

Measurement of Neutron Multiplicity in Charged Current Neutrino Interactions on Oxygen

By

JULIE HE
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Physics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Robert Svoboda, Chair

Michael Mulhearn

Eric Prebys

Committee in Charge

2024

Abstract

The Accelerator Neutrino Neutron Interaction Experiment (ANNIE) is a 26-ton gadolinium-doped water Cherenkov detector located 100 meters downstream in the Booster Neutrino Beam (BNB) at the Fermi National Accelerator Laboratory (Fermilab). Its primary goals are to (1) measure the neutron yield from ν_μ interactions as a function of momentum transfer Q^2 so that neutrino-nucleus interaction models can be better constrained, and (2) demonstrate the power of novel, fast-timing detectors with the first deployment of Large Area Picosecond PhotoDetectors (LAPPDs).

Current knowledge of neutrino-nucleus interactions fall short in modeling the topologies of such interactions, leading to inaccurate reconstruction of event kinematics such as particle energy, direction, and vertex. As a consequence, the accuracy of cross section measurements is impacted, which is necessary for precise physics measurements. Neutrons are an indication of inelasticity and affect the determination of the energy of the parent neutrino. Quantifying the neutron yield is a step towards reducing the associated uncertainties, and thus improving our understanding of these complex interactions and benefiting the next generation of long-baseline neutrino experiments. ANNIE will make use of LAPPDs to measure neutron multiplicity of CC- 0π ν_μ interactions on oxygen, making it the first experiment to deploy an array of these photodetectors. Because the LAPPD is a novel photodetection technology, much customization is required to integrate it into existing electronics.

The first half of this thesis covers the significant progress made towards the first deployment of the LAPPD system. From its test stand at Fermilab, the LAPPD system was systematically tested and put together until deployment readiness was achieved. I present my contributions to the design, fabrication, and testing of the waterproof housing and cables, and the commissioning of the LVHV board that powers the LAPPD and its readout electronics. These efforts brought the LAPPD system significantly closer to deployment, and eventually first data.

The second half of this thesis presents the vertex and energy reconstruction algorithms developed to analysis the beam data with PMT-only information. While much progress has been made towards the deployment of LAPPDs, with several in the detector tank, efforts to integrate the LAPPD datastream are in progress. Thus, I developed a ring edge detection technique using PMT data to fit the muon vertex and determine its energy. The analysis in this thesis finds average neutron yields of $\bar{n}_{data} = 0.452 \pm 0.039(\text{stat}) \pm 0.27(\text{sys})$ for a selection of muon neutrino candidates in the fiducial volume of ANNIE and corrected with an averaged neutron detection efficiency. An equivalent analysis for simulated beam data results in an average neutron yield of $\bar{n}_{MC} = 0.582 \pm 0.018(\text{stat}) \pm 0.25(\text{sys})$. Future work includes application of efficiency corrections at a positional level, quantification of all systematic uncertainties, and neutron multiplicity measurements with other event topologies.

To my mom and dad,

who left their lives in a familiar place and built new ones in a foreign country. Without their sacrifices, I would not be where I am and who I am today. Thank you from the bottom of my heart. Dad, thanks for being my dad. You can't be here to see me graduate, but I hope I made you proud. I miss you every day.

Contents

Contents	v
List of Figures	vii
List of Tables	xix
1 Introduction	1
2 Neutrinos	6
2.1 Neutrinos in the Standard Model	12
2.2 Neutrino Mass	16
2.3 Neutrino Oscillation	21
2.4 Neutrino Interactions with Matter	28
2.5 Open Questions and Current Parameters	38
3 ANNIE	43
3.1 ANNIE Phase I	44
3.2 ANNIE Phase II Detector	48
3.3 Booster Neutrino Beam (BNB)	70
3.4 Physics Program	74
4 The Large Area Picosecond Photodetector	80
4.1 Conventional Photodetectors	81
4.2 The Design & Technology Behind the LAPPD	86
4.3 Why Are LAPPDs Important in ANNIE?	89
4.4 Full LAPPD System: Overview & Challenges	90
4.5 LAPPD System Components	95
4.6 Characterization & Testing	115
4.7 Deployment & Integration	122
4.8 My Work	126
4.9 First LAPPD Data	127

5 Calibration	128
5.1 Detector Calibration	128
5.2 Timing Calibration	140
5.3 Water Transparency and Quality	142
5.4 Neutron Capture and Detection Calibration	145
6 Neutron Multiplicity Measurement	151
6.1 Processing Raw Data Into ANNIE Events	152
6.2 Tank and MRD Cluster Finding	158
6.3 MRD Track Reconstruction	160
6.4 Muon Candidate Selection (Prompt Window)	164
6.5 Imaging the Cherenkov Ring	166
6.6 Muon Vertex Fit and Energy Reconstruction	175
6.7 Neutron Candidate Selection (Delayed Window)	182
6.8 Neutron Multiplicity as a Function of Momentum Transfer	187
6.9 Brief Overview of Systematic Uncertainties	191
7 Conclusion	195
7.1 Summary	195
7.2 Future Directions	197
A List of Modifications to LAPPD System Components	199
A.1 Re-design of the LVHV Board	200
A.2 Modifications to the ACDC Boards	203
A.3 Modifications to the Watertight Housing	203
A.4 Re-splicing Waterproof Cables	204
B Slow Controls Routine	207
B.1 PIC Programming	207
B.2 Diagram of the Slow Controls Routine	209
C Postdoc Performance in Fitting Tank Track Lengths	210
D Graduate Student Julie at Work	211
Bibliography	213

List of Figures

2.1	A periodic table of the Standard Model particles, highlighting the different generations (or families) of fermions, the gauge bosons that mediate the strong, weak, and electromagnetic interactions, and the Higgs boson that gives the W and Z bosons their mass. Figure from Wikimedia Commons.	15
2.2	Feynman diagrams comparing $2\nu\beta\beta$ decay (left) and $0\nu\beta\beta$ decay (right). If neutrinos are Majorana, the two neutrinos would annihilate each other and no neutrinos would be observed. This process violates lepton number by two. Feynman diagrams adapted from [40].	19
2.3	A pictorial representation of the normal and inverted mass orderings. The colors represent the flavor composition of each mass eigenstate. The subscripts for Δm^2 terms indicate the type of neutrino experiments that are able to determine the mass squared differences. Note that Δm_{sol}^2 is equivalent to Δm_{12}^2 and Δm_{atm}^2 to Δm_{23}^2 . Figure from [46].	27
2.4	For both types of scattering: the left Feynman diagram is a charged current interaction in which a $W^{+/-}$ boson is exchanged, and the right is a neutral current interaction in which a Z^0 boson is exchanged. The α subscript can be any of the three neutrino flavors.	29
2.5	The famous plot of neutrino cross section as a function of energy. Quasi-elastic interactions dominate in the low to medium energy range. At medium energies, resonant processes turn on and contribute a significant portion of events. At higher energies, deep inelastic scattering is the main interaction. Figure from [50].	30
3.1	The Neutron Capture Volume (NCV) is an acrylic vessel filled with EJ-335, a mixture of liquid scintillator and gadolinium. A PMT coupled to the NCV is wrapped in black to prevent light generated in the tank from triggering the data acquisition system. The array of PMTs at the bottom of the Phase I detector can be seen facing upward in this image.	45
3.2	The NCV was deployed at various locations in the ANNIE Phase I detector. These background measurements show that background neutron rates are sufficiently low in a fiducial volume located in a central location slightly upstream of the tank center, indicated by the purple circle in the diagram above. Figure from [65].	47

3.3	ANNIE Phase II detector diagram. The tank is 3-m in diameter and 4-m tall, with PMTs and LAPPDs lining its walls. The Front Muon Veto and Muon Range Detector sit upstream and downstream of the tank, respectively. The readout electronics and power supplies are located on the level above the detector. In this diagram, the direction of the beam is from left to right. Figure from [66].	48
3.4	LEFT: Fishbowl perspective of the inner structure, wrapped in tape made of polytetrafluoroethylene (PTFE). RIGHT: 3D model of the inner structure. The different colors indicate different PMT types. Unistrut rails between two panels are used for LAPPD deployment.	50
3.5	Sections of a rectangular steel pipe soaking in 1% gadolinium-loaded water. The right pipe is wrapped in teflon tape. It is clear that any presence of rust will quickly degrade the water transparency, so precautions must be taken to reduce the risk of rust contamination.	51
3.6	The inner structure is lowered into the experimental hall while photo shoots among collaborators take place. Plastic sheets protect the inner structure against dust during transport. The SBND (Short-Baseline Near Detector, as part of the Short-Baseline Neutrino Program) experimental hall is the large building to the right.	52
3.7	A view of the tank lid from the top level of ANNIE Hall. The five green circles indicate the locations of the ports used for source deployment. The central port is part of the larger removable hatch. The red rectangle highlights one of the eight “mail slots” used for LAPPD deployment.	53
3.8	The blue glow of Cherenkov radiation can be seen in the core of a nuclear reactor. This is a consequence of the phenomenon’s λ^{-2} dependence, resulting in higher intensities of light at shorter wavelengths.	57
3.9	Examples of spectra for different materials soaking in 1% gadolinium-loaded water. Samples of the soak solution are analyzed frequently in the first few weeks of testing, and gradually reduce in frequency. The increasing absorbance at wavelengths below 300 nm will not affect light collection in ANNIE, but may indicate that something is leaching into the water.	59
3.10	Materials degrading in gadolinium loaded water. The blue arrows and circles indicate that the material is deteriorating in the solution. This eliminates the material from use in the detector.	60
3.11	LEFT: Diagram of water flow through the ANNIE’s custom water filtration system. Black arrows mark points where the flow can be manually stopped for maintenance and testing. RIGHT: Testing the water system at Fermilab before installation in the hall. Blue tubing connect the pump, filters, ion exchanger, and lamps.	61

3.12	An image of the Front Muon Veto (FMV) taken from [77], before the steel tank was placed in the hall. The two layers are placed in opposite orientations and are slightly offset in height from each other to make up for the decreasing efficiency of the scintillator paddles at far end and cover any gaps between layers.	64
3.13	The MRD is an iron-scintillator sandwich held by an aluminum support frame. The scintillator paddles are coupled to PMTs and alternate between horizontal and vertical orientations. Diagram from [77] and picture from [67].	65
3.14	The monitoring toolchain used to monitor the subsystems of the ANNIE detector. The toolchain runs continuously to post up-to-date conditions for shifters. . . .	69
3.15	A schematic of Fermilab accelerators. The Pre-Accelerator, Linac, and Booster make up the Proton Source, which converts a beam of H^- ions into an 8 GeV proton beam. The protons are sent to several destinations, one of them to make the Booster Neutrino Beam.	71
3.16	Diagram of the Booster Neutrino Beam target system at Fermilab. The Booster delivers an 8 GeV beam of protons to a beryllium target, creating pions and kaons that are focused with a magnetic focusing horn. The mesons decay into muons and muon neutrinos. The muons are absorbed and what is left is a muon (anti)neutrino beam.	71
3.17	The Booster Neutrino Beam (BNB) energy spectrum shows that peak neutrino production occurs around 700 MeV. Figure from [83].	72
3.18	Neutron multiplicity measurements by Super-K [86] (left) and SNO [87] (right) experiments. Measurements were made with atmospheric neutrinos, whose sign and direction are not well-known. In addition, neutron yields are characterized as a function of visible energy.	76
3.19	A typical event in ANNIE. (a) A neutrino-nucleus interaction creates a muon that undergoes Cherenkov radiation as it travels out of the water tank and ranges out in the MRD. (b) Final state neutrinos are created and thermalize (few μs). (c)-(d) Neutrons capture on gadolinium and produces flashes of light to be observed by photodetectors.	77
4.1	TOP: A selection of Hamamatsu photomultiplier tubes (PMTs) of varying shapes and sizes [93]. BOTTOM: PMTs are vacuum tubes that can detect a single photon through the photoelectric effect. A layer of photocathode material, typically a bialkali, is deposited on the underside of the glass window. When a photon is incident on photosensitive area, it may be converted into an electron. This electron is promptly accelerated through the vacuum to a chain of dynodes, metal objects that are coated with a layer of emissive material. The dynodes emit additional electrons when struck by incident electrons, thus creating an avalanche effect. Charge is collected on an anode at the end of the chain and is read out [94].	82

4.2	A diagram of a microchannel plate (MCP) and its principle of operation [97]. A MCP is a matrix of through-going, closely-packed channels that are 10-100 μm in diameter. The interior of each channel is coated in an emissive material so that when a particle of sufficient energy is incident on the channel, an avalanche of electrons follows. A voltage is applied across the plate to accelerate the electrons towards a charge collection plane. A single plate has a gain of 10^3 - 10^4 and several plates can be stacked to achieve higher gain values. MCPs can be 10s to 100s of millimeters in size and can be cut into any shape (often round disks).	85
4.3	A cross section of the Chevron microchannel plates (MCPs) in a Large Area Picosecond Photodetector (LAPPD) [99]. When an incoming photon is incident on the top window, the photocathode on the inner surface of the window absorbs the photon and emits an electron. Like the PMT and the MCP detectors, the electron is accelerated through the vacuum volume by a strong electric field. The photoelectron initiates a cascade of secondary electrons after it enters the first MCP. The cascade then enters the second MCP after traversing the inter-plate gap and electrons are further multiplied before collection by the anode plane. The many micron-sized channels of the MCPs provide sub-centimeter spatial resolution compared to conventional PMTs. The thickness of the entire stack comes out to about 2 cm.	87
4.4	A plot depicting the transit-time spread (TTS) of a LAPPD. The y -axis is the number of pulses. Transit-time is the amount of time that passes from the creation of the initial photoelectron to the collection and readout of the resulting charge avalanche by the anode. For a LAPPD, this time is affected by the paths that electrons can take inside the MCP pores and noise from the readout electronics [101]. This TTS plot is made by shining a PILAS laser with < 45 ps root-mean-square resolution at an LAPPD, and taking data with a 10 gigahertz (GHz), 10 GSample/s oscilloscope and PSEC4 waveform digitizing electronics [100]. A timing resolution of 64 ps was observed in the main peak. The slight secondary peak to the right is due to after-pulsing.	89
4.5	LEFT: A distribution of the Δr values. RIGHT: A cumulative distribution of the Δr , normalized to a percentage of successfully reconstructed events [66]. Blue curves represent reconstructed events using 128 PMTs, while red curves represent reconstructed events using 128 PMTs + 5 LAPPDs. Calculations were performed on the same sample of 1800 events. The addition of just 5 LAPPDs improves the vertex reconstruction resolution by a factor of about three, from 38 cm to 12 cm (gray lines). Vertex reconstruction efficiency is also improved to 90% from 85%.	90

4.6	A diagram of the full LAPPD system. The LAPPD, along with a trigger board, two ACDC boards, and a LVHV board are encased in a watertight housing. The housing is machined from a block of gray PVC and has an acrylic front window and stainless steel backplate. Two waterproof cables that carry power, data, communication, and alarm lines connect the LAPPD package to a breakout box that manages power and signal distribution. Data signals are directed to an ANNIE Central Card (ACC), which sends the beam gate, PPS (pulse-per-second), and clock signals to the ACDC boards downstream. From the ACC, data is saved into event stores by the ANNIE DAQ.	91
4.7	This LAPPD (#63) is still attached to a circuit board with 60 SMA connectors (one highlighted by a green arrow): 56 for both ends of the 28 anode striplines + four spares for trigger signal inputs. In this picture, an SMA cable (light pink) is connected to one of the 28 output connectors on one half of the circuit board. If all 56 SMA connectors were used, it would be a cabling nightmare! Red arrows indicate power inputs.	93
4.8	The ANNIE Analog Pickup Board, designed by Dr. Bernhard Adams. This pickup board replaces the 56 SMA connectors with two high channel density SAMTEC connectors (white arrows) for signal readout. Traces for each stripline can be seen on the two halves of the board. Four SMA connectors (red boxes) remain as spare input/output channels for trigger signals. Mounting holes in the LAPPD frame are indicated by the blue arrows. Bernhard was quite fond of photographing dragonflies, hence the dragonfly on the pickup board.	94
4.10	A picture of the ANNIE Central Card (ACC). A +5 V DC input enters from the top of the board (white box) and powers the Altera Arria-V FPGA (yellow box) and other peripherals on the board. The ACC can control up to eight ACDC boards via CAT6 cables that connect to the block of RJ45 connectors at the bottom of the board (blue box). The LVDS signals are processed by microchips (green box). External signals such as the beam gate, PPS, and clock signals are input via the connector in the bottom left corner (red box).	99
4.11	LEFT: A diagram depicting how the LAPPD electronics boards are connected to each other. Data flows from the ACDCs to the LVHV board (passive), over the waterproof cables and through the breakout box, and finally to the ACC. A Raspberry Pi with LAPPD data acquisition (DAQ) software acts as an interface between the ACC and the main ANNIE DAQ. Note the diagram is not to scale. RIGHT: An image of the ACC board with two ACDC boards (not pictured) connected via ethernet cables. The red box highlights the input port for external clock, PPS, and beam gate signals.	100

4.12	A trigger card (1) connected to an LAPPD's anode terminals via spring-loaded contact pins ("pogo pins"). The LAPPD frame must be opened up before installing or removing the card. The output (2) of the trigger board is sent to two of the four spare SMA channels on the Analog Pickup Board and is readout with the data by the ACDC boards. Power and communication lines (3) are found at one end of the board.	101
4.13	A basic circuit of one of the channel comparators on the LAPPD trigger board. The (+) input is biased to 1.25 V so that a DAC0 output of 1.24 V is consistent with a waveform threshold of -10 mV. Figure from [108].	102
4.14	Version 2 of the LVHV board for LAPPD-40. Some important peripherals are labeled. The LVHV board provides the high and low voltages needed to power the LAPPD and the trigger card and ACDC boards attached to it. It also monitors the temperature, humidity, and light levels inside the housing. The PIC microcontroller serves as the main point of communication between the housing and surface-level electronics. Users communicate with the PIC via CAN communication.	104
4.15	The LAPPD watertight housing comprises a 13-inch by 17-inch by 2-inch PVC frame, a UV-transparent acrylic window, and stainless steel backplate. O-rings sit in grooves under the front window and backplate to provide the watertightness needed for the LAPPD system. The backplate offers some level of heat dissipation.	106
4.16	A PVC frame with a LVHV board mounted inside. The LAPPD Ultem frame can be bolted to the four protrusions ("nubs") indicated by the white arrows. The MacArtney SubConn socket connectors screw into the holes indicated by the red arrows. The top of the housing is where the LVHV board is mounted. One of the six nubs (underneath the LVHV board) needed to be removed as it was in the way of the power module of the LVHV.	107
4.17	A 10-ft long deployment panel machined from gray PVC, ready for deployment testing. The housing is attached to the panel at the stainless steel backplate with washers, nuts, and bolts. The waterproof cables are secured with stainless steel zip ties. Materials that make up LAPPD components (e.g., PVC, cable jacket) were carefully selected for chemical compatibility with Gd-loaded water. Before deployment in the tank, everything needs to be dust- and contaminant-free. . .	109
4.18	Four polypropylene sliders (shown by white arrows) attach to the back of the LAPPD deployment panel. Rectangular grooves were machined into the sides of each block so fit into the eight rails of the octagonal inner structure.	110
4.19	LEFT: A MacArtney SubConn Power Ethernet Circular 13-pin connector (blue section) is spliced to the Falmat XtremeNet™ Underwater Cable (neon green). The spliced region (indicated by the white line) is made with crimp connectors, heat shrink tubing, electrical tape, and splicing tape. RIGHT: A waterproof cable spliced with the original MacArtney SubConn 12-pin connector. The cable is about 10 m, or 33 ft, in length.	112

4.20	Inside the breakout box are two sets of electronics and components that service two LAPPDs: (1) an input voltage of +12 V is fanned out to the relay HATs on the Raspberry Pis before making it down to the LVHV board underwater; (2) the Raspberry Pi power supply also receives the +12 V input to power the Pis with +5 V; (3) a relay HAT and a PiCAN 2 HAT sits atop a Raspberry Pi that runs a routine to fetch slow control measurements; (4) power, CAN bus, data, sine wave, and alarm lines are routed through terminal blocks; (5) the LVDS data signals exit the breakout box via RJ45 jacks; (6) communication with the Raspberry Pis can occur over an Ethernet connection.	113
4.21	A single LAPPD is mounted inside a dark box. A motor-driven apparatus holds the LED and laser used to perform quantum efficiency and timing studies. . . .	115
4.22	The results from a quantum efficiency (QE) scan for LAPPD-63. Note that the QE is fairly uniform across the photocathode and comparable to the QE of convention PMTs. The dark regions at the perimeter of the LAPPD and that form an “X” are where the frame and X-spacer sit. Over time, it was realized that performing QE scans before and after extensive testing is recommended to study how the photocathode performance changes over time. After extensive studies with the laser, the average QE had dropped by about a percent.	117
4.23	LEFT: An empty watertight housing soaking in deionized water. Note that this soak test was carried out with the old 12-pin connectors. The waterproof cable is wrapped in PTFE tape where the MacArtney SubConn connectors mate, so that no rubber insulation is exposed due to incompatibility with Gd-loaded water. The black O-ring on the acrylic window side of the housing is indicated by the red arrow. Three humidity indicator strips were taped to the backplate inside the housing to get rough estimates of humidity levels. RIGHT: The full LAPPD system. An LED (green arrow), wrapped in plastic to waterproof it, is placed in front of the acrylic window to confirm that the LAPPD still sees light even in its housing and underwater.	121
4.24	A commercial power supply (on top) provides the +12 V to the breakout box (on bottom), which distributes the power accordingly. (1) The input voltage enters the breakout box at one port on the front panel. (2) Waterproof cables connect to Amphenol receptacle connectors. The RJ45 ports to the right of the cables are where the ACC connects. (3) Another set of connectors and ports for a second LAPPD.	123
4.25	An empty LAPPD housing attached to a deployment panel. Testing of the deployment mechanism took place in a clean tent, before the octagonal inner structure was transported to ANNIE Hall. The panel slides along the rail via polypropylene sliders bolted to its back. All eight panels were tested to make certain that they slid down the rails smoothly. Sliders that did not go down a track smoothly naturally had pieces of polypropylene shaved off by the track. An ANNIE collaborator can be seen behind the inner structure holding the deployment panel.	125

4.26	An increase in events with a width of $1.6 \mu\text{s}$ confirms that the LAPPD does in fact see the neutrino beam.	127
5.1	Distributions of PMT operating voltages (left) and peak-to-valley ratios (right). Several PMTs in the purple distribution, representing ETEL type PMTs, exhibit lower operating voltages compared to the green distribution of WATCHBOY PMTs. ETEL PMTs also have lower peak-to-valley ratios. Given that the ETEL PMTs were more sensitive to operating conditions, they were installed under the tank lid and thus experienced less pressure underwater.	130
5.2	Distributions of dark noise (left) and current (right). WATCHBOY PMTs (green) have low dark noise compared to ETEL PMTs (purple), which have a wide range of dark noise levels. The current draw for ETEL PMTs also vary more than WATCHBOY PMTs.	131
5.3	An example pulse charge distribution curve fitted with a double Gaussian. The first visible peak is the pedestal and is slightly greater than zero charge. The second visible peak is the single photoelectron peak. This charge represents the gain-amplified photoelectron.	132
5.4	Six LED models are installed on the inner structure of the detector. Their locations cover all of the PMTs so that gain calibration and water transparency monitoring can be performed while the experiment is running.	134
5.5	A distribution of initial gain values on the left [112] and updated gain values on the right [113]. The PMT gains were recalculated after an update to the hit finding procedure. The average gain has a value of 9×10^6 compared to the previous value of 7×10^6	135
5.6	An event with a through-going muon. Note the missing paddle hit. The long and short (dashed) lines represent the horizontal or vertical orientations of the scintillator paddles, depending on the perspective view. Double dashed lines represent the two halves of the paddle layers. The iron plate absorbers fill the gaps between scintillator layers. Figure from [77].	137
5.7	Most MRD paddles have efficiency values of at least 80%, as high as 95%. Paddles with low efficiency or are turned off will be replaced in future maintenance during beam downtime. Figure from [114].	138
5.8	Front Muon Veto (FMV) paddle efficiencies. The average paddle efficiency is about 80%, ranging from 65% to 90%. Figures from [114].	139
5.9	Initial timing calibration of the tank PMTs using LED light sources. Note the different types of PMTs are grouped together due to similar cable lengths and positions in the tank. Measurements are accurate to about 10 ns. Figure from [112].	141
5.10	Timing calibration of the tank PMTs using a laser light source. Figure from [115]. Different colors represent PMTs at different locations.	141

5.11	With the timing offsets measured using the laser diffuser ball, the structure of the Booster Neutrino Beam is now visible. The top left plot is a distribution of hit clusters whose cluster times fall in the prompt $2\ \mu\text{s}$ acquisition window. The $1.6\ \mu\text{s}$ beam spill can be seen atop a flat background (above the dashed red line). Before the timing corrections, the beam structure is not visible. The top right plot shows the timing distribution of a sample of dirt muon clusters after the timing corrections from laser calibration were applied. The bottom two plots are zoomed in versions of the top right plot to highlight the beam structure. Figures from [115].	143
5.12	A record of transparency measurements. Significant changes in the transparency can often be attributed to an external event, usually maintenance on the water filtration system.	144
5.13	LEFT: AmBe source in acrylic housing. A deployment mechanism was being tested here. CENTER: AmBe source housing wrapped in teflon (PTFE) tape to maximize light collection. RIGHT: AmBe source deployment winch being tested. This was the method selected for deployment.	146
5.14	TOP: A schematic of source deployment and port locations. BOTTOM: Neutron capture detection efficiency as a function of tank position. This map was created when the BGO crystal was coupled to two SiPMs. A new map using data taken with a PMT is in progress. Diagram from [112].	150
6.1	This diagram shows the tools that are involved in various analysis toolchains (blue). The tools relevant to event building starts at <code>LoadGeometry</code> and ends at <code>ANNIEEventBuilder</code> . This toolchain is where data from each subsystem are matched to the same physics event. Other toolchains can produce event displays, root files, and run validation plots.	153
6.2	A schematic of the timestamp matching steps of the <code>ANNIEEvent</code> building process. Raw data from each subsystem is matched according their timestamps to form event entries that belong to the same physical process. Any unmatched data go to an <code>OrphanStore</code> . Figure from [118].	155
6.3	The pre-trigger region of the above waveform is divided into M regions with N samples. The darker regions were used to estimate the baseline, while the lighter regions are where the baseline subtraction is applied, in addition to the darker regions. Note the sinusoidal aspect of this particular waveform that comes from power supply noise. For the future, a sinusoid may be a better fit to estimate the baseline. Figure from [112].	156
6.4	Cellular automation method for MRD track reconstruction. (a) Hits in the same layer are grouped into clusters. (b) Clusters in neighboring layers form cells. (c) Each cell is assigned a status value. (d)-(e) Combinations of cells are made according to status until the end of a track is reached. (f) Linear fits are made to track candidates. Figure from [123].	161

6.5	Difference between reconstructed and true muon track angle for events that pass the FV cut. The mean difference is -0.026 rad with a σ of 0.16 rad. The error on the reconstructed track angle will be propagated through to the muon energy reconstruction.	162
6.6	One hundred muons with initial kinetic energy of 500 MeV were fired into a MRD-like geometry in Geant4 . Typical paths are not linear and demonstrate the effects of scattering. Actual path lengths are longer than the path length determined by fitting a line to the muon entry and stopping points, which is the current method in existing analysis tools.	164
6.7	Ratios of a linear (“straight”) track to the actual track lengths as a function of initial muon kinetic energy. The linear track is calculated by taking the distance between the entry point and the stopping point. For lower energy muons entering the MRD, the track actual track length is longer than the track length determined by connecting the entry and stopping points. Results generated from Geant4 simulations. Corrections are made to the track length depending on initial estimates of the muon energy. Data points beyond 800 MeV did not have a muon stopping in the MRD geometry in the simulation.	165
6.8	LEFT: MRD cluster times vs. tank cluster times. One can see a correlation between hit clusters that come from beam events. RIGHT: Difference between MRD and tank cluster times. Correlated activity has a time difference of ~ 750 ns. The accepted time window if (700 ns, 800 ns) is marked with black lines. Note the logarithmic scale in the number of clusters. Figures from [112].	167
6.9	Given a muon vertex, we can define the variables above. a_i represents the tank segment between the location of the muon and where it exits the tank. R_i is the vector from the tank exit point to PMT i . b_i is the vector from the vertex to PMT i	168
6.10	A muon that travels some distance Δx will emit a constant number of photons. These photons will make up a section of the Cherenkov cone called a frustum, whose surface area is imaged as a ring on the photodetectors. Left figure from [125].	170
6.11	This is a photon density η vs. tank track segment a_i plot with truth information shown. To the left and right of the true tank track length (vertical blue line), there is a clear distinction in the average photon density at shorter track segments (left horizontal red line) and at longer track segments (right horizontal green line).	172
6.12	For events with a clear break in the photon density, a manual fit is easily made. For this event, a tank track of 150 cm will be the fit. The true tank track length is 148 cm.	173
6.13	TOP LEFT: An event display of a <i>beam</i> event. TOP RIGHT: A η vs. a_i plot. Here one would estimate that the track segment is about 190 cm. BOTTOM LEFT: An event display of a <i>simulated</i> beam event. BOTTOM RIGHT: A η vs. a_i plot. Here one would estimate that the track segment is about 250 cm.	174

6.14	A distribution of differences between the manually fitted and true track lengths in water. The mean difference is -12 cm, with a standard deviation of 30 cm and RMS of 32 cm.	175
6.15	A Recurrent Neural Network (RNN) was trained on η vs. a_i data and truth information to develop a model to fit the track length in the water tank. Compared to Figure 6.14, this model performs well.	176
6.16	Once the track length L_{track} in the detector tank is determined, the muon vertex (blue star) is found by moving backwards into the tank from the tank exit point (red star) by a distance equal to the fitted track length.	177
6.17	Reconstructed muon vertices and energy in beam data (no FV cut). The x and y distributions are centered around 0 cm, while the z distribution has a slight shift towards the upstream region of the tank. The reconstructed muon energy is centered around 800 MeV, which agrees with expectations based on the BNB energy spectrum.	178
6.18	A fiducial volume (FV) cut is made on the reconstructed vertices. The left plot is a topdown view of the detector, while the right is a side view of the detector. The red outlines the main FV cut that is shifted upstream from the center. The blue outlines a downstream FV cut used for neutron containment studies in a later section. The green outlines the active volume.	179
6.19	The top two graphs demonstrate the performance of vertex reconstruction in ANNIE simulations with and without LAPPDs. The bottom two graphs demonstrate the performance of the vertex reconstruction method in this thesis. The Δr distribution corresponds to the blue curve in the distribution above it. The vertex reconstruction method performs similar to the scenario of 128 PMTs without LAPPDs deployed. In the cumulative distribution, the ANNIE reconstruction methods indicate a spatial resolution of about 40 cm, while the method presented in this thesis has a spatial resolution of 60 cm.	180
6.20	A comparison of reconstruction energy to the true energy of the muon. A Gaussian curve is fitted to the sample of events after a fiducial cut is made on the reconstructed muon vertices. The mean energy difference is about -15 MeV, meaning the energy reconstruction algorithm is slightly underestimating the muon's kinetic energy. The width of this fit is about 66 MeV. Events with large differences are mainly attributed to pions or events with partially missing truth information, e.g., the track length in the MRD was zero.	182
6.21	This plot visualizes the cuts made to select neutron cluster candidates. The lower left quadrant is where candidate neutron clusters will be taken. The cuts were selected based on the first full AmBe source calibration campaign. AmBe data is available for a more recent campaign, but analysis is ongoing. Once completed, new selection cuts should provide a better sample of neutron cluster candidates. Figure from [114].	184

6.22	A distribution of neutron candidate cluster times is fitted with an exponential function. The time constant of this exponential is $(33 \pm 2) \mu\text{s}$, which is close to the expected value of $\tau \approx 30 \mu\text{s}$, but is slightly higher. Further investigations are needed to understand this shift to a longer time constant.	185
6.23	Distributions of events according to number of neutrons in the final state, following a neutrino-nucleus interaction. The blue curve represents all neutron candidates, the black curve are the number of candidates with a stopping MRD track, and the red curve are the remaining candidates with a FV cut. The majority of events have no neutrons in the final state. The right plot is the left plot in logarithmic scale.	187
6.24	The number of neutrons as a function of visible energy of the muon for data with a stopping MRD track (left) and with a fiducial volume cut (right). Most events produce zero to two neutrons. Note that all plots are in logarithmic scale in the z-axis.	189
6.25	The average neutron multiplicity as a function of visible muon energy (top two plots) and momentum transfer (bottom two plots). Compared with results from Super-K and SNO, ANNIE sees a lower neutron yield. The left plots are measurements without any selection cuts applied. The right plots have the FV cut and neutron detection efficiency applied.	190
6.26	Efficiency maps for MC (left) and data (right), the latter of which was reproduced from Figure 5.14. Simulations of the AmBe source deployment were performed to produce the efficiency map on the left. The efficiencies at Ports 1 and 5 between -50 cm and 50 cm heights were averaged and applied to the neutron yield measurements in this work. The average efficiency is 65% for data and 76% for MC. Future studies can apply efficiencies dependent on location. Figure from 114.	192
6.27	Comparison of upstream (left) and downstream (right) neutron multiplicity as a function of muon energy E_μ . Both plots are efficiency-corrected.	194
6.28	Average neutron multiplicity as a function of Q^2 , after combining the upstream and downstream fiducial volumes. All cuts and efficiency corrections have been applied.	194
A.1	Modifications to the housing (left) were needed for the new LAPPD waterproof connectors due to different threading and socket diameter as seen in the right image.	204
A.2	TOP: Falmat cables (neon green) re-spliced with new connectors (blue). BOTTOM: Heat shrink tubing is placed over the splice to waterproof it.	206
B.1	Slow controls routine. Users communicate with the LVHV board via Raspberry Pis. Checks for relative humidity, temperature, and voltages are performed.	209

List of Tables

2.1	A table of the Standard Model particles, with their representations in the gauge groups $SU(3)_C$ and $SU(2)_L$, and their respective weak hypercharge Y_W and electric charge Q . The top half of the table shows the quantum numbers for fermions and the bottom half shows the quantum numbers for bosons. Here, the definition $Q = T_3 + \frac{1}{2}Y_W$ is used.	14
2.2	Current values of neutrino oscillation parameters with errors for both the normal (NO) and inverted (IO) mass orderings [58]. For Δm_{3l}^2 , $l = 1$ for NH and $l = 2$ for IH.	38
3.1	Examples of media of different refractive indices n and Cherenkov angles θ_C .	56
3.2	The different types of PMT installed in the ANNIE Phase II detector. Listed for each PMT type is the PMT model, diameter, average quantum efficiency, and number installed.	63
6.1	A summary of the average neutron multiplicity with various cuts applied.	190
A.1	A list of known issues of LVHV board v1 and proposed solutions.	201
A.2	A list of issues with v2 of the LVHV board and their solutions.	202

Acknowledgments

Many people have entered my life and made it all the more an enriching experience. Whether it was on campus, on the plane, or at the Users Center, thanks for stopping by. Please excuse my fried brain, if I forget to include you, for these past months have been a wild ride.

First, I want to show my gratitude and appreciation for my advisor Bob Svoboda. I am fortunate to have found a place in your research group as I was figuring out what I wanted to do after undergrad. Thank you for taking me in and showing me the possibilities of grad school. You are always joyful, witty, and full of grand ideas, and I will miss the stories you tell during long car rides to collaboration meetings. You gave me incredible opportunities to develop skills I never thought I would have and gave me enough independence to do my own thing. The last few years have been a struggle, but I am thankful for your infinite patience and a simple enough analysis task so I could graduate faster.

Next, I want to thank the rest of my committee, Eric Prebys and Mike Mulhearn, for taking time out of their busy schedules and agreeing to review this thesis. Thank you Eric, for your patience, feedback, and contributions to this work. I appreciate the fun tidbits left in your comments. Continuing on with professors and others in the department, I want to thank Robin Erbacher and Emilija Pantic. Both of you are inspirational powerhouses to women in physics. I am grateful for the concern that you have shown me and other students. It is much needed in the stresses of grad school. I also extend my gratitude

to Tracy Lade, Amy Folz, Steve Carlip, and Jaroslav Trnka for supporting my efforts in creating an undergraduate peer mentoring program through TEAM-UP (for which I thank Brian Lemaux for believing in me to carry on). Your guidance and participation motivated me to keep going. And of course, thank you to Dr. Randy Harris, whose gentle instruction and solid foundation in physics encouraged me to pursue a degree in this field. You made physics accessible to us all. On a similar note, I am grateful to the McNair Scholars Program for showing me grad school as a potential career path. Because of McNair, I got to work with Maruša Bradač, whom I thank for her patience as I learned to conduct research and for sending me to Hawai‘i.

There are many talented and wonderful people I met during my time as an ANNIE collaborator whom I am fond of and will miss dearly. First, thank you Emrah Tiras for starting all our on-site meetings on ET (Emrah Time). I appreciate your constant positivity and infectious laughter despite having to work late hours preparing the inner structure for deployment. Thanks to Lena Korkeila for joining me on an adventurous road trip to Chicago, delivering PMTs and saving them from heists, and Michael Nieslony, David Maksimović, and Pratyush “Reik” Das (honorary ANNIE member) for a fantastic summer at Fermilab, full of shenanigans, bowling, and escape rooms. Thank you Carrie McGivern, for all the ANNIE swag that you’ve given us. I will rock the stickers and beanie forever. I am glad to have met and worked with Gian Caceres and Paul Hackspacher, both great postdocs and quite silly at times. I wish you both the best on your job searches! Also from Bob’s group is Tianqi Zhang. Thanks for the ML script for this thesis. Hope I get to meet your cat one

day. I welcome the newest additions to ANNIE: Jaydip Singh (thank you for helping with my analysis) and Dylan Fleming (thanks for offering to fight during my breaks). Thank you Amanda Weinstein, for coralling the LAPPD cats and being a huge and important part in getting the LAPPDs to where they are today. I will always be in awe of your brilliance and meticulous attention to detail. I hope, one day, all our ducks will be in a row. And thank you Matt Wetstein for showing me that professors can still be silly and do great work. Your excitement and passion are unmatched. Thanks to Marvin Ascencio, for being such a great postdoc and helping me take pictures and make measurements in the hall. You deserve the longest vacation ever. To Emily Pottebaum (5/7 memes), Marc Breisch (thanks for all the side conversations during Zoom meetings), Rory Edwards (we will roll one day!), Rob Foster, Yue Feng, James Minock, Steven Doran, Franklin Lemmons, Noah Goehlke, Amala Augusthy, Dhaval Ajana, and Adil Aman: though our time together was short, I am glad I got to meet you all and I am excited to see what is in store for you next.

Moving a bit out of the ANNIE collaboration, many thanks to Bernhard Adams and Paul Stucky for teaching me all I know about electronics. It still amazes me how such tiny devices can do so much. Both of you are a fountain of knowledge and I hope I will have even just a fraction of it. Thank you Dave Hemer for letting me hang out with you at the machine shop and feed peanuts to the squirrels and birds while you work on parts for us.

In the time between undergrad and grad school, I had the pleasure and honor of meeting Lenore. I am glad I got to spend the last few years of your life with you, watching movies and sharing thoughtful chats about the going-ons in the world. I am now an amateur movie

enthusiast because of you. Thank you for the stories of your young adulthood. Your sons, Paul and Marc, are equally delightful and I am glad I got to meet them too.

Anyone who knows me knows that I have been training in the martial arts for a while. Thank you Master Bob, Louman, Kitty, and Constante for teaching me lethal, but fun, joint locks, throws, and especially, the kicks. I look forward to training with you all soon, and becoming a fraction of the badasses that you all are. Thank you to Coach Hector for shaping me into a better fighter with fast reflexes, and to MaeLynn for being the best sparring partner a boxer could ask for and testing the limits of my cardio. I'd also like to give a shoutout to Pedro and the BJJ team at the ARC. To all the wonderful people I got to roll with at B1, thank you for making Illinois feel like a second home. I look forward to sparring with you all again. Thanks to Professor Lalo Ortega and my team at Element Martial Arts for pushing me to the next levels. Although this PhD has kicked my butt, you will all get to continue kicking mine (and vice versa!).

I want to thank the people I met at the Energy Conservation Office on campus, especially to Daniel Imperiale, Hiroko Masuda, Nico Fauchier-Magnan, and the Energy Group Grads. Thank you for supporting me in my endeavors to enter the energy industry. Daniel, I have never felt such genuineness from anybody in wanting to see me succeed than I have from you, and thank you for taking the time to help me with my interviews. For that, I am forever grateful to you. To my R&D team at SMUD, thank you for showing me the innovative side of the energy industry. Thanks Denver, Deepak, Carol, Aaron, Adrian, Tyler, Dago, Josh, Tasha, Terry, James, Matt, and Biju (phew!) for making my time at SMUD well worth it.

I'm looking forward to working with you all in the future. I extend my deepest gratitude to my coworkers at Energeia for their patience and understanding as I finish my PhD.

