2003, ISSN: 0954-3899. DOI: 10.1088/0954-3899/29/8/343.