To my "OG Peeps" Linda, Judy, and Vu: it's been a struggle to get us all in one location, but it's always fun when we do manage to get together. I look forward to the girls trips we've put on hold. Shoutout to my "OG Roommates" Sally and Sherry. Thanks for letting me rant to y'all about anything and everything and for having my back. Let's travel together soon! I'd like to express my immense appreciation for Ricardo, for being a pillar of support and making me a better person. We've had our rollercoaster rides, but I am glad you are still in my life.

Of course, you can't do grad school alone. I am proud to say that I have the best cohort a grad student could ever ask for. I want to especially thank Karthik Prabhu, Tyler Erjavec, and Yash Bezawada for all the silly and thoughtful conversations and tea and for making grad school a little more bearable. And to Patty Bolan: hey buddy! We didn't get to hang out often, but I'm glad you got to be part of our cohort. You are an amazing inspiration and always unapologetically yourself. I hope to affect a fraction of the lives that you touched in your time here.

To my brother David, thanks for all the pictures and videos of your cats. I can't wait to meet them. Lastly, I cannot begin to repay my parents for all that they've done for me and my brother. Thank you for pushing us to pursue higher education so that we could have better opportunities in life. Although you never wanted to see us struggle, I admire both of your work ethics and I hope to be just as dedicated and resilient. 多谢爸爸妈妈.

Chapter 1

Introduction

Following the postulation of its existence, more than twenty years elapsed until the discovery of the neutrino. The search for these elusive, ghost-like particles was enabled by new technology in photosensors and electronics. We witnessed the innovation of large-scale neutrino detectors and detection media, such as water Cherenkov detectors and water-based liquid scintillator. Since its discovery, this new field of neutrino physics has been quite exciting and rewarding as neutrinos threw at us one surprise after another: there are several generations – or “flavors” – of neutrinos, neutrinos can oscillate between different flavors, they travel great distances before interacting with matter, and so on. Neutrinos can also be used to understand other interesting topics in physics, such as supernovae, proton decay, the Diffuse Supernova Neutrino Background, the matter-antimatter asymmetry in the universe, and our very own Sun.

Although much knowledge about this fundamental particle has been gained, questions

surrounding the neutrino remain and are challenging to answer. For instance, the absolute masses of the neutrinos, the ordering of the mass eigenstates, and the charge conjugation and parity (CP) violating phase are all currently unknown. Today, there exist many neutrino experiments that attempt to answer these questions with various approaches and technologies. One such experiment is the Accelerator Neutrino Neutron Interaction Experiment (ANNIE), a modestly sized water Cherenkov detector located at the Fermi National Accelerator Laboratory (Fermilab) in Batavia, Illinois. The experiment sits on-axis of a muon neutrino beam with and utilizes Cherenkov radiation to detect neutrino interactions in the 26-ton water target. ANNIE aims to improve the understanding of interactions between neutrinos and nuclei by measuring the neutron multiplicity (or neutron yield) of such interactions. Current models feature large uncertainties regarding the primary neutrino interaction as well as the secondary hadronic interactions that may follow. These large uncertainties affect cross section measurements which are needed for precise physics measurements, such as the extent of CP symmetry violation in the neutrino sector that may explain the matter-antimatter asymmetry of the universe. Therefore, it is crucial to develop more accurate models to describe these interactions. Along the way, ANNIE serves as a testbed for novel technologies such as gadolinium-doped water, water-based liquid scintillator (WbLS), and the Large Area Picosecond PhotoDetectors (LAPPD). Given the relatively small size of ANNIE to other water Cherenkov detectors, the fast-timing capabilities of LAPPDs are needed to achieve its physics measurements with high precision. The work at hand documents the neutron multiplicity analysis and the progress made towards the first-ever deployment of the LAPPD in

a neutrino physics experiment.

This thesis utilizes the ANNIE detector to measure the neutron yield in neutrino interactions with oxygen in exclusive interactions in the 100-500 MeV range. Oxygen is a common target in neutrino experiments, yet few measurements of the interaction cross section exist. In a typical interaction, a neutrino interacts with a neutron in the nucleus and creates a lepton in the process. Since ANNIE sits on a muon neutrino beam with a peak energy of ~ 700 MeV, the resulting lepton is a muon. Neutrons are also expected in the final state, due to nuclear effects and other final state interactions (FSIs). Without a complete understanding of these nuclear effects and a good handle on the number of neutrons and other final state particles, it is difficult to accurately reconstruct the incident neutrino energy, which is important for constraining uncertainties in cross section measurements. To better characterize the neutron yield of neutrino-nucleus interactions, the water target is loaded with gadolinium (Gd), due to its high average neutron capture cross section. Equipped with novel photodetectors and detection media, ANNIE can make a measurement of the neutron production rate as a function of muon momentum transfer.

This work starts with an overview of the current knowledge of neutrinos within and beyond their role in the Standard Model of particle physics. Chapter [2](#) also covers neutrino and nucleus interactions models and takes a deeper dive into the motivations for improving these models. This chapter ends with a summary of current values of neutrino parameters and some of the remaining challenges in the field.

As mentioned earlier, ANNIE serves as a technology testbed for a new Gd-loaded wa-

ter purification and filtration system, water-based liquid scintillator (WbLS), and LAPPDs. Water purification and filtration systems often remove contaminants and impurities from water, so ANNIE demonstrates how gadolinium can be kept in solution while removing unwanted particles. WbLS shows promise as a detection medium that can be used to distinguish Cherenkov radiation from scintillation for particle identification and ANNIE has demonstrated this capability. LAPPDs offer picosecond timing resolution that is needed for accurate vertex and energy reconstruction of the muon. A detailed description of the ANNIE detector setup and its physics goals follows in Chapter [3](#). My work in conducting material compatibility tests is also presented in this chapter.

The technology of LAPPDs explored in the hardware chapter, Chapter [4](#), along with a discussion of how they aid ANNIE in achieving its physics goals. In this chapter, I describe the work I performed in testing LAPPD electronics, putting together the LAPPD system from its subcomponents, and conducting full system tests and troubleshooting during my year at Fermilab when the COVID-19 pandemic was in full swing. A critical piece of the LAPPD system is the LVHV board that serves as a power source for the LAPPD and its auxiliary electronics and as a monitoring system. Developing and testing the initial LVHV board and designing the next board revisions constitutes much of my thesis work. My contributions to the design and fabrication of the waterproof housing and waterproof cables are also presented in this chapter.

The remainder of this thesis then covers software and analysis. Chapter [5](#) outlines the calibration of the ANNIE detector. A description of the photomultiplier tube (PMT) tests

that I performed as a junior specialist, prior to my time as a graduate student, can be found here. The method for measuring the neutron yield with PMTs and a Gd-loaded water target in ANNIE is outlined in Chapter [6](#). Here, the algorithms for data cleaning and reduction, event classification, vertex and energy reconstruction, and neutron multiplicity are described. For this thesis, an alternative method for vertex reconstruction was developed for PMT-only data. The technique involves imaging the edge of the Cherenkov disk that results from the passage of a muon through the detector. The method is included in the reconstruction paper that is in review in the collaboration at the time of this writing. This chapter ends with a measurement of neutron multiplicity as a function of momentum transfer and compares the findings to existing measurements.

Finally, this thesis closes with a summary of the work presented and an outlook for the future of ANNIE in Chapter [7](#). While the LAPPD system is discussed extensively, data from the photodetectors is not included in the analysis. Future work includes the addition of the LAPPD datastream in event reconstruction, absolute cross section measurements, and filling the detector with WbLS and Gd-loaded WbLS.

Chapter 2

Neutrinos

The neutrino is one of the fundamental particles of our universe. It is a neutral fermion of extremely light mass and there exist three flavors¹ of neutrinos, much like their lepton family counterparts, the electron (e), muon (μ), and tau (τ). Although much is known about the neutrino today, with a whole branch of physics dedicated to the study of it, this was not the case several decades ago. Neutrinos have quite the history and it all starts with the beta decay reaction:

$$n \rightarrow p + e^- + \bar{\nu}_e \tag{2.1}$$

In 1914, James Chadwick reported a measurement of a continuous energy distribution while studying beta decay with a magnetic spectrometer and Geiger counter [1]. Since the α decay spectrum was found to be a discrete energy distribution, an indication of a simple two-body decay, it was a surprise that the beta decay spectrum was not equally discrete. Some

¹As far as we know.

postulated that the beta particle (electron or positron) was in fact emitted with discrete energies, but lost varying amounts of its energy from final state interactions [2]. Neils Bohr even suggested that the energy conservation could be violated and that energy in beta decays only needs to be conserved on average [3].

As a “desperate remedy”, Wolfgang Pauli postulated the existence of a new particle, which he dubbed the “neutron”, in his famous letter² to a nuclear physics conference in Tübingen, Germany, in 1930 [4]. In this same letter, Pauli mentions the spin statistics problem of nitrogen and lithium, and that this theoretical particle would be the solution. Before the discovery of the modern neutron, scientists were perplexed that nitrogen had an integer spin, despite having an odd number of nuclear constituents (14 nucleons plus 7 electrons³). Pauli postulated his “neutron” to be a fermion as well, with spin- $\frac{1}{2}$ to resolve this issue of nuclear magnetic moment, and a mass similar to that of the electron.

Pauli’s “neutron” did not become the neutron that we know today, which was discovered by Chadwick in 1932 while examining electrically neutral radiation from beryllium after irradiation with alpha particles. This discovery got Chadwick the Nobel Prize in 1935. The modern neutron was far heavier than the neutral particle needed to avoid detection in beta decay, prompting Enrico Fermi to adopt a new name given to Pauli’s hypothesized particle, the “neutrino”. While it was only a hypothesis, Fermi turned the idea into a successful theory of beta decay in 1934, in analogy to quantum electrodynamics (QED) [5, 6]. To

²An English translation of the original letter [4] (in German) can be found [2].

³At the time, “nuclear electrons” were considered as part of the nucleus to explain why nuclei were heavier than the sum of the protons that make up its charge.

explain the one-unit change of nuclear spin in some β decays, George Gamow and Edward Teller extended Fermi's theory by introducing axial vectors in a way that still conserved parity [7, 8]. At the time, parity violation was still unthinkable.

As successful and useful as Fermi's 1934 theory on weak processes was, there were still some issues that needed addressing, such as the lack of renormalizability. Much of the theoretical framework for our modern understanding of neutrino interactions (and weak processes in general) was set up by Sheldon Glashow [9, 10], Abdus Salam [11, 12], and Steven Weinberg [13] in the late 1950s and 1960s. The Glashow-Weinberg-Salam (GWS) model, which won them the 1979 Nobel Prize in Physics, provides a unified treatment of electromagnetic (EM) and weak interactions. Under this model, the Lagrangian governing the electromagnetic and weak interactions is invariant under symmetry transformations of the local gauge group $SU(2) \otimes U(1)$.

It took 26 years after its postulation until scientists confirmed the existence of the neutrino. The first indication was at the Hanford Site in 1953 and the final discovery was made at the Savannah River Plant in 1956 as part of "Project Poltergeist" [14]. The newly discovered neutrino was actually the electron antineutrino and it was detected via Inverse Beta Decay (IBD)

$$\bar{\nu}_e + p \rightarrow e^+ + n \tag{2.2}$$

using a cadmium-loaded liquid scintillator target and a nuclear reactor as a source of antineutrinos. The scintillator was used to detect flashes of light coming from positron annihilation

and neutron absorption by cadmium. The discovery of the neutrino won Clyde L. Cowan and Frederick Reines the 1995 Nobel Prize in Physics.

Later, it was discovered that there were other flavors of neutrinos. The muon neutrino was discovered in 1962 at the Brookhaven National Laboratory with the brand new Alternating Gradient Synchrotron (AGS) [15]. The AGS used a beam of protons to produce pions (pions) that traveled 70 feet toward a 5,000-ton steel wall made of old battleship plates. The pions decayed into muons and neutrinos, but only the neutrinos could pass the steel wall into a neon-filled spark chamber detector. The neutrinos interacted with the aluminum plates and produced muons that left spark trails that could be detected and photographed, thus providing evidence for the existence of muon neutrinos. Leon Lederman, Melvin Schwartz, and Jack Steinberger were awarded the 1988 Nobel Prize in Physics for their work on demonstrating the difference between neutrino flavors. When the tau lepton was discovered in 1975 [16], a corresponding neutrino was expected to exist. However, experimental challenges prevented its discovery until the year 2000, when the DONUT (Direct Observation of the Nu Tau) collaboration reported the observation of four tau neutrino interaction events in a nuclear emulsion detector [17]. Their results confirmed the results of earlier experiments at the Stanford Linear Collider (SLC) at SLAC⁴ and Large Electron-Positron Collider (LEP) at CERN⁵, which all indirectly pointed to the existence of three active neutrino flavors with masses lighter than half the Z boson mass [18, 19].

⁴Stanford Linear Accelerator Center

⁵Conseil européen pour la Recherche nucléaire (European Organization for Nuclear Research)

Just as we thought we were beginning to understand the neutrino, new puzzles popped up. For one, experiments consistently reported a ν_e flux from the sun that was much less than the number predicted by solar models at the time. In 1968, the Homestake experiment found the neutrino flux of the sun (at which point it was known that the sun produced neutrinos) to be one-third of the expected amount [20]. Similar deficits were reported by KamiokaNDE, SAGE, and GALLEX experiments⁶, and came to be known as the “solar neutrino problem”. Studies of atmospheric neutrinos also saw deficits in neutrino flux, such as in the IMB, NUSEX, and Super-Kamiokande experiments⁷. These experiments found an unexpectedly low ratio of atmospheric ν_μ to ν_e events [21]. This was labeled the “atmospheric neutrino problem”.

Although neutrino oscillations were first anticipated theoretically in the late 1950s and 1960s, direct evidence for neutrino changing between flavors would not come until the 1990s and 2000s. Evidence that convinced the neutrino community that flavor change and oscillations were the cause of the solar and atmospheric neutrino anomalies came from the Super-Kamiokande (Super-K) and Sudbury Neutrino Observatory (SNO) experiments. In 1998, the Super-K collaboration published an analysis of the atmospheric neutrino flux for both electron and muon flavors and found a directional asymmetry in the number of detected atmospheric ν_μ events, establishing a disappearance of muon neutrinos if they passed through the earth before making their way to the detector [22].

⁶KamiokaNDE: Kamioka Nucleon Decay Experiment; SAGE: Soviet-American Gallium Experiment (sometimes Russian-American Gallium Experiment); GALLEX: Gallium Experiment

⁷IMB: Irvine-Michigan-Brookhaven; NUSEX: Nucleon Stability Experiment

Starting in 1999, the SNO collaboration investigated the solar neutrino problem and found that the flux from the sun was consistent with predictions if all neutrino flavors were taken into account [23]. SNO reported a measurement in 2002 of the relative rates of the charged current (CC) and the neutral current (NC) reactions induced by the solar neutrinos within their heavy water-based detector.

$$\begin{aligned}\nu_e + d &\rightarrow p + p + e^- \text{ (CC)} \\ \nu_x + d &\rightarrow p + n + \nu_x \text{ (NC)}\end{aligned}\tag{2.3}$$

Because only the ν_e component of the flux could participate in the CC reaction and all neutrino flavors ($x = e, \mu, \tau$) could participate in the NC reaction, a simultaneous measurement of the two reaction rates allowed a determination of the total solar neutrino flux and the ν_e fraction. The results of Super-K and SNO collaborations offered compelling evidence, under the assumption of neutrino flavor change taking place, and provided an elegant solution to these neutrino anomalies. To add to their findings, the Kamioka Liquid Scintillator Antineutrino Detector (KamLAND) experiment measured the energy and distance dependency of all neutrino flavors from various nuclear reactors [24]. The observations from these experiments indicated that neutrinos possess mass, albeit very small, which leads to quantum mechanical effects that facilitate flavor transitions due to the differing mass and flavor eigenstates. This results in phenomena such as flavor oscillations, confirming that neutrinos do in fact undergo oscillations and that these oscillations are responsible for the atmospheric neutrino deficit and adiabatic conversions driven by Mikheyev-Smirnov-Wolfenstein (MSW) effect, causing the solar neutrino deficit [25]. The 2002 Nobel Prize was awarded to Homestake's Raymond

Davis, Jr. and KamiokaNDE’s Masatoshi Koshihara for finding the deficits, while the 2015 Nobel Prize was awarded to the leaders from the Super-K and SNO experiments, Takaaki Kajita and Arthur McDonald, for the solution to these deficits.

The scientific community has made significant progress in characterizing neutrinos, like their spectra, quantum numbers, and other properties. From claiming that they would be impossible to detect to using them to study the sun [26, 27] and probe supernovae [28], neutrinos are promising in their potential to answer some of our most challenging questions, such as why there is more matter than antimatter in the universe. However, neutrinos also push us to rethink what we currently know. In fact, their very mass raises questions and tensions about the Standard Model (SM), in which neutrinos are unable to obtain their mass like other SM particles do. The remainder of this chapter provides a brief summary of our current understanding of neutrinos, before we tackle the mysteries and challenges that lie ahead.

2.1 Neutrinos in the Standard Model

The Standard Model (SM) describes the strong, electromagnetic, and weak interactions between the fundamental particles of the universe in the framework of quantum field theory (QFT). It is a gauge theory based on the local symmetry group

$$SU(3)_C \times SU(2)_L \times U(1)_Y \tag{2.4}$$

where the subscripts denote the color charge (C), weak isospin (L), and weak hypercharge (Y) gauge fields. The theory is Lorentz-invariant and renormalizable, since its Lagrangian is dimension four. The SM governs the behavior of eight massless gluons, corresponding to the eight generators of $SU(3)_C$, and that mediate strong interactions; three massive (W^\pm and Z bosons) and one massless (the photon γ) gauge bosons, corresponding to the three generators of $SU(2)_L$ and one generator of $U(1)_Y$, and that mediate the electroweak interactions. In addition to the bosons, the SM governs the behavior of the standard model fermions: six quarks (u, d, c, s, t, b) divided into three generations, three families of charged leptons (electron, muon, tau), their accompanying neutrinos (ν_e, ν_μ, ν_τ), and their antiparticle counterparts. Finally, the SM also includes the scalar Higgs boson, which gives the particles (except, perhaps, for the neutrino; see later sections) their masses. Figure 2.1 shows the standard model particles, organized by generations and families. Table 2.1 shows the particles and their quantum numbers.

The masses of fermions result from the interactions of the fermion fields with the Higgs doublet. These are referred to as Yukawa interactions or couplings, named after Hideki Yukawa who developed the theory. A fermion mass term must involve a coupling of left-handed (LH) and right-handed (RH) fields, the strength of which depends on the Yukawa coupling factors. Since all fermions are Weyl fermions, neutrinos can be included in the SM via a LH doublet, which still maintains the renormalizability of the Lagrangian. This is enough to describe the LH neutrino and RH antineutrino. The weak interaction only couples to chirally left-handed particles whereas right-handed charged leptons can still couple to

1st Gen.	2nd Gen.	3rd Gen.	$SU(3)_C$	$SU(2)_L$	Y_W	Q
$\begin{pmatrix} u \\ d \end{pmatrix}_L$	$\begin{pmatrix} c \\ s \end{pmatrix}_L$	$\begin{pmatrix} t \\ b \end{pmatrix}_L$	3	2	1/3	$\begin{pmatrix} 2/3 \\ -1/3 \end{pmatrix}$
u_R	c_R	t_R	$\bar{3}$	1	4/3	2/3
d_R	s_R	b_R	$\bar{3}$	1	-2/3	-1/3
$\begin{pmatrix} \nu_e \\ e \end{pmatrix}_L$	$\begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}_L$	$\begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}_L$	1	2	-1	$\begin{pmatrix} 0 \\ -1 \end{pmatrix}$
e_R	μ_R	τ_R	1	1	-2	-1
Bosons			$SU(3)_C$	$SU(2)_L$	Y_W	Q
$H = \begin{pmatrix} H^+ \\ H^0 \end{pmatrix}$			1	2	1	$\begin{pmatrix} 1 \\ 0 \end{pmatrix}$
$\begin{pmatrix} W^+ \\ W^- \end{pmatrix}$			1	2	0	$\begin{pmatrix} +1 \\ -1 \end{pmatrix}$
$\begin{pmatrix} Z^0 \\ \gamma \end{pmatrix}$			1	2	0	0
g			8	1	0	0

Table 2.1: A table of the Standard Model particles, with their representations in the gauge groups $SU(3)_C$ and $SU(2)_L$, and their respective weak hypercharge Y_W and electric charge Q . The top half of the table shows the quantum numbers for fermions and the bottom half shows the quantum numbers for bosons. Here, the definition $Q = T_3 + \frac{1}{2}Y_W$ is used.




Figure 2.1: A periodic table of the Standard Model particles, highlighting the different generations (or families) of fermions, the gauge bosons that mediate the strong, weak, and electromagnetic interactions, and the Higgs boson that gives the W and Z bosons their mass. Figure from Wikimedia Commons.

the electromagnetic force, so their inclusion in the SM gives rise to a Yukawa term in the Lagrangian⁸

$$\mathcal{L}_{\text{Yukawa}} = -y\overline{\psi}_L\phi\psi_R + \text{h.c.} \quad (2.5)$$

that describes the interaction between LH and RH lepton fields ψ_L and ψ_R with scalar field ϕ (Higgs field) in the SM. The mass terms for the leptons arise after spontaneous symmetry-breaking. The absence of right-handed neutrino fields in the SM results in massless neutrinos,

⁸“h.c.” stands for Hermitian conjugate.

since there are no right-handed fields to complete this interaction for the neutrino. However, we have seen thus far and have overwhelming evidence that this is not the case. We will see in the next section (Section [2.2](#)) how neutrinos may obtain their mass. This discrepancy tells us that the SM is incomplete and may be an effective theory of a more general theory that includes gravity.

2.2 Neutrino Mass

In the current Standard Model, neutrinos are massless. However, for neutrinos to oscillate and mix their flavors, they must have mass. Currently, there are a few mass models that give different predictions of the neutrino mass, depending on the nature of the neutrino's identity. If neutrinos are distinct from their antiparticle counterparts, then they are considered Dirac fermions. But if the neutrino and its antineutrino partner can be described by the same field, then they are Majorana fermions. As such, answering questions about the mass of a neutrino means exploring the nature of the neutrino's identity.

From cosmological observations, the sum of the three neutrino masses has an upper limit of 0.17 eV, set with Planck data [\[29\]](#). While neutrino mixing tells us the values of the mass squared differences (see Section [2.3](#)), we do not know the absolute scale of their masses. Experiments currently looking to determine the absolute mass scale include KATRIN (Karlsruhe Tritium Neutrino) and Project 8. Both use tritium to try to determine the absolute masses of neutrinos. In 2011, data from the Troitsk experiment was reanalyzed and set an

upper bound of $m(\bar{\nu}_e) \sim 2.05$ eV [30]. KATRIN set limits on the effective electron neutrino mass to $m(\nu_e) < 1.1$ eV/c² at the 90% confidence interval [31] in 2019 and lowered this limit to $m(\nu_e) < 0.7$ eV/c² in 2022 [32]. Both experiments aim to improve this limit by one (KATRIN) or two (Project 8) orders of magnitude.

2.2.1 Dirac Masses

The simplest way for neutrinos to obtain their masses is by coupling with the Higgs field like the other massive standard model fermions. In this treatment, neutrinos are Dirac fermions, named after Paul Dirac, who formulated the Dirac equation to describe spin- $\frac{1}{2}$ massive particles that are distinct from their antiparticles. In addition to the left-handed neutrinos that we observe, there would be a right-handed counterpart with a left-handed antineutrino. We can introduce right-handed Weyl spinors that behave as singlets and these would couple to the Higgs via a term in the Lagrangian with a Yukawa coupling factor:

$$\mathcal{L}_{\text{Dirac}} = -y^{\alpha i} \overline{L_{L\alpha}} \tilde{H} \nu_{Ri} + \text{h.c.} \quad (2.6)$$

where $y^{\alpha i}$ represents the Yukawa coupling matrices, $L_{L\alpha}$ is the left-handed Dirac doublet that combines the Weyl spinors for left-handed charged leptons and left-handed neutrinos into one object, \tilde{H} representing the Higgs field H via $\tilde{H} = i\sigma_2 H^*$, and ν_{Ri} being the right-handed neutrinos.

While the generation of the neutrino Dirac masses involves a straightforward and minimal extension of the Standard Model, it might be considered theoretically unsatisfying due to

the smallness of the neutrino masses, $\mathcal{O}(1 \text{ eV})$ at most, compared to the largeness of the Higgs vacuum expectation value (vev), $v \approx 246 \text{ GeV}$. Since the Dirac neutrino masses are proportional to v via

$$m_D^{\alpha i} = y^{\alpha i} v \quad (2.7)$$

this implies that the coupling constants y_k must be extremely small, $\mathcal{O}(10^{-13})$. For comparison, the couplings are $\mathcal{O}(10^{-6} \text{ to } 10^{-2})$ for other charged leptons and $\mathcal{O}(10^{-5} \text{ to } 1)$ for quarks [33].

2.2.2 Majorana Masses

Another way to give neutrinos mass is to treat neutrinos as Majorana fermions, named after Ettore Majorana⁹, who in 1937, came up with the formulation in which a neutrino gains mass by coupling to its own charge conjugate. This means that the neutrino is its own antiparticle and its mass is generated from the same field. Given a chirally left-handed field Ψ_L , its charge conjugate $\Psi_L^c = -i\gamma^2\Psi_L^*$ is a chirally right-handed field. The Lagrangian generated from this field can then be written as

$$\mathcal{L}_{\text{Majorana}} = -\frac{m_M}{2}\overline{\Psi}_L\Psi_L^c + \text{h.c.} \quad (2.8)$$

Charge is not conserved in this treatment because the fields couple to their charge conjugates. Since neutrinos are chargeless, they are the only fermions that are allowed to be Majorana. Another noteworthy feature of the Majorana mass term is that it violates lepton

⁹Majorana mysteriously disappeared in 1938 at the age of 31.




Figure 2.2: Feynman diagrams comparing $2\nu\beta\beta$ decay (left) and $0\nu\beta\beta$ decay (right). If neutrinos are Majorana, the two neutrinos would annihilate each other and no neutrinos would be observed. This process violates lepton number by two. Feynman diagrams adapted from [40].

number symmetry by two. This makes neutrinos a possible candidate for generating the matter-antimatter asymmetry in the Universe through leptogenesis [34].

One way to confirm if neutrinos are Majorana is to study reactions in which no neutrinos are produced and lepton number is violated by two, also called neutrinoless double beta ($0\nu\beta\beta$) decay (Figure 2.2 depicts this). This nuclear process involves two neutrons transforming into two protons and two electrons without the emission of any neutrinos since they annihilate each other. Experiments searching for $0\nu\beta\beta$ decay include GERDA, EXO-200, MAJORANA, SNO+ and KamLAND-Zen. Current limits on the half-life of $0\nu\beta\beta$ decay are on the order of $T_{1/2}^{0\nu} \sim 10^{25}$ years [35–39].

2.2.3 Other Neutrino Mass Models

While the above models view Dirac masses and Majorana masses as distinct possibilities, both of these mechanisms can be present in the Lagrangian. Sterile right-handed neutrinos would couple to active left-handed states through the Dirac mass, while coupling to their self-conjugate via the Majorana mass. This introduces off-diagonal entries in the mass matrix, which, when diagonalized, adds contributions that give a compelling explanation for the discrepancy in the observed neutrino masses. This model is referred to as the *seesaw mechanism*, as the right-handed Majorana particles would need to be heavy to compensate for the lightness of the left-handed neutrinos, hence the name “seesaw”. The Lagrangian for this looks like

$$\mathcal{L}_M = -\frac{m_M^{\alpha\beta}}{2}\bar{\nu}_{L\alpha}\nu_{L\beta}^c + \text{h.c.}; \quad m_M^{\alpha\beta} = \frac{\lambda^{\alpha\beta}v^2}{\Lambda} \quad (2.9)$$

Λ represents a high energy scale where new physics responsible for the Majorana nature of neutrinos to arise, similar to how the vacuum expectation value v is representative of the electroweak energy scale. This seesaw mechanism is referred to as seesaw type-I. A variation of the seesaw mechanism – called type-III seesaw – includes new right-handed neutrinos being $SU(2)_L$ triplets instead of singlets. After spontaneous symmetry breaking, this results in a Majorana mass matrix giving mass to both types of neutrinos in the form of

$$\mathcal{L}_M = -\frac{1}{2} \begin{pmatrix} \bar{\nu}_L & \bar{\nu}_R^c \end{pmatrix} \begin{pmatrix} 0 & m_D \\ m_D^T & M_R \end{pmatrix} \begin{pmatrix} \nu_R^c \\ \nu_R \end{pmatrix} \quad (2.10)$$

The mass of the neutrino is then determined by

$$m_\nu = -m_D M_R^{-1} m_D^T \quad (2.11)$$

Other mass models include seesaw type-II, which involves adding a boson triplet for the neutrinos to couple to that serves the same purpose as the Higgs. In this type, the boson compensates for the lightness of the SM neutrinos instead of the right-handed heavy neutrino as in type-III. More details of the different seesaw mechanisms and other mass models can be found in [41–44].

2.3 Neutrino Oscillation

It is difficult to talk about neutrino oscillations without talking about their masses, since the existence of massive neutrinos (regardless of which mass model holds true) leads to neutrino mixing. Oscillations were not new physics at the time the first neutrino was discovered. In the quark sector, the unitary transformation between the weak eigenstates and the mass eigenstates of quarks is described by the Cabbibo-Kobayashi-Maskawa (CKM) matrix, first introduced by Nicola Cabbibo in 1963 and later extended by Makoto Kobayashi and Toshihide Maskawa in 1973 [45]. The mismatch of the quantum states of quarks results in quarks mixing and changing flavors when participating in weak interactions.

The equivalent to the CKM matrix for neutrinos is the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, which describes the mixing of their mass and flavor eigenstates. The PMNS matrix was first introduced in 1962 by Ziro Maki, Masami Nakagawa, and Shoichi

Sakata to describe the neutrino oscillations predicted by Bruno Pontecorvo, who was aware of the oscillations from K^0 into its antiparticle \bar{K}^0 . Neutrinos oscillations are dependent on neutrinos having two distinct eigenstates: in weak interactions, neutrinos interact via their flavor eigenstates, while they propagate through space via their mass eigenstates due to gravity. Each of the flavor states can be written as a weighted superposition of the three mass states and vice versa

$$|\nu_\alpha\rangle = \sum_{i=1}^3 U_{\alpha i}^* |\nu_i\rangle \quad (2.12)$$

$$|\nu_i\rangle = \sum_{\alpha=e,\mu,\tau} U_{\alpha i} |\nu_\alpha\rangle \quad (2.13)$$

where ν_α is a flavor eigenstate, ν_i is a mass eigenstate, and $U_{\alpha i}$ are elements of the PMNS matrix.

In the three flavor case, the PMNS matrix is generally expressed as the product of three rotation matrices, each parameterized by one of the mixing angles θ_{12} , θ_{13} , and θ_{23} , along with a complex phase δ_{CP} . Noting that $s_{ij} \equiv \sin \theta_{ij}$ and $c_{ij} \equiv \cos \theta_{ij}$,

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (2.14)$$

which when written out, becomes

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{13}s_{23}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{13}s_{23}e^{i\delta_{CP}} & c_{13}s_{23} \\ s_{12}s_{23} - c_{12}s_{13}c_{23}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}s_{13}c_{23}e^{i\delta_{CP}} & c_{13}c_{23} \end{pmatrix} \quad (2.15)$$

The complex phase, δ_{CP} , takes on a value other than zero or π in the case that charge conjugation and parity (CP) symmetry violation occurs in the neutrino sector, and generates differences in oscillation probabilities for neutrinos and antineutrinos. Two additional phases η_1 and η_2 are added to the PMNS matrix via

$$U_M = \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\eta_1/2} & 0 \\ 0 & 0 & e^{i\eta_2/2} \end{pmatrix} \quad (2.16)$$

in the scenario that neutrinos are Majorana. These two phases can be rewritten into multiplicative factors and do not contribute to neutrino oscillations when calculating the oscillation probabilities as they cancel out.

To see how mass is tied to oscillations, let us first calculate the oscillation probability. Neutrinos propagate through space in their mass states as a result of interacting with gravitational potentials. In vacuum, mass states are eigenstates of the Hamiltonian \mathcal{H} , with their time evolution and total energy described by the Schrödinger equation:

$$\begin{aligned} i \frac{d}{dt} |\nu_i(t)\rangle &= \mathcal{H} |\nu_i(t)\rangle \\ &= E_i |\nu_i(t)\rangle \\ &= E_i \cdot e^{-iE_i t} |\nu_i(t)\rangle \end{aligned} \quad (2.17)$$

Using Eqs. [2.12](#) and [2.13](#), we can rewrite the above equation as

$$|\nu_\alpha(t)\rangle = \sum_{i=1}^3 U_{\alpha i}^* e^{-iE_i t} |\nu_i\rangle \quad (2.18)$$

The probability of a neutrino with initial flavor α at time $t = 0$ to be observed as a neutrino with flavor β at a later time t is given by

$$\begin{aligned}
P_{\nu_\alpha \rightarrow \nu_\beta} &= |\langle \nu_\beta | \nu_\alpha(t) \rangle|^2 \\
&= \left| \sum_{i=1}^3 U_{\alpha i}^* U_{\beta i} e^{-iE_i t} \right|^2 \\
P_{\nu_\alpha \rightarrow \nu_\beta} &= \sum_{i,j=1}^3 U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* e^{-i(E_i - E_j)t}
\end{aligned} \tag{2.19}$$

We can rewrite the exponent in Eq. [2.19](#), using the notation $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$ and with the following Taylor expansion approximation

$$E_i = \sqrt{p_i^2 + m_i^2} \approx E_i + \frac{m_i^2}{2E} \tag{2.20}$$

$$\begin{aligned}
E_i - E_j &\approx \left(E_i + \frac{m_i^2}{2E} \right) - \left(E_j + \frac{m_j^2}{2E} \right) \\
E_i - E_j &\approx \frac{\Delta m_{ij}^2}{2E}
\end{aligned} \tag{2.21}$$

since neutrinos are ultra-relativistic particles with negligible masses compared to their total energies and momenta, so $E \approx p$. We can also assume that all mass eigenstates of the same neutrino have the same total energy $E_i = E_j = E$, so that the energy terms cancel each other out. In addition, we can also approximate $t \approx L$ in natural units due to the ultra-relativistic nature of the neutrino. Now we can write the probability in Eq. [2.19](#) as

$$\begin{aligned}
P_{\nu_\alpha \rightarrow \nu_\beta}(L) &= \sum_{ij} U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* \exp \left(-i \frac{\Delta m_{ij}^2 L}{2E} \right) \\
P_{\nu_\alpha \rightarrow \nu_\beta}(L) &= \delta_{\alpha\beta} - 4 \sum_{i>j} \Re(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin^2 \left(\frac{\Delta m_{ij}^2 L}{2E} \right) \\
&\quad + 2 \sum_{i>j} \Im(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \cos^2 \left(\frac{\Delta m_{ij}^2 L}{2E} \right)
\end{aligned} \tag{2.22}$$

where we have taken the real and imaginary components of the appropriate sums. In units that are convenient for an experimental physicist, the mixing argument in Eq. [2.22](#) can be simplified to

$$\left(\frac{\Delta m_{ij}^2 L}{2E}\right) \Rightarrow \left(\frac{\Delta m_{ij}^2 c^4}{4 \hbar c} \frac{L}{E}\right) = 1.27 \Delta m_{ij}^2 [eV^2] \frac{L[km]}{E[GeV]} \quad (2.23)$$

so that the oscillation probability takes on its more familiar form

$$P_{\nu_\alpha \rightarrow \nu_\beta}(L) = \sin^2 2\theta \sin^2 \left(1.27 \Delta m_{ij}^2 \frac{L}{E} \right). \quad (2.24)$$

A similar equation can be derived for antineutrinos. Since observations show that neutrinos oscillate, the mass term Δm_{ij}^2 must be non-zero, confirming that neutrinos have masses (at least two of them) and thus relating their masses to oscillations.

Using Eq. [2.24](#) as a starting point, we can study neutrino oscillations and determine the mass squared differences and mixing angle parameters. In general, once we know the energy of the neutrino beam E and design a baseline with known length L , we can determine Δm_{ij}^2 , θ_{ij} , and δ_{CP} . Current experiments studying CP violation and neutrino oscillations include the Deep Underground Neutrino Experiment (DUNE) and NuMI Off-axis ν_e Appearance ($\text{NO}\nu\text{A}$).

2.3.1 Matter Effects on Neutrino Oscillation

The probability calculations made in the previous section are for neutrino oscillations in vacuum. When neutrinos travel through matter, there are several important effects that affect the oscillation probability. In general, matter presents itself as an effective potential

for propagating neutrinos. In addition, the Hamiltonian is not the same in matter as it is in vacuum. The Hamiltonian in matter includes an additional potential caused by coherent forward elastic scattering with surrounding particles and dependent on the density of the material constituents such as electrons, protons, and neutrons. Because of this, neutrinos have a different set of mass eigenstates, while keeping the same flavor eigenstates (cf. Eq [2.17](#)). This results in new mixing angles. The modified mixing angles results in a modification of the probability of flavor change, also known as the Mikheyev-Smirnov-Wolfenstein (MSW) effect. In high density materials, like the core of the sun, mixing is almost completely suppressed because the mixing angles are modified in such a way that interference between the different mass eigenstates does not take place. The MSW effect is seen especially for electron neutrinos, since regular matter contains electrons. Muon and tau neutrinos have less channels of interaction, so their oscillation probabilities are unaffected when traveling through matter. The MSW effect allowed the determination of the sign of Δm_{12}^2 , which is positive, meaning $m_2^2 > m_1^2$.

2.3.2 Mass Ordering

As shown in Eq. [2.22](#), the oscillation probability depends on the non-zero Δm_{ij}^2 terms. In the standard three neutrino flavor framework, only two of these are independent. While we have set some limits on these squared differences, we have yet to determine the absolute ordering of m_1 , m_2 , and m_3 , in other words, the signs of the mass differences. From solar

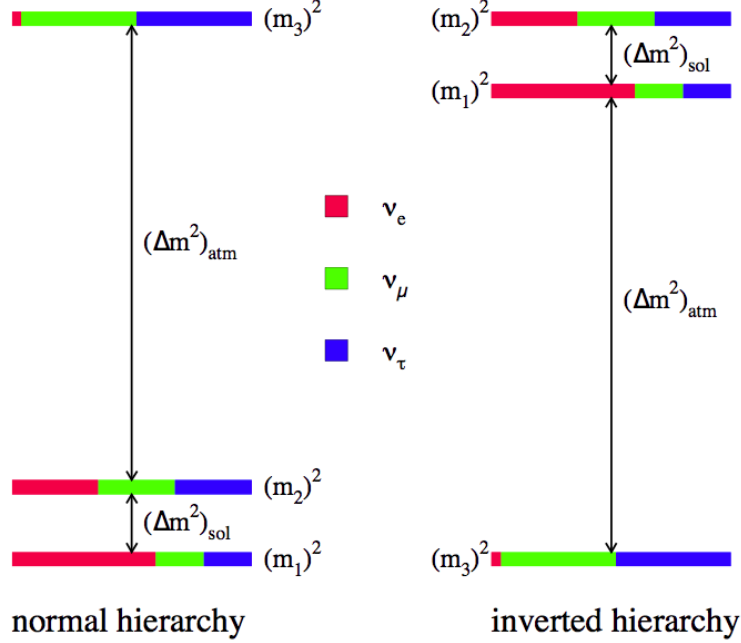


Figure 2.3: A pictorial representation of the normal and inverted mass orderings. The colors represent the flavor composition of each mass eigenstate. The subscripts for Δm^2 terms indicate the type of neutrino experiments that are able to determine the mass squared differences. Note that Δm_{sol}^2 is equivalent to Δm_{12}^2 and Δm_{atm}^2 to Δm_{23}^2 . Figure from [46].

neutrino experiments, we have determined the sign of Δm_{12}^2 due to the MSW effect, which sets m_2 to be heavier than m_1 . However, we do not know the sign of Δm_{23}^2 , which can be determined from atmospheric neutrinos and currently gives us two possible mass orderings. The scenario where $m_3 > m_2$ is referred to as the *normal mass ordering*, while $m_2 > m_3$ is referred to as the *inverted mass ordering*. See Figure 2.3 for a depiction of the mass ordering.

Identifying the correct mass ordering is one of the important goals in neutrino physics. Knowing the sign of Δm_{31}^2 will determine the total sum of all three neutrinos. It is also important in CP violation searches, as the sign will affect the expected signal. Experiments

such as T2K¹⁰ and NO ν A are trying to determine the correct mass ordering. At the time this writing, the latest fits to existing data lean towards the normal ordering, although there is a small tension between the NO ν A dataset and the combination of T2K and reactor data. The significance is not enough to fully exclude the inverted mass ordering [47–49].

2.4 Neutrino Interactions with Matter

Despite their elusiveness, neutrinos can be detected. As we’ve seen earlier in this chapter, neutrinos interact with matter through the weak and gravitational forces, the latter of which is negligible in typical experimental situations. Because we can usually ignore interactions through gravity, neutrino interactions can be described with two categories: charged current and neutral current. We detect neutrinos from the byproducts of these in weak interactions. In charged current (CC) interactions, a neutrino exchanges a W boson (either W^+ or W^-) with another particle and is transformed into the charged lepton of the same flavor, e.g., ν_e becomes e^- and ν_μ becomes μ^- . In neutral current (NC) interactions, a neutrino exchanges a Z^0 boson with another particle, recoils, and emerges with its flavor unchanged. Figure 2.4 depicts these CC and NC processes in Feynman diagrams. Derivations of the aforementioned exchange currents terms in the Lagrangian can be found in [8].

Across the energy spectrum low, medium, and high energy scales, CC and NC interactions can be further described as such interactions as

¹⁰T2K: Tokai to Kamioka




Figure 2.4: For both types of scattering: the left Feynman diagram is a charged current interaction in which a $W^{+/-}$ boson is exchanged, and the right is a neutral current interaction in which a Z^0 boson is exchanged. The α subscript can be any of the three neutrino flavors.

- **Quasi-Elastic (QE) Scattering:** at low (~ 10 s of MeV) to medium (~ 1 GeV) neutrino energies, the simplest hadronic interaction in which a charged weak interaction causes a transition in weak isospin space of the nucleon;
- **Resonant Pion Production (RES):** above ~ 0.5 GeV, neutrinos can induce the production of mesons, often pions, when interacting with a nucleus; and
- **Deep Inelastic Scattering (DIS):** starting around 2 GeV, the neutrino is increasingly able to interact with the quark substructure of the nucleon and generate hadronic cascades.

Referring to Figure [2.5](#), the Accelerator Neutrino Neutron Interaction Experiment (ANNIE) is mainly concerned with QE and resonant interactions, since the Booster Neutrino




Figure 2.5: The famous plot of neutrino cross section as a function of energy. Quasi-elastic interactions dominate in the low to medium energy range. At medium energies, resonant processes turn on and contribute a significant portion of events. At higher energies, deep inelastic scattering is the main interaction. Figure from [50].

Beam (BNB) that the detector sits on is peaked at an energy of ~ 700 MeV. In the medium-energy regime, $\mathcal{O}(1 \text{ GeV})$, the interaction length is hadronic ($\sim 1 \text{ fm}$), so nuclear effects become important. These are the interactions of great interest to atmospheric and accelerator-based neutrino oscillation experiments due to the medium energies of atmospheric and accelerator neutrinos [51].

The rate at which an interaction is expected to occur is the cross section. With neutrinos, their interaction cross sections are quite small and make them difficult to detect. Almost all the major questions in physics involve counting the number of interactions. Thus, it is

crucial that we have a good handle of cross sections, which is a function of the neutrino energy, as depicted in Figure 2.5. Early cross section measurements were done with light nuclei bubble chambers, which had low statistics but minimal nuclear effects, both of which are due to the low A of the nuclei. However, in recent years, neutrino experiments have moved towards larger detectors with heavier nuclei, like argon and oxygen. Oxygen is the main nuclear target in water Cherenkov detectors, yet it has few cross section measurements.

2.4.1 Neutrino-Nucleus Interaction Models

A particular type of interaction that is of interest is charged-current quasi-elastic (CCQE) scattering. A few examples are

$$\begin{aligned}\nu_\mu + n &\rightarrow \mu^- + p \\ \bar{\nu}_\mu + p &\rightarrow \mu^+ + n\end{aligned}\tag{2.25}$$

which have the largest cross sections at ~ 1 GeV (Figure 2.5), and are important for neutrino oscillation experiments that operate in this energy range. CCQE interactions are important because they allow us to estimate the neutrino energy and momentum transfer of the interaction from simple kinematics of the detected lepton. As mentioned previously, many neutrino parameters depend on a precise understanding of interaction cross sections, which in turn depend on the neutrino energy.

The simplest description of neutrino-nuclear scattering is built upon two ingredients: a characterization of neutrino-nucleon scattering and a model for nucleons in the nucleus.

For quasi-elastic scattering in charged current neutrino interactions, the former is well-established, while the latter has been the subject of much theoretical and experimental work in the context of lepton scattering over the past 40 years [51]. In the simplest and most common scenario, neutrino-nucleus scattering is treated as the incoherent sum of scattering from free nucleons, that is, we treat the interaction of the neutrino with just one nucleon with no further effects from traveling through the nucleus. This is the so-called *impulse approximation approach*. In reality, as we will see, the nucleons are not independent particles, and as a result, more complex nuclear dynamics are involved.

To understand the impulse approximation model, we can look at the case of an electron scattering elastically off a free nucleon. The amplitude for this scenario can be written as

$$M = \frac{e^2}{q^2} \bar{u}_e(k') \gamma^\mu u_e(k) j_\mu \quad (2.26)$$

where $e = \sqrt{2\pi\alpha}$, α being the fine-structure constant; $u_e(k)$ and $\bar{u}_e(k')$ represent the initial and final states of the electron; and j_μ is the hadronic vector current. The hadronic vector current captures the underlying nucleon structure and is written in terms of the well-known Dirac and Pauli nucleon-vector form factors, $F_1(q)$ and $F_2(q)$, which are functions of a single variable, the magnitude of the three-momentum transfer, q . Conserved vector current (CVC) and isospin conservation allow QE electron scattering (QEES) to establish the vector current contributions of both charged and neutral current scattering of neutrinos.

In analyzing electron scattering, it is useful to separate the scattering yields into longitudinal and transverse contributions. The longitudinal contribution originates from scattering,

in which the polarization of the virtual photon is along the direction of momentum transfer (Coulomb scattering), whereas for transverse scattering, it is perpendicular to the momentum transfer. The differential cross section for electron-nucleon scattering is

$$\frac{d\sigma}{d\Omega_e} = \sigma_{Mott} \frac{E'}{E_0} \left[\frac{G_E^2(Q^2) + \tau G_M^2(Q^2)}{1 + \tau} + 2\tau G_M^2(Q^2) \tan^2 \left(\frac{\theta}{2} \right) \right] \quad (2.27)$$

At medium energies, the transverse response dominates, particularly at large scattering angles.

For neutrinos scattering off nucleons, an axial current enters the picture. The total nucleon current coupling to the charged weak leptonic current is an isovector one-body nucleon current with both vector and axial-vector components: $j^\mu(Q^2) = j_V^\mu(Q^2) + j_A^\mu(Q^2)$. The full nucleon weak current was written down by Llewellyn-Smith [\[52\]](#), but it is sufficient to write the axial current of the nucleon as

$$j_A^\mu(Q^2) = \bar{u}(p') \left(G_A(Q^2) \gamma^\mu + \frac{1}{2M} G_P(Q^2) q^\mu \right) \gamma^5 u(p) \quad (2.28)$$

where the induced psuedoscalar form factor $G_P(Q^2) = 4m_N^2 G_A / (m_\pi^2 + Q^2)$ is determined by partially conserved axial current (PCAC) and the axial-vector form factor $G_A(Q^2)$ is established from experiment.

Finally, the differential cross section for neutrino QE scattering off free nucleons can be expressed as

$$\frac{d\sigma}{dQ^2} = \frac{G_F^2 M^2 \cos^2 \theta_C}{8\pi E_\nu^2} \left[A \mp \frac{(s-u)}{M^2} B + \frac{(s-u)^2}{M^4} C \right] \quad (2.29)$$

The $(-)+$ refers to (anti)neutrino scattering, $(s-u) = 4ME_\nu - Q^2 - m^2$, and m is the lepton mass. A , B , and C are functions of the Q^2 -dependent vector, axial-vector, and pseudoscalar form factors.

For neutrino-nuclear scattering models, the vector form factors were modeled after electron scattering and small contributions from the pseudoscalar form factor for ν_μ scattering. Thus, early studies of neutrino QE scattering focused on investigating the axial-vector form factor of the nucleon. Spectral functions are needed to account for the overlap between wavefunctions of the nuclear ground state and all other possible final states associated with the observed interaction. The simplest is the Fermi Gas model initially used by Smith and Moniz [53]. This model assumes a uniform distribution of nucleon momenta up to the Fermi surface.

Early experiments involving spark chambers and bubble chambers provided the first event samples of neutrino QE scattering. The primary aim of these experiments was to measure free nucleon form factors. By 1990s, scientists concluded that neutrino QE cross sections could be accurately and consistently described by $V-A$ theory, assuming dipole axial vector form. The axial mass M_A values determined from these experiments were on the higher end than what was predicted with the impulse approximation model.

As neutrino experiments moved to targets with higher A values, discrepancies popped up between prediction and experiment that pushed the neutrino physics community to revisit the model describing neutrino QE scattering interactions. For example, the MiniBooNE experiment made the first differential cross section measurement for QE scattering on carbon-

12 and observed a value $\sim 45\%$ higher than the cross section predicted based on the impulse approximation approach [54, 55]. These discrepancies can also lead to impacts on neutrino energy reconstruction [56]. Attempts to reconcile theory and experiment include introducing a scaling factor, y . Other models tried to incorporate nuclear effects such as short-range correlations (SRCs). With these new considerations, we must also consider how a QE event is defined for an experiment. The extent to which event selection criteria are inclusive affects the sample of true QE events, and thus the calculation of cross sections and energies. Overall, it is clear that our understanding of nuclear effects in neutrino-nucleus interactions falls short.

2.4.2 Nuclear Effects on Neutrino-Nucleus Interactions

While QE electron scattering is a good model to start with in describing neutrino QE scattering off nuclei, there are some shortcomings. First, in electron scattering experiments, the electron beam flux and energies are known to a very high accuracy and precision. Meanwhile, neutrino fluxes and energies are not very well-characterized and contribute to errors in cross sections upwards of 20% and by factors greater than two. Second, while role of nuclear effects in electron scattering were known, they were not considered when describing neutrino QE scattering. Some nuclear effects are described in the following sections.

2.4.2.1 Multi-Nucleon Correlations

The impulse approximation model treats the nucleus as a sum of single nucleons that act independently of each other. However, that is not what's really happening here. Nucleons can exert short-range forces on each other (*short-range correlations*) and this is not accounted for in the average potential used to describe the nucleus. When two nucleons are close to each other ($r_{ij} < 1.5f$), the strong short range forces greatly increase their relative momentum, propelling the nucleons far off-shell. These SRCs produces $2p$ - $2h$ (2-particle-2-hole) fluctuations in Fermi momentum distribution and causes individual particle momenta involved in fluctuation to acquire momentum well above Fermi momentum and impacting the total cross section. Thus, SRCs and meson exchange must be taken into consideration when attempting to make a realistic calculation of neutrino CCQE.

Nucleons can also experience *long-range correlations* (LRCs). LRCs occur when strong interactions between nucleons generate a polarization of the nuclear medium which modifies the weak charge of nucleons relative to those in vacuum [57]. Short and long range correlations are taken into account in Benhar's spectral function.

2.4.2.2 Final State Interactions

In addition to mutli-nucleon correlations, Final State Interactions (FSI) also contribute to the higher neutrino CCQE cross sections. The definition of FSIs can be a bit arbitrary, but in general, they are the subsequent strong interactions, that follow a neutrino-nucleus

interaction, between the product of the electroweak vertex and the other nucleons in the nucleus. As some examples, the outgoing nucleon can knock out additional nucleons on its way out, charge exchange interactions can take place, pion rescattering, and so on. At the moment, there is little experimental work done in the study of FSIs.

All in all, it is important to model these nuclear effects accurately, since it impacts the number of signal and background events, the observed final state particles, and neutrino energy inferred. Events with FSIs can mimic CCQE signals, which affect the determination of the incident neutrino energy as it relies on using a simple kinematic model. Thus, it is critical that we study these effects to improve models implemented in event generators and get a better handle on systematic uncertainties needed for the next generation of neutrino experiments.

2.4.3 Towards a Neutron Multiplicity Measurement

Now that we have a sense of the challenges that exist in studying neutrino-nucleus interactions, it is clear how important it is to improve interaction cross section measurements. Experiments like T2K and DUNE that are attempting to measure the CP violating phase rely on precise cross section measurements, as most of the uncertainty stems from them. As we will see in the following chapters, ANNIE attempts to make sense of the nuclear effects that take place in neutrino-nucleus interactions by measuring the neutron multiplicity (or neutron yield), since neutrons hint at inelastic interactions. This is why it's important

that ANNIE makes a measurement of neutron multiplicity to characterize how much energy neutrons take away and improve accuracy in event reconstruction.

2.5 Open Questions and Current Parameters

While much has been learned about neutrinos over the past few decades, there are still mysteries that puzzle us. This section will review the current values of the relevant neutrino parameters and give a brief overview of some of the questions we are trying to answer.

2.5.1 Current Parameter Values

Table 2.2 shows the most current values at the time of this writing for the neutrino parameters, both under the normal mass ordering and inverted mass ordering.

Parameter	Normal Ordering (best fit)	Inverted Ordering ($\Delta\chi^2 = 4.14$)
$\sin^2 \theta_{12}$	$0.304^{+0.013}_{-0.012}$	$0.304^{+0.013}_{-0.012}$
$\theta_{12} [^\circ]$	$33.55^{+0.78}_{-0.75}$	$33.45^{+0.78}_{-0.75}$
$\sin^2 \theta_{23}$	$0.570^{+0.018}_{-0.024}$	$0.575^{+0.017}_{-0.021}$
$\theta_{23} [^\circ]$	$49.0^{+1.1}_{-1.4}$	$49.3^{+1.0}_{-1.2}$
$\sin^2 \theta_{13}$	$0.02221^{+0.00068}_{-0.00062}$	$0.02240^{+0.00062}_{-0.00062}$
$\theta_{13} [^\circ]$	$8.57^{+0.13}_{-0.12}$	$8.61^{+0.12}_{-0.12}$
$\delta_{CP} [^\circ]$	195^{+51}_{-25}	286^{+27}_{-32}
$\Delta m_{21}^2 [10^{-5} \text{eV}^2]$	$7.42^{+0.21}_{-0.20}$	$7.42^{+0.21}_{-0.20}$
$\Delta m_{3l}^2 [10^{-3} \text{eV}^2]$	$+2.514^{+0.028}_{-0.027}$	$-2.497^{+0.028}_{-0.028}$

Table 2.2: Current values of neutrino oscillation parameters with errors for both the normal (NO) and inverted (IO) mass orderings [58]. For Δm_{3l}^2 , $l = 1$ for NH and $l = 2$ for IH.

2.5.2 Sterile Neutrinos

Up until now, the above discussions of neutrinos have mainly involved three active neutrino generations. This number has been confirmed and is strongly constrained by studies of the decay width of the Z^0 gauge boson via electron-positron collisions [18, 19, 59]. However, the mass of the Z^0 provides an upper bound for this number, as it was the limit of the beam energy, and there is a possibility that additional neutrinos exist. Additional neutrinos could be heavy, or light but sterile, meaning they would not participate in weak interactions and so would not appear in gauge boson decay. However, they could exist through oscillations. One motivation for the existence of at least a fourth neutrino comes from the Liquid Scintillator Neutrino Detector (LSND) [60] anomaly. The experiment showed an excess of $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations compared to what was expected from three-flavor predictions [61]. The MiniBooNE experiment confirmed the LSND results with reports of their own excess of ν_e [55]. Efforts to address anomalies include the Short-Baseline Neutrino Program (SBN) at Fermilab and PROSPECT at Oak Ridge National Laboratory (ORNL). If confirmed, anomalies may signal the existence of a fourth sterile neutrino which does not participate in weak interactions.

The existence of additional sterile neutrinos would mean expanding the PMNS matrix from a 3×3 matrix into a unitary $N \times N$ matrix, with $(N - 3)$ sterile neutrinos. This new

matrix would look like

$$U' = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} & \dots & U_{en} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} & \dots & U_{\mu n} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} & \dots & U_{\tau n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ U_{s_{N-3}1} & U_{s_{N-3}2} & U_{s_{N-3}3} & \dots & U_{s_{N-3}n} \end{pmatrix} \quad (2.30)$$

The original PMNS matrix in the upper left corner would no longer be unitary on its own, and the oscillation probabilities from Eq. [2.19](#) would include additional terms that may affect active-active flavor transitions along with active-sterile and sterile-active flavor transitions. Current data significantly restrict sterile neutrino contributions, so searches for sterile neutrinos are mainly conducted in short baseline experiments, before active flavor transitions have a higher chance of occurring.

2.5.3 CP Violation and Matter-Antimatter Asymmetry

Andrei Sakharov proposed that CP violation is necessary to explain the observed imbalance of matter and antimatter [\[62\]](#). Because the extent of CP violation in the quark sector is not enough to explain the observed asymmetry, we must look to the lepton sector, specifically the neutrino sector. As seen in the PMNS matrix in Eq. [2.15](#), the complex phase δ_{CP} quantifies the extent of CP violation in neutrinos. The magnitude of CP violation is proportional to the Jarlskog invariant

$$J_{CP,l} = \frac{1}{8} \cos \theta_{13} \sin(2\theta_{23}) \sin(2\theta_{13}) \sin \delta_{CP} \quad (2.31)$$

which is equal to $0.33 \sin \delta_{CP}$ with current oscillation parameter values. It has the potential to be three orders of magnitude larger than its value in the quark sector. CP violation can be measured in muon neutrino to electron neutrino oscillations and the corresponding antineutrino oscillations. Experiments such as T2K and NO ν A are studying these oscillations using accelerator-produced neutrino beams and the long-baseline. Current limits on the CP phase are given in Table [2.2](#).

2.5.4 Diffuse Supernova Neutrino Background

One of the most efficient sources of neutrinos are core-collapsed supernovae [\[63\]](#), [\[64\]](#). On average, a single collapse happens every second in the observable universe and produces 10^{58} neutrinos with $\mathcal{O}(\text{MeV})$ energies. Combined with the fact that neutrinos rarely interact, it is reasonable to expect that there is a steady flux of neutrinos and antineutrinos that contain information about the history of supernovae. This flux is known as the Diffuse Supernova Neutrino Background (DSNB).

Current searches for the DSNB are constrained by background noise so only upper limits can be made. At the time of this writing, KamLAND has placed a limit of $98.1 \text{ cm}^{-2}\text{s}^{-1}\text{MeV}^{-2}$ for antineutrino flux in the 8.3–13.3 MeV energy range. In 2006, SNO placed an upper limit on the electron antineutrino flux of $2.7 \text{ cm}^{-2}\text{s}^{-1}$ for the neutrino energies between 22.9 and

36.9 MeV.

Similar to the Cosmic Microwave Background (CMB), we expect a low energy neutrino background of 10^4 to 10^{-6} eV called the Cosmic Neutrino Background (CNB). These are also called relic neutrinos.

Chapter 3

ANNIE

With the advent of neutrino beams and photodetection technologies, neutrino detectors fall into a few categories including liquid argon time projection chambers (LArTPCs), liquid scintillator (LS), and water Cherenkov (WC) detectors. In LArTPCs, like MicroBooNE and DUNE, a neutrino interaction with an argon nucleus produces a lepton that ionizes the liquid argon as it travels through the rest of the TPC. The resulting charge is collected onto wire planes and the event topology can be reconstructed from the time and position information. In liquid scintillator detectors, such as KamLAND, Daya Bay, and JUNO¹, the daughter particles deposit their energy in the scintillator, creating flashes of light that can be detected. In water Cherenkov detectors, like SNO and Super-K, the traveling lepton produces Cherenkov radiation, which is collected with photodetectors such as photomultiplier tubes (PMTs). For all detector types, large quantities of the target medium are used so that there

¹JUNO: Jiangmen Underground Neutrino Observatory

is a sufficient number of nuclear targets for neutrinos to interact with. Coupled with neutrino beam, a significant number of neutrino interaction events can be collected. We focus on the water Cherenkov detector, which has the following main advantages: (1) the relatively cheap cost of water, (2) the environmental safety of water, (3) the fast timing of the signal, (4) the ability to load neutron absorbers, and (5) the lack of natural radioisotopes.

The Accelerator Neutrino Neutron Interaction Experiment (ANNIE) is a water Cherenkov neutrino detector located on-axis in the Booster Neutrino Beam (BNB) at Fermi National Accelerator Laboratory (FNAL or Fermilab, for short). ANNIE has grown from a small experiment to a collaboration between 19 institutions and two national labs across five countries. The experiment is designed to measure the final state neutron multiplicity of neutrino-nucleus interactions, determine the charged current inclusive cross section on water, and serve as a testbed for new technologies. This chapter details the detector and the physics goals of ANNIE.

3.1 ANNIE Phase I

ANNIE is divided into several phases that mark different experimental geometries and measurement goals. The first phase of ANNIE demonstrated the feasibility of the experimental concept. The main goal was to assess the level of background neutrons and decide whether the experiment could move forward. The Phase I detector consisted of a cylindrical tank lined with white PVC, a front veto of scintillator paddles (Section [3.2.5](#)), and a partially



Figure 3.1: The Neutron Capture Volume (NCV) is an acrylic vessel filled with EJ-335, a mixture of liquid scintillator and gadolinium. A PMT coupled to the NCV is wrapped in black to prevent light generated in the tank from triggering the data acquisition system. The array of PMTs at the bottom of the Phase I detector can be seen facing upward in this image.

instrumented Muon Range Detector (Section [3.2.6](#)). At the bottom of the tank sat an array of 58 Hamamatsu R5912 photomultiplier tubes with 8-inch diameters. Above the tank lid was a set of six plastic scintillator paddles, each coupled to a 2-inch PMT via a light guide and secured to a metal frame. The front veto, MRD, and the top scintillator paddles were used to tag muons.

A large 50 cm x 50 cm acrylic cylinder filled with EJ-335, a 0.25% by weight gadolinium-loaded liquid scintillator from Eljen Technology, was deployed into the water-filled tank to measure the rate of beam-correlated background neutrons (see Figure [3.1](#)). Sources of

background neutrons include:

- **sky-shine neutrons:** neutrons that make their way out of the beam line and undergo scattering and enter the detector from above, and
- **dirt neutrons:** neutrons that are produced in the dirt.

Although not shown in Figure 3.1, this Neutron Capture Volume (NCV) was wrapped in layers of reflective material to maximize light collection and an outer layer of black plastic to optically isolate it from the rest of the tank. Two 3-inch PMTs were coupled to the NCV to trigger on scintillation light and initiate data acquisition.

The NCV was moved to different locations in the tank as indicated by the blue and red markers in Figure 3.2. After a neutron gets captured by gadolinium (Gd) in the NCV, the de-excitation photons emitted from Gd are amplified by the liquid scintillator, resulting in thousands of photons that make their way to the NCV PMTs. Their waveforms are recorded along with any activity in the tank following a beam trigger. In addition to the physics data runs, calibration runs of the NCV were carried out with a Cf-252 source. More details about the Phase I experimental setup, calibration, and data analysis can be found in [65].

Referring to Figure 3.2, we can see that an active fiducial volume placed lower in the tank (indicated by the blue curve) and bit upstream of the tank center (indicated by the red curve) sees very few background neutrons. The average beam-induced background neutron flux in this volume measures less than 0.02 neutrons/beam-spill/m³ [65]. The level of background neutrons was deemed sufficiently low to move forward with a full ANNIE detector. Since the




Figure 3.2: The NCV was deployed at various locations in the ANNIE Phase I detector. These background measurements show that background neutron rates are sufficiently low in a fiducial volume located in a central location slightly upstream of the tank center, indicated by the purple circle in the diagram above. Figure from [65].

Phase II installation, modifications to the area surrounding the experimental hall prompted an updated analysis of the beam-correlated neutron background. These changes include a large mound of dirt, called a berm, placed upstream of the ANNIE detector and a new experimental hall (SBND) built downstream. An updated analysis is ongoing at the time of writing this thesis.




Figure 3.3: ANNIE Phase II detector diagram. The tank is 3-m in diameter and 4-m tall, with PMTs and LAPPDs lining its walls. The Front Muon Veto and Muon Range Detector sit upstream and downstream of the tank, respectively. The readout electronics and power supplies are located on the level above the detector. In this diagram, the direction of the beam is from left to right. Figure from [66].

3.2 ANNIE Phase II Detector

The second phase involved fully commissioning the ANNIE detector. Referring to Figure 3.3, the main subsystems of the Phase II detector are the

- **0.1% gadolinium-doped water target:** active medium consisting of deionized water loaded with 0.1% gadolinium by mass, giving the target an enhanced sensitivity to neutrons;

- **132 photomultiplier tubes:** conventional photomultipliers of various models and sizes used to detect Cherenkov radiation for energy and vertex reconstruction;
- **5 Large Area Picosecond Photodetectors:** novel photodetectors with picosecond timing and sub-centimeter spatial resolution for improved vertex reconstruction;
- **steel cylindrical tank and inner structure:** encapsulates the water and holds photodetectors;
- **water filtration/purification system:** network of pumps, pipes, tubes, filters, lamps, and ion exchange resins that clean and recirculate the water;
- **Front Muon Veto:** two layers of scintillator paddles for tagging charged particles created upstream of the detector;
- **Muon Range Detector:** iron-scintillator sandwich used to tag downstream muons and reconstruct their tracks;
- and **electronics rack:** system of high voltage power supplies, digitization and readout modules, and circuit boards that power the detector subsystems and take data.

The Fermilab Booster Neutrino Beam encounters the Front Muon Veto, the water tank, and Muon Range Detector in the order listed. A detailed paper describing the ANNIE Phase II detector is in the process for publication [\[67\]](#).




Figure 3.4: LEFT: Fishbowl perspective of the inner structure, wrapped in tape made of polytetrafluoroethylene (PTFE). RIGHT: 3D model of the inner structure. The different colors indicate different PMT types. Unistrut rails between two panels are used for LAPPD deployment.

3.2.1 Detector Tank

One of the main components of the Phase II detector is a 3-m (10-ft) diameter, 4-m (13-ft) tall right cylindrical tank constructed from 7-gauge (4.5 mm) steel, which was also used in Phase I. To prevent the formation of and contamination by rust, the inside and outside of the tank are coated with epoxy and several layers of plastic separate the water target and the steel walls of the tank. The first two layers are 6 mil polyethylene plus a 30 mil sheet at the bottom to shield the main layer from rough areas of the tank. The main layer is a




Figure 3.5: Sections of a rectangular steel pipe soaking in 1% gadolinium-loaded water. The right pipe is wrapped in teflon tape. It is clear that any presence of rust will quickly degrade the water transparency, so precautions must be taken to reduce the risk of rust contamination.

40 mil white PVC liner manufactured by Kentain [68]. The presence of rust would quickly degrade the water transparency so it is crucial that the water target does not make contact with the steel tank. Figure 3.5 compares sections of steel pipe with and without a teflon² tape wrap as a precautionary measure.

Within the liner is an octagonal stainless steel inner structure designed to hold 132 PMTs and approximately 40 LAPPDs. PMTs are mounted on one of the eight side panels, while LAPPDs are deployed along Unistrut rails that are installed where two panels meet. See Figure 3.4 for a photo and 3D representation of the inner structure. While the inner structure is made of stainless steel, the welding required to put it together compromises its resistance

²Teflon is a common brand name for polytetrafluoroethylene (PTFE)-based materials.



Figure 3.6: The inner structure is lowered into the experimental hall while photo shoots among collaborators take place. Plastic sheets protect the inner structure against dust during transport. The SBND (Short-Baseline Near Detector, as part of the Short-Baseline Neutrino Program) experimental hall is the large building to the right.

to corrosion and rust. To restore the protective passive layer, the entire inner structure was electropolished by Dr. Emrah Tiras³ and me. This electrochemical finishing process removes a thin layer of the metal with a mild acid and electrical current to repassivate its surface. As an additional preventative measure, the inner structure is wrapped with 2-inch wide PTFE tape secured with metal zip ties to slow down the rust process and leaching of contaminants. Black low-density polyethylene (LDPE) plastic sheets are wrapped around the inner structure to optically isolate the encapsulated volume. The tank lid and inner structure are seen being lowered into the hall in Figure 3.6.

³Erciyes Üniversitesi




Figure 3.7: A view of the tank lid from the top level of ANNIE Hall. The five green circles indicate the locations of the ports used for source deployment. The central port is part of the larger removable hatch. The red rectangle highlights one of the eight “mail slots” used for LAPPD deployment.

The tank lid is welded to the top of the inner structure to form a single unit. The lid has several circular and rectangular ports, with a large removable hatch in the center (cf. Figure 3.7). The small circular ports permit the easy deployment and removal of calibration sources at different locations within the tank. The central hatch was designed for larger apparatuses, like the NCV of Phase I. Four PMTs are secured to the hatch, as seen in Figure 3.4, and thus, are removable. The rectangular ports, referred to as “mail slots”, are where LAPPDs will be deployed *in situ* so that commissioning of the rest of the detector can be completed in parallel with the development and testing of the LAPPD systems.

3.2.2 0.1% Gadolinium Water Target

The detector tank is filled with 26 tons of 0.1% by mass of gadolinium-loaded water. Because one of ANNIE's physics goals is to measure the rate of neutron production from neutrino interactions in the tank, gadolinium (Gd) was added to the water due to its high average neutron capture cross section of about 49,000 barns (b), compared to hydrogen's 0.33 b and oxygen's 0.19 mb [69]. Natural gadolinium consists of 15.7% Gd-155 and 14.5% Gd-157, which have capture cross sections of 61,000 b and 254,000 b, respectively. With such high cross sections, the addition of just 0.1% natural Gd by weight results in 90% of neutron captures happening on Gd, with the remaining 10% on hydrogen [70]. The time to capture is also reduced to $\sim 30 \mu s$ compared to $\sim 200 \mu s$ in pure water [71]. Any more Gd would significantly reduce the capture time constant, making it difficult to distinguish neutron captures from muons traversing through the detector.

Another benefit of gadolinium for neutron detection is the increase of visible energy in its de-excitation gammas. Following a neutron capture by Gd-155 and Gd-157, the isotopes decay through a cascade of energy levels, emitting a total of 8 MeV in energy shared among 3-4 gammas⁴ [70]. This unique signal is much more discernible than the 2 MeV gammas that hydrogen emits.

Other experiments that have used gadolinium-doped water as their neutron-tagging

⁴The peak energy is 5 MeV and ranges from 3-8 MeV as the Compton-scattered electrons eventually go below the Cherenkov threshold for detection.

medium include WATCHBOY⁵ [72] and GADZOOKS!⁶ [70]. At the time of this writing, the Super-Kamiokande experiment is beginning its Gd phase, Super-K-Gd (SK-Gd). While WATCHBOY uses GdCl_3 as its source of Gd, GADZOOKS! and ANNIE use $\text{Gd}_2(\text{SO}_4)_3$ due to the less reactive nature of SO_4^{2-} ions. The first milestone of the GADZOOKS! project is the EGADS⁷ demonstrator [73].

3.2.2.1 Cherenkov Radiation

While gadolinium is an important component, the majority of the target is pure, deionized water. When light travels through a dielectric medium such as water, its speed is reduced to $c' = c/n$, where c is the speed of light in vacuum and n is the index of refraction of the medium and $n > 1$. When this happens, it is then possible for energetic, charged particles to travel faster than light in the same medium, much like a sonic boom.

A charged particle with sufficient energy traveling at a velocity $\beta > 1/n$ through water leaves a trail of polarized molecules. Photons carry away the excess energy as the molecules relax, creating the phenomenon known as *Cherenkov radiation*, which can be detected in transparent media. The photons constructively interfere, producing a wavefront that defines a cone with an opening angle that satisfies

$$\cos \theta_C = \frac{1}{n\beta} \quad (3.1)$$

⁵A 2-ton prototype for WATCHMAN (WATER CHerenkov Monitor for AntiNeutrinos).

⁶Gadolinium Antineutrino Detector Zealously Outperforming Old Kamiokande, Super!

⁷Evaluating Gadolinium's Action on Detector Systems

Medium	Index of Refraction (n)	Cherenkov Angle (θ_C)
Air	1	$< 1.4^\circ$
Water	1.33	42°
Mineral Oil	1.47	47°
Diamond	2.42	63°

Table 3.1: Examples of media of different refractive indices n and Cherenkov angles θ_C .

For ultra-relativistic particles with $\beta \approx 1$ traveling through water ($n = 1.33$), the Cherenkov angle is constant and has the maximum value of 42° . Table 3.1 lists examples of other media used in physics.

The number of Cherenkov photons per unit distance traveled per wavelength emitted by a particle with charge Ze can be described by the Frank-Tamm formula⁸

$$\frac{d^2 N}{dx d\lambda} = \frac{2\pi\alpha Z^2}{\lambda^2} \left(1 - \frac{1}{\beta^2 n^2(\lambda)} \right) = \frac{2\pi\alpha Z^2}{\lambda^2} \sin^2 \theta_C \quad (3.2)$$

where α is the fine structure constant. The β^{-2} dependence indicates that the rate of photon emission is weakly dependent on the particle's energy, so the number of photons emitted per unit distance is roughly constant. The λ^{-2} dependence results in a greater intensity at shorter wavelengths, hence why the “blue glow” is often seen in nuclear reactor cores (cf. Figure 3.8).

For an ultra-relativistic particle of unit charge, the average number of photons emitted per cm is ~ 390 in the wavelength region of 300-700 nm. Thus, it is important that we keep

⁸Named after physicists Ilya Frank and Igor Tamm, who developed the theory in 1937 and won the 1958 Nobel Prize for their work.



Figure 3.8: The blue glow of Cherenkov radiation can be seen in the core of a nuclear reactor. This is a consequence of the phenomenon's λ^{-2} dependence, resulting in higher intensities of light at shorter wavelengths.

the transparency of the water target in this wavelength region as high as possible so that we can collect as much light as possible.

3.2.2.2 Material Compatibility

Maintaining high transparency of the water target is important if we want to detect as many Cherenkov photons as possible for accurate event reconstruction. Materials can break down or leach contaminants when submerged in solution, decreasing target transparency. As such, all materials are soaked in a 1% gadolinium (III) sulfate solution and must meet compatibility requirements before installation in the detector tank. The intent of using a

solution with a ten times higher concentration of Gd is to accelerate the appearance of any detrimental effects.

I led the efforts in the compatibility studies at UC Davis, which were later continued by others. My test procedure is as follows:

1. A batch of 1% Gd sulfate solution is prepared by dissolving 10 g of $\text{Gd}_2(\text{SO}_4)_3$ into ~ 1 L of deionized water. This is referred to as the soak solution.
2. Material samples are cleaned with ethanol and rinsed in deionized water to remove oils and dust and placed in glass jars filled with the soak solution.
3. A small volume sample of the Gd sulfate solution (no contact with any materials yet) is placed in a spectrophotometer and its UV-vis spectrum is taken. This is the control.
4. A sample of the soak solution is taken from the jar that a material is soaking in for spectrophotometry measurements.

Regular measurements were taken with a Shimadzu UV-1800 UV/Visible Spectrophotometer to monitor water transparency over the course a few weeks to several months. Materials tested include entire photomultiplier tubes, acrylic, PVC, and stainless steel.

The criteria for deciding whether a material makes the cut are that the absorbance remains nearly unchanged and no additional features show up compared to the control sample. As shown in Figure [3.9](#), the absorbance at shorter wavelengths increases over the course of a few days and weeks for some materials. However, since the photodetectors




Figure 3.9: Examples of spectra for different materials soaking in 1% gadolinium-loaded water. Samples of the soak solution are analyzed frequently in the first few weeks of testing, and gradually reduce in frequency. The increasing absorbance at wavelengths below 300 nm will not affect light collection in ANNIE, but may indicate that something is leaching into the water.

are opaque to wavelengths below 300 nm, this is not a concern, though it may indicate that something is leaching into the water. To make an assessment based on the ANNIE detector, the ratio of exposed surface area of the material and the soak volume is scaled up to extrapolate absorbance and transparency values and attenuation lengths for the actual detector size.

Another factor that would disqualify a material from installation in the detector is if it physically starts to degrade. Some examples are in Figure 3.10. Clumps of material can be seen suspended in the solution, scattering light and reducing the attenuation length. Less light would be observed by the photodetectors.

The main record of materials and their properties can be found here [74].




Figure 3.10: Materials degrading in gadolinium loaded water. The blue arrows and circles indicate that the material is deteriorating in the solution. This eliminates the material from use in the detector.

3.2.3 Water Filtration System

While we do our best to prevent contaminants from entering the water target, that is not enough to maintain the water transparency. Over time, dust, iron ions, and other particulates will make their way into the water target. If no actions are taken, they will accumulate in the tank, absorbing and scattering light and attenuating the signal strength by the time it reaches a photodetector. This prevents us from seeing events that otherwise would have been seen by the ANNIE detector.

In addition to electropolishing the inner structure, wrapping it with PTFE tape, and taking care to select compatible materials, the water target is circulated through a custom filtration system. The system is scaled down from the larger filtration system of the EGADS demonstrator [75]. Referring to Figure 3.11 for a diagram of the custom filtration system

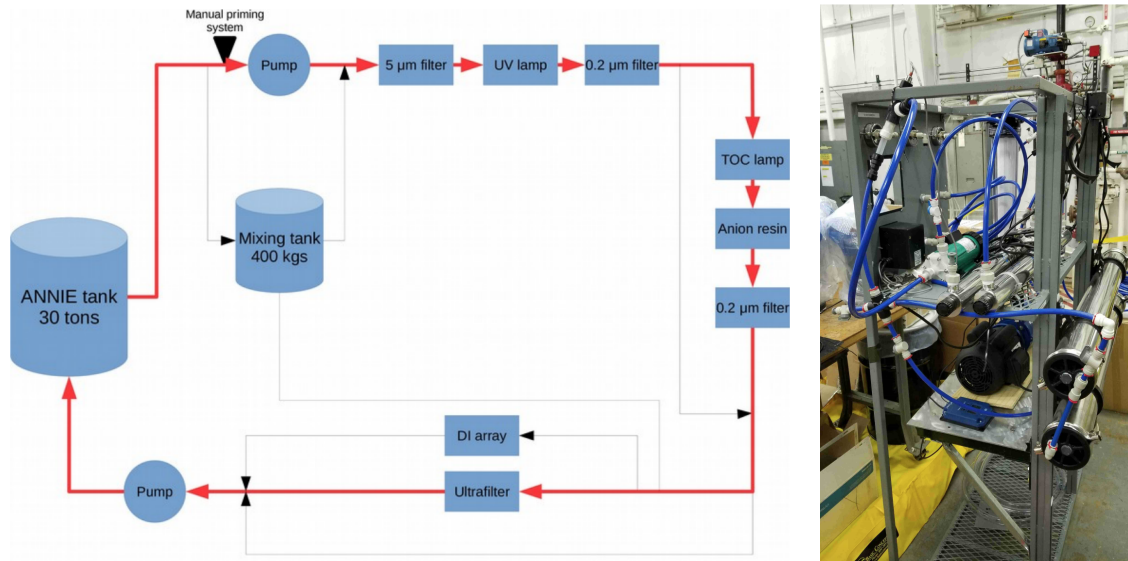


Figure 3.11: LEFT: Diagram of water flow through the ANNIE’s custom water filtration system. Black arrows mark points where the flow can be manually stopped for maintenance and testing. RIGHT: Testing the water system at Fermilab before installation in the hall. Blue tubing connect the pump, filters, ion exchanger, and lamps.

developed at UC Davis [76], the water from the main tank first encounters a 2-micron (μm) filter that is intended to remove large particles such as dust. Next the water is filtered with a ion exchange resin that removes negative ions while keeping Gd^{3+} ions in the water. This step in the filtration system removes any contaminant ions that may degrade other materials and produce more contaminants in the detector tank. Ref. [76] contains more details on selecting and producing the best ion exchange resin for ANNIE. After passing through the ion exchanger, the water then passes a 0.1-micron filter that removes even smaller particulates, such as the ion exchange beads that may come out of the ion bed. Finally, the water is sent through UV and TOC⁹ lamps that are used to stop microorganism

⁹Total Organic/Oxidizable Carbon

growth and break down any carbon compounds.

Inside the tank itself, a nitrogen blanket sits above the surface of the water to displace any oxygen molecules, further preventing microorganisms from growing and contaminating the water. Nitrogen is also circulated throughout the water.

The water system was used to load gadolinium, a process in which I was involved. The mixture was created by dissolving 52 kg of $\text{Gd}_2(\text{SO}_4)_3$ in pure water to make batches of 1% gadolinium solution. The concentrated solution is loaded into a 110-gallon mixing tank and circulated in the filtration system before it is mixed with the tank water. Following each loading is a wait period to monitor the transparency and quality of the water. When no significant changes are perceived, the next batch is loaded until 100% loading is achieved.

3.2.4 Photodetectors

ANNIE Phase II is instrumented with 132 conventional photomultiplier tubes (PMTs) and 5 novel Large Area Picosecond Photodetectors (LAPPDs), providing about 20% photocathode coverage. Of the 132 PMTs, 44 of them were purchased new from Hamamatsu, 48 are from the WATCHBOY experiment, 4 borrowed from WATCHMAN, 20 are Hamamatsu PMTs from LUX, and 20 are from Electron Tube Enterprises Ltd (ETEL). Table 3.2 lists the specifications of each PMT.

The LUX PMTs are installed in an array at the base of the inner structure. The WATCHBOY, WATCHMAN, and ANNIE PMTs are distributed throughout the side panels of the

PMT Type	Model	Diameter	QE	# in tank
ETEL	ETEL D784UKFLB	11"	25%	20
LUX	Hamamatsu R7081	10"	25%	20
WATCHBOY	Hamamatsu R7081	10"	25%	48
WATCHMAN	Hamamatsu R7081-100	10"	35%	4
ANNIE	Hamamatsu R5912-100	8"	35%	44

Table 3.2: The different types of PMT installed in the ANNIE Phase II detector. Listed for each PMT type is the PMT model, diameter, average quantum efficiency, and number installed.

inner structure. The ETEL PMTs are installed at the top, underneath the tank lid, since their ability to withstand the water pressure is uncertain.

Before their deployment in the detector tank, a portion of the total PMTs were tested at UC Davis and Fermilab to characterize their features and properties under lab settings. The WATCHBOY and ETEL PMTs were tested in a dark box at UC Davis with an LED to measure their gain values (see Chapter 5). All PMTs were tested in a dark box at Fermilab to determine or confirm gain values. More information about the lab testing of PMTs can be found in [77].

Unlike the PMTs which need to be installed before water fill, LAPPDs can be deployed at any time along one of the eight Unistrut rails located under a mail slot on the tank lid (see Figure 3.7). Chapter 4 provides a deeper discussion on LAPPDs.




Figure 3.12: An image of the Front Muon Veto (FMV) taken from [77], before the steel tank was placed in the hall. The two layers are placed in opposite orientations and are slightly offset in height from each other to make up for the decreasing efficiency of the scintillator paddles at far end and cover any gaps between layers.

3.2.5 Front Muon Veto (FMV)

Upstream of the tank is the Front Muon Veto (FMV): two layers of plastic scintillator paddles stacked horizontally on top of each other (Figure 3.12). Each layer has 13 paddles for a total of 26 paddles, all recycled from the CDF¹⁰ experiment at Fermilab [78]. Each paddle is 322 cm long, 31 cm wide, and 2.5 cm thick, and is coupled to a 2-inch (5-cm) diameter EMI-9815 PMT via a tapered glass light guide. Typical operating voltages of the PMTs are in the 2,000-2,500 V range. The two layers are vertically offset from each other by 0.5 cm to cover any gaps between paddles. Overall, the FMV is about 3.2 m wide and 4 m tall.

¹⁰Collider Detector at Fermilab




Figure 3.13: The MRD is an iron-scintillator sandwich held by an aluminum support frame. The scintillator paddles are coupled to PMTs and alternate between horizontal and vertical orientations. Diagram from [77] and picture from [67].

The FMV tag muons and other charged particles that were produced upstream and therefore did not originate in the tank. Tagging upstream muons allows their removal from the event sample.

3.2.6 Muon Range Detector (MRD)

Similar to the FMV, the Muon Range Detector (MRD) tags muons that make their way downstream of the tank. These are most muons since the ANNIE tank is relatively small, so they lose most of their energy and decay outside of the tank. The MRD is a legacy of the SciBooNE experiment [79], which ran in the same experimental hall that ANNIE now occupies. It comprises alternating layers of iron plate absorbers, which slow down the muons,

and scintillator paddles, which mark their passage through the MRD. Each iron layer has an area of 274 cm x 305 cm with a thickness of 5 cm and an average density of 7.841 g/cm³ (Figure 3.13). The scintillator layers consist of paddles that alternate in orientation and have various dimensions. The horizontally oriented paddles, all originally from SciBooNE, are 20 cm wide, 155 cm long, and 6 mm thick. The original vertical paddles are 20 cm wide, 138 cm long, and 6 mm thick. Two vertical layers were refurbished with components of the paddles from the KTeV experiment. These are 15 cm wide, 150 cm long, and 14 mm thick. The alternating orientation allows for xy -position reconstruction of the particle's track. Information from the MRD is also used to reconstruct the momentum and direction. In total, there are 306 MRD paddles spread over 6 horizontal and 5 vertical layers, with a total of 11 iron slabs and an overall dimension of 3.5 m wide and 4 m tall. The paddle efficiencies range from 80-95%. More details about the MRD in ANNIE such as PMT model and refurbishment process can be found in [67] and [77].

3.2.7 DAQ and HV

ANNIE's electronics system handles the collection of waveforms, timestamps, and so on, so that data analysis can be performed offline. It also powers the photodetectors in the tank. The FMV and MRD PMT signals are collected by Time-to-Digital Converters (TDCs), so only time information is saved when it comes to these PMTs; there is no charge information from the FMV and MRD. Their signals are discriminated by LeCroy 4413 CAMAC discrim-

inators before readout by LeCroy 3377 TDCs. High voltage for these PMTs are supplied by LeCroy 1440 HV crates.

On the other hand, readout of the tank PMT waveforms are managed by a VME system, and both the time and waveform information are saved so that charge information can be extracted. These signals are digitized by Analog-to-Digital Converter (ADC) cards and may be amplified if the PMT gain is too low to meet thresholds set at the discriminator level. The tank PMTs are powered by CAEN SY527 crates and their waveforms are recorded with VME digitizer cards originally designed for the KOTO experiment [80]. Both the CAEN and LeCroy HV crates can be remotely controlled by experts and monitored by shifters.

The integration of LAPPDs to the existing DAQ and HV systems require some modifications to the hardware and firmware, which is currently underway. Readout of the LAPPDs is performed by custom made electronics boards, which are discussed in Chapter 4.

3.2.7.1 Trigger Scheme

ANNIE has several trigger schemes for different data acquisition scenarios. Various data-taking runs include

- **Beam:** data acquisition with physics quality beam;
- **Cosmic:** data acquisition when a cosmic muon triggers a readout;
- **LED:** data acquisition with LED light sources; and
- **AmBe:** data acquisition with an AmBe run deployed in the water tank.

When the Booster Neutrino Beam is on, Fermilab's Accelerator Division sends a signal informing experiments that beam is on its way. ANNIE uses this beam signal to prepare the electronics for data acquisition. A $2\text{-}\mu\text{s}$ window is open and is extended to $70\text{-}\mu\text{s}$ if there is activity in the tank. At least one PMT with a registered hit equivalent to 5 photoelectrons is required to extend the acquisition window. These conditions make up the beam runs. Occasionally, random beam trigger windows are extended for studies on beam-correlated activity. In cosmic runs, the trigger scheme is set up to readout all subdetector systems if the first and last layers of the MRD register coincident hits from cosmic muons passing through the detector.

In LED runs, the LEDs that are installed on the inner structure are fired and a readout of the entire tank occurs. No FMV or MRD PMT data is recorded in the acquisitions. The goal of these runs is to monitor the transparency of the water. The LEDs were also used to get a first estimate of cable lengths and delays. However, a laser source was later used to better quantify timing offsets.

AmBe runs are used when an AmBe source [\[5\]](#) is deployed in the tank to characterize the neutron detection efficiency of the detector. When the PMTs coupled to the source register a hit that meet a threshold, waveforms from all the tank PMTs are recorded.

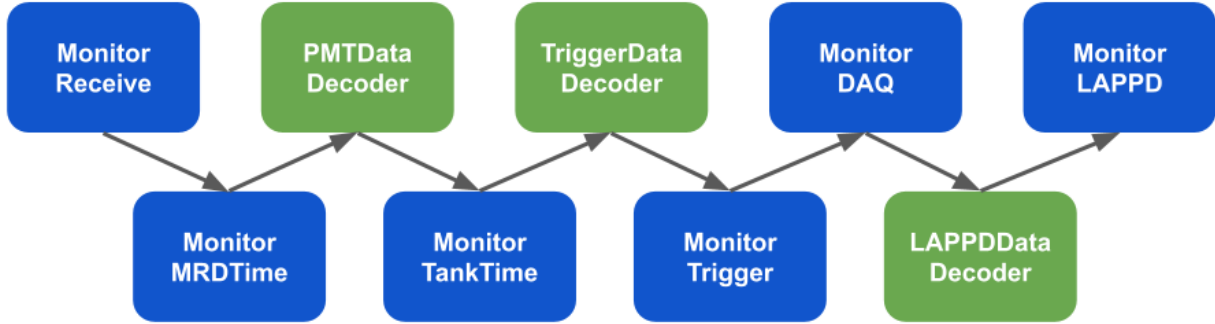


Figure 3.14: The monitoring toolchain used to monitor the subsystems of the ANNIE detector. The toolchain runs continuously to post up-to-date conditions for shifters.

3.2.8 ToolAnalysis and ToolDAQ

ToolAnalysis is a modular analysis program developed for data analysis in ANNIE. It is based on the ToolDAQ [81] framework written by Dr. Ben Richards¹¹, which is used in the DAQ systems of ANNIE and Hyper-Kamiokande [82]. Users can create tools that make up software toolchains to process and analyze data. The data and other variables are stored in `BoostStores` accessible to all tools. The diagram in Figure 3.14 depicts an example toolchain.

Each tool has an `Initialise` function that retrieves variables from configuration files and sets up input parameters, geometry, output files, and so on, that usually remain unchanged throughout the analysis chain. After initialization, any calculations and analysis techniques that need to be performed on multiple events are placed under the `Execute` function. This function is looped over for each tool until all entries are processed. Finally, the `Finalise`

¹¹University of Warwick

function closes out files, clears variables from memory, and writes information to file if the user specifies.

ANNIE’s monitoring system makes up one of the toolchains used. Monitoring tools running in a continuous tool chain produce and post plots to monitoring webpages for shifters to identify when things go wrong. More information about monitoring can be found in Chapter 5 on calibration of the ANNIE detector. Other toolchains are in charge of readout and control of other components of the DAQ system.

3.3 Booster Neutrino Beam (BNB)

The Booster Neutrino Beam (BNB) is one of the active beamlines at the Fermi National Accelerator Laboratory (Fermilab) in Batavia, Illinois. It is the primary source of neutrinos for current and past experiments such as ANNIE, SBND, ICARUS, SciBooNE, MiniBooNE, and MicroBooNE. The other neutrino beamline is the Neutrinos from the Main Injector (NuMI) beam. NuMI delivers neutrinos to the MINERvA, MINOS, and NO ν A experiments.

For the BNB to provide a steady beam of muon neutrinos, several key steps are required. First, a source of protons is created with the *Proton Source*, which consists of the Pre-Accelerator, Linac, and Booster (see Figure 3.15). Two hydrogen (H^-) ion beams of 35 keV are focused, chopped¹², and accelerated to 750 keV with magnets and radio-frequency quadrupoles (RFQs) of the Pre-Accelerator. The Linac segment of the Proton Source then

¹²An Einzel lens is used as a beam chopper. Beam chopping provides more control over the length and refines portions of the beam pulse that move on to the next stages of the accelerator complex.

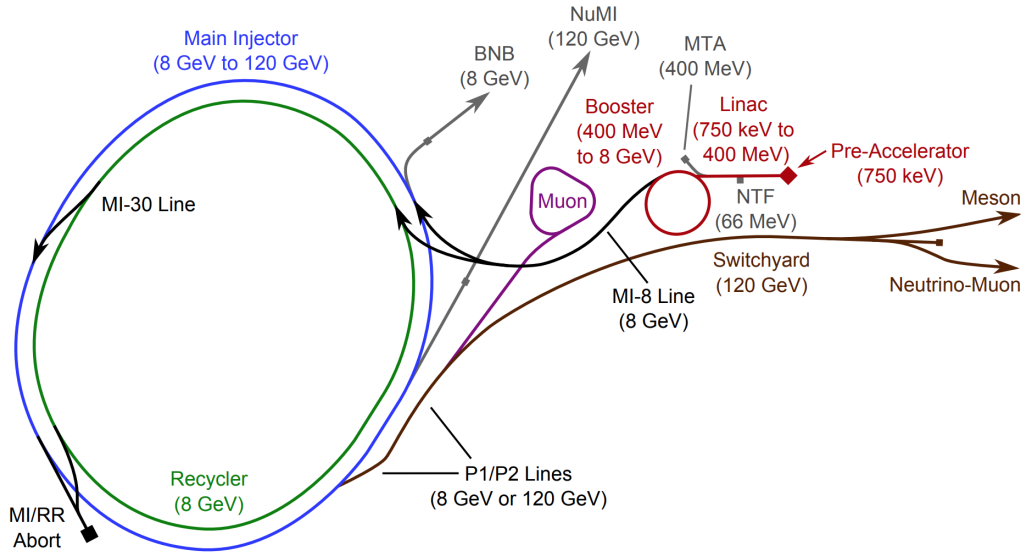


Figure 3.15: A schematic of Fermilab accelerators. The Pre-Accelerator, Linac, and Booster make up the Proton Source, which converts a beam of H^- ions into an 8 GeV proton beam. The protons are sent to several destinations, one of them to make the Booster Neutrino Beam.




Figure 3.16: Diagram of the Booster Neutrino Beam target system at Fermilab. The Booster delivers an 8 GeV beam of protons to a beryllium target, creating pions and kaons that are focused with a magnetic focusing horn. The mesons decay into muons and muon neutrinos. The muons are absorbed and what is left is a muon (anti)neutrino beam.




Figure 3.17: The Booster Neutrino Beam (BNB) energy spectrum shows that peak neutrino production occurs around 700 MeV. Figure from [83].

accelerates the H^- ion beam to 116.5 MeV in the low energy section and to 400 MeV in the high energy section. The 400 MeV beam is then directed to several destinations, one of them being the Booster, a 74.5 m radius (474.2 m circumference), rapid-cycling synchrotron. Before injection into this final stage of the Proton Source, the H^- ions are passed through a stripping foil to remove the electrons, resulting in a beam of protons. The Booster accelerator, using a series of magnets and RF cavity stations, focuses and boosts the 400 MeV protons to 8 GeV. The resulting beam has 84 proton bunches and a revolution period of $1.6 \mu s$. The MI-8 Line delivers the proton beam from Proton Source to the Booster dump, the Main Injector (MI), the Recycler, or the BNB beamline.

To create the BNB, the protons are directed at a beryllium (Be) target, creating positive and negative pions and kaons. The pions decay into muons and neutrinos. The kaons would also decay into muons and pions, which in turn decay into more muons and neutrinos. The mesons are sent through an electromagnetic focusing horn, originally built for MicroBooNE, where magnets are used to select the desired charge, or “sign”, of the pions. Flipping the magnetic field produces a neutrino beam of the opposite sign. The beam of pions is collimated about 259 cm downstream of the Be target and sent to a decay region of about 50 m in length, where they can decay into muons and muon neutrinos. A steel and concrete absorber terminates the decay region, stopping the muons while the muon neutrinos continue forward. The resulting neutrino beam is about 93% ν_μ and 6.4% $\bar{\nu}_\mu$ from $p + Be \rightarrow \pi^- \rightarrow \mu^- + \bar{\nu}_\mu$ and 0.6% ν_e from decays of μ^+ produced in the main π^+ decay channel and other channels. More details about the accelerator, BNB, and other beamlines at Fermilab can be found in [84] and [85].

ANNIE sits about 100 m downstream from the beryllium target, placing it in a on-axis of the BNB. About 5×10^{12} protons-on-target (POT) per spill are delivered to the BNB with an average rate of 5 Hz. With these numbers, about 10,000 charged-current events per m^3 per year in the ANNIE detector [66]. ANNIE is expected to see $\mathcal{O}(10^{20})$ POT per year. This corresponds to a flux of $\sim 5.19 \cdot 10^{-10} \nu/\text{cm}^2/\text{POT}$ [83]. The flux of this beam peaks at about 700 MeV (Figure 3.17), which is in the energy region of interest of long baseline and atmospheric neutrino experiments. The close proximity to the target provides a high flux of neutrinos, which means more statistics!

Monitoring the beam conditions is necessary for precise physics measurements. When the accelerator is not producing physics quality beam, then very few physics events can be studied. Four systems are in place upstream of the target to monitor the BNB: two toroids (TOR875, TOR860) that measure beam intensity (protons-per-pulse), beam position monitors (BPM) and a multiwire chamber to determine the beam width and position, and a resistive wall monitor (RWM) to measure the time and intensity of beam spills. Measurements from these beam monitoring devices are used to identify periods of acceptable beam quality for data analysis.

3.4 Physics Program

ANNIE's physics goals are to measure the final state neutron multiplicity of neutrino interactions in water and the charged current (CC) inclusive neutrino cross section. These measurements will benefit next generation, long-baseline neutrino experiments that are studying neutrino oscillation, the Diffuse Supernova Neutrino Background (DSNB), and proton decay (PDK), by characterizing relevant backgrounds and reducing systematic uncertainties.

The ANNIE detector also functions as a testbed to demonstrate new technologies. For example, as will be discussed in a later chapter, ANNIE is the first physics experiment on a neutrino beam to deploy LAPPDs. By demonstrating the capabilities of this picosecond technology, production of these photodetectors can scale up, reducing the cost of one unit. Other experiments can also take advantage of these devices, along with the custom electronics

developed for the LAPPD system. ANNIE will demonstrate that even with just a few LAPPDs, experimental measurements can be improved significantly. With LAPPDs, ANNIE may also demonstrate multi-track reconstruction, distinguish between overlapping events, and better resolve the structure of electromagnetic showers.

ANNIE is also a testbed for new detection media: gadolinium-doped water and water-based liquid scintillator (WbLS). It is also the first Gd-doped water Cherenkov detector on a neutrino beam. Being able to identify neutrons more easily complements experiments that are unable to see neutrons but can see protons. Combining Gd-doped water and WbLS allows for the separation of Cherenkov radiation (to detect the muon) and scintillation light (to detect the proton), so that we can better reconstruct neutrino events. Coupled with the picosecond timing capabilities of the LAPPD, ANNIE will be able to demonstrate the power of these new technologies.

3.4.1 Neutron Multiplicity Measurement

Recall in section [2.4](#) that it is important to quantify the number of neutrons that are produced in neutrino-nucleus interactions because they can take away a significant fraction of the energy from the outgoing lepton, thus biasing event reconstruction. Existing neutron multiplicity measurements carried out by the SNO and Super-K experiments (see Figure [3.18](#)) were made with atmospheric neutrinos, which are a mixture of all flavors and a significant fraction of antineutrinos. The direction and sign of the neutrinos and antineutrinos were

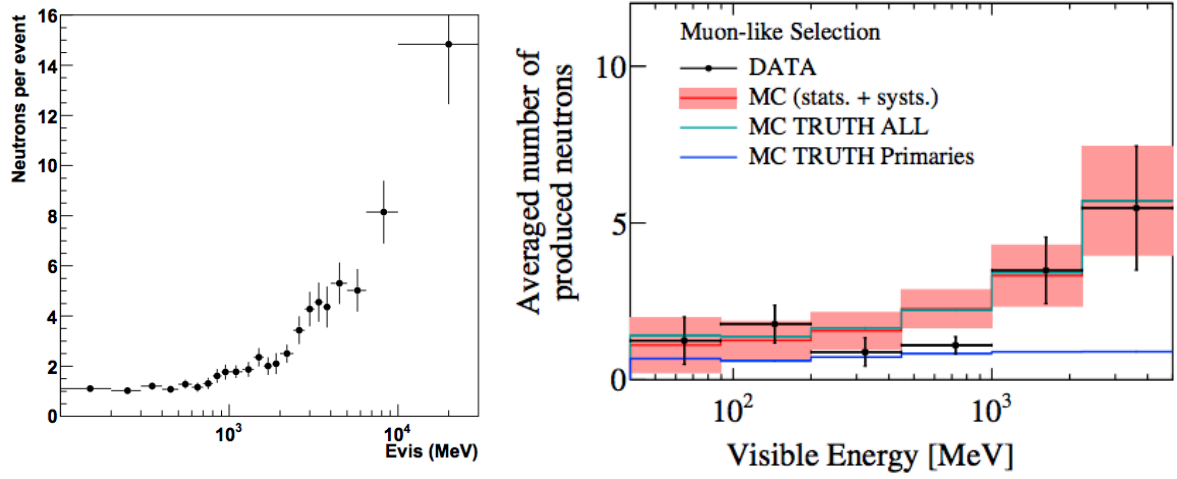


Figure 3.18: Neutron multiplicity measurements by Super-K [86] (left) and SNO [87] (right) experiments. Measurements were made with atmospheric neutrinos, whose sign and direction are not well-known. In addition, neutron yields are characterized as a function of visible energy.

also not known. ANNIE will improve upon this measurement due to its unique position. For one, ANNIE sits on a neutrino beam, so the statistics are significantly larger, the energy of the beam is constrained to a narrower range, and the neutrino flavor composition is purer (single flavor and sign). Being on a neutrino beam also gives us the knowledge of the neutrino direction. Secondly, the Super-K and SNO measurements of neutron multiplicity by were made as a function of visible energy. This can be made more precise by presenting the measurement as a function of muon momentum, or even neutrino energy. Finally, Super-K's and SNO's measurements were performed in pure water and heavy water, respectively, where the neutron detection efficiency is very low. ANNIE (and SK-Gd) has a much improved neutron detection efficiency due to its Gd loading.




Figure 3.19: A typical event in ANNIE. (a) A neutrino-nucleus interaction creates a muon that undergoes Cherenkov radiation as it travels out of the water tank and ranges out in the MRD. (b) Final state neutrinos are created and thermalize (few μs). (c)-(d) Neutrons capture on gadolinium and produces flashes of light to be observed by photodetectors.

The neutrino interactions of interest are:

$$\begin{aligned}\nu_{\mu} + n &\rightarrow \mu^{-} + p \\ \bar{\nu}_{\mu} + p &\rightarrow \mu^{+} + n\end{aligned}\tag{3.3}$$

Since the BNB is 97% ν_{μ} , the former reaction will make up the majority of events. However, there is some contamination from $\bar{\nu}_{\mu}$ so the latter reaction will also take place in the detector. The daughter muon travels out of the detector and may stop in the MRD. A typical ANNIE event is shown in Figure [3.19](#). First, a muon neutrino interacts with a nucleus and a muon is created. The muon emits Cherenkov radiation as it exits the water tank and ranges out in the MRD, providing a prompt signal. Any final state neutrons resulting from nuclear processes thermalize until they are captured by gadolinium. The de-excitation of gadolinium produces a delayed signal in the tank.

To reconstruct the muon energy, charge and timing information from the detector tank

along with track information from the MRD are used to reconstruct the muon vertex and direction. Using simple kinematics, the Q^2 value of the neutrino interaction can be extracted. The number of neutrons can be determined by identifying clusters of hits that fall into the delayed/extended window, and can be plotted against the muon's momentum transfer. While various final state topologies exist for neutrino-nucleus interactions, this thesis focuses on the CC-0 π event topology (CC events with no pions) so that a simple kinematic model can be used to reconstruct the neutrino energy from lepton kinematics. Future studies will look at CC-1 π (CC events with one pion), NC, and other event topologies.

3.4.2 Cross Section Measurement

As mentioned in the previous chapter, oxygen has very few cross section measurements despite its common use as a target in water Cherenkov neutrino detectors. A precise measurement of the CC inclusive cross section in water is beneficial for the next generation of neutrino oscillation experiments. These experiments require low systematic uncertainties to be able to achieve their physics goals. ANNIE is in a good position to improve the cross section measurements in water due to its high statistics of neutrino interaction events. In addition, since the SBN¹³ near detector is downstream of the ANNIE hall, there is an opportunity to perform a joint argon and oxygen cross section analysis and a combined proton and neutron multiplicity measurement. Although efforts to perform this measurement are ongoing, this thesis will focus on the CC-0 π neutron multiplicity measurement.

¹³Short-Baseline Neutrino Program

3.4.3 Water-based Liquid Scintillator Technology

There is growing interest in combining the capabilities of water Cherenkov detectors with the capabilities of liquid scintillators, and being able to detect and separate scintillation light and Cherenkov radiation. The motivation for this is the fact that scintillation light can provide additional calorimetric information to improve energy resolution and the ability to detect neutral particles with energies that don't meet the Cherenkov threshold. Various methods such as spectral photon sorting, time separation with fast photodetectors, and development of novel detection medium (e.g., slow scintillators) were explored. ANNIE would use fast photodetectors, namely the LAPPDs, in combination with a WbLS medium to demonstrate this capability.

In fact, ANNIE has already done so. Details of this technical demonstration can be found in [\[88\]](#). In short, ANNIE deployed a small acrylic vessel filled with WbLS in a program called Scintillator in ANNIE for Neutrino Detection Improvement (SANDI). The demonstration was successful and the tools developed to distinguish scintillation from Cherenkov radiation create a basis on which to build and develop vertex reconstruction and particle identification tools.

As we can see, there are many physics opportunities with ANNIE. This rest of this work focuses on the development and testing of LAPPDs and a $CC-0\pi$ neutron multiplicity measurement.

Chapter 4

The Large Area Picosecond Photodetector

The Large Area Picosecond Photodetector (LAPPDTM) is a 20 cm by 20 cm photodetector that takes advantage of breakthrough technologies to deliver fast timing and high spatial resolution. This novel photodetector was developed by Incom Inc., in partnership with the U.S. Department of Energy, Argonne National Laboratory, and University of Chicago. LAPPDs exhibit excellent timing and localization of signals, making them great for applications such as particle tracking, nuclear science, and medical imaging. Refs. [\[89\]](#), [\[90\]](#), and [\[91\]](#) list potential imaging applications that could change the way we detect particles.

4.1 Conventional Photodetectors

4.1.1 Photomultiplier tubes

Let's first begin with the conventional photomultiplier tube (PMT). The PMT is a vacuum tube designed to detect low levels of light by converting photons into measurable electric currents. PMTs have been around since the 1930s [92] and are used in a plethora of experiments today. They are quite robust and may operate for decades, with the same PMT being used in multiple experiments throughout its lifetime. In general, a photomultiplier tube consists of

- a **glass window** to admit light,
- a **photocathode** made of a photoemissive material (typically a bialkali) deposited onto the inner surface of the glass window that converts a photon into an electron,
- **several dynodes** which are electrodes coated in an emissive material and that release secondary electrons in response to impinging electrons, and
- an **anode** to collect the resulting electron avalanche.

Additionally, there are electronic structures that accelerate and focus the electrons. Some PMTs of varying sizes and shapes are shown in Figure 4.1.

PMTs operate based on the principle of the photoelectric effect. When a photon of sufficient energy is transmitted through the glass window, it is absorbed by the photocathode




Figure 4.1: TOP: A selection of Hamamatsu photomultiplier tubes (PMTs) of varying shapes and sizes [93]. BOTTOM: PMTs are vacuum tubes that can detect a single photon through the photoelectric effect. A layer of photocathode material, typically a bialkali, is deposited on the underside of the glass window. When a photon is incident on photosensitive area, it may be converted into an electron. This electron is promptly accelerated through the vacuum to a chain of dynodes, metal objects that are coated with a layer of emissive material. The dynodes emit additional electrons when struck by incident electrons, thus creating an avalanche effect. Charge is collected on an anode at the end of the chain and is read out [94].

and an electron is emitted in response. This electron, referred to as the “photoelectron” (PE), is then accelerated through the vacuum volume by a strong electric field. Since the charge of an electron is quite small ($1.602 \cdot 10^{-19}$ Coulombs), amplification of the photoelectron signal is a necessity. This amplification is done using metal structures, called dynodes, that are coated with an emissive material. As the photoelectron accelerates down the PMT, it strikes a dynode resulting in the emission of several secondary electrons. These secondary electrons make their way down to other dynodes through a potential of about 100 V, knocking off more electrons in the process and creating an avalanche effect. Eventually a large enough current (on the order of 10^{-12} Coulombs, or picoCoulombs) can be detected at the anode of the PMT. The number of electrons that result from the initial photoelectron is called the gain of the PMT and is dependent on the operating voltage. A typical gain value is 10^7 . Another parameter commonly used to characterize a PMT is the quantum efficiency (QE), which is the likelihood that the incident photon is converted into a photoelectron. For more information about the different types of PMTs and their specifications, refer to [95] and [96].

Because PMTs typically have one amplification stage and one charge collection anode¹, they are considered “single pixel” detectors and their position resolution is limited by the area of the photosensitive surface. In addition, since the electrons need to travel through the vacuum tube and down a chain of dynodes, the timing response is largely determined by the myriad paths that the electrons can take [96] and thus the size of the PMT. Typical time resolutions are on the order of nanoseconds.

¹There are PMTs that have multiple amplification stages and charge collection anodes.

4.1.2 Microchannel plates

A microchannel plate (MCP) detector is a particle and photon detector that amplifies signals by the emission of secondary electrons, much like the PMT. However, unlike the PMT, the MCP provides better spatial resolution. The development of higher resolution MCPs has been driven largely by night vision systems.

In its most basic design, a MCP is an array of closely-packed channels, or “pores”. The interior of each channel is coated in a layer of emissive material, like the PMT dynodes, so that each one functions as an electron multiplier. Each channel acts as a continuous dynode and leads from one face of the MCP to the other (refer to Figure [4.2](#)). A voltage is applied across the plate to accelerate the electrons towards a charge collection plane. Channels typically enter the plate in a normal orientation or at a small angle (8° - 15°) from the normal. The slight angle reduces the effects of positive ion-feedback, a phenomenon in which gaseous atoms or adsorbed atoms in the channel surface are ionized by the electron avalanche and accelerated in the opposite direction. If the ion gains enough energy, it can start a second electron avalanche and produce after-pulses.

A plate is fabricated from a resistive material, typically lead glass. Heat treatment follows in a hydrogen-reducing environment to give the pores their secondary emissive properties. One plate has a typical gain of 10^3 - 10^4 , so stacking two MCPs on top of each other and in a Chevron configuration (in which the plates are rotated 180° from each other and the channels create a shallow “v”) can produce gains of 10^6 - 10^7 . MCP channels have diameters




Figure 4.2: A diagram of a microchannel plate (MCP) and its principle of operation [97]. A MCP is a matrix of through-going, closely-packed channels that are 10-100 μm in diameter. The interior of each channel is coated in an emissive material so that when a particle of sufficient energy is incident on the channel, an avalanche of electrons follows. A voltage is applied across the plate to accelerate the electrons towards a charge collection plane. A single plate has a gain of 10^3 - 10^4 and several plates can be stacked to achieve higher gain values. MCPs can be 10s to 100s of millimeters in size and can be cut into any shape (often round disks).

between 10 and 100 microns (μm) and lengths between 400 and 1000 μm [96, 98]. Spatial resolution is limited by the diameter of the pores and the spacing between them, typically 15 μm . Because the total height of a MCP or stack of MCPs is on the order of millimeters, MCPs can deliver timing resolutions that are <100 picoseconds (ps). For more information about MCPs and their properties, see Refs. [96] and [98].

Generally, MCPs are very expensive to produce due to the manufacturing steps. Slabs of lead glass need to be chemically etched, which is an expensive process. In addition, during the hydrogen reduction step to form the emissive layer, lead particles tend to agglomerate at

high temperature, reducing their effectiveness in the electron conduction process [98]. The lifetime of a MCP is affected by what is known as “electron scrubbing”, a phenomenon in which the electron cloud hits the walls of the MCP channels, changes its surface properties, and affects the secondary electron yield. Consequently, the gain drops and the MCP does not perform optimally.

4.2 The Design & Technology Behind the LAPPD

LAPPDs are classified as MCP-PMTs and were designed to detect light with maximal sensitivity to the position and arrival time of photons. They operate similarly to conventional photomultiplier tubes (PMTs) in that both rely on the photoelectric effect to detect photons. However, rather than use a chain of dynodes to multiply the photoelectron, LAPPDs rely on a pair of microchannel plates (MCPs) stacked in a Chevron configuration with a sub-millimeter gap between the plates. Applying a voltage across this inter-plate gap reduces the rise time of pulses to 100s of picoseconds [98]. Figure 4.3 shows a cross-section of a LAPPD.

LAPPDs are an attempt to make MCP photodetectors cheaper. Two key technological breakthroughs enabled the creation of the LAPPD: (1) Incom Inc.’s ability to produce large blocks of hollow, glass capillary arrays (GCAs) with micron-sized channels using low cost glass, and (2) the emergence of Atomic Layer Deposition (ALD), a method of coating objects with thin, uniform layers of material [100]. Recall that in conventional methods, MCP channels are made by chemical etching and a single lead sub-oxide resistive and emissive layer




Figure 4.3: A cross section of the Chevron microchannel plates (MCPs) in a Large Area Picosecond Photodetector (LAPPD) [99]. When an incoming photon is incident on the top window, the photocathode on the inner surface of the window absorbs the photon and emits an electron. Like the PMT and the MCP detectors, the electron is accelerated through the vacuum volume by a strong electric field. The photoelectron initiates a cascade of secondary electrons after it enters the first MCP. The cascade then enters the second MCP after traversing the inter-plate gap and electrons are further multiplied before collection by the anode plane. The many micron-sized channels of the MCPs provide sub-centimeter spatial resolution compared to conventional PMTs. The thickness of the entire stack comes out to about 2 cm.

with variable composition is made during a hydrogen-forming reduction process. Incom Inc.’s proprietary hollow core GCA process eliminates the need for chemical etching to remove core material. ALD allows for the optimization of both the resistive and emissive layers of the MCP, allowing for better performance, stability, and durability for longer periods of time.

What sets LAPPDs apart from traditional MCP-PMTs is their large photocathode area, which is comparable to an 8-inch diameter PMT. However, LAPPDs have better spatial resolution and are less bulky than PMTs that offer similar photocoverage, due to their use

of MCPs and a stripline anode geometry. The glass window, the two MCPs, and the anode plane add up to a height of about 2 cm, including the space between layers [100]. This small form factor gives the LAPPD its fast, picosecond-level timing response since electrons travel a much shorter distance before readout. One way to characterize the timing limitations is to look at the full-width at half-maximum of the transit-time spread (TTS). Transit-time is the amount of time between the emission of the photoelectron from the photocathode and the collection of the resulting pulse's charge. In the case of LAPPDs, the standard deviation of the TTS for a single photoelectron is ~ 50 ps [101]. Figure 4.4 shows a plot of the TTS for a LAPPD.

To accelerate the photoelectron through the vacuum, high voltage is applied at the photocathode and is divided down by a resistor voltage divider circuit for a total of five high voltage points: the photocathode (PC), the top and bottom faces of the first MCP, and the top and bottom faces of the second MCP. At the PC, typical operating voltages are 2,000-2,500 V. The voltage applied across the MCPs is typically 900-1,000 V. The voltage drops in the gaps between the PC, MCPs, and anodes (Figure 4.3) are typically 100-200 V. After the photoelectron is multiplied by the two MCPs, charge is collected on one or more anode striplines of the 28 available. One stripline is about 20 cm long and 5.2 mm wide, with a gap of 1.7 mm between strips [100]. Each can be read out at both ends to determine where along the stripline the photon hit.




Figure 4.4: A plot depicting the transit-time spread (TTS) of a LAPPD. The y -axis is the number of pulses. Transit-time is the amount of time that passes from the creation of the initial photoelectron to the collection and readout of the resulting charge avalanche by the anode. For a LAPPD, this time is affected by the paths that electrons can take inside the MCP pores and noise from the readout electronics [101]. This TTS plot is made by shining a PILAS laser with < 45 ps root-mean-square resolution at an LAPPD, and taking data with a 10 gigahertz (GHz), 10 GSample/s oscilloscope and PSEC4 waveform digitizing electronics [100]. A timing resolution of 64 ps was observed in the main peak. The slight secondary peak to the right is due to after-pulsing.

4.3 Why Are LAPPDs Important in ANNIE?

LAPPDs are necessary for ANNIE to achieve its physics goals with high precision. Thanks to the photodetector’s sub-centimeter spatial and picosecond timing resolution, we can reconstruct the neutrino interaction vertex, the outgoing muon’s momentum, and thus the incident neutrino’s energy with higher accuracy and precision. In a simulation of neutrino events inside the ANNIE detector, one can see in Figure 4.5 that, with the addition of just 5

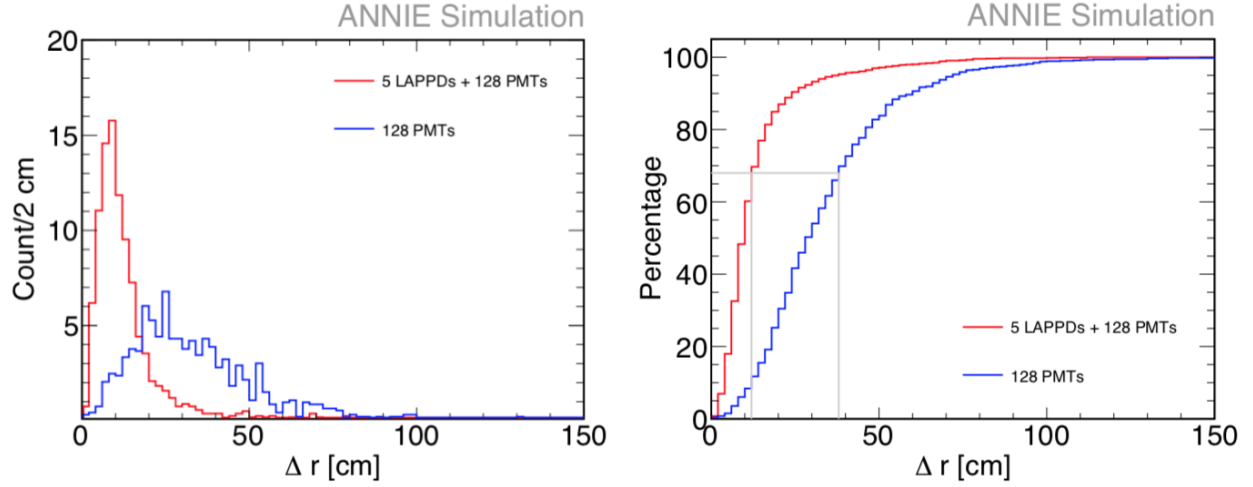


Figure 4.5: LEFT: A distribution of the Δr values. RIGHT: A cumulative distribution of the Δr , normalized to a percentage of successfully reconstructed events [66]. Blue curves represent reconstructed events using 128 PMTs, while red curves represent reconstructed events using 128 PMTs + 5 LAPPDs. Calculations were performed on the same sample of 1800 events. The addition of just 5 LAPPDs improves the vertex reconstruction resolution by a factor of about three, from 38 cm to 12 cm (gray lines). Vertex reconstruction efficiency is also improved to 90% from 85%.

LAPPDs, the position resolution improves by roughly a factor of three, from 38 cm to 12 cm [66]. In addition, the reconstruction efficiency—the efficiency of being able to reconstruct the interaction vertex in a neutrino event—is improved from 85% to 90%.

4.4 Full LAPPD System: Overview & Challenges

The LAPPD by itself is not enough to achieve ANNIE’s physics goals. Auxiliary components are needed for waveform digitization and readout, data transmission, clock and trigger signals, slow controls, and power. Figure 4.6 shows a diagram of the full LAPPD system.

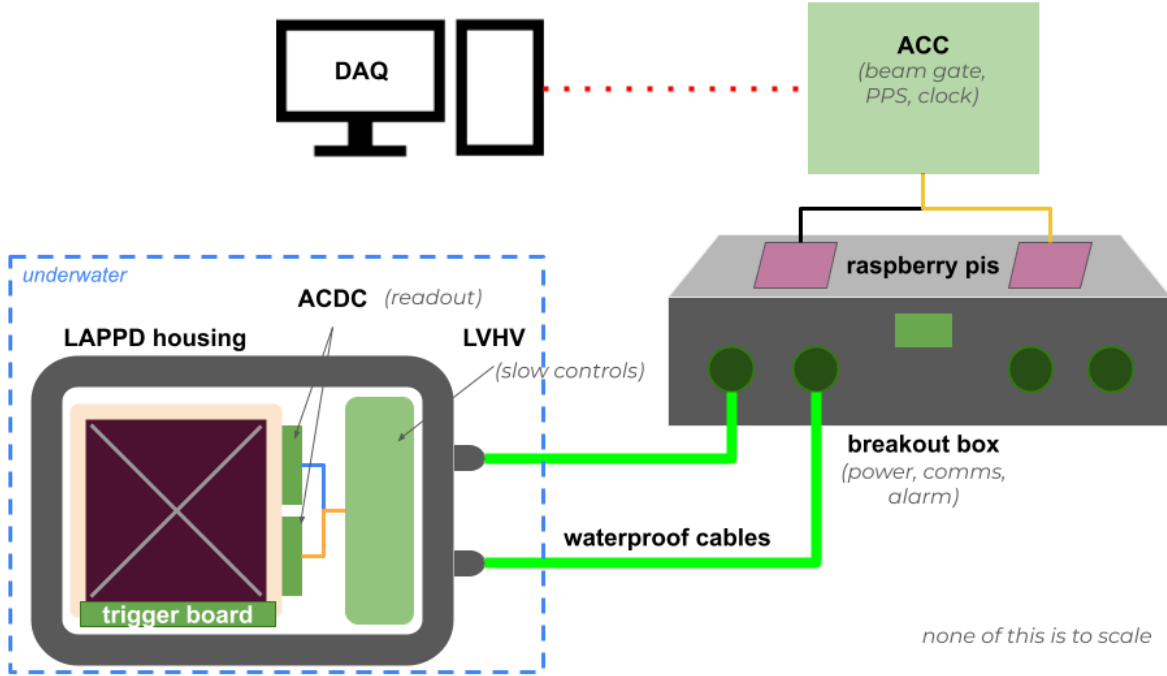


Figure 4.6: A diagram of the full LAPPD system. The LAPPD, along with a trigger board, two ACDC boards, and a LVHV board are encased in a watertight housing. The housing is machined from a block of gray PVC and has an acrylic front window and stainless steel backplate. Two waterproof cables that carry power, data, communication, and alarm lines connect the LAPPD package to a breakout box that manages power and signal distribution. Data signals are directed to an ANNIE Central Card (ACC), which sends the beam gate, PPS (pulse-per-second), and clock signals to the ACDC boards downstream. From the ACC, data is saved into event stores by the ANNIE DAQ.

The LAPPD and its accompanying electronics (described in later sections) are sealed in a watertight housing. Power to and data from the LAPPD are transferred over waterproof cables that are meters long and which connect to a breakout box. This breakout box manages distribution of power and signals and interfaces with an ANNIE Central Card (ACC), which in turn communicates with the main ANNIE Data Acquisition (DAQ) System.

Deploying a LAPPD is easier said than done. There are challenges that come with using this novel photodetector. For one, because the LAPPD is a new technology, many of the system components listed above need to be custom-made. These include the cables, the housing, and the readout electronics. This is also the first time a LAPPD will be deployed remotely in a physics experiment. In its original design, the LAPPD is attached to a circuit board with 60 SMA connectors (cf. Figure [4.7](#)), one for each end of the 28 anode striplines, plus four spare connectors for trigger signal input. The SMA connectors allowed users to read out signals with digitizers in the immediate vicinity. However, this design is impractical for most physics experiments that need to deploy many photodetectors remotely and in a medium like water. Cable management would be quite cumbersome if all striplines are to be read out, and it is necessary to mitigate issues that arise with watertightness and heat dissipation. For small experiments that employ one LAPPD [\[102\]](#), reading out 60 outputs is manageable. Some CAEN digitizers have 8-32 input channels, so one would need 2-4 digitizers to read out both ends of each stripline. For larger experiments, however, one can see how messy it can quickly get.

In the current state of ANNIE, there is not enough electronics rack space or capacity to




Figure 4.7: This LAPPD (#63) is still attached to a circuit board with 60 SMA connectors (one highlighted by a green arrow): 56 for both ends of the 28 anode striplines + four spares for trigger signal inputs. In this picture, an SMA cable (light pink) is connected to one of the 28 output connectors on one half of the circuit board. If all 56 SMA connectors were used, it would be a cabling nightmare! Red arrows indicate power inputs.

manage this many cables and digitizers, especially if there will be 5 LAPPDs deployed in the detector tank. As a result, the ANNIE Analog Pickup Board was designed and made by Dr. Bernhard Adams, so that data signals from the LAPPD can be read out with two high channel density SAMTEC connectors² (see Figure 4.8). This eliminates the need for cables coming out of 60 SMA connectors per LAPPD, and waveform digitization and readout can be accomplished by two electronics boards attached to the back of the LAPPD (refer to the

²Part number: QRF8-78-05.0L-D-A or QRM8-78-05.0L-D-A




Figure 4.8: The ANNIE Analog Pickup Board, designed by Dr. Bernhard Adams. This pickup board replaces the 56 SMA connectors with two high channel density SAMTEC connectors (white arrows) for signal readout. Traces for each stripline can be seen on the two halves of the board. Four SMA connectors (red boxes) remain as spare input/output channels for trigger signals. Mounting holes in the LAPPD frame are indicated by the blue arrows. Bernhard was quite fond of photographing dragonflies, hence the dragonfly on the pickup board.

next section on ACDCs).

In addition to reading out signals, the LAPPD and its electronics must be remotely powered while underwater in the detector tank. This is done with two custom-made, waterproof cables. With many electronics and high voltage packed into a compact housing, we must also be able to monitor the internal conditions once the LAPPD system is deployed underwater. Power and monitoring of temperature and relative humidity are accomplished with a custom LVHV board.

The following section will discuss in more detail the different components of the full LAPPD system. Note that because the LAPPD is a new technology, many parts were manufactured for the first time. Each component must be tested as much as possible to understand how they will function in the long-term.

4.5 LAPPD System Components

4.5.1 ACDC Card (Waveform Digitization & Readout)

The LAPPD is a photodetector capable of achieving picosecond-level timing resolutions. For this reason, the need for ultra-fast waveform digitizing electronics capable of preserving and recording the fast timing responses from multiple channels of the LAPPD led to the development of the PSEC4 chip, an application-specific integrated circuit (ASIC). This “oscilloscope-on-a-chip”, designed by Eric Oberla, is capable of reading out six channels,




Figure 4.9: A picture of an unmodified^a ACDC board. The ACDC board is a front-end waveform digitizing printed circuit board that houses five PSEC4 chips (red box), capable of preserving the fast, picosecond-level rise times of pulses. A single PSEC4 chip can record 25 ns of waveform data from six channels, at a sampling rate of 4-15 GSamples/s and 1.5 GHz bandwidth. The white arrow points to the SAMTEC connector that mates with one of the two connectors on the ANNIE Analog Pickup Board attached to the LAPPD. After waveform digitization, data is transferred over LVDS lines that connect to the RJ45 port, indicated by the blue box. The yellow box shows the Altera Cyclone-IV GX FPGA that controls the five PSEC4 chips. The green boxes highlight the linear regulators that power the PSEC4 chips, clock, and FPGA.

^aUnmodified for power. See Appendix [A](#).

with an adjustable sampling rate of 4-15 GSamples/s and a 1.5 GHz bandwidth. On each of the six analog channels, a switched capacitor array (SCA) that is 256 samples deep, corresponding to 25 ns of ring-buffer, samples waveforms at relatively high rates to preserve timing information but with some latency in readout. Some dead-time (on the order of μs , but this has not been precisely measured) is acceptable since many high energy physics experiments like ANNIE rely on trigger schemes to signal the start of potentially interesting events. For the sake of brevity, the reader is referred to [103] and [104] for more about the specifications of the PSEC4 ASIC.

Five PSEC4 chips are mounted on a front-end readout board called the ACDC (ACquisition and Digitization with pseC4) board [105], so that one ACDC can read up to 30 channels. Thus, two ACDC boards are needed to digitize the 28 anode striplines of a LAPPD. A SAMTEC connector installed on the board connects the ACDC to the Analog Pickup Board attached to the back of the LAPPD. Data transfer occurs over eight twisted-pair LVDS (low-voltage differential signaling) lines that are connected to RJ45 ports. Figure 4.9 shows a picture of one ACDC board.

The Altera Cyclone-IV GX FPGA (field-programmable gate array) controls the five PSEC4 chips and other peripherals on the ACDC board. On-board linear regulators power the clock-generating chip, PSEC4 chips, and FPGA. In the original design, a single +5 V DC input was distributed to the linear regulators. However, this proved to generate a significant amount of heat, which is a cause for concern since the ACDC boards are adjacent to the LAPPD. Consequently, the power distribution was re-designed so that the LVHV board

(discussed later in this section) provided the necessary voltages [106].

Several revisions of the ACDC have been made, with newer ones in development and testing. ANNIE uses the Rev B ACDC boards. For more information on the ACDC board, see Refs. [105] and [107].

4.5.2 ACC (Beam Gate, PPS, Clock, & Readout)

Since the ACDC boards were designed to match the speed of the signals coming from the LAPPD, the ANNIE Central Card (ACC) was designed to manage the ACDC boards and their large data output. Designed by Mircea Bogdan of the University of Chicago Electronics Design Group, the ACC can control up to eight ACDC boards, depending on the firmware flashed on the Altera Arria-V FPGA on-board. Firmware was modified by David Green-shields of the University of Warwick to suit ANNIE's needs. In addition to data acquisition, the ACC sends the beam gate, PPS (pulse per second), and clock signals to the ACDCs for waveform alignment between channels and synchronization with the rest of the ANNIE DAQ.

The communication protocol between the ACC and ACDC boards follows the LVDS standard. The boards can be connected via a CAT6 cable where two LVDS pairs are used for transmitting and receiving data, one LVDS pair for the clock signal, and the remaining LVDS pair is for the trigger signals. For more information about the ACC, refer to [105] and [107]. An ACC is shown in Figure 4.10.




Figure 4.10: A picture of the ANNIE Central Card (ACC). A +5 V DC input enters from the top of the board (white box) and powers the Altera Arria-V FPGA (yellow box) and other peripherals on the board. The ACC can control up to eight ACDC boards via CAT6 cables that connect to the block of RJ45 connectors at the bottom of the board (blue box). The LVDS signals are processed by microchips (green box). External signals such as the beam gate, PPS, and clock signals are input via the connector in the bottom left corner (red box).

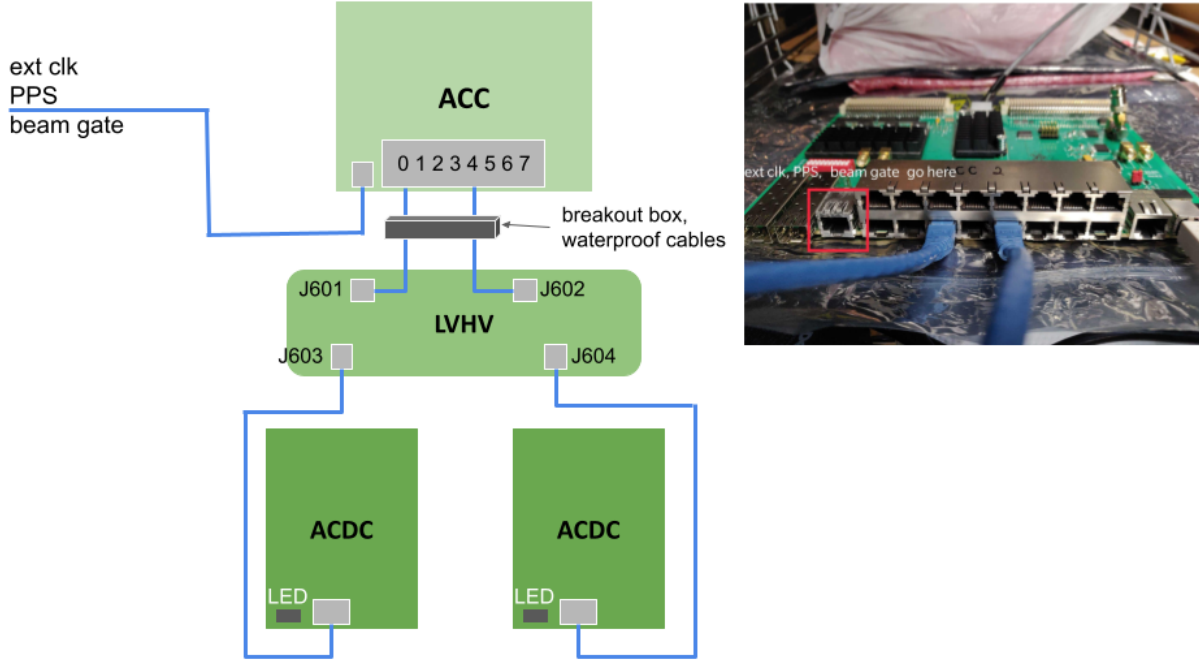


Figure 4.11: LEFT: A diagram depicting how the LAPPD electronics boards are connected to each other. Data flows from the ACDCs to the LVHV board (passive), over the waterproof cables and through the breakout box, and finally to the ACC. A Raspberry Pi with LAPPD data acquisition (DAQ) software acts as an interface between the ACC and the main ANNIE DAQ. Note the diagram is not to scale. RIGHT: An image of the ACC board with two ACDC boards (not pictured) connected via ethernet cables. The red box highlights the input port for external clock, PPS, and beam gate signals.

4.5.3 Trigger Card

The trigger card is a 205 mm by 18 mm circuit board also designed by Bernhard Adams to initiate digitization of waveforms when desired trigger conditions are met [108]. The board comprises identical circuits of capacitors, resistors, and operational amplifiers (op-amps) to monitor each of the 28 anode striplines and provide multi-channel coincidence logic. The trigger board is attached to the anode terminals on one side of the LAPPD, with spring-




Figure 4.12: A trigger card (1) connected to an LAPPD’s anode terminals via spring-loaded contact pins (“pogo pins”). The LAPPD frame must be opened up before installing or removing the card. The output (2) of the trigger board is sent to two of the four spare SMA channels on the Analog Pickup Board and is readout with the data by the ACDC boards. Power and communication lines (3) are found at one end of the board.

loaded contact pins (“pogo pins”), and signals are routed, also via the pins, to the Analog Pickup Board behind the LAPPD.

Using a digital-to-analog converter (DAC) of type MCP4725, users can set a threshold from -1,250 mV to +1,750 mV for individual channels. This DAC (DAC0) is common to all of the “channel comparator” circuits and provides the input to the $(-)$ input of the comparator. The $(+)$ input is biased to 1.25 V so that a DAC0 output of 1.24 V corresponds to a waveform threshold of -10 mV (the output of photodetectors are typically negative-going pulses). Figure [4.13](#) shows an example circuit diagram of one comparator. Channel comparators determine how large a pulse must be to trigger an event. The first trigger board




Figure 4.13: A basic circuit of one of the channel comparators on the LAPPD trigger board. The (+) input is biased to 1.25 V so that a DAC0 output of 1.24 V is consistent with a waveform threshold of -10 mV. Figure from [108].

has a threshold of -30 mV, above the noise level of the board itself.

Output from the individual channel comparators are summed and sent as an input to the “multiplicity comparator”, which sets a multiplicity threshold for an event. This sets the minimum number of pulses that meet the channel threshold before data acquisition is stopped and samples are digitized. Users can set the multiplicity threshold using a DAC (DAC1) also of type MCP4725. The output of the multiplicity comparator is a 200-ns wide, 700-mV negative-going square wave that is converted from a LVDS signal to a RF (radio frequency) signal with a wideband transformer³. This square wave is routed to two of the four spare SMA channels on the Analog Pickup Board so that the trigger signal is also digitized with the stripline channels and time alignments can be made offline.

³Part number: Coilcraft WB2-1-2W 2026D


4.5.4 LVHV Board (Slow Controls)

With the LAPPD and fast readout electronics closely packed in a watertight enclosure (see section on Watertight Housing), air circulation to help dissipate heat is low. On top of that, the PVC⁴ housing is nearly flush with the PVC deployment panel, so water circulation between the backplate and panel is also low. While the stainless steel backplate does offer some level of heat transfer out of the housing, its thermal conductivity is modest, and overheating is a serious concern for the LAPPD system. Placing electronics underwater also puts them at risk for water damage, and because the LAPPD is very light-sensitive, it is at risk for damage by excessive amounts of light. Thus, it is crucial that we monitor the conditions inside the housing. This is referred to as “slow controls” and includes monitoring of the temperature, humidity, high and low voltages, and light levels.

The LVHV board, designed by Seth Hillbrand, is dedicated to monitoring the aforementioned slow controls variables. It also provides the high (HV) and low voltages (LV) needed to power the LAPPD, trigger card, and ACDC boards. The DSPIC33EV256GM102 (hereafter known as the “PIC”) by Microchip Technology serves as the main point of communication between the surface and the peripheral devices on-board. Users communicate with the PIC via the CAN bus communication protocol, which is serviced by the MCP2515 chip.

A +12 V input voltage to the LVHV board is stepped down to +5 V and +3.3 V power rails needed by other modules on-board. Step-down converters and linear regulators provide the +3.8 V, +3.0 V, and +1.8 V rails that power the chips on the ACDC boards. HV

⁴Polyvinyl chloride



- | | | |
|------------------------------------|---|----------------------------|
| 1 - Power input (+12V, GND) | 6-9 - LV lines (+3.8V, +3.5V, +1.8V) | 14 - RJ45 port |
| 2 - UV/OV protection circuit | 10 - TPS562200 step-down converter | 15 - HV output circuit |
| 3 - PIC microcontroller | 11 - ADP7158 linear regulator | 16 - LTC2631-HZ DAC |
| 4 - HIH6030-021 RH&T sensor | 12 - Trigger board power and comms | 17 - HV divider |
| 5 - CAN bus | 13 - LVDS circuit | |

Figure 4.14: Version 2 of the LVHV board for LAPPD-40. Some important peripherals are labeled. The LVHV board provides the high and low voltages needed to power the LAPPD and the trigger card and ACDC boards attached to it. It also monitors the temperature, humidity, and light levels inside the housing. The PIC microcontroller serves as the main point of communication between the housing and surface-level electronics. Users communicate with the PIC via CAN communication.

is distributed by an EMCO C40N power module, which is programmed with a digital-to-analog converter (DAC) of type LTC2631-HZ and an Omron G5Q-1A4 DC12 relay. Under a software routine, the PIC fetches measurements from a relative humidity and temperature (RH&T) sensor⁵, a photodiode, and its internal analog-to-digital converters (ADC), and sends humidity, temperature, light level, and voltage readings to the surface for monitoring [109].

Software protections in the main DAQ routines warn users of undesirable conditions

⁵Part number: Honeywell HIH6030-021 (SPI version)

in the LAPPD system, while a salt bridge installed inside the housing and monitored at the surface level detects emergency leaks. An activated salt bridge prompts an automated shutdown if the LAPPD and electronics are on, or prevents devices from being powered if a leak is detected immediately after the system is deployed in the water.

After extensive testing and debugging, the LVHV board design underwent revisions by Seth Hillbrand and Paul Stucky to include more protections for the LAPPD system and the LVHV itself. These include an undervoltage/overvoltage (UV/OV) protection circuit for the +12 V input and high voltage read-back. Although this second version was an improvement from the first version, there were still some issues that were found and needed to be addressed. See Appendix [A](#) for a list of the changes between the first and second versions and modifications to LVHV-v2.

4.5.5 Watertight Housing & Deployment Panel

Although the LAPPD has a frame to protect the fragile vacuum tile (see Figure [4.7](#)), there is no protection from water as it stands. The Ultem[®] frame designed by Incom Inc. [\[110\]](#) protects the LAPPD from mechanical forces, distributes high voltage to electrodes that penetrate the vacuum tile, and provides mounting holes (shown in Figure [4.8](#)) to secure the LAPPD to a structure for testing or deployment. Thus, a form of housing is needed to protect the LAPPD and its electronics from water.

The LAPPD watertight housing and deployment panel were designed by me and David



Figure 4.15: The LAPPD watertight housing comprises a 13-inch by 17-inch by 2-inch PVC frame, a UV-transparent acrylic window, and stainless steel backplate. O-rings sit in grooves under the front window and backplate to provide the watertightness needed for the LAPPD system. The backplate offers some level of heat dissipation.

Hemer from the UC Davis Math and Physical Sciences Machine Shop. This housing consists of a 13-inch by 17-inch by 2-inch PVC frame, an acrylic front window so the LAPPD can detect light generated within the detector volume, and a stainless steel backplate⁶ for heat dissipation and mounting (Figure 4.15). Grooves for O-rings and 60 screw holes can be found on both sides of the PVC frame, for a grand total of 120 bolts and two O-rings for each LAPPD. Short protrusions, or “nubs”, in the interior of the frame provide a shelf for the LAPPD and LVHV board to rest on and be securely mounted in the housing (see Figure 4.16). The front window is made of UV-transparent (UVT) acrylic and is 0.25 inches thick. When installed in the housing, the LAPPD frame nearly touches the front window. Two holes were drilled into the top edge of the housing (Figure 4.16) for the MacArtney SubConn Power

⁶The stainless steel backplate was machined by an external company using laser cutting technology.




Figure 4.16: A PVC frame with a LVHV board mounted inside. The LAPPD Ultem frame can be bolted to the four protrusions (“nubs”) indicated by the white arrows. The MacArtney SubConn socket connectors screw into the holes indicated by the red arrows. The top of the housing is where the LVHV board is mounted. One of the six nubs (underneath the LVHV board) needed to be removed as it was in the way of the power module of the LVHV.

Ethernet Circular 13-pin socket connectors⁷ to screw into. The socket connectors are where the waterproof cables plug into the housing.

The LAPPD deployment panel is a 10-ft long, 11-inch wide, and 0.25-inch thick PVC board with four polypropylene “sliders” attached to its back. The sliders have a cross-section in the shape of an “I” (cf. Figure 4.18) and fit into the six rails of the octagonal inner structure. At the top of each panel is a rectangular hole used for lowering and retrieving the panel into and out of the tank. The LAPPD housing attaches to the deployment panel via the holes in the stainless steel protruding from the backplate. Once a position is chosen, holes can be drilled into the panel where the housing will be bolted and secured. A set of smaller holes are drilled in the left and right halves of the panel so that the waterproof cables can be secured with metal zip ties. This is to prevent the cables from drifting into the inner volume of the detector and obstructing other photodetectors. Figure 4.17 shows a deployment panel fitted with an empty housing unit and a set of waterproof cables, ready to be deployed in the water.

4.5.6 Waterproof Cables

The LAPPD waterproof cables manage the power, CAN bus communication, data, sine wave, and alarm lines. It was necessary to design them from scratch since the LAPPD is new detector technology and there were no commercially available parts. The cable design accommodates

⁷Part number: DBH13F (socket connector)




Figure 4.17: A 10-ft long deployment panel machined from gray PVC, ready for deployment testing. The housing is attached to the panel at the stainless steel backplate with washers, nuts, and bolts. The waterproof cables are secured with stainless steel zip ties. Materials that make up LAPPD components (e.g., PVC, cable jacket) were carefully selected for chemical compatibility with Gd-loaded water. Before deployment in the tank, everything needs to be dust- and contaminant-free.




Figure 4.18: Four polypropylene sliders (shown by white arrows) attach to the back of the LAPPD deployment panel. Rectangular grooves were machined into the sides of each block so fit into the eight rails of the octagonal inner structure.

- 2 lines for power and ground,
- 8 pairs of lines (16 total) for data transfer and communication with two ACDC boards,
- 2 lines of CAN bus communications,
- 2 lines for a sine wave or other calibration signal, and
- 2 lines for an alarm

for a total of 24 lines needed. The Falmat XtremeNet™ Underwater Cable⁸ best suited our needs with four single conductors of 18 AWG and four twisted, jacketed pairs of CAT6 lines of 23 AWG, for a total of 12 lines per cable. Two cables serve one LAPPD and are not identical, since power, CAN bus, sine wave, and alarm need to be divided between the cables. Thus, it is important that a cable connects to the appropriate connector, so that signals are not delivered to the wrong location and damage electronics.

To complete the cable design⁹, pin connectors are needed to terminate the wet end (the end that goes underwater) and plug into the housing. The MacArtney SubConn Power Ethernet Circular 13-pin¹⁰ connector¹¹ was selected for cable termination, as its pins are specialized for paired data signals like LVDS. The MacArtney SubConn pin connectors were spliced to the Falmat XtremeNet cables, making the cable a little over 10 m (33 ft) in length. On the dry end of the LAPPD cable, data lines were terminated with RJ45 jacks, while the single conductor lines (for power, CAN, etc.) were directly routed inside the breakout box.

⁸Part number: FM022208-14

⁹At the time, Falmat was unable to terminate the cables for us, but we also needed the connectors to be waterproof.

¹⁰Unfortunately, the 12-pin connector (P/N: IL12M) initially selected did not function to our standards and needed replacement. The data lines extending from the splicing side of the connector were not twisted and picked up significant levels of noise. This affected communications between the ACC and ACDC boards. All cables needed to be re-spliced with the new connectors and the housings modified for the new socket connectors.

¹¹Part number: DIL13M (pin connector)



Figure 4.19: LEFT: A MacArtney SubConn Power Ethernet Circular 13-pin connector (blue section) is spliced to the Falmat XtremeNet™ Underwater Cable (neon green). The spliced region (indicated by the white line) is made with crimp connectors, heat shrink tubing, electrical tape, and splicing tape. RIGHT: A waterproof cable spliced with the original MacArtney SubConn 12-pin connector. The cable is about 10 m, or 33 ft, in length.

4.5.7 Breakout Box

The breakout box (BoB), also designed and assembled by Seth Hillbrand, serves as a distribution hub for power, data, and other important signals. Each BoB services two¹² LAPPDs with a set of terminal blocks, a Raspberry Pi and Pi HATs¹³, and power module for the Pis, as shown in Figure 4.20.

¹²The prototype breakout box serviced only one LAPPD and was designed to take up one rack unit (1U). However, upon realizing that the Pi HATs would take up more space in the vertical direction, the breakout box was expanded so that it could service two LAPPDs instead, making the height 2U.

¹³Raspberry Pi Model 3B. A Pi HAT (Hardware Attached on Top) is an expansion board that connects to the Raspberry Pi's GPIO pins and adds extra functionality to the Pi, such as lights, motors, and CAN bus.




Figure 4.20: Inside the breakout box are two sets of electronics and components that service two LAPPDs: (1) an input voltage of +12 V is fanned out to the relay HATs on the Raspberry Pis before making it down to the LVHV board underwater; (2) the Raspberry Pi power supply also receives the +12 V input to power the Pis with +5 V; (3) a relay HAT and a PiCAN 2 HAT sits atop a Raspberry Pi that runs a routine to fetch slow control measurements; (4) power, CAN bus, data, sine wave, and alarm lines are routed through terminal blocks; (5) the LVDS data signals exit the breakout box via RJ45 jacks; (6) communication with the Raspberry Pis can occur over an Ethernet connection.

An input voltage of +12 V enters the breakout box and is fanned out to two Raspberry Pi Power Relay HATs before powering the electronics downstream. The relay serves as one point of protection from accidental power-ons: users must command the Raspberry Pi to switch the relays to power the LVHV board, else the electronics remain off. The +12 V input also powers the +5 V power supplies for the Raspberry Pis. Another Pi HAT, the PiCAN 2 HAT, adds CAN bus communication functionality to the Raspberry Pi. Users have the capability to send commands to the LVHV board via CAN messages sent by the Raspberry Pi. Communication with the Raspberry Pis themselves occurs via `ssh` but can also be done with standard Ethernet cables.

Terminal blocks route signals in and out of the breakout box. In Figure [4.20](#), twisted pairs of wires coming from a waterproof cable are routed to a terminal block, and then to RJ45 ports. CAT5/6 Ethernet cables connect to the BoB on the other side of these ports and deliver the LVDS signals to the ACC.

Amphenol receptacle connectors on the front panel of the BoB proved unnecessary as testing of the LAPPD system components progressed and added noise to the LVDS signals transferred between the ACC and ACDC boards. Thus, the connectors were bypassed and data signals are transferred over the waterproof cables directly to the ACC, while power and CAN bus were still managed by the breakout box.

Although not shown in Figure [4.20](#), the Yoctopuce Yocto-MaxiThermistor was installed to measure resistances of the salt bridge and thermistor that were added to the LAPPD system in the housing. For future designs of the breakout box, reorganization of the internal




Figure 4.21: A single LAPPD is mounted inside a dark box. A motor-driven apparatus holds the LED and laser used to perform quantum efficiency and timing studies.

components and removal of unnecessary parts should be taken into consideration.

4.6 Characterization & Testing

LAPPDs exhibit unit-to-unit variability in performance, so it is necessary to thoroughly characterize each one. Some characterization measurements include the quantum efficiency (QE) and timing capabilities for each LAPPD. Additionally, since many components are designed specifically for ANNIE, they must be individually and extensively tested before integration into the full LAPPD system.

4.6.1 Measuring Quantum Efficiency with an LED

Recall that quantum efficiency (QE) is the likelihood that a photon is converted into a photoelectron by the photocathode. Since LAPPD production is relatively new, manufacturing steps may not necessarily result in identical units, particularly the steps involving deposition of photocathode material on the front window and sealing of the vacuum tile. This makes QE a major source of variability among LAPPD tiles. Early LAPPD tiles exhibited modest photocathode QEs [100], but more recent tiles have QEs of at least 25%, comparable to that of conventional PMTs.

To measure the QE of the entire photocathode area, an LED of 420 nm wavelength was used with a NIST¹⁴ photodiode serving as a reference. First, a motor-driven apparatus that holds the light sources is moved to the reference photodiode. The LED shines onto the photodiode and current is measured. After this reference current is measured, the motor-driven apparatus then shines the LED at points across the entire area of the LAPPD in increments of 5 mm in both the vertical and horizontal directions, measuring a current at each location. The scan takes about 8 hours to complete. Once complete, the data is analyzed using Python scripts written by Bernhard Adams to produce a QE plot. Figure 4.22 shows the results for LAPPD-63, one of ANNIE's better performance LAPPDs.

As extensive characterization studies progressed, the average QE across the face of the LAPPD was observed to decrease over time. Thus, QE scans before and after extensive testing is recommended for to study how photocathode performance changes over time.

¹⁴National Institute of Standards and Technology




Figure 4.22: The results from a quantum efficiency (QE) scan for LAPPD-63. Note that the QE is fairly uniform across the photocathode and comparable to the QE of convention PMTs. The dark regions at the perimeter of the LAPPD and that form an “X” are where the frame and X-spacer sit. Over time, it was realized that performing QE scans before and after extensive testing is recommended to study how the photocathode performance changes over time. After extensive studies with the laser, the average QE had dropped by about a percent.

4.6.2 Timing Performance Characterization

LAPPDs are known for their picosecond-level timing. Naturally, measuring the timing capabilities is another way to characterize a LAPPD. Timing studies are performed with a PILAS PiL040X laser from NKT Photonics with a wavelength range of $405 \text{ nm} \pm 15 \text{ nm}$ and a $<45 \text{ ps}$ resolution rather than with an LED because lasers are capable of producing pulses with rise times and widths on the order of picoseconds. Sending well-defined, narrow

pulses allows us to determine how well the LAPPD is able to capture the pulse’s timing information.

Mounted on the same motor-driven apparatus as the LED, the laser is pointed at a fiber-optic collimator, which is coupled to an optical fiber. The fiber reduces the intensity of the laser so that single photoelectron (SPE) levels of light can be reached. SPE levels of light is key to understanding the response of a LAPPD because we want to make sure that we are able to detect a single photon. An Advanced Laser Diode Systems (ALS) EIG2000DX Controller powers and controls the laser. Settings were such that the laser fired at a rate of 100 Hz (the lowest setting unless externally triggered) and at a tune¹⁵ of 70-90%.

The PILAS laser was pointed at the LAPPD using the same Python software that controls the LED and data was taken at several locations on the photocathode. Studies include horizontal scans to study timing differences between the ends of the striplines, vertical scans to study charge spread and cross-talk between striplines, and high statistics scans. A laptop computer with LAPPD data acquisition software¹⁶ installed connects to the ACC to take data.

4.6.3 Performance Stability Tests

Understanding all the features of a LAPPD is important, but understanding its supporting electronics and parts is equally as important. Since most of the components will be under-

¹⁵Tune is a setting on the ALS controller box that influences the intensity of the laser. A larger tune (i.e., more tuning) value means a lower intensity, and vice versa.

¹⁶The LAPPD DAQ software is written in C/C++ and can be found at <https://github.com/lappd-daq/acdc-revc-daq>.

water and not easily accessible, operation of each system component needs to be reliable.

To test the performance stability of the ACC and ACDC boards, high statistics scans often ran for at least 8 hours to a day. Any crash or interruption in a run was debugged to ensure long-term performance.

Burn-in tests of the LVHV board were carried out to study the robustness of the different communication protocols (I²C, SPI, CAN) on-board the LVHV. The slow controls routine of fetching humidity, temperature, low voltage output, etc., ran for several hours to find any software bugs that may come up during long periods of operation. In addition to testing the slow controls, burn-in tests with HV were also conducted to ensure that the LAPPD is powered without interruption and data quality is not affected. Sudden changes in applied voltage can damage the LAPPD and render it inoperable. For these tests, the LVHV board is powered and the HV output is gradually ramped up until the nominal voltage is reached. Dummy loads with resistances close to that of the microchannel plates (MCPs) were attached to the LVHV to get as close to actual operating conditions as possible. After HV tests were deemed successful, the LAPPD itself was connected to the LVHV board and powered for several hours to confirm performance stability.

4.6.4 Soak Test of Housing

To test the capability of the watertight housing in keeping the interior dry, soak tests were performed with and without electronics installed. O-rings provide the watertightness and

each housing receives two O-rings: one under the acrylic front window and one under the stainless steel backplate. In these soak tests, the waterproof cables need to be plugged in to prevent water from entering through the pin holes of the socket connectors¹⁷. In effect, the MacArtney SubConn pin and socket connectors were also tested in their ability to keep water out, along with the waterproof integrity of the cables. Because the rubber that the MacArtney SubConn connectors were manufactured with is not compatible with Gd-loaded water, those sections of the cable and the socket connectors (see Figure 4.19) were wrapped in PTFE tape, like the inner structure of the detector tank.

A large bin was lined with plastic and filled with deionized water. The housing with the cables attached are then submerged and held down since the entire unit is buoyant. Soak tests with an empty housing ran for several days since there were no electronics to worry about. Humidity indicator strips were taped to the backplate inside the housing to get a rough measure of humidity levels inside, as can be seen in Figure 4.23. In actual deployment, when all electronics are installed, some (six to eight) silicone gel desiccant packets were distributed throughout the housing to absorb moisture and keep humidity levels to a minimum. Soak tests were repeated with LAPPD components installed, first with the electronics off and then powered on. When the LVHV board is on, humidity can be measured with the sensor in addition to the indicator strips.

A successful soak test demonstrated that no water entered the housing and retrieval from

¹⁷Small O-rings between the socket connectors and the housing to provide watertightness where the connectors are screwed into the housing frame.




Figure 4.23: LEFT: An empty watertight housing soaking in deionized water. Note that this soak test was carried out with the old 12-pin connectors. The waterproof cable is wrapped in PTFE tape where the MacArtney SubConn connectors mate, so that no rubber insulation is exposed due to incompatibility with Gd-loaded water. The black O-ring on the acrylic window side of the housing is indicated by the red arrow. Three humidity indicator strips were taped to the backplate inside the housing to get rough estimates of humidity levels. RIGHT: The full LAPPD system. An LED (green arrow), wrapped in plastic to waterproof it, is placed in front of the acrylic window to confirm that the LAPPD still sees light even in its housing and underwater.

the water. Although the humidity indicator strips did show an increase in humidity levels, it was not cause for concern because it is known that acrylic is slightly permeable to moisture. As for the cables, continuity tests and resistance measurements of each wire, using a digital multimeter, showed that no water penetrated the cables. It is noted here that frequent opening and closing of the housing requires replacement of O-rings as they may fail with frequent compression and decompression.

4.7 Deployment & Integration

Once the LAPPD components have been individually and extensively tested, they must be integrated and tested as a whole system in water before deployment in the ANNIE tank. During early development and testing, four commercial power supplies provided the high voltages necessary for the LAPPD photocathode and microchannel plates (MCPs) described in Section 4.2. The last voltage drop between the lower MCP and the anode was established by a +200 V zener diode. Power to the ACDC boards also came from external power supplies, before integration of the LVHV board.

After collecting sufficient characterization data and performing burn-in tests, the LVHV board was the first component integrated into the full system. In this configuration, the LAPPD and ACDC boards received power from the LVHV board instead of commercial power supplies, while the LVHV and ACC boards were powered externally. The LVHV board was encased by a plastic enclosure¹⁸ to mitigate the dangers of exposed high voltage. Some data was taken with the laser to confirm that this configuration of the system worked and that the LAPPD could see pulses.

Next, the waterproof cables and breakout box were integrated into the system together¹⁹. Rather than power the LVHV board directly, the power supply delivered +12 V to the front of the breakout box, which then distributed power to the rest of the system over the waterproof

¹⁸Designed and assembled by an electrical engineer at Iowa State University.

¹⁹At the time of integration testing, the cables were terminated with Amphenol connectors and only the breakout box had the correct mating connector and pinout. Pinout refers to the assignment of wires and pins to their respective functions.



Figure 4.24: A commercial power supply (on top) provides the +12 V to the breakout box (on bottom), which distributes the power accordingly. (1) The input voltage enters the breakout box at one port on the front panel. (2) Waterproof cables connect to Amphenol receptacle connectors. The RJ45 ports to the right of the cables are where the ACC connects. (3) Another set of connectors and ports for a second LAPPD.

cables and to other modules inside the breakout box. This is shown in Figure [4.24](#). As integration testing progressed, it was observed that the waterproof cables were long enough to cause the ground path to drift and cause issues in the LVDS signals exchanged between the ACC and ACDC boards. In turn, the altered LVDS signals damaged some of the buffer chips on-board the ACC. Fortunately, a redesign of the LVHV board mitigated the grounding issues (refer to Appendix [A](#) for more details).

After much testing and debugging in the dark box, the nearly complete system operated to our satisfaction. The final step was to move the LAPPD and all its electronics into the

watertight housing and test the system in water. As a sanity check, the fully assembled housing was placed inside a small portable dark box with a green LED: the LAPPD system still operated as expected. Several burn-in tests in water were carried out with the LAPPD at nominal voltage to establish standard operating conditions.

4.7.1 Deployment Mechanism

As briefly described in Section [4.5.5](#), the watertight housing is bolted onto the deployment panel. The waterproof cables are secured to the panel by stainless steel zip ties at points along the panel to prevent the cables from drifting towards the center of the tank and obstructing the other photodetectors. Bolted to the back of the panel are four polypropylene blocks with rectangular grooves machined into the sides so that the panel can slide up and down one of the eight tracks. The bottommost block also acts as a stopper, to prevent the panel from sliding all the way through to the bottom of the inner structure and hitting other PMTs.

Before the inner structure was transported to the experimental hall, the deployment mechanism was tested in air, in a clean tent located at the D0 Assembly Building (DAB) at Fermilab. Figure [4.25](#) shows the panel with an empty housing in one of the eight rails. A cross-like pole with a hook at the end was assembled out of PVC and bolts to retrieve the panel out of the tank. This method of retrieval was tested as well.

Following the installation of the inner structure and the filling of the tank with Gd-loaded water, the LAPPD deployment and retrieval mechanisms were tested again to establish the



Figure 4.25: An empty LAPPD housing attached to a deployment panel. Testing of the deployment mechanism took place in a clean tent, before the octagonal inner structure was transported to ANNIE Hall. The panel slides along the rail via polypropylene sliders bolted to its back. All eight panels were tested to make certain that they slid down the rails smoothly. Sliders that did not go down a track smoothly naturally had pieces of polypropylene shaved off by the track. An ANNIE collaborator can be seen behind the inner structure holding the deployment panel.

best way to deploy the system in water and in the hall. Given the limited space in the experimental hall, any issues that may put the LAPPD system at risk for damage needed to be identified and mitigated before actual deployment. Buoyancy is also a concern, so a stainless steel rod was attached to the bottom of the PVC panel. The added weight made lifting the panel over the guard rails above the tank lid slightly more awkward. Regardless, the deployment and retrieval mechanisms were demonstrated with great success.

4.8 My Work

For this thesis, I spent a year at Fermilab debugging, developing, and testing the individual components of the LAPPD system leading up to and including the full system integration into the ANNIE detector. During subsequent visits to Fermilab, I tested the redesigned LVHV board, modified the housing to accept the replacement MacArtney connectors, and helped address the heating issues inside the housing. I spliced the original 12-pin MacArtney connector to the 10 Falmat cables, and respliced a pair of the cables with the 13-pin connector. I also helped design the plugs needed to modify the housing to accept these new connectors.

I wrote and tested the firmware that programs the PIC microcontroller, so that it can communicate with peripheral devices on the LVHV board such as the humidity and temperature sensor, photodiode, and HV power module. I also developed a user-friendly software (on which the slow controls portion of the LAPPD DAQ toolchain software is based) that

allows users to set outlines a routine for monitoring slow controls variables [109]. When the temperature and/or humidity reach(es) a certain threshold, the software starts an automated ramp-down of the high voltage. This way, the LAPPD and the other electronics are protected from severe damage.

4.9 First LAPPD Data

Finally, on March 29, 2022, the ANNIE Collaboration deployed its first LAPPD. Figure 4.26 presents the first LAPPD data taken with the ANNIE detector. The beam spill is clearly visible. Efforts to integrate this datastream are in progress.




Figure 4.26: An increase in events with a width of $1.6 \mu\text{s}$ confirms that the LAPPD does in fact see the neutrino beam.

Chapter 5

Calibration

A good and thorough understanding of how a detector responds to a physics event is crucial for any precision measurement. To measure the response of each system component, we must perform calibration tests, both *ex situ* and *in situ*, to ensure performance stability during actual experimental runs. No detector is perfect, so it is pertinent to assess its overall response to events and understand how well it can detect particles.

5.1 Detector Calibration

The realization of the physics goals of ANNIE requires measurements of particle type, energy, vertex, and direction. The detector must also be able to veto interactions that occur in the upstream rock and efficiently detect the 8 MeV gammas from neutron captures on gadolinium. To extract this information from the detector, we need a good handle on pho-

todetector gains, timing, and efficiencies. Setting the absolute energy scale of the detector is also necessary, that is, determining how much energy is represented by some amount of charge. If the charge-energy conversion is unknown, then particle energies cannot be determined accurately. Similarly, if the timing of signal is not well constrained, uncertainties in energy and position reconstruction will be large.

For these reasons, photodetectors were tested and characterized using test stands in lab and custom software. The efficiencies of the Front Muon Veto (FMV) and Muon Range Detector (MRD) were determined by looking at a sample of through-going muons in the data.

5.1.1 Target Tank PMT Gains

In a conventional photomultiplier tube (PMT), dynodes amplify a single photoelectron (SPE) into a detectable current. This amplification factor—the ratio of the anode current to the photocathode current—is the gain of the PMT. Gain varies with supply voltage. If the supply voltage is too low, the result is low gain and poor single photoelectron (PE) detection efficiency since the anode current may be too small. If the supply voltage is too high, then dark noise (electrons spontaneously emitted from photocathode), after-pulsing (correlated delayed pulses caused by ions from residual gases drifting into the photocathode and freeing additional electrons) rates, and current draw increase, potentially damaging the PMT and reducing its lifetime. Finding the right gain and operating voltage that sufficiently amplifies




Figure 5.1: Distributions of PMT operating voltages (left) and peak-to-valley ratios (right). Several PMTs in the purple distribution, representing ETEL type PMTs, exhibit lower operating voltages compared to the green distribution of WATCHBOY PMTs. ETEL PMTs also have lower peak-to-valley ratios. Given that the ETEL PMTs were more sensitive to operating conditions, they were installed under the tank lid and thus experienced less pressure underwater.

a photoelectron is important if we want a uniform detector response. PMTs are not exactly identical to one another, so each will have its own operating voltage for the same gain value. ANNIE has several PMT types, namely the ETEL, Hamamatsu, LUX, WATCHBOY, and WATCHMAN PMTs, so gain-voltage curves are needed for each of these PMTs.

I characterized the WATCHBOY and ETEL PMTs at UC Davis before transporting them to Fermilab. A large darkbox was set up with a green LED at one end and one PMT at a time directly facing it. The LED intensity was set to a level such that one in every ten pulses would be detected by the PMT. This ensures that the signals collected come from a single PE and not multiple PEs. For each PMT, waveform data was taken with an Tektronix AFG 3022B oscilloscope at different voltages increasing in steps of 100 V. Charge values were




Figure 5.2: Distributions of dark noise (left) and current (right). WATCHBOY PMTs (green) have low dark noise compared to ETEL PMTs (purple), which have a wide range of dark noise levels. The current draw for ETEL PMTs also vary more than WATCHBOY PMTs.

extracted from the data by integrating the waveforms via Ohm's Law

$$V(t) = I(t)R$$

$$Q_{total} = \int_{t_i}^{t_f} I(t) dt = \frac{1}{R} \int_{t_i}^{t_f} V(t) dt \quad (5.1)$$

where $V(t)$ is the measured voltage at each time bin and R is the resistance of the digitizer.

Because the light intensity is low, the majority of events will not have pulses in them. These form the first peak called the pedestal. Ideally, the pedestal is centered at zero charge, arising from events where no photon is detected. However, variations in the baseline due to noise and hardware can shift it to non-zero values. The next obvious peak consists of the single PE events, whose central value represents the charge amplified by the gain. To




Figure 5.3: An example pulse charge distribution curve fitted with a double Gaussian. The first visible peak is the pedestal and is slightly greater than zero charge. The second visible peak is the single photoelectron peak. This charge represents the gain-amplified photoelectron.

extract the pedestal and gain values, Gaussian curves were fit to each peak, as shown in Figure 5.3. If the pedestal differs from zero, which is often the case since hardware is not perfect, then its central value is subtracted from the overall charge distribution. Once the pedestal-subtracted single PE charge value is extracted, dividing by the electron's charge gives the gain of the PMT. This procedure was repeated for all of the WATCHBOY and ETEL PMTs.

Once all PMTs were relocated to Fermilab, characterization tests were performed again to check the stability of PMT performance. Up to four PMTs were tested in a dark box with an LED light and a custom test stand built from CAMAC electronics to measure cool

down times, gain, operating voltage, and after-pulsing rates. Refer to [77] for more details of these characterization tests. This second round of testing offered an independent way of determining the gains for ANNIE PMTs and confirmed the results of testing at UC Davis. Gains were determined using an algorithm from the DEAP-3600 collaboration [111]. In addition to the two Gaussian fits to the pedestal and the SPE peaks, a third Gaussian was fit to the double PE peak along with an exponential for events where the electron does not follow the typical path down the dynode chain. Overall, the operating voltages were selected so that the PMT gains were relatively uniform and had a mean value of 7×10^6 (cf. Figure 5.5).

After PMT installation, *in situ* measurements were performed to monitor the stability of the PMTs. These measurements were carried out with six 468 nm LED optical fiber modules installed at various locations on the inner structure, and include measuring PMT response and cable propagation delays. In addition, the LEDs are used for water transparency monitoring. For example, the light output from the top LED can be measured by the bottom PMTs (Figure 5.4). If the transparency is relatively constant, then we can conclude that water transparency remains stable.

Gain calibration is done periodically with the LED light sources to monitor the stability of the PMTs. Using the DEAP-3600 method, a combination of functions were used to model




Figure 5.4: Six LED models are installed on the inner structure of the detector. Their locations cover all of the PMTs so that gain calibration and water transparency monitoring can be performed while the experiment is running.

the different contributions to the PMT's charge distribution:

$$\begin{aligned}
 Q(x) = & A_{ped} \cdot \text{Gauss}(x, \mu_{ped}, \sigma_{ped}) + B e^{(x-\sigma)/\tau} H(Q - \mu_{ped}) \\
 & + A_{1PE} \cdot (\text{Gauss}(x, \mu, \sigma) + \text{Gauss}(x, f_{\mu} \cdot \mu, f_{\sigma} \cdot \sigma)) \\
 & + A_{2PE} \cdot (\text{Gauss}(x, 2\mu, \sqrt{2}\sigma) + \text{Gauss}(x, (1 + f_{\mu}) \cdot \mu, \sqrt{1 + f_{\sigma}} \cdot \sigma))
 \end{aligned} \tag{5.2}$$

where $\text{Gauss}(x, \mu, \sigma)$ is the Gaussian function $\text{Gauss}(x, \mu, \sigma) = e^{-(x-\mu)^2/(2\sigma^2)}$, A_x represents the amplitude of the Gaussian associated with the pedestal, single PE, and double PE peaks,




Figure 5.5: A distribution of initial gain values on the left [112] and updated gain values on the right [113]. The PMT gains were recalculated after an update to the hit finding procedure. The average gain has a value of 9×10^6 compared to the previous value of 7×10^6 .

and $H(x)$ is the Heaviside function with B as its amplitude. The fractions f_x model the incomplete photoelectron amplification by the PMT dynodes. These more recent gain values with an updated hit finding procedure are used for the analysis in this thesis. A distribution of the new gain values is shown in Figure 5.5.

5.1.2 LAPPDs

Similar to conventional PMTs, the gain and operational voltage are determined for each LAPPD. In addition, since LAPPDs are new photodetection technology, the quantum efficiency needs to be characterized since production processes were not perfected at the time ANNIE purchased its first LAPPDs. Although LAPPD data will not be used in this analysis, details of LAPPD calibration can be found in Chapter 4.

5.1.3 FMV and MRD Efficiencies

Both the FMV and MRD consist of scintillator paddles coupled to PMTs. Plastic scintillators age and may degrade over time, so paddle performance must be monitored periodically.

5.1.3.1 Muon Range Detector Efficiency

The Muon Range Detector (MRD) records the passage of muons through its iron-scintillator sandwich structure. For the accurate reconstruction of the muon vertex and energy, it is crucial that the MRD can detect the muon entry and stopping points. As described in Section [3.2.6](#), the MRD was inherited from the SciBooNE experiment and some layers needed refurbishment. For the replacement paddles, the efficiency was determined in lab using threefold coincidence measurements at three different locations along the length of the paddle. These serve as reference values for the *in situ* efficiency measurements.

The *in situ* measurements are made by using data and selecting events with through-going muons, events in which the first and last layers of the MRD register hits. Because the end layers are used to determine whether a muon is through-going, they will not have efficiency values assigned to them. For the interior layers, their efficiencies are determined by reconstructing the muon track inside the MRD using paddle hits positions (cf. Section [6.3](#)) and looking for missing paddle hits. Each layer is checked for a hit that is in proximity of the fitted track. If no hit is found, then the efficiency of that particular paddle is reduced. For each paddle i , the efficiency is calculated by comparing the number of observed hits to

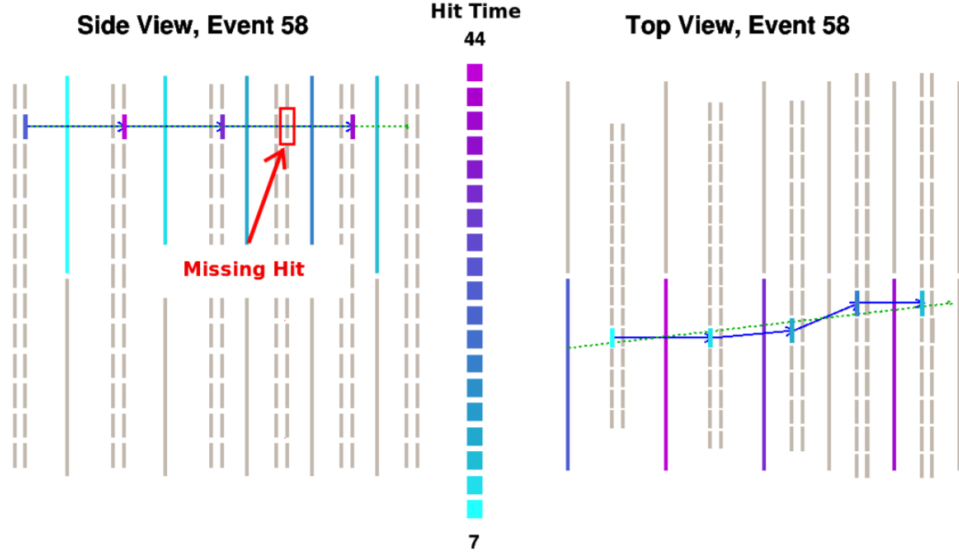


Figure 5.6: An event with a through-going muon. Note the missing paddle hit. The long and short (dashed) lines represent the horizontal or vertical orientations of the scintillator paddles, depending on the perspective view. Double dashed lines represent the two halves of the paddle layers. The iron plate absorbers fill the gaps between scintillator layers. Figure from [77].

the number of expected hits:

$$\epsilon_i = \frac{N_{obs,i}}{N_{exp,i}} \quad (5.3)$$

Figure 5.6 shows an example of an event with missing paddle hits. The resulting MRD efficiencies are shown in Figure 5.7. For most of the paddles, efficiencies of at least 80% are observed.

5.1.3.2 Front Muon Veto Efficiency

The Front Muon Veto (FMV) tags events in which muons are created upstream. This lets us remove events where the muon originated outside of the tank and reject entering




Figure 5.7: Most MRD paddles have efficiency values of at least 80%, as high as 95%. Paddles with low efficiency or are turned off will be replaced in future maintenance during beam downtime. Figure from [114].

backgrounds. External events would otherwise bias our vertex and energy reconstruction, and it must do so with high efficiency. The efficiency of the FMV was measured in the lab using a threefold coincidence setup and cosmic muons. Test results showed that the efficiency was about 90% at the end nearest the PMT and decreased to about 50% towards the other end of the paddle. By orienting the layers such that the PMTs are on opposite ends, lower efficiency on one end of the layer can be made up by higher efficiency on the other end of the second layer and vice versa. A muon that doesn't register a hit in either FMV layer would be missed.

In situ efficiency measurements of the FMV scintillator paddles are determined in a similar manner as the MRD paddles. Because there are only two layers of scintillator paddles




Figure 5.8: Front Muon Veto (FMV) paddle efficiencies. The average paddle efficiency is about 80%, ranging from 65% to 90%. Figures from [114].

that make up the FMV, we must also rely on the MRD to determine the efficiency of the FMV. Events with a track in the tank and MRD are selected. The reconstructed is then extrapolated back to the FMV, where each layer is checked for hits and compared with the other layer. Using Eq. 5.3, efficiency values were determined for each of the paddles. The efficiency was determined to be in the range of 60-80% for the paddles, comparable to what was measured in lab settings. Combined, the efficiency of the FMV in vetoing dirt muons is about 96%.

More details for the procedure for determining the efficiency of the FMV and MRD paddles can be found in [114]. The technique was applied to MC data and the efficiency values found in MC show agreement with data.

5.2 Timing Calibration

To achieve its physics goals, ANNIE uses optical phenomena, which occurs at very fast time scales. As a result, it is important to understand the timing response of the detector and detector system. Otherwise, poor detector performance can limit ANNIE’s capabilities or lead to misunderstandings in the acquired data. For ANNIE, accurate timing information is needed to reconstruct the muon vertex. Accurate knowledge of when a photon was registered by a PMT can help constrain the vertex position so that the muon energy can be determined by track length in the tank and MRD. To perform a timing calibration for the ANNIE detector, a laser is fired into the tank and the hit times of PMTs are recorded. Any deviation from the actual laser trigger is a correction that needs to be made before data analysis. In addition, the differing cable lengths of all the PMTs need to be taken into account as well. Cable delays were programmed into the firmware of the DAQ system, so that they do not need to be corrected for in offline analysis.

As a first measurement, LED optical fiber modules were used to determine the delays between each PMT. Six of these LED light sources are installed at various locations of the inner structure to ensure complete coverage of the tank [112]. Because LED pulses are wider than laser pulses, the measured timing offsets are only accurate to ~ 10 ns, resulting in lower spatial resolution due to the uncertainty in where photon emission occurred. Figure 5.9 shows the initial estimates.

Timing calibration was performed using a PILAS picosecond pulsed diode laser of 400 nm




Figure 5.9: Initial timing calibration of the tank PMTs using LED light sources. Note the different types of PMTs are grouped together due to similar cable lengths and positions in the tank. Measurements are accurate to about 10 ns. Figure from [112].




Figure 5.10: Timing calibration of the tank PMTs using a laser light source. Figure from [115]. Different colors represent PMTs at different locations.

into the water tank. Each pulse is 30 ps wide with a full width at half maximum (FWHM) value of 19 ps with a jitter of 3 ps. The laser is coupled to a single-mode fiber with a teflon diffuser ball at the end that disperses the light. A distribution of the hit times fitted with a Gaussian function was used to determine the timing offsets for each tank PMT. Overall, the PMTs have uncertainties between 1-1.5 ns and the BNB bunch structure is now visible in the ANNIE neutrino beam data, representing 81 53 MHz bunches from the Booster (Figure 5.11) [116]. These offsets are now incorporated into the event building analysis procedure (discussed in Chapter 6) for new runs and can be applied retroactively for older datasets. Figure 5.10 shows the distribution of the relative timing offsets.

5.3 Water Transparency and Quality

While many precautions were taken to reduce the possibility of contamination (see section 3), there is still a chance that something might happen to reduce the transparency of the Gd water. In addition, the Gd dissolved in the water might precipitate out of solution or be removed from the water filtration system. ANNIE uses LEDs to monitor the water transparency and quality, as changes indicate changes in the amount of light that can be seen by the photodetectors.

In an LED run, the LEDs installed at different positions on the inner structure are fired and only tank PMT waveforms are collected. More specifically, the bottom LED is fired at the same high intensity at the top PMTs (of ETEL type) and the amount of charge

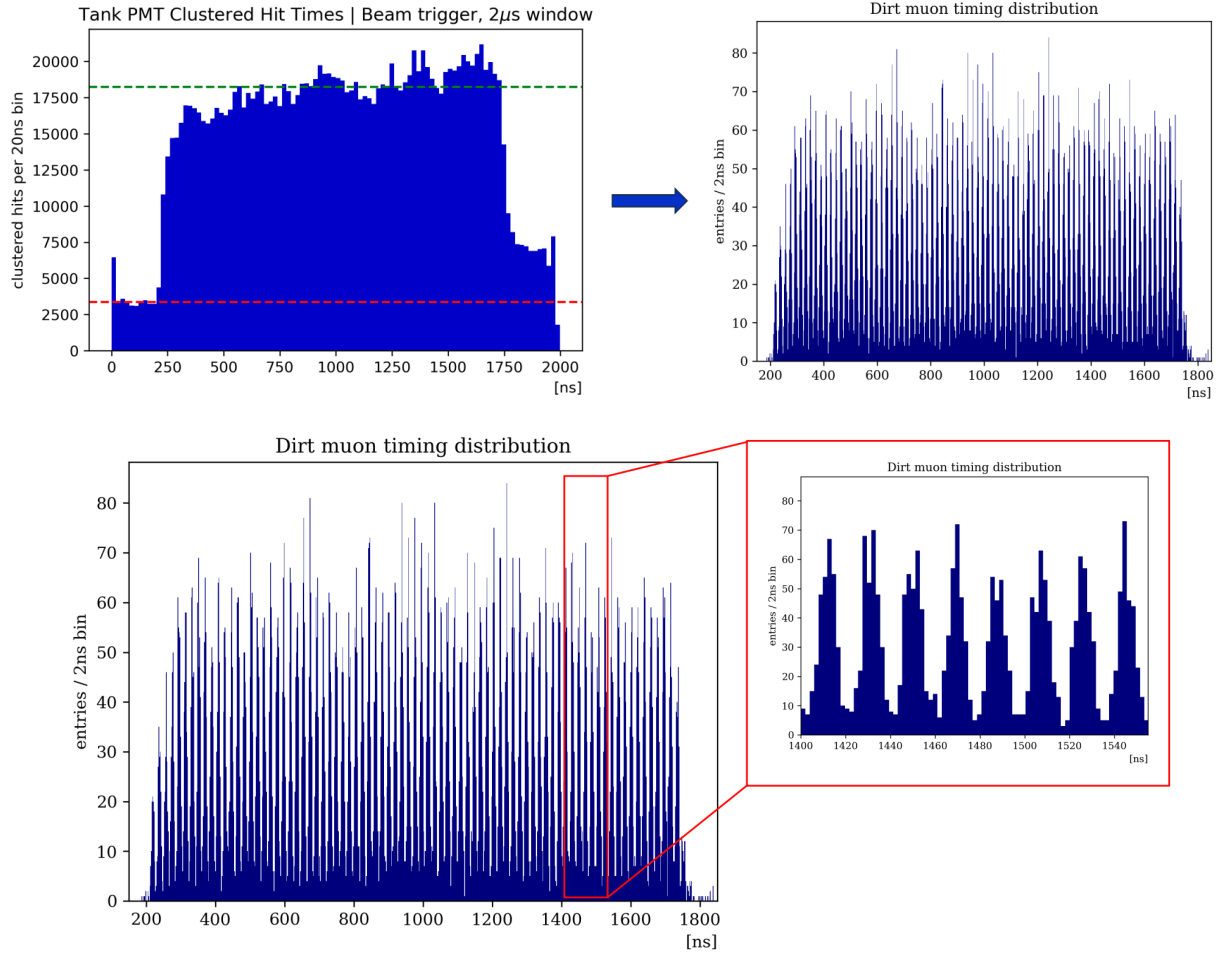


Figure 5.11: With the timing offsets measured using the laser diffuser ball, the structure of the Booster Neutrino Beam is now visible. The top left plot is a distribution of hit clusters whose cluster times fall in the prompt 2 μ s acquisition window. The 1.6 μ s beam spill can be seen atop a flat background (above the dashed red line). Before the timing corrections, the beam structure is not visible. The top right plot shows the timing distribution of a sample of dirt muon clusters after the timing corrections from laser calibration were applied. The bottom two plots are zoomed in versions of the top right plot to highlight the beam structure. Figures from [115].




Figure 5.12: A record of transparency measurements. Significant changes in the transparency can often be attributed to an external event, usually maintenance on the water filtration system.

collected is monitored over time. The initial intensity of the LEDs is compared to the light seen by the PMTs to determine whether transparency is still stable. Figure 5.12 shows a record of transparency measurements and shows that transparency has remained relatively stable since the start of running. As seen in the figure, decreases and increases in water transparency can usually be attributed to an external event, such as the replacement of filters or resins. Changes in transparency inform collaborators when maintenance of the water system is needed. Note that these are relative transparency measurements.

In addition to the transparency monitoring, the water quality is also checked by moni-

toring the conductivity, temperature, and pH of the water. Conductivity is important for monitoring the concentration of Gd and contamination from other ions. An increase in conductivity is an indicator for unwanted ions leaching into the water target. Acidity is also important because Gd requires a slightly acidic environment to remain dissolved in solution. Otherwise, a more basic environment encourages the formation of $\text{Gd}(\text{OH})_3$, which will precipitate Gd out of solution, reducing the neutron capture time. Water samples are taken from the circulation system regularly to check their spectra, similar to the material compatibility tests described in Section [3.2.2](#).

5.4 Neutron Capture and Detection Calibration


One of ANNIE's physics goals is to measure the neutron production rate following a neutrino-nucleus interaction. As such, it is important that the detector can detect neutron captures. By loading the water target with gadolinium (Gd), efficiency in capturing neutrons is increased: $\sim 90\%$ of neutron captures happen on Gd, while the rest happen on hydrogen. However, not all neutron captures will be detected due to fluctuations in the amount of light detected and the fact that neutrons capture on hydrogen only produced 2.2 MeV of light. This sections covers the characterization of the detector's neutron detection efficiency.



Figure 5.13: LEFT: AmBe source in acrylic housing. A deployment mechanism was being tested here. CENTER: AmBe source housing wrapped in teflon (PTFE) tape to maximize light collection. RIGHT: AmBe source deployment winch being tested. This was the method selected for deployment.

5.4.1 AmBe Source

An Americium-Beryllium (AmBe) source is a neutron source comprised of americium (Am-241) oxide with beryllium (Be-9) metal powder. Am-241 emits an α particle, which is absorbed by Be-9, which then emits a neutron. The following depicts this process:



The neutron is almost always produced with a 4.44 MeV gamma ray (γ), which comes from the de-excitation of carbon-12. The γ functions as a trigger to initiate readout of the photodetectors in the water tank and AmBe source setup, shown in Figure [5.13](#).

The 100 μCi AmBe source that ANNIE uses is housed in a 2-inch diameter cylindrical acrylic vessel with a closed end and a conical cap. Two O-rings secured under the cap provide watertightness, in addition to the bolts that attach the lid to the rest of the housing. Grooves are made along the length and through the base of the acrylic vessel for wires to sit flush and to hold the housing during deployment. Inside the housing, the AmBe source sits at the base with a 2-inch by 2-inch cylindrical bismuth germanium oxide (BGO) scintillator crystal on top of it. In initial designs, a stainless steel plate separates the source and the crystal to stop any X-rays that may come from neptunium-237, but was later removed due to size constraints. The crystal is wrapped in white teflon tape to maximize light collection through reflection. Attached to the underside of the cap are two silicon photomultipliers (SiPMs) for light detection. In later AmBe source calibration campaigns, the SiPMs are replaced by a 2-inch PMT for better signal-to-noise discrimination.

After the process in Eq. 5.4 takes place, the 4.44 MeV γ enters and Compton scatters in the BGO crystal, causing it to scintillate. BGO has a typical light yield of about 8,500 photons per MeV of energy deposited, with a peak in the 375-650 nm range. The scintillation photons make their way to the other end of the BGO crystal to be detected by the coupled PMT. When the PMT registers a hit, an acquisition window is opened to capture any activity in the tank resulting from a neutron capture. In the two SiPM configuration, a coincidence between the two SiPMs is required to trigger a readout and ensure that hits were not due to noise. This requirement reduced the number of background events by a factor of 25. Additionally, the AmBe source is optically isolated to prevent the scintillation photons from

contaminating the neutron data.

5.4.2 Neutron Detection Efficiency

With a 100 μCi source and a 10^{-4} efficiency for the (α, n) interaction channel shown above, a rate of approximately 100 tagged neutrons per second is expected. The emitted neutron from the AmBe can have energies of up to ~ 10 MeV [117], providing them with enough energy to make their way out of the acrylic housing and into the water to be captured by gadolinium (Gd). Following a neutron capture, Gd de-excites by emission of gamma rays, which in turn excite electrons that will go on to produce Cherenkov radiation that can be detected by the tank PMTs. To determine the neutron capture efficiency at a particular location in the detector, the number of triggers from the PMT is compared to the number of hit clusters found in the acquisition window. If the AmBe PMT fires but there is no tank activity, then the efficiency at that position is reduced. Note that other factors affect this efficiency.

Two calibration campaigns have taken place so far, both during periods of no beam. The trigger scheme for these runs relies on the PMT/SiPMs registering a hit above threshold (two consecutive samples in the waveform of at least 10 ADC counts above baseline, corresponding to a threshold of about 5.9 mV). In the event that this occurs, a 67 μs window is opened to record all tank waveforms, including the ones from the AmBe SiPMs/PMT. The AmBe source was deployed into the tank at one of the five calibration ports in the tank lid using

a custom built winch (see Figure 5.13). Data was taken at different heights to get a map of neutron capture efficiency with positional dependence in the tank. As shown in Figure 5.14, the neutron capture efficiencies range from $\sim 55\text{-}70\%$. These efficiency values will be used to correct the neutron production rate at different positions in the tank in the next chapter.

As discussed in Chapter 3, ANNIE conducted a neutron background measurement to determine the rate of neutrons entering the detector tank from outside. Other sources of background include cosmic muons (that may clip the detector), radioactive isotopes, and spallation. These sources of background will need to be characterized in future studies.

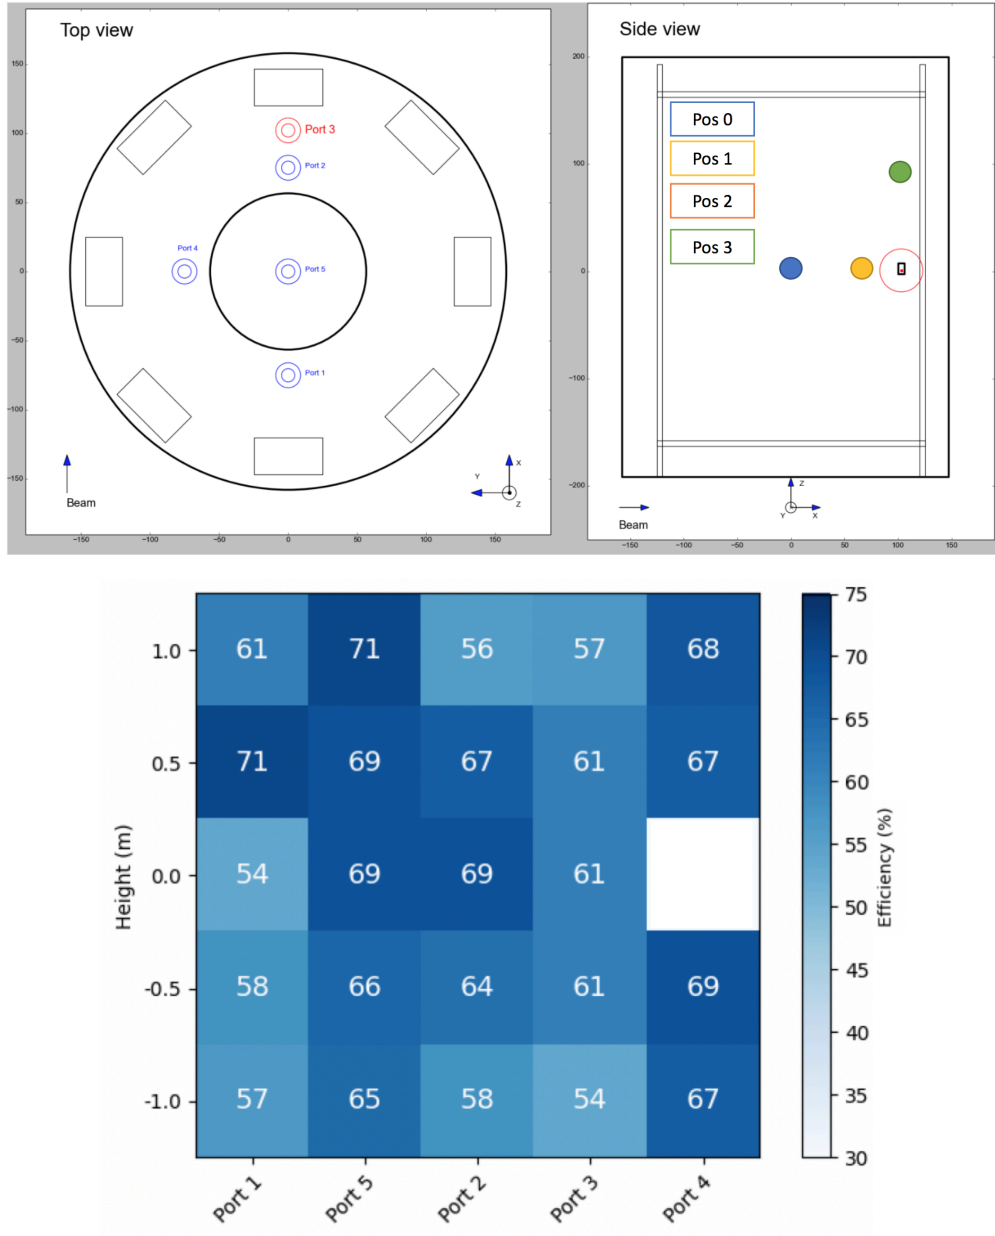


Figure 5.14: TOP: A schematic of source deployment and port locations. BOTTOM: Neutron capture detection efficiency as a function of tank position. This map was created when the BGO crystal was coupled to two SiPMs. A new map using data taken with a PMT is in progress. Diagram from [112].

Chapter 6

Neutron Multiplicity Measurement

Neutron multiplicity refers to the number of neutrons that are produced in the final state of neutrino-nucleus interactions. A measurement of neutron multiplicity is needed to improve the nuclear models that neutrino event generators use. While purely charged current quasi-elastic (CCQE) interactions do not involve any neutrons, the use of heavy nuclei as targets in neutrino experiments increases the likelihood of neutrons appearing in the final state due to significant nuclear effects. For example, an excited nucleus can de-excite by emitting neutrons (“boil-off” neutrons), a resonant pion can knock out a neutron, and so forth. Neutrons are an indicator of inelasticity, since they carry away energy from the initial interaction and often go undetected due to their lack of charge. This obfuscates the true neutrino energy and introduces biases in energy reconstruction. As mentioned in Chapter [3](#), the neutron multiplicity measurement performed by the Super-K and SNO experiments can be improved upon since the direction and sign (neutrino or antineutrino) of the incident neutrino were

unknown, and measurements were made as a function of visible energy. Since ANNIE sits on a muon neutrino beam with known sign and direction, and the energy of the incident neutrinos is known to some extent, a measurement of neutron multiplicity of an exclusive reaction can be made. More specifically, we can make a measurement of neutron multiplicity as a function of momentum transfer of CC- 0π interactions.

In general, the following steps outline the procedure to obtain a measurement of neutron multiplicity in ANNIE:

- Perform data reduction on the raw data, extracting only necessary information;
- Identify muons and calculate the relevant kinematic variables;
- Identify candidate neutron captures;
- Account for background neutrons; and
- Correct for neutron detection efficiency.

The rest of this chapter covers these steps in more detail.

6.1 Processing Raw Data Into ANNIE Events

Before any analysis takes place, the raw data must be cleaned and reduced to more manageable and user-friendly files. This section covers the data processing procedures developed by the ANNIE collaboration. Figure [6.1](#) depicts the toolchain associated with the processing of

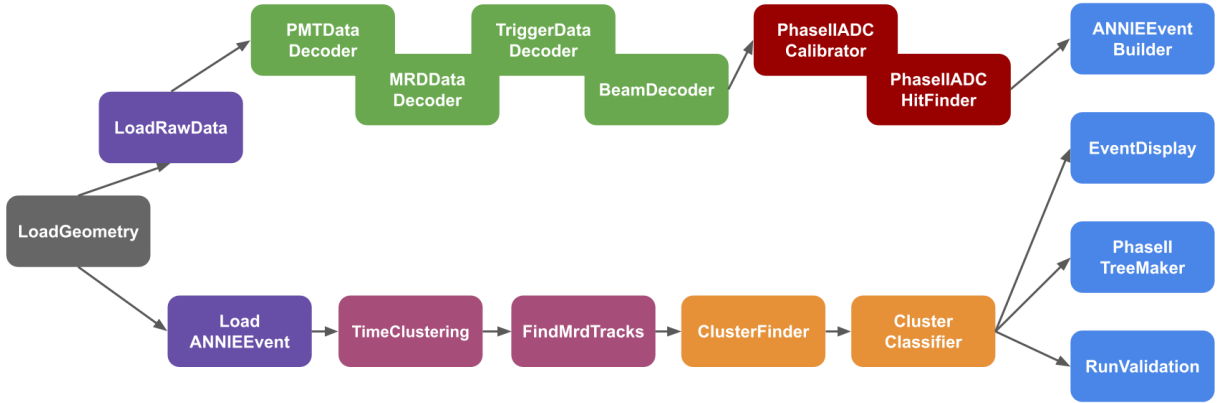


Figure 6.1: This diagram shows the tools that are involved in various analysis toolchains (blue). The tools relevant to event building starts at **LoadGeometry** and ends at **ANNIEEventBuilder**. This toolchain is where data from each subsystem are matched to the same physics event. Other toolchains can produce event displays, `root` files, and run validation plots.

raw data, which will be explained in more detail in the following sections. Beam data from runs 2506-2643 from the 2021 and runs 3774-4188 from the 2023¹ beam years are used for this thesis.

6.1.1 Event Building

For each beam signal from the resistive wall monitor (RWM) in the hall, 2- and 4- μs windows are open for the tank and MRD PMTs, respectively. The width of these time windows is set long enough to capture the 1.6- μs spread of the neutrino beam spill. When at least one tank PMT registers a hit of at least 5 PE, the acquisition window is extended to 70 μs to

¹It is noted that few events from the 2023 runs passed the selection cuts due to the presence of multiple reconstructed tracks in the MRD, which have known issues.

include tank activity that occurs after the first two microseconds. Data from each of the detector subsystems are recorded and saved separately and at times that differ from the global timestamp logged by the Central Trigger Card (CTC) due to the nature of electronics and signal propagation and readout times. As a consequence, the raw output from each subsystem must be timestamp decoded and matched to form an entry that belongs to the same physics event before performing any analysis, as depicted in Figure [6.2](#). The activity in the different subsystems must fall within certain time tolerances of each other for them to be assigned to the same physical event. This process is referred to as “event building” and the processed data is grouped into `ANNIEEvent` objects stored in `BoostStores` with raw waveforms, hit times, and so on. Any data entries that are not matched are sent to an `OrphanStore`. Additionally, beam information is saved alongside each event so that beam quality cuts can be made. This information includes the beam horn current and intensity.

6.1.2 Baseline Subtraction

Following event building and PMT gain calibration, baseline subtraction is performed for each PMT waveform due to the imperfect and varying nature of electronics. Given a fixed threshold for all PMTs (which could in principle be made unique for each PMT), differing baselines across PMTs would affect the pulse-finding algorithm discussed next.

Baseline subtraction in ANNIE is based on the baseline characterization method found in the `ZE3RA` (`ZEPLIN-III Reduction and Analysis`) software developed for the `ZEPLIN-III`

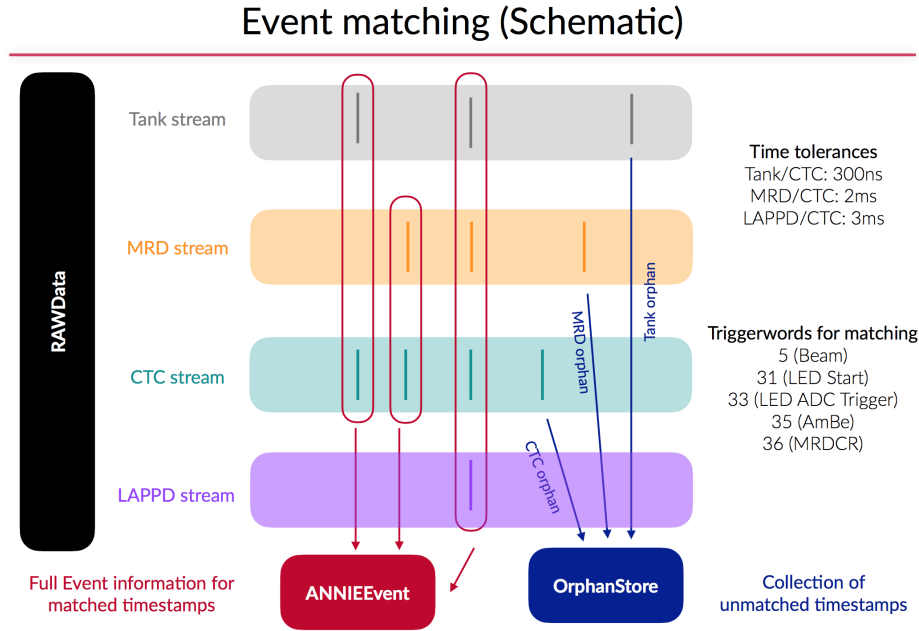


Figure 6.2: A schematic of the timestamp matching steps of the `ANNIEEvent` building process. Raw data from each subsystem is matched according their timestamps to form event entries that belong to the same physical process. Any unmatched data go to an `OrphanStore`. Figure from [118].

collaboration [119, 120]. The algorithm works as follows:

- Split the pre-trigger section of a PMT waveform in M regions with N samples each;
- Calculate the mean and standard deviation for region m ;
- Define a ratio R for neighboring regions m and $m + 1$ using their standard deviations;
- Remove regions that do not pass a F -test (probability threshold and hypothesis test);
- Estimate the baseline mean and standard deviation using the remaining regions.




Figure 6.3: The pre-trigger region of the above waveform is divided into M regions with N samples. The darker regions were used to estimate the baseline, while the lighter regions are where the baseline subtraction is applied, in addition to the darker regions. Note the sinusoidal aspect of this particular waveform that comes from power supply noise. For the future, a sinusoid may be a better fit to estimate the baseline. Figure from [112].

The F -test was selected based on the principle that if the baseline ADC values of regions m and $m + 1$ contain independent and random samples drawn from populations of the same variance, then the ratio R should follow an F -distribution. The probability P_m of observing a deviation of R_m from 1 is compared to a selected threshold P_{crit} , such that R_m passes the test if $P_m < P_{crit}$ is satisfied. Failed tests mean that regions m and $m + 1$ do not share the same mean and standard deviation, and region $m + 1$ is rejected from the final baseline estimation. More details about the algorithm can be found in [120].

The ZE3RA algorithm works well in excluding regions that may have an unexpected pulse

in them, so as to not bias the estimated baseline value. In turn, adjusting the baseline so that it is as close to zero as possible prevents any baseline shifts from affecting the hit-finding and charge integration that occurs later. Note that a sinusoidal component can be seen in Figure [6.3](#) due to noise from a PMT power supply. In the future, a sinusoidal fit could be used to estimate the baseline.

6.1.3 Hit Finding

The next crucial steps are to look for pulses, or hits, and to convert their ADC counts into voltages. For the tank PMTs, the conversion of ADC counts to voltage is 2^{12} ADC counts to 2.415 V, or 1,696 counts per volt. PMTs in the Front Muon Veto (FMV) and Muon Range Detector (MRD) do not have charge information associated with them, so do not need conversion.

After baseline subtraction, a waveform is scanned for pulses that meet the 7 ADC count threshold above the baseline. The integration window begins five samples before the point in the pulse where the threshold was crossed and ends after the waveform returns to a value within the uncertainty of the baseline determined in the baseline characterization step. This hit finding method was also applied to the PMT calibration procedure described in Section [5.1.1](#) to update the gain values from previous calibrations.

The isolated pulse is integrated to get the total charge Q of the hit via

$$Q = \frac{1}{R} \int_{t_i}^{t_f} v(t) dt \quad (6.1)$$

where $v(t)$ is the voltage of the waveform in a time bin defined by the sampling rate, R is the $50\text{-}\Omega$ impedance of the digitizer, and t_i and t_f define the start and end of the identified hit. This charge represents the amount of energy the PMT registered. In the past, the algorithm for finding hits used a different threshold requirement. Users should note that older runs may have been processed with this method and may require reprocessing.

The `PhaseIIADCCalibrator` and `PhaseIIADCHitFinder` tools in the event building toolchain are used to perform these steps, shown in Figure [6.1](#).

6.2 Tank and MRD Cluster Finding

6.2.1 Tank PMT Clusters

Now that the raw data has been timestamp-matched and processed into `ANNIEEvent` entries with identified hits with charge and timing information, it is time to start identifying interesting physics events. When a muon traverses and exits the detector tank, it is expected that a group of PMTs will register hits that close in time.

The clustering tools in the `ToolAnalysis` framework takes this assumption and looks for clusters of PMT hits that are within 50 ns of each other. This time period was chosen based on the initial studies of signal propagation delays. A sliding 50 ns window scans across the full $70\text{ }\mu\text{s}$ readout window and searches for the position that encapsulates the largest number of hits. These hits are formed into a cluster and the search is repeated with the remaining

hits until the largest number of hits in the sliding window falls below 10 hits. Each cluster is assigned a cluster time, defined as the average of all PMT hit times, and a cluster charge, the total charge observed by the PMTs in the cluster. The `ClusterFinder` tool perform clusterization for the tank PMTs.

An additional tool (`ClusterClassifier`) is applied to the tank clusters as a way to label clusters that may belong to different physical processes or particles. The tool assigns each cluster a charge balance parameter, CB , defined as

$$CB = \sqrt{\frac{\sum_i^N Q_i^2}{Q_{sum}^2} - \frac{1}{N}} \quad (6.2)$$

where N is the total number of PMTs in the cluster, Q_i is the total amount of charge seen by the i -th PMT, and Q_{sum} is the total charge observed in the entire cluster. This classifier is used to flag clusters that may arise due to instrumentation effects and noise. If the cluster charge is evenly distributed among the cluster PMTs, CB approaches zero. However, if all the charge is registered by a single PMT, CB approaches one. This method for classifying clusters was adapted from the WATCHBOY experiment to identify neutron captures on gadolinium [121].

6.2.2 MRD PMT Clusters

Similarly, time clusters are found for the MRD PMTs and used for track reconstruction. This is done by the `TimeClustering` tool. For these clusters, a track will be fit to the hits to determine the track direction and energy of the muon in the MRD, which we will be used to

fit the muon vertex in the tank described in Section [6.6.1](#). No cluster classifiers are applied to the MRD clusters, but selection cuts can be made based on the track angle, number of layers hit, whether the muon stopped in the MRD, and so on.

6.3 MRD Track Reconstruction

MRD data comes from TDC units, which record hit timing information but not charge information. Some positional information is available due to the way the scintillator paddles are setup: in every horizontal and vertical layer, there are two paddles in each row so that each layer consists of two halves.

The MRD track reconstruction algorithm is based on the cellular automation method used by K2K [\[122\]](#) and T2K [\[123\]](#). It takes the MRD time clusters found using the method in the previous section and fits tracks to the hits. Horizontal and vertical views of the track are reconstructed independently and then combined in the end. First, for each layer, hits in adjacent paddles are grouped into a “cluster” (different from the clusters described earlier). Each cluster is then grouped with all clusters in the next two downstream layers to form a “cell”. These cells represent a segment of the MRD track and are assigned an upstream neighbor that is sufficiently aligned. This is repeated until the end of a track is reached. A visualization of the track reconstruction algorithm can be found in Figure [6.4](#).

The result of this procedure are 2D tracks for the horizontally- and vertically-oriented layers that are combined to construct a complete 3D track. As a final step, a linear fit is

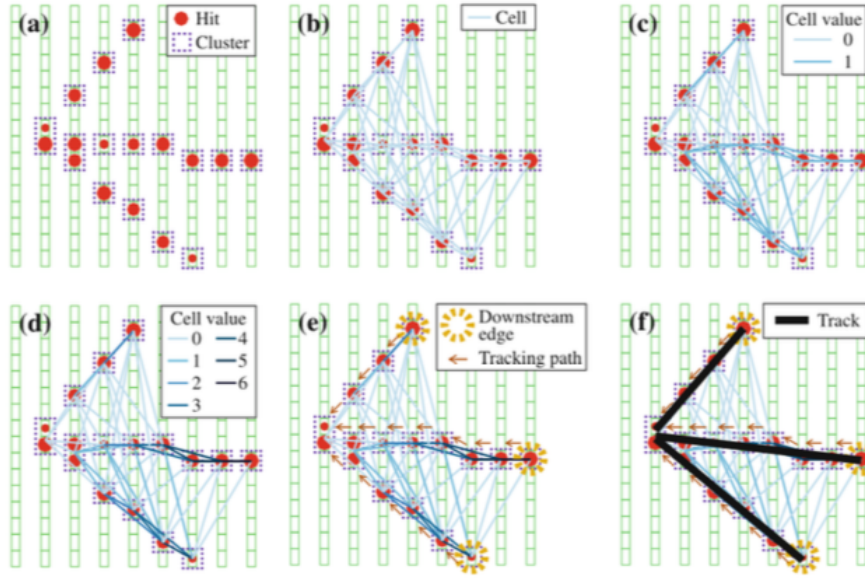


Figure 6.4: Cellular automation method for MRD track reconstruction. (a) Hits in the same layer are grouped into clusters. (b) Clusters in neighboring layers form cells. (c) Each cell is assigned a status value. (d)-(e) Combinations of cells are made according to status until the end of a track is reached. (f) Linear fits are made to track candidates. Figure from [123].

performed on all the clusters to determine properties such as total length and angle with respect to the beam direction. The performance of the track angle reconstruction is shown in Figure 6.5. The mean difference between the reconstructed and true track angle is about -0.03 radians (-1.4°), with a width of about $\sigma_\theta \sim 0.16$ radians (9.4°). The uncertainty on the reconstructed track angle will be propagated through to the reconstructed muon energy.

In addition to the outlined procedure, certain requirements for the MRD topology need to be met. Hits in the MRD must be close in time (coincident) with each other form a cluster. A cluster requires a minimum of four paddle hits within a 30-ns window, more specifically,




Figure 6.5: Difference between reconstructed and true muon track angle for events that pass the FV cut. The mean difference is -0.026 rad with a σ of 0.16 rad. The error on the reconstructed track angle will be propagated through to the muon energy reconstruction.

at least two hits in vertical paddles and two hits in horizontal paddles. More details of the MRD track-fitting algorithm can be found in [77] and [114].

6.3.1 Effective Track Length in MRD

While the above algorithm performs fairly well in reconstructing the MRD portion of the muon track, a separate track reconstruction method was developed for energy reconstruction purposes. The high density of iron increases the chances of scattering of the muon as it traverses the MRD (see Figure 6.6). Scattering adds to the total track length and results in more energy deposited compared to a straight linear track [124]. Thus, the MRD track

reconstruction method was adjusted to account for these effects. We do so by counting the number of scintillator layers that were crossed, which has nearly a direct correspondence to the number of iron layers. This number is converted into an effective amount of iron traversed by the muon by multiplying by 5 cm (thickness of iron slab) and dividing by the cosine of the track angle, determined with the above reconstruction technique. Since it is uncertain where the muon stops, half an iron layer is added to the final effective track length to be used in energy reconstruction later. This way, we can take advantage of the minimum ionizing nature of the muon and use the Bethe-Bloch formula to determine the energy deposited in the ANNIE detector.

The impact of scattering on the total track length can be seen in Figure [6.7](#). For muons of all energies in the range of 100-1,000 MeV, the actual track length traversed is longer than the track length estimated by calculating the distance between the MRD entry and stopping or exit points in the current ANNIE reconstruction methods. Impacts of scattering are taken into account by using a percentage correction based on **Geant4** simulations. The track length in the MRD is corrected based on the muon's energy when entering the MRD. For example, referring to Figure [6.7](#), muons with entering energies less than 200 MeV would have tracks corrected by a factor of 80-90%.




Figure 6.6: One hundred muons with initial kinetic energy of 500 MeV were fired into a MRD-like geometry in `Geant4`. Typical paths are not linear and demonstrate the effects of scattering. Actual path lengths are longer than the path length determined by fitting a line to the muon entry and stopping points, which is the current method in existing analysis tools.

6.4 Muon Candidate Selection (Prompt Window)

An important step towards an improved neutron multiplicity measurement is to select a pure sample of muon candidates. The muons give us the Q^2 dependence of the rate of neutron production. Since we know approximately when neutrinos are incoming (due to the signal from the Accelerator Division), we can reasonably expect muons to be created within the first $2\ \mu\text{s}$ of the acquisition window. To find the perfect muon candidate, the following cuts were made:




Figure 6.7: Ratios of a linear (“straight”) track to the actual track lengths as a function of initial muon kinetic energy. The linear track is calculated by taking the distance between the entry point and the stopping point. For lower energy muons entering the MRD, the track actual track length is longer than the track length determined by connecting the entry and stopping points. Results generated from **Geant4** simulations. Corrections are made to the track length depending on initial estimates of the muon energy. Data points beyond 800 MeV did not have a muon stopping in the MRD geometry in the simulation.

- Select BEAM trigger events;
- Clusters whose cluster times were within 2000 ns of the event start time;
- If there is more than one cluster in the prompt window, select the one with the most amount of charge;
- No hits in the FMV to ensure events originate within the water target;
- One and only one reconstructed track that stops in the MRD so we are not dealing

with multi-track events;

- Extending the MRD track backwards into the tank should cross a one cubic meter spherical volume around the tank center;
- Tank activity and MRD activity should be close in time to each other (time difference of $750 \text{ ns} \pm 50 \text{ ns}$ for data and $0 \pm 50 \text{ ns}$ for simulations).

Tracks that stop in the MRD are required, since it would not be possible to reconstruct the energy of those muons that do not end in the MRD. The requirement of a time coincidence between tank and MRD activities ensures that hit clusters from each subsystem are correlated with each other and belong to the same physical process. Figure [6.8](#) plots the times of tank and MRD hit clusters. There is a clear correlation between certain clusters and these result in a time coincidence of $\sim 750 \text{ ns}$. The accepted time window is $700 \text{ ns} < \Delta t < 800 \text{ ns}$. Once the above cuts have been applied, we move on with a ring imaging technique to fit the muon vertex. This technique is described in the next section.

6.5 Imaging the Cherenkov Ring

While LAPPDs have been extensively tested and deployed in the ANNIE tank, the relevant tools to integrate this datastream are currently in development. Thus, this thesis is not able to take advantage of the picosecond timing that LAPPDs have to offer. As a result, I developed a method to fit the vertex at which the muon is created with the data available. The




Figure 6.8: LEFT: MRD cluster times vs. tank cluster times. One can see a correlation between hit clusters that come from beam events. RIGHT: Difference between MRD and tank cluster times. Correlated activity has a time difference of ~ 750 ns. The accepted time window if (700 ns, 800 ns) is marked with black lines. Note the logarithmic scale in the number of clusters. Figures from [112].

vertex reconstruction will be done with the following ring imaging technique that involves detecting the ring edge by taking into account two ideas explained below.

First, for each PMT that fired, we can find a segment of the muon track in water, constrained along the MRD track direction and by the Cherenkov angle in water, θ_C . We call this segment of the tank track a_i . The location of photon emission can be inferred from the track segments a_i , and the amount of light emitted at points along the track can be determined. This is the second part of the technique and will be described later. Referring to Figure [6.9], the following definitions are as follows:

- $\hat{\mu}$: opposite of the muon track direction;
- \vec{R}_i : a vector that starts at a reconstructed tank exit point (the star in Figure [6.9]) and points to the i^{th} PMT;




Figure 6.9: Given a muon vertex, we can define the variables above. a_i represents the tank segment between the location of the muon and where it exits the tank. R_i is the vector from the tank exit point to PMT i . b_i is the vector from the vertex to PMT i .

- \vec{b}_i : a vector that starts at a point along the muon track and ends at the i^{th} PMT;
- \vec{a}_i : a vector that starts at a point along the muon track and ends at the reconstructed tank exit point, with

$$\vec{a}_i = a_i \hat{\mu} \quad (6.3)$$

We can write

$$\vec{a}_i + \vec{R}_i = \vec{b}_i \Rightarrow \vec{a}_i = \vec{b}_i - \vec{R}_i \quad (6.4)$$

Then we dot $\hat{\mu}$ into the equation above

$$\hat{\mu} \cdot \vec{a}_i = (\hat{\mu} \cdot \vec{b}_i) - (\hat{\mu} \cdot \vec{R}_i)$$

$$a_i = b_i \cos \theta_C + (\hat{\mu} \cdot \vec{R}_i) \quad (6.5)$$

Next we cross Eq. [6.4](#) with $\hat{\mu}$

$$\begin{aligned} \hat{\mu} \times \vec{a}_i &= 0 = (\hat{\mu} \times \vec{b}_i) - (\hat{\mu} \times \vec{R}_i) \\ 0 &= -b_i \sin \theta_C - (\hat{\mu} \times \vec{R}_i) \\ \vec{b}_i &= -\frac{\hat{\mu} \times \vec{R}_i}{\sin \theta_C} \end{aligned} \quad (6.6)$$

Now taking Eq. [6.6](#) and substituting it into Eq. [6.5](#), we get

$$a_i = -\frac{(\hat{\mu} \times \vec{R}_i)}{\tan \theta_C} + (\hat{\mu} \cdot \vec{R}_i) \quad (6.7)$$

Thus, for every PMT that fired, there is a track segment a_i associated with it, indicating the point where the photon was emitted.

The second part of this ring imaging method is based on the fact that muons with high energy are taken to be minimum ionizing particles when traveling through a medium such as water. This means that they do not lose much energy to ionization and continue to radiate a constant amount of Cherenkov photons at a fixed angle per unit distance traveled. Let us call this η with units of # photons/cm.

Referring to Figure [6.10](#), let's say the muon travels some distance Δx , emitting a constant number of photons. In the resulting ring of Cherenkov light, the total number of photoelectrons (PE) should be the same as the total number of PE emitted in the next ring that the




Figure 6.10: A muon that travels some distance Δx will emit a constant number of photons. These photons will make up a section of the Cherenkov cone called a frustum, whose surface area is imaged as a ring on the photodetectors. Left figure from [125].

muon creates after traveling the same amount of distance Δx . More accurately, this ring is the surface area of a frustum – a portion of a cone (see left diagram in Figure 6.10) – and we can calculate the area of this frustum, which we call ΔA_{ring} . This area ΔA_{ring} overlaps with the surface area of some PMTs, and even for those overlapping PMTs, it may not overlap completely with the entire photocathode. Thus, we define a fraction f_i for the fraction of ΔA_{ring} that the PMT takes up. Then, knowing the total amount of PE that falls in the ring, divided by the total fraction of photocathode overlapping with the ring area, we can get the total number of PE that the muon radiated in Δx and thus the photon density, η . In summary, we can get an η value for each ring, as we take incremental steps (Δx) back into the tank.

The fraction of the photocathode area of the i -th PMT that overlaps with the frustum can be written as

$$f_i = \frac{A_{PMT,i}}{2} \frac{(1 + \cos \theta_i)}{\Delta A_{ring}} \quad (6.8)$$

so that

$$f_{total} = \sum f_i.$$

Since we know that η should be a constant photon density, we can extract the actual η from

$$\eta = \frac{\text{total \# p.e. in ring}}{f_{total} \Delta x} \quad (6.9)$$

Combining the two ideas above, we can plot η as a function of track segment a_i . At each a_i , there is a ring associated with it and we can extract an η value from the PMTs that fall inside this ring area. As a_i increases, that is, as we move away from the tank exit point and deeper into the tank, we expect η to start at high values since we are inside the Cherenkov cone/disk. Once we are outside the cone/disk, η values should drop to a lower value. Figure 6.11 shows an example Monte Carlo (MC) event with truth information shown to highlight some notable features. The vertical blue line indicates the truth track length in water. To the left and right of the blue line are green and red horizontal lines that show the average η values, respectively. The left average η value (inside the disk) is higher than the right average η (outside the disk), which is what we expect. Looking at the event displays and η vs. a_i plots for data and MC (Figure 6.13), this method is effective in finding the ring edge.




Figure 6.11: This is a photon density η vs. tank track segment a_i plot with truth information shown. To the left and right of the true tank track length (vertical blue line), there is a clear distinction in the average photon density at shorter track segments (left horizontal red line) and at longer track segments (right horizontal green line).

After going through several of these η vs. a_i plots with truth information, the truth information is then removed and the track length is manually fitted by eye. I scanned through about a thousand events and guessed where the true track length might be. For example, in Figure [6.12](#), I would guess a tank track length of about 150 cm. The true track length was 148.0 cm, which is not too far off. Figure [6.14](#) shows a distribution of the differences between the fitted and true track lengths in water. The mean difference is about 12 cm, with a standard deviation of 30 cm and RMS of 32 cm. For cross checks,




Figure 6.12: For events with a clear break in the photon density, a manual fit is easily made. For this event, a tank track of 150 cm will be the fit. The true tank track length is 148 cm.

postdoc helpers were recruited to manually fit the events in data and in MC to determine the systematic error for this technique. See Appendix [C](#).

For additional statistics, a machine learning script was developed by fellow graduate student Tianqi Zhang [\[126\]](#) to train a Recurrent Neural Network (RNN) model to fit the tank track lengths. Figure [6.15](#) shows the performance of the model. The majority of the events were fit with the RNN model.

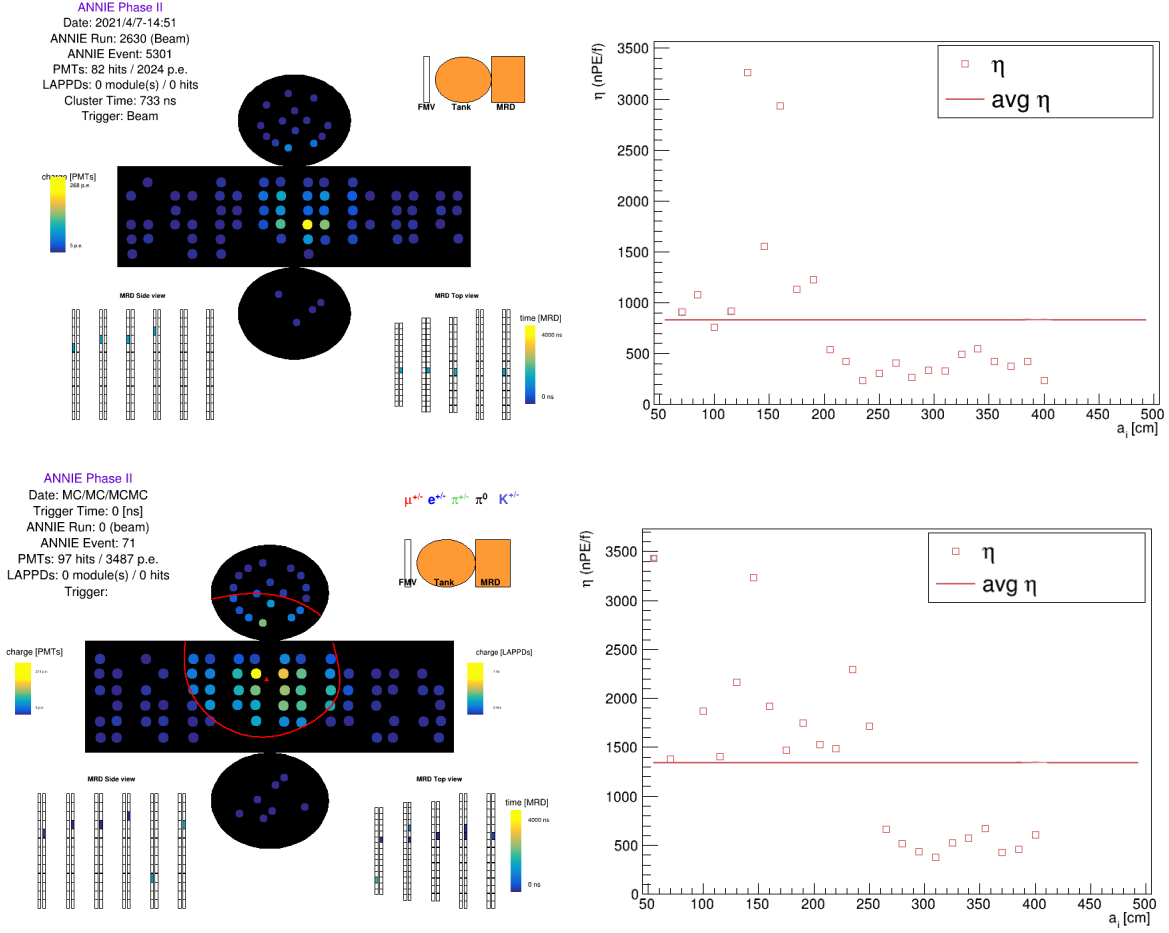


Figure 6.13: TOP LEFT: An event display of a *beam* event. TOP RIGHT: A η vs. a_i plot. Here one would estimate that the track segment is about 190 cm. BOTTOM LEFT: An event display of a *simulated* beam event. BOTTOM RIGHT: A η vs. a_i plot. Here one would estimate that the track segment is about 250 cm.




Figure 6.14: A distribution of differences between the manually fitted and true track lengths in water. The mean difference is -12 cm, with a standard deviation of 30 cm and RMS of 32 cm.

6.6 Muon Vertex Fit and Energy Reconstruction

6.6.1 Fitting the Vertex

Once a manual fit to the tank track length is performed, the muon vertex is found by starting at the reconstructed tank exit point and moving backwards in the tank by a distance equal to the fitted track length. Figure [6.16](#) depicts this process in a diagram.

A fiducial volume (FV) cut is then made on the muon vertices to select the final sample of muons. The Phase I background measurement from section [3.1](#) indicate that with enough water shielding, the neutron background rate is reduced to low enough levels for physics analyses. Given these results, a 36 cm overburden and a 20-27 cm water burden from the




Figure 6.15: A Recurrent Neural Network (RNN) was trained on η vs. a_i data and truth information to develop a model to fit the track length in the water tank. Compared to Figure 6.14, this model performs well.

upstream side of the tank is placed around an approximately 2.5 ton fiducial volume. This FV is also slightly offset upstream from the center of the tank as simulation showed that, setting the FV upstream from the tank center increases the neutron detection acceptance and encapsulates more muons that range out in the MRD. The resulting FV is a half cylinder with a radius of 1 m and a height of 1 m. As will be discussed in a later section, measurements using a downstream FV cut (closer to the MRD) is compared to the current upstream FV. Figure 6.18 highlights the two volumes. After making a FV cut, this technique performs to similar expectations as in [66] in reconstructing the muon vertex without LAPPD data (see




Figure 6.16: Once the track length L_{track} in the detector tank is determined, the muon vertex (blue star) is found by moving backwards into the tank from the tank exit point (red star) by a distance equal to the fitted track length.

Figure 6.19).

6.6.2 Reconstructing Muon Energy

Above ~ 200 MeV, muons lose little energy to ionization and can deposit a fixed amount of energy by Cherenkov radiation per unit distance traveled in water. Their energy loss can be characterized by the Bethe-Bloch equation [124, 127]:

$$\left\langle -\frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \quad (6.10)$$

For a muon with $\mathcal{O}(10^2)$ MeV kinetic energy, Eq. 6.10 is roughly $2 \text{ MeV}/(\text{g}/\text{cm}^2)$ in water.

This means a muon will deposit at minimum $2 \text{ MeV}/\text{cm}$ in the water. Since we know the track length in the water tank and in the MRD, we can determine how much energy was

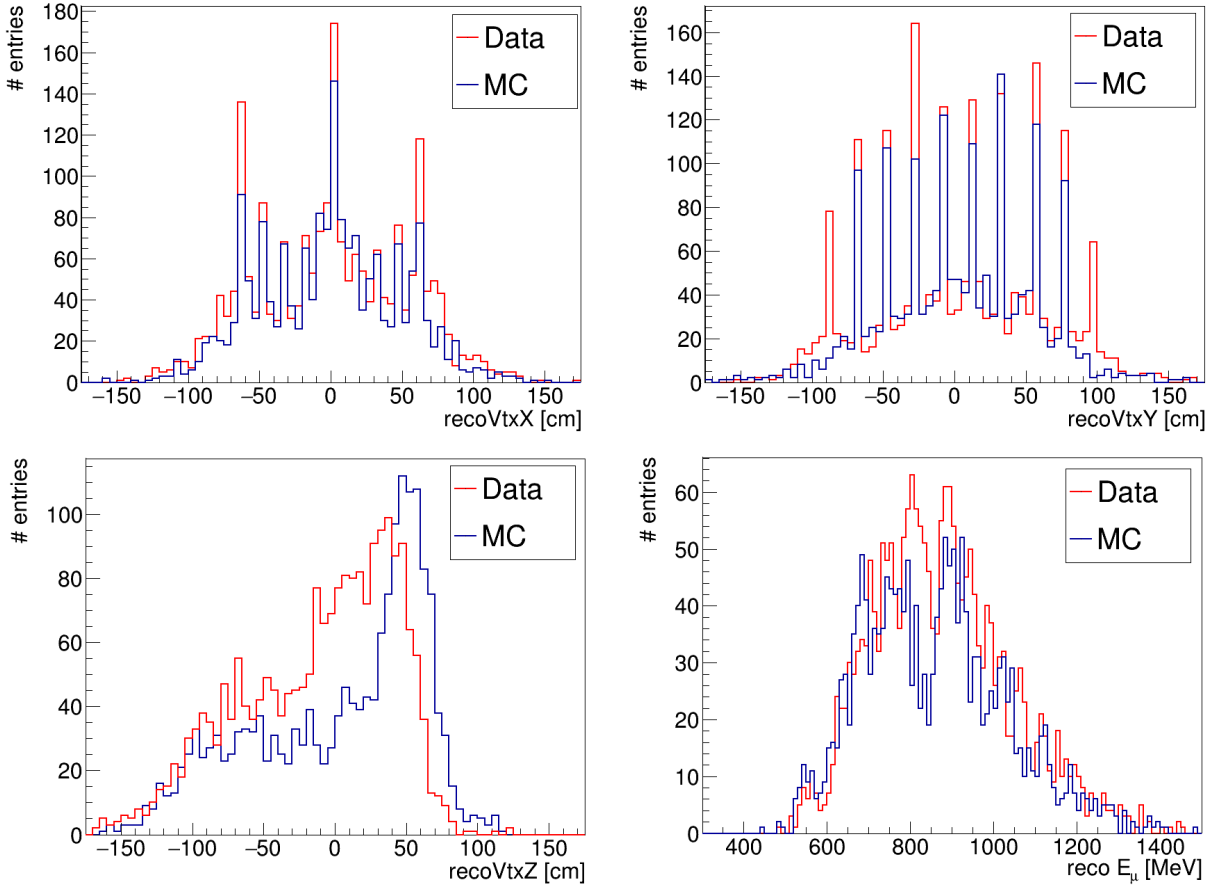


Figure 6.17: Reconstructed muon vertices and energy in beam data (no FV cut). The x and y distributions are centered around 0 cm, while the z distribution has a slight shift towards the upstream region of the tank. The reconstructed muon energy is centered around 800 MeV, which agrees with expectations based on the BNB energy spectrum.




Figure 6.18: A fiducial volume (FV) cut is made on the reconstructed vertices. The left plot is a topdown view of the detector, while the right is a side view of the detector. The red outlines the main FV cut that is shifted upstream from the center. The blue outlines a downstream FV cut used for neutron containment studies in a later section. The green outlines the active volume.

deposited in the water ($\rho = 1 \text{ g/cm}^3$). For example, a track length in the tank of 140 cm and an average dE/dx of 2 MeV/cm results in 280 MeV of energy deposited in the water, while an effective track length of 40 cm and an average dE/dx of 11 MeV (assuming iron is the main source of energy loss) results in 440 MeV deposited in the MRD. In total, this muon is estimated to have 720 MeV of kinetic energy. This is very simple approximation. To make a more accurate determination of the muon's energy, we take the initial determination of the energy and input it in the Bethe-Bloch equation to update the dE/dx value as the muon travels through the detector. The procedure is outlined below:

1. Use average dE/dx factors to make an initial estimate E_0 of the energy deposited in the tank and in the MRD.
2. Take the initial energy E_0 and calculate a new dE/dx value from the Bethe-Bloch

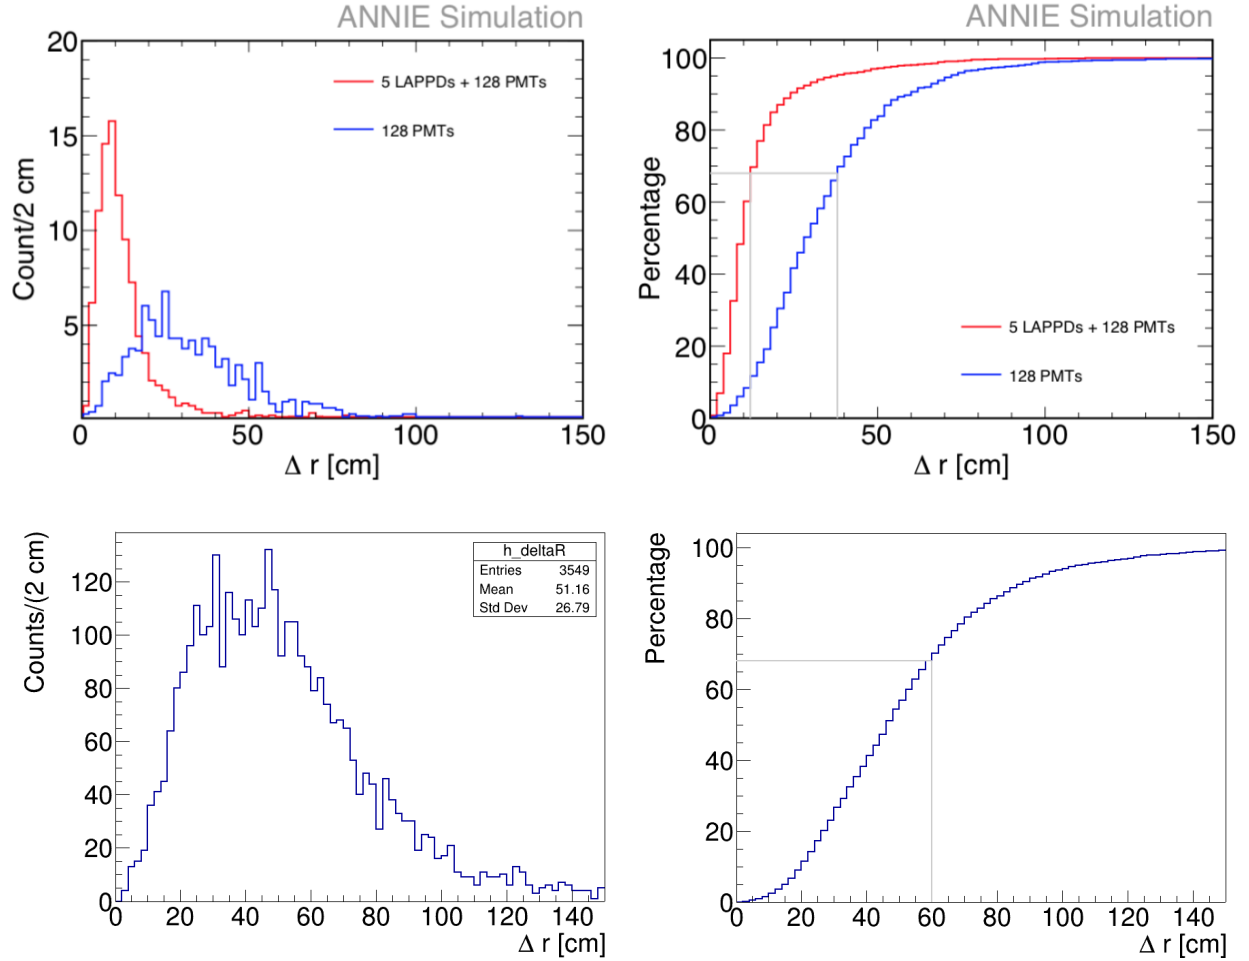


Figure 6.19: The top two graphs demonstrate the performance of vertex reconstruction in ANNIE simulations with and without LAPPDs. The bottom two graphs demonstrate the performance of the vertex reconstruction method in this thesis. The Δr distribution corresponds to the blue curve in the distribution above it. The vertex reconstruction method performs similar to the scenario of 128 PMTs without LAPPDs deployed. In the cumulative distribution, the ANNIE reconstruction methods indicate a spatial resolution of about 40 cm, while the method presented in this thesis has a spatial resolution of 60 cm.

equation.

3. Starting from the vertex, multiply the new dE/dx value by the fixed amount of distance traveled by the muon, e.g., 2 cm. Add the deposited energy to a cumulative sum.
4. Subtract the deposited energy from E_0 and input the updated energy E'_0 into the Bethe-Bloch formula to get a new dE/dx value.
5. Repeat the above two steps until the remaining input energy falls below 20 MeV.
6. Add the remaining input energy (~ 20 MeV) to the cumulative sum.

The final sum is the new reconstructed energy of the muon. Initial efforts to reconstruct the MRD energy were based on a look-up table generated from simulations. Once the track angle is determined, the dE/dx value is found from the look-up table and used to determine the energy loss. The outlined method provides a more realistic picture of the energy loss.

Using the effective MRD track length found in the MRD track reconstruction section, we can implement the Bethe-Bloch method outlined above. This relies on the assumption that muons deposit most of their energy in the iron. Recall that muons whose tracks stop in the MRD are reconstructed, since the energy of a muon that exits the MRD cannot be determined. Figure [6.20](#) shows a distribution of the difference between the reconstructed muon energy and the true muon energy for events that pass the FV cut. A width of about 65 MeV corresponds to an energy resolution $\Delta E/E$ of about 10% (assuming muon energy peak is ~ 800 MeV according to Figure [6.17](#)).




Figure 6.20: A comparison of reconstruction energy to the true energy of the muon. A Gaussian curve is fitted to the sample of events after a fiducial cut is made on the reconstructed muon vertices. The mean energy difference is about -15 MeV, meaning the energy reconstruction algorithm is slightly underestimating the muon’s kinetic energy. The width of this fit is about 66 MeV. Events with large differences are mainly attributed to pions or events with partially missing truth information, e.g., the track length in the MRD was zero.

6.7 Neutron Candidate Selection (Delayed Window)

Recall that when a beam trigger is received by the DAQ electronics, a $2 \mu\text{s}$ window is open to capture any additional activity in the detector. If at least one PMT registers a hit consistent with at least 5 PE, then the acquisition window is extended to $70 \mu\text{s}$. The first $2 \mu\text{s}$ of the event is the prompt component, while the remainder of the $70 \mu\text{s}$ is the delayed component

of the event. With the addition of gadolinium, the neutron capture time is reduced to about $30 \mu s$ in ANNIE compared to $200 \mu s$ in pure water. Clusters of hits found in the prompt window are usually associated with muons, whereas those found in the delayed window are associated with neutrons. It is in the delayed window where we will find and count our neutrons. Section [6.2](#) outlined the clustering algorithm.

To select neutron candidate clusters, first a time cut is made. While anything after the first $2 \mu s$ is considered delayed, any activity between 2 and $10 \mu s$ is excluded due to after-pulsing associated with the initial energetic muon. Delayed clusters with times greater than $10 \mu s$ are selected, to avoid these effects.

Next are charge cuts. Clusters are selected based on a charge balance (CB) classifier to get rid of any noise clusters and isolated peculiarities. Charge balance is defined in Equation [6.2](#). Based on previous studies, clusters with $CB < 0.4$ are selected. An additional charge cut limiting the total amount of charge to $Q_{tot} < 120$ PE. These cuts are summarized below and visualized in Figure [6.21](#).

$$t_{cluster} > 10 \mu s$$

$$Q_{tot} < 120 \text{ PE}$$

$$CB < 0.4$$

$$CB < 0.5 - \frac{Q_{tot}}{300 \text{ PE}}$$




Figure 6.21: This plot visualizes the cuts made to select neutron cluster candidates. The lower left quadrant is where candidate neutron clusters will be taken. The cuts were selected based on the first full AmBe source calibration campaign. AmBe data is available for a more recent campaign, but analysis is ongoing. Once completed, new selection cuts should provide a better sample of neutron cluster candidates. Figure from [114].

While a data for the most recent AmBe source calibration campaign is available, analysis of the data is ongoing. Once completed, new selection cuts should be defined with the latest AmBe calibration dataset and new neutron capture efficiency values.

Figure 6.22 shows a timing distribution of the neutron candidates that passed the above cuts. An exponential fit² is performed on the data to determine the capture time:

$$f(t) = e^{A+Bt} = A' \cdot e^{Bt}$$

²Using the built-in exponential function in ROOT




Figure 6.22: A distribution of neutron candidate cluster times is fitted with an exponential function. The time constant of this exponential is $(33 \pm 2) \mu\text{s}$, which is close to the expected value of $\tau \approx 30 \mu\text{s}$, but is slightly higher. Further investigations are needed to understand this shift to a longer time constant.

$$\tau = \frac{1}{B}$$

An exponential fit to the capture times returns a time constant of about $(33 \pm 2) \mu\text{s}$, which is close to the expected value of $\tau_{exp} = 30 \mu\text{s}$.

Since a time cut of $t_{cluster} > 10 \mu\text{s}$ is made when selecting neutron candidates, there will be missed neutrons in the final neutron multiplicity measurement. This systematic can be

quantified by the fraction of missed neutrons:

$$f = \frac{\int_{0.2}^{10} e^{-t/\tau} dt + \int_{70}^{\infty} e^{-t/\tau} dt}{\int_{0.2}^{\infty} e^{-t/\tau} dt} \quad (6.11)$$

where the first integral in the numerator represents the missed events from the time it takes the neutron to slow down ($0.2 \mu s$) to the time cut, and the second integral are the neutrons that were captured beyond the acquisition window. The integral in the denominator is used for normalization. For the time constant determined above, Eq. [6.11](#) returns a fraction f of 0.378, that is, about 38% of neutron captures are missed by applying the time cut of $t_{cluster} > 10 \mu s$. The efficiency in the time cut is then given by

$$\epsilon_t = 1 - f \quad (6.12)$$

and returns a value of 0.622, or $\sim 62\%$. This correction is applied to the average neutron yields presented in later sections and treated as a systematic uncertainty.

The distribution of the number of neutrons with varying event selection cuts applied is shown in Figure [6.23](#). The majority of events have no neutrons in the final state, and the total number of events is significantly reduced with each cut.




Figure 6.23: Distributions of events according to number of neutrons in the final state, following a neutrino-nucleus interaction. The blue curve represents all neutron candidates, the black curve are the number of candidates with a stopping MRD track, and the red curve are the remaining candidates with a FV cut. The majority of events have no neutrons in the final state. The right plot is the left plot in logarithmic scale.

6.8 Neutron Multiplicity as a Function of Momentum Transfer

The energy of the neutrino can be determined using the kinematics of the outgoing lepton, which is the muon in this case:

$$E_{rec}^{\nu} = \frac{2(M - E_b)E_{\mu} - (E_b^2 - 2ME_b + m_{\mu}^2 + \Delta M^2)}{2(M - E_b - E_{\mu} + |\vec{p}_{\mu}| \cos \theta_{\mu})} \quad (6.13)$$

The momentum transfer Q^2 can then be calculated by

$$Q^2 = 2E_{rec}^{\nu}(E_{\mu} - p_{\mu} \cos \theta_{\mu}) - M_{\mu}^2 \quad (6.14)$$

For each event with a muon candidate, the number of neutrons is counted and a Q^2 value calculated. Figure [6.24](#) displays neutron multiplicity across ranges of visible energy

E_μ , momentum transfer Q^2 , and track angle $\cos\theta_\mu$. An increase in the neutron yield can be seen at higher Q^2 values.

It is useful to look at the average neutron multiplicity in each energy bin. For each 200 MeV wide energy bin i , the average neutron multiplicity and uncertainties are determined by

$$\bar{n}_i = \frac{1}{N_{total,i}} \sum_{j=1}^{N_{bins,y}} j \cdot n_{ij} \quad (6.15)$$

$$\sigma_i = \sqrt{\frac{1}{N_{total,i} - 1} \sum_{j=1}^{N_{bins,y}} (j - \mu_i)^2 \cdot n_{ij}} \quad (6.16)$$

where j is the neutron multiplicity, n_{ij} is the number of events in the i -th energy bin with j neutrons, $N_{total,i}$ is the total number of events in the same bin, and μ_i is the averaged neutron multiplicity in energy bin i . Figure 6.25 presents the averages as a function of muon energy and momentum transfer. Compared with the Super-K and SNO measurements (Figure 3.18), ANNIE sees a lower neutron multiplicity for events with similar visible energies.

The overall average neutron multiplicity across all muon energies and Q^2 values is $\bar{n}_{data} = 0.294 \pm 0.015(\text{stat}) \pm 0.18(\text{sys})$ and $\bar{n}_{MC} = 0.442 \pm 0.007(\text{stat}) \pm 0.27(\text{sys})$ for all muons that have a stopping track, and $\bar{n}_{FV, data} = 0.452 \pm 0.039(\text{stat}) \pm 0.27(\text{sys})$ and $\bar{n}_{FV, MC} = 0.582 \pm 0.018(\text{stat}) \pm 0.38(\text{sys})$ after a FV cut and neutron detection efficiency correction. The systematic uncertainty comes from the timing cut efficiency from an earlier section. These are summarized in Table 6.1. The high average neutron multiplicity seen in simulations compared to data is currently not understood and suggests that our nuclear models need to be further investigated and updated.

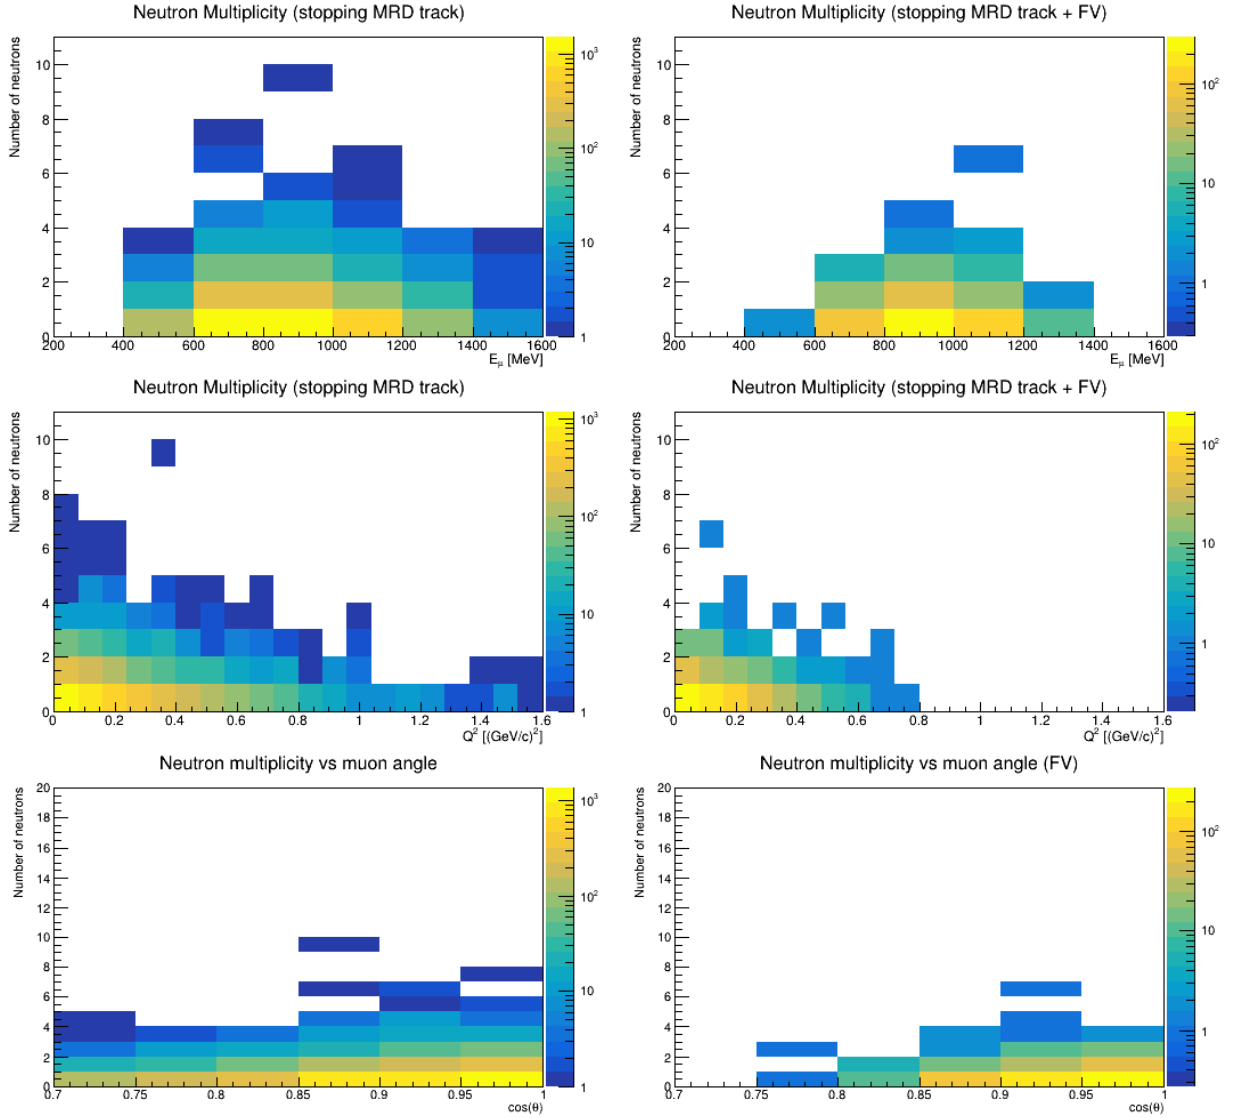


Figure 6.24: The number of neutrons as a function of visible energy of the muon for data with a stopping MRD track (left) and with a fiducial volume cut (right). Most events produce zero to two neutrons. Note that all plots are in logarithmic scale in the z -axis.

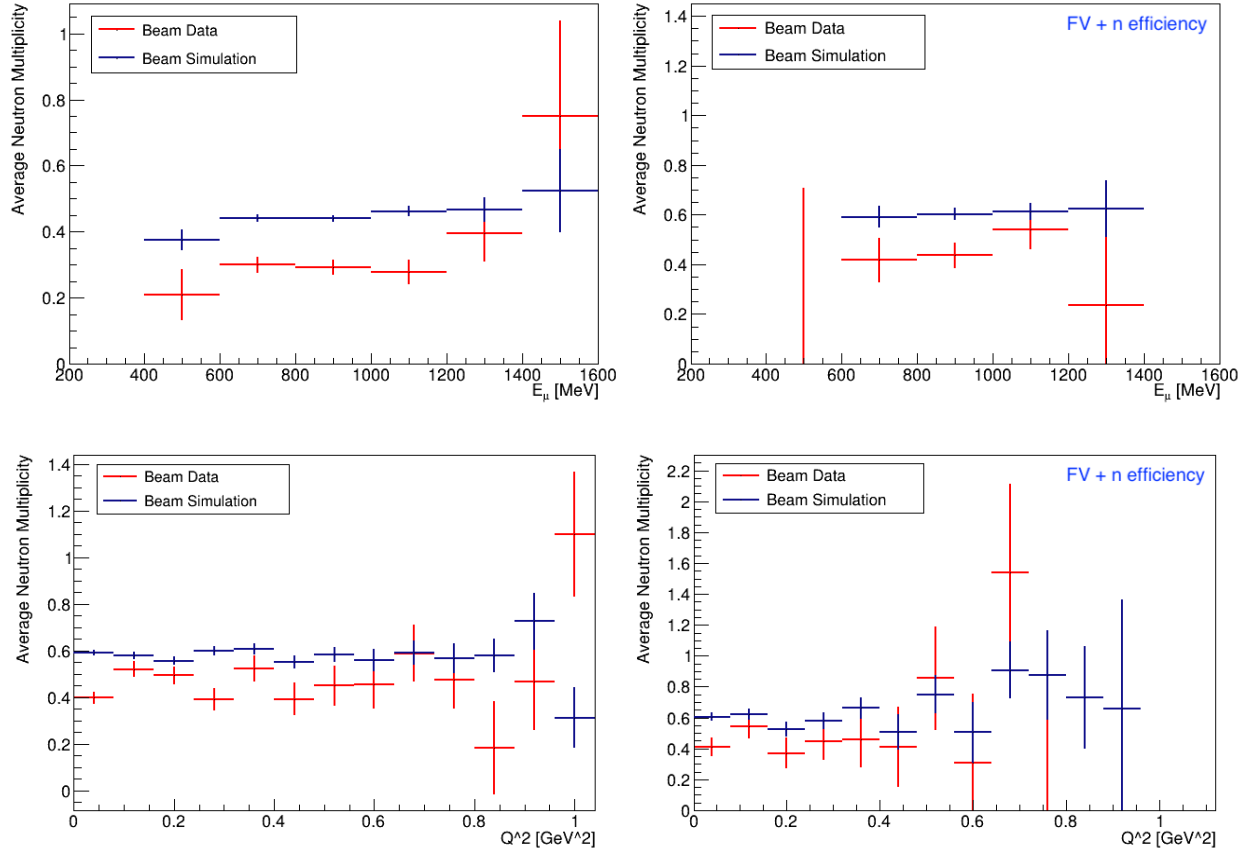


Figure 6.25: The average neutron multiplicity as a function of visible muon energy (top two plots) and momentum transfer (bottom two plots). Compared with results from Super-K and SNO, ANNIE sees a lower neutron yield. The left plots are measurements without any selection cuts applied. The right plots have the FV cut and neutron detection efficiency applied.

Cut	Data	Simulation
Stopping track	$0.294 \pm 0.015(\text{stat}) \pm 0.18(\text{sys})$	$0.442 \pm 0.007(\text{stat}) \pm 0.27(\text{sys})$
+ FV cut	$0.294 \pm 0.039(\text{stat}) \pm 0.18(\text{sys})$	$0.449 \pm 0.018(\text{stat}) \pm 0.27(\text{sys})$
+ n efficiency	$0.452 \pm 0.039(\text{stat}) \pm 0.27(\text{sys})$	$0.582 \pm 0.018(\text{stat}) \pm 0.38(\text{sys})$

Table 6.1: A summary of the average neutron multiplicity with various cuts applied.

6.9 Brief Overview of Systematic Uncertainties

While there is a number of statistical and systematic uncertainties to address for the above analysis, this thesis covers only a few. Studies to cover the remaining sources of error can be found in [77], [112], and [114], or are ongoing.

6.9.1 Neutron Background

This analysis uses the neutron background rates discussed in Section 3.1 and found in [65], since an updated analysis is in progress. From Phase I, the neutron background rate was measured to be 0.02 neutrons/beam-spill/m³ in the fiducial volume. This number is applied as an overall correction to the number of neutrons for each muon vertex in the FV after scaling with the volume.

6.9.2 Neutron Capture Efficiency

With the AmBe analysis discussed in Chapter 5, we can correct for the actual number of neutrons produced. Since the exact location of the neutron capture is not known, the efficiencies in the FV are averaged. Referring to the schematic of the AmBe source deployment positions in Figure 5.14, Ports 1 and 5 are on the same upstream side of the FV. Efficiencies at these locations and with heights between -50 cm and 50 cm (relative to the tank center) are averaged for a value of 65% for data and 76% for MC. These values were used to correct the calculations above. While this is a simple correction, much work is needed to quantify




Figure 6.26: Efficiency maps for MC (left) and data (right), the latter of which was reproduced from Figure 5.14. Simulations of the AmBe source deployment were performed to produce the efficiency map on the left. The efficiencies at Ports 1 and 5 between -50 cm and 50 cm heights were averaged and applied to the neutron yield measurements in this work. The average efficiency is 65% for data and 76% for MC. Future studies can apply efficiencies dependent on location. Figure from [114].

the uncertainty that comes from neutron detection, as it is one of the largest sources of uncertainty.

6.9.3 Neutron Containment

As noted previously, neutrons can be created with tens of MeV of energy [117]. This can be enough energy for the neutron to escape detection in the tank since ANNIE is a small detector. To quantify this uncertainty, we compare the neutron multiplicity measured in the

upstream region of the tank to the downstream region (cf. Figure 6.27). The difference in the neutron production rate can be quantified and serve as our systematic. The average neutron multiplicity for the upstream half cylindrical FV is $\bar{n}_{\text{FV-up, data}} = 0.294 \pm 0.039(\text{stat})$, while the downstream half cylindrical FV is $\bar{n}_{\text{FV-down, data}} = 0.292 \pm 0.026(\text{stat})$. The systematic error is $\sigma_{\text{sys,data}} = 0.002$ and thus negligible for the beam data. In simulations, the average $\bar{n}_{\text{FV-up, MC}} = 0.459 \pm 0.018(\text{stat})$, while the downstream half cylindrical FV is $\bar{n}_{\text{FV-down, MC}} = 0.439 \pm 0.010(\text{stat})$, resulting in a systematic error of $\sigma_{\text{sys,MC}} = 0.02$, which is also negligible for MC data. Combining the datasets (cf. Figure 6.28) from both volumes gives an overall average of $\bar{n}_{\text{FV-comb, data}} = 0.293 \pm 0.022(\text{stat})$ and $\bar{n}_{\text{FV-comb, MC}} = 0.293 \pm 0.009(\text{stat})$. With efficiency corrections, the combined average neutron multiplicities are $\bar{n}_{\text{FV-comb, data}} = 0.450 \pm 0.022(\text{stat})$ and $\bar{n}_{\text{FV-comb, MC}} = 0.583 \pm 0.009(\text{stat})$. However, a thorough investigation of escaping should be performed to get a better idea of the actual fraction of neutrons that make it out of the detector tank. A planned study with the GENIE event generator is in the works.




Figure 6.27: Comparison of upstream (left) and downstream (right) neutron multiplicity as a function of muon energy E_μ . Both plots are efficiency-corrected.




Figure 6.28: Average neutron multiplicity as a function of Q^2 , after combining the upstream and downstream fiducial volumes. All cuts and efficiency corrections have been applied.

Chapter 7

Conclusion

7.1 Summary

From its initial postulation to its discovery to its production in beams, the physics community has come a long way in the study of neutrinos. Many technological advancements were made and surely more to come. Much has been learned about the neutrino, but still, there are questions that remain.

This thesis gave an overview of the current knowledge and understanding of neutrinos, and presented some of the challenges in answering the remaining questions. One challenge pertains to current models of neutrino-nucleus interactions, as they do not paint the full picture of what happens within the nucleus. A better understanding of these nuclear models is necessary if precise measurements are to be made of neutrino mixing angles and the CP violating phase, for example. Among these first steps, a neutron multiplicity measurement

will serve to improve our understanding, as neutrons are indicators of inelasticity. This measurement can constrain nuclear models and improve accuracy in event reconstruction. Although Super-Kamiokande and SNO have published first results, their neutron yields are functions of visible energy and the sign and direction of the neutrino are unknown. The Accelerator Neutrino Neutron Interaction Experiment (ANNIE) is in a unique position to make a neutron multiplicity measurement as a function of momentum transfer, among other physics goals.

A detailed description of the ANNIE detector and its subsystems follows, along with an outline of the experiment's physics goals. ANNIE, a 26-ton water Cherenkov detector equipped with 132 conventional photomultiplier tubes, serves as a testbed for new technology: a gadolinium-doped detection medium, water-based liquid scintillator, and Large Area Picosecond Photodetectors (LAPPDs). The hardware portion of this thesis discusses the work performed in bringing the LAPPD system to a deployment-ready state. I spent a year at Fermilab testing and preparing the LAPPD system for remote deployment in the water tank. I also did much of the work on commissioning the LVHV board and identifying key areas for improvement of the board for future revisions.

Following my lab residence, I developed a method for reconstructing the muon vertex and energy without the fast-timing benefits of LAPPDs. Because ANNIE is a small detector compared to Super-Kamiokande and SNO, vertex reconstruction techniques developed by these collaborations cannot be applied in the same manner in ANNIE. This thesis ends with a presentation of neutron multiplicity as a function of momentum transfer and compares the

results of Super-Kamiokande and SNO. What we found was that the multiplicity measured with a pure ν_μ , low energy beam was less than that for atmospheric neutrino studies.

7.2 Future Directions

While this thesis contains an extensive discussion of LAPPDs, it does not contain a neutron multiplicity measurement with LAPPD data. The analysis tools that will integrate this datastream are in development at the time of this writing. ANNIE may also install more LAPPDs in addition to the nominal five. Additionally, the LVHV board that powers the LAPPD and its auxiliary electronics can be improved in future revisions.

Several sources of uncertainties still need to be accounted for to improve the neutron multiplicity measurement presented in this thesis, especially those that come from neutron detection efficiency. Additional ANNIE goals, such as a joint-analysis on the neutrino cross section on water with the neighboring SBND experiment, will also be achieved. This thesis only covered the CC- 0π event topology, but other topologies such as CC- 1π and NC events can also be studied.

7.2.1 ANNIE Phase III / ANNIE WURST

ANNIE has already demonstrated the feasibility of separating Cherenkov radiation from scintillation light using timing. This study was carried out using the SANDI vessel filled with 5% water-based liquid scintillator (WbLS). This opens the door to improved event

reconstruction capabilities. To build upon this initial demonstration, ANNIE could fill a larger volume of the tank with WbLS and even load it with gadolinium. This would model the KamLAND detector in that a balloon-like structure filled with liquid scintillator would be installed in ANNIE. This will be known as the ANNIE WbLS Underwater Retrievable Scintillator Tank (WURST).

It is an exciting time for ANNIE. I leave it to the next generation of graduate students to perform this important and interesting work.

Appendix A

List of Modifications to LAPPD System Components

The LAPPD is a new photodetection technology, making commercial parts for functions like power and readout nearly non-existent. For deployment in the ANNIE detector, many components that make up the LAPPD system were designed specifically for the experiment. As with most prototypes, the system components did not work perfectly right away. This appendix documents some of the modifications necessary for satisfactory operation.

A.1 Re-design of the LVHV Board

A.1.1 From v1 to v2

As a reminder, the functions of the LVHV board are to

- supply high voltage to the LAPPD;
- supply low voltages to the ACDC boards;
- monitor humidity, temperature, light levels, voltage outputs; and
- power the trigger board and set threshold values.

Testing on a bench top is quite different from testing in an actual system. Early on in testing when the LVHV board was first integrated into the LAPPD system, a bad power supply damaged a chip on-board. This incident prompted the need for protection circuitry for the input voltage and high voltage. Although we can read back the low voltage outputs, there was nothing in place for reading back the high voltage, which is crucial for LAPPD operation. Due to a miscommunication, the LVHV board initially designed to communicate with the trigger board via SPI protocol, needed in-field modifications so that it could communicate via I²C. There were also issues with the LVDS DC vs. AC coupling.

Thus, the second revision of the LVHV board took into account all the issues that arose during testing and implemented the fixes that were made in-field. A list of all the changes can be found in Table [A.1](#).

Issue	Proposed Solution
Linear regulator (step-down converter) chip damaged by bad power supply	Implement UV/OV protection circuit to input voltage
No HV read-back	Implement HV read-back circuit
No HV protection	Install a flyback diode to protect relay
Concerns about voltage drops between photocathode and MCP, across MCPs, and between MCPs	Installed zener diodes to prevent voltage gaps from exceeding set limits
Low voltage outputs (+3.3 V, +2.5 V, +1.8 V) too low by the time it reaches ACDC	Increased outputs to +3.8 V, +3.0 V, +1.8 V
Incorrect communication protocol with trigger card	Switch to I ² C from SPI and pull lines up to +3.3 V instead of +5 V
No dedicated power to trigger card	Added +3.3 V utility
RHT sensor interfering with I ² C comms	Switch to SPI version of RHT sensor
Two RH&T sensors with same I ² C address	Remove second RH&T sensor
Concerns about ground path drifts affecting beam gate, PPS, and clock signals on one LVDS pair	Install AC-coupled LVDS line for these signals

Table A.1: A list of known issues of LVHV board v1 and proposed solutions.

A.1.2 LVHV-v2 Changes

Since no one double checked the second revision carefully, some minor modifications were needed for LVHV board v2. They can be found in Table [A.2](#).

CAN bus lines were swapped (a mislabeling). The UV/OV window was too narrow and pushed us to operate the LVHV board at a higher voltage, too high for our comfort. This needed a change of programming resistors. Programming resistors for the HV read-back were incorrect. After some trial and error, it was ruled that the set of programming resistors divided down the HV by a factor of 500.

Issue	Fix
CAN bus lines swapped	Re-labeled
Incorrect linear regulator output for trigger card (+1.8 V instead of +3.3 V)	Replace with correct linear regulator (P/N: ADP7158ARDZ-3.3)
UV/OV window too narrow	Selected different set of programming resistors (replace resistors: R202 = 1 MOhm, R206 = 263.8 kOhm, R207 = 36.2 kOhm)
Incorrect programming resistors for HV read-back (HV divider for reading back the HV is not outputting a reasonable voltage (e.g. when HV = 2500V, HV divider outputs 0.36V when we expect 2.5V))	Replaced with correct programming resistors
HV readings occasionally dropped out (pattern varied between boards) due to floating SHDN pin (pin 5) on U506 (ADM8828). This caused the -5V output of the chip to go to 0V, and as it went back to -5V, the op-amp output would drop	Grounded the shutdown pin (pin 5) on U506 to ground (pin 4) with solder
Zener diodes (75V, 150V, 75V) preventing us from establishing the voltage gaps between MCPs that Matt wants (200V, etc)	Replace with zeners with higher voltage rating
3V8 line outputs 4.20 V (too high)	Replace programming chip resistor (R401) with a 34-kOhm resistor
U203 is incorrect: originally labeled 3.3V and has the 1.8V part (due to copy-paste in schematics), but should actually be 2.5V	Replace LR (U203) with 2.5V LR
AC-coupled beam gate path is not working as expected	Switch to DC-coupled beam gate

Table A.2: A list of issues with v2 of the LVHV board and their solutions.

A.2 Modifications to the ACDC Boards

The majority of ACDC boards fabricated are designed to operate in test stands that are not compact like the watertight housing. In addition, not many test experiments (see [107]), required a completely dark setup.

An initial fit test revealed that some of the RJ45 jacks were too tall. The LEDs on-board were removed since they introduce light into an environment that we want dark. The interface between the ACDC boards and the LVHV board for power did not match so they were converted into wires that plugged into the LVHV board with some terminal blocks. Some linear regulators also needed to be replaced to follow a new power scheme.

A.3 Modifications to the Watertight Housing

During fit tests of the LAPPD and the electronics inside the housing, one of the nubs prevented the LVHV board from sitting flush with the other nubs. The EMCO HV module was located exactly where a nub was. To resolve this, this nub was cut off. For future housings, this nub should be excluded if the LVHV board design remains the same.

Additionally, a different set of waterproof connectors were selected, so the housing needed modifications since the threading and socket diameter were different. To resolve this issue, David Hemer and I designed circular plugs with the correct diameter and threading and machined the plugs from PVC. The plugs were glued to the housing frame where the existing holes were with PVC cement.




Figure A.1: Modifications to the housing (left) were needed for the new LAPPD waterproof connectors due to different threading and socket diameter as seen in the right image.

A.4 Re-splicing Waterproof Cables

The MacArtney SubConn Circular 12-pin pin and socket connectors initially selected for termination of the LAPPD waterproof cables and installation in the housing did not meet our requirements for signal integrity. All of the wires coming out of the pin connector were straight, so there was a one to two feet section of cable where the LVDS (low voltage differential signal) pairs were not twisted, which is not great for LVDS. We believe this played a significant part in why the ACDC boards were unable to interpret some signals coming from the ACC. Ultimately, the decision to purchase the MacArtney Ethernet SubConn Circular 13-pin connectors (pin and socket) and replace the existing connectors was our saving grace.

Of course, the ten LAPPD cables had already been spliced and terminated at the wet end by the time they were transported to Fermilab. When the first set of the 13-pin connectors arrived, I re-spliced a pair of waterproof cables selected for LAPPD deployment. During

the re-splicing procedure, some of the solder butt connections were tenuous and sometimes came loose, resulting in several redos of the splice. Eventually, we made the switch to crimp connectors and double, triple checked that the splice was sturdy. In addition, we added splints, fashioned from thick, industrial plastic zip ties, to prevent the cable from bending too much where the splice is made. That way, there is a reduced risk of undoing a crimped connection. The remaining eight cables will be re-spliced with this new procedure.




Figure A.2: TOP: Falmat cables (neon green) re-spliced with new connectors (blue).
BOTTOM: Heat shrink tubing is placed over the splice to waterproof it.

Appendix B

Slow Controls Routine

B.1 PIC Programming

The main player in the slow controls is the Microchip DSPIC33EV256GM102 microcontroller (the “PIC”). This PIC microcontroller serves as the interface between users at the surface and the electronics underwater. To power the LAPPD, the ACDC boards, and the trigger board, users must send a command over CAN bus via a Raspberry Pi in the breakout box (BoB) to the PIC. The PIC then sends its own command to the appropriate modules to turn them on. Without the PIC, we also wouldn’t have readings of humidity, temperature, light levels, and voltage settings.

In this repository are the source files for programming the PIC:

<https://github.com/jujuhub/LVHVPIC-files>. In an infinite loop, it checks for CAN messages from the surface-level Raspberry Pi. It matches the CAN message ID to the appropriate

block of code to execute.

The source code for the slow controls can be found here:

<https://github.com/jujuhub/lvhv>

B.2 Diagram of the Slow Controls Routine

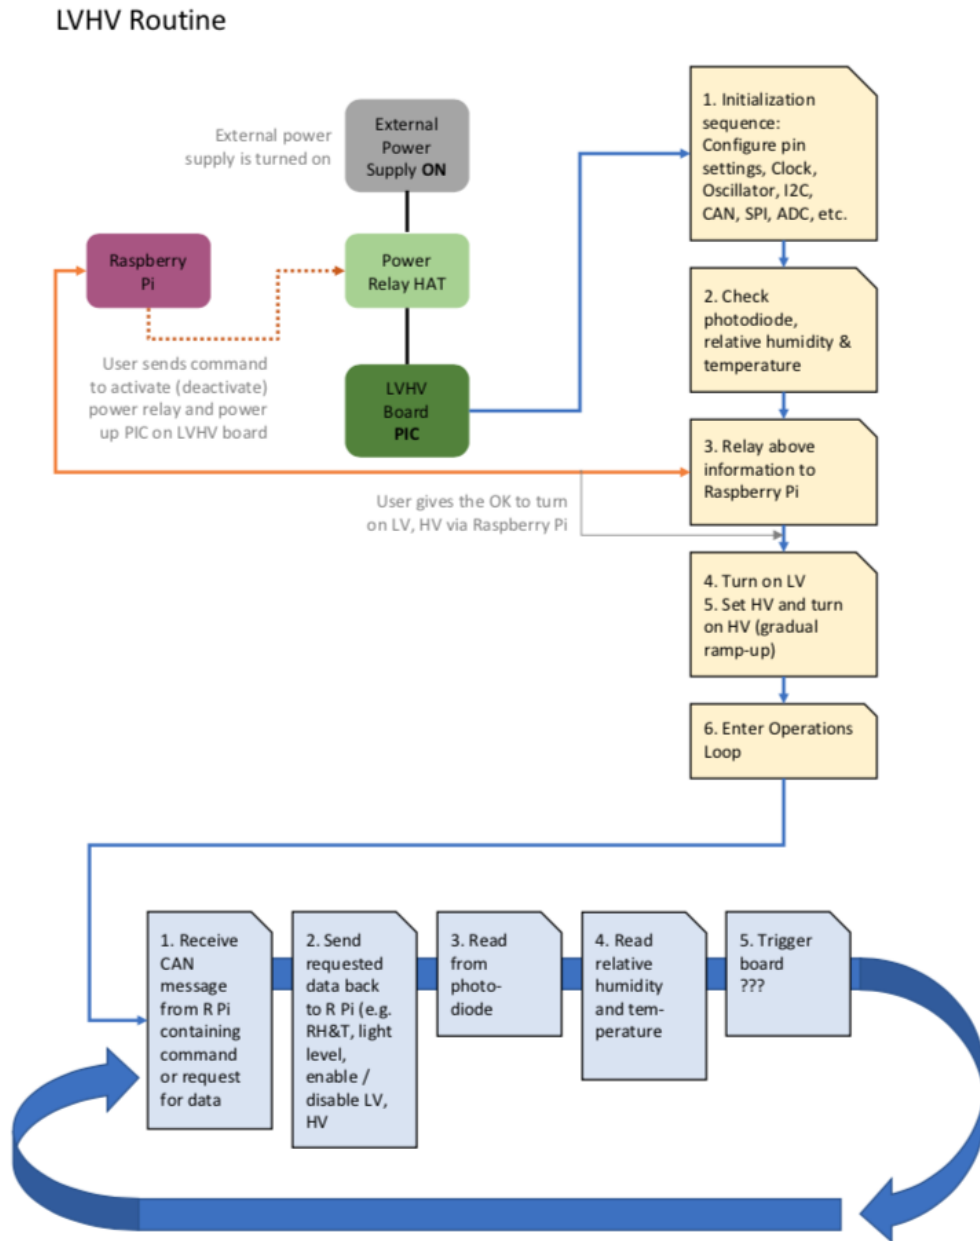



Figure B.1: Slow controls routine. Users communicate with the LVHV board via Raspberry Pis. Checks for relative humidity, temperature, and voltages are performed.

Appendix C

Postdoc Performance in Fitting Tank Track Lengths

In early stages of analysis, I recruited postdocs Paul Hackspacher and Gian Caceres to manually fit tank tracks by eye. Below are their results. Paul seems to be a great BI (biological intelligence, in reference to artificial intelligence, since imaging techniques can be done with AI). Good job to you both!



Appendix D


Graduate Student Julie at Work



Holding a LAPPD



PVC is no match against me!



Shifters watch the ANNIE detector, but who is watching the shifters?

Bibliography

- [1] J. Chadwick. “Intensitätsverteilung im magnetischen Spectrum der β -Strahlen von radium B + C”. In: *Verhandl. Dtsc. Phys. Ges.* 16 (1914), pp. 383–391.
- [2] L. M. Brown. “The idea of the neutrino”. In: *Physics Today* 31.9 (Sept. 1978), pp. 23–28. ISSN: 0031-9228. DOI: [10.1063/1.2995181](https://doi.org/10.1063/1.2995181). eprint: https://pubs.aip.org/physicstoday/article-pdf/31/9/23/8285438/23_1_online.pdf. URL: <https://doi.org/10.1063/1.2995181>.
- [3] A. de Gouvêa. *2004 TASI Lectures on Neutrino Physics*. 2004. arXiv: [hep-ph/0411274](https://arxiv.org/abs/hep-ph/0411274) [hep-ph].
- [4] W. Pauli. *Aufsatz und Vorträge über Physik und Erkenntnistheorie*. 1961. URL: <https://www.math.utah.edu/~beebe/talks/2015/qtm/pdf/pauli-1930-ltc.pdf>.
- [5] E. Fermi. “An attempt of a theory of beta radiation. 1.” In: *Z. Phys.* 88 (1934), pp. 161–177. DOI: [10.1007/BF01351864](https://doi.org/10.1007/BF01351864).
- [6] E. Fermi. “Tentativo di una Teoria Dei Raggi β ”. In: *Nuovo Cim* 11 (1934), pp. 1–19. DOI: [10.1007/BF02959820](https://doi.org/10.1007/BF02959820).
- [7] G. Gamow and E. Teller. “Selection Rules for the β -Disintegration”. In: *Phys. Rev.* 49 (12 June 1936), pp. 895–899. DOI: [10.1103/PhysRev.49.895](https://doi.org/10.1103/PhysRev.49.895).
- [8] C. Giunti and C. W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. 2007. ISBN: 978-0-19-850871-7.
- [9] S. L. Glashow. “The renormalizability of vector meson interactions”. In: *Nuclear Physics* 10 (1959), pp. 107–117. ISSN: 0029-5582. DOI: [https://doi.org/10.1016/0029-5582\(59\)90196-8](https://doi.org/10.1016/0029-5582(59)90196-8).
- [10] S. L. Glashow. “Partial-symmetries of weak interactions”. In: *Nuclear Physics* 22.4 (1961), pp. 579–588. ISSN: 0029-5582. DOI: [https://doi.org/10.1016/0029-5582\(61\)90469-2](https://doi.org/10.1016/0029-5582(61)90469-2).
- [11] A. Salam and J. C. Ward. “Weak and electromagnetic interactions”. In: *Nuovo Cim* 11 (1959), pp. 568–577. DOI: <https://doi.org/10.1007/BF02726525>.

- [12] A. Salam and J. Ward. “Electromagnetic and weak interactions”. In: *Physics Letters* 13.2 (1964), pp. 168–171. ISSN: 0031-9163. DOI: [https://doi.org/10.1016/0031-9163\(64\)90711-5](https://doi.org/10.1016/0031-9163(64)90711-5).
- [13] S. Weinberg. “A Model of Leptons”. In: *Phys. Rev. Lett.* 19 (21 Nov. 1967), pp. 1264–1266. DOI: [10.1103/PhysRevLett.19.1264](https://doi.org/10.1103/PhysRevLett.19.1264).
- [14] C. L. Cowan et al. “Detection of the Free Neutrino: a Confirmation”. In: *Science* 124.3212 (1956), pp. 103–104. DOI: [10.1126/science.124.3212.103](https://doi.org/10.1126/science.124.3212.103).
- [15] G. Danby et al. “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”. In: *Phys. Rev. Lett.* 9 (1 July 1962), pp. 36–44. DOI: [10.1103/PhysRevLett.9.36](https://doi.org/10.1103/PhysRevLett.9.36).
- [16] M. L. Perl et al. “Evidence for Anomalous Lepton Production in e^+e^- Annihilation”. In: *Phys. Rev. Lett.* 35 (22 Dec. 1975), pp. 1489–1492. DOI: [10.1103/PhysRevLett.35.1489](https://doi.org/10.1103/PhysRevLett.35.1489).
- [17] K. Kodama et al. “Observation of tau neutrino interactions”. In: *Physics Letters B* 504.3 (2001), pp. 218–224. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/S0370-2693\(01\)00307-0](https://doi.org/10.1016/S0370-2693(01)00307-0).
- [18] R. L. Workman et al. “Review of Particle Physics”. In: *PTEP* 2022 (2022), p. 083C01. DOI: [10.1093/ptep/ptac097](https://doi.org/10.1093/ptep/ptac097).
- [19] F. Dydak. “Results from LEP and the SLC”. In: *Conf. Proc.* (1991). Ed. by K. K. Phua and Y. Yamaguchi, pp. 3–32.
- [20] B. T. Cleveland et al. “Measurement of the solar electron neutrino flux with the Homestake chlorine detector”. In: *Astrophys. J.* 496 (1998), pp. 505–526. DOI: [10.1086/305343](https://doi.org/10.1086/305343).
- [21] E. Kearns. “Experimental measurements of atmospheric neutrinos”. In: *Nuclear Physics B - Proceedings Supplements* 70.1 (1999). Proceedings of the Fifth International Workshop on topics in Astroparticle and Underground Physics, pp. 315–323. ISSN: 0920-5632. DOI: [https://doi.org/10.1016/S0920-5632\(98\)00441-1](https://doi.org/10.1016/S0920-5632(98)00441-1).
- [22] Y. Fukuda et al. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Phys. Rev. Lett.* 81 (8 Aug. 1998), pp. 1562–1567. DOI: [10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562).
- [23] Q. R. Ahmad et al. “Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by ^8B Solar Neutrinos at the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 87 (7 July 2001), p. 071301. DOI: [10.1103/PhysRevLett.87.071301](https://doi.org/10.1103/PhysRevLett.87.071301).
- [24] T. Araki et al. “Measurement of Neutrino Oscillation with KamLAND: Evidence of Spectral Distortion”. In: *Phys. Rev. Lett.* 94 (8 Mar. 2005), p. 081801. DOI: [10.1103/PhysRevLett.94.081801](https://doi.org/10.1103/PhysRevLett.94.081801).
- [25] A. Y. Smirnov. “Solar neutrinos: Oscillations or No-oscillations?” In: (2017). arXiv: [1609.02386 \[hep-ph\]](https://arxiv.org/abs/1609.02386).

- [26] C. Arpesella et al. “First real time detection of ^7Be solar neutrinos by Borexino”. In: *Physics Letters B* 658.4 (Jan. 2008), pp. 101–108. ISSN: 0370-2693. DOI: [10.1016/j.physletb.2007.09.054](https://doi.org/10.1016/j.physletb.2007.09.054). URL: <http://dx.doi.org/10.1016/j.physletb.2007.09.054>.
- [27] G. Bellini et al. “First Evidence of pep Solar Neutrinos by Direct Detection in Borexino”. In: *Physical Review Letters* 108.5 (Feb. 2012). ISSN: 1079-7114. DOI: [10.1103/physrevlett.108.051302](https://doi.org/10.1103/physrevlett.108.051302). URL: <http://dx.doi.org/10.1103/PhysRevLett.108.051302>.
- [28] K. Sato and H. Suzuki. “Analysis of neutrino burst from the supernova 1987A in the Large Magellanic Cloud”. In: *Phys. Rev. Lett.* 58 (25 June 1987), pp. 2722–2725. DOI: [10.1103/PhysRevLett.58.2722](https://doi.org/10.1103/PhysRevLett.58.2722).
- [29] F. Couchot et al. “Cosmological constraints on the neutrino mass including systematic uncertainties”. In: *Astronomy & Astrophysics* 606 (Oct. 2017), A104. ISSN: 1432-0746. DOI: [10.1051/0004-6361/201730927](https://doi.org/10.1051/0004-6361/201730927). URL: <http://dx.doi.org/10.1051/0004-6361/201730927>.
- [30] V. N. Aseev et al. “Upper limit on the electron antineutrino mass from the Troitsk experiment”. In: *Phys. Rev. D* 84 (11 Dec. 2011), p. 112003. DOI: [10.1103/PhysRevD.84.112003](https://doi.org/10.1103/PhysRevD.84.112003). URL: <https://link.aps.org/doi/10.1103/PhysRevD.84.112003>.
- [31] M. Aker et al. “Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN”. In: *Phys. Rev. Lett.* 123 (22 Nov. 2019), p. 221802. DOI: [10.1103/PhysRevLett.123.221802](https://doi.org/10.1103/PhysRevLett.123.221802).
- [32] M. Aker et al. “Direct neutrino-mass measurement with sub-electronvolt sensitivity”. In: *Nat. Phys.* 18 (2022), pp. 160–166. DOI: [10.1038/s41567-021-01463-1](https://doi.org/10.1038/s41567-021-01463-1).
- [33] H. E. Logan. *TASI 2013 lectures on Higgs physics within and beyond the Standard Model*. 2022. arXiv: [1406.1786 \[hep-ph\]](https://arxiv.org/abs/1406.1786).
- [34] W. Buchmüller et al. “Leptogenesis for pedestrians”. In: *Annals of Physics* 315.2 (2005), pp. 305–351. ISSN: 0003-4916. DOI: <https://doi.org/10.1016/j.aop.2004.02.003>.
- [35] M. Agostini et al. “Final Results of GERDA on the Search for Neutrinoless Double- β Decay”. In: *Physical Review Letters* 125.25 (Dec. 2020). ISSN: 1079-7114. DOI: [10.1103/physrevlett.125.252502](https://doi.org/10.1103/physrevlett.125.252502). URL: <http://dx.doi.org/10.1103/PhysRevLett.125.252502>.
- [36] I. J. Arnquist et al. “Final Result of the Majorana Demonstrator’s Search for Neutrinoless Double- β Decay in ^{76}Ge ”. In: *Phys. Rev. Lett.* 130 (6 Feb. 2023), p. 062501. DOI: [10.1103/PhysRevLett.130.062501](https://doi.org/10.1103/PhysRevLett.130.062501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.062501>.

- [37] D. Q. Adams et al. “New Direct Limit on Neutrinoless Double Beta Decay Half-Life of ^{128}Te with CUORE”. In: *Phys. Rev. Lett.* 129 (22 Nov. 2022), p. 222501. DOI: [10.1103/PhysRevLett.129.222501](https://doi.org/10.1103/PhysRevLett.129.222501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.129.222501>.
- [38] G. Anton et al. “Search for Neutrinoless Double- β Decay with the Complete EXO-200 Dataset”. In: *Phys. Rev. Lett.* 123 (16 Oct. 2019), p. 161802. DOI: [10.1103/PhysRevLett.123.161802](https://doi.org/10.1103/PhysRevLett.123.161802). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.161802>.
- [39] S. Abe et al. “Search for the Majorana Nature of Neutrinos in the Inverted Mass Ordering Region with KamLAND-Zen”. In: *Phys. Rev. Lett.* 130 (5 Jan. 2023), p. 051801. DOI: [10.1103/PhysRevLett.130.051801](https://doi.org/10.1103/PhysRevLett.130.051801). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.051801>.
- [40] F. T. Avignone et al. “Double beta decay, Majorana neutrinos, and neutrino mass”. In: *Reviews of Modern Physics* 80.2 (Apr. 2008), pp. 481–516. ISSN: 1539-0756. DOI: [10.1103/revmodphys.80.481](https://doi.org/10.1103/revmodphys.80.481).
- [41] A. de Gouvêa. “Neutrino Mass Models”. In: *Annual Review of Nuclear and Particle Science* 66 (2016), pp. 197–217. DOI: <https://doi.org/10.1146/annurev-nucl-102115-044600>.
- [42] Z.-Z. Xing and S. Zhou. *Neutrinos in Particle Physics, Astronomy and Cosmology*. Springer, 2011.
- [43] K. Babu. “Model of “calculable” Majorana neutrino masses”. In: *Physics Letters B* 203.1 (1988), pp. 132–136. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(88\)91584-5](https://doi.org/10.1016/0370-2693(88)91584-5).
- [44] E. Ma. “Verifiable radiative seesaw mechanism of neutrino mass and dark matter”. In: *Phys. Rev. D* 73 (7 Apr. 2006), p. 077301. DOI: [10.1103/PhysRevD.73.077301](https://doi.org/10.1103/PhysRevD.73.077301).
- [45] M. Kobayashi and T. Maskawa. “CP-Violation in the Renormalizable Theory of Weak Interaction”. In: *Progress of Theoretical Physics* 49.2 (Feb. 1973), pp. 652–657. ISSN: 0033-068X. DOI: [10.1143/PTP.49.652](https://doi.org/10.1143/PTP.49.652). eprint: <https://academic.oup.com/ptp/article-pdf/49/2/652/5257692/49-2-652.pdf>. URL: <https://doi.org/10.1143/PTP.49.652>.
- [46] A. de Gouvêa. *On Determining the Neutrino Mass Hierarchy*. FNAL Theory Seminar. 2006.
- [47] P. Dunne. *Latest Oscillation Results from T2K*. 2020. URL: <https://zenodo.org/records/4154355>.
- [48] A. Himmel. *New Oscillation Results from the NO ν A Experiment*. 2020. URL: <https://zenodo.org/records/3959581>.

- [49] P. B. Denton et al. “CP-Violating Neutrino Nonstandard Interactions in Long-Baseline-Accelerator Data”. In: *Physical Review Letters* 126.5 (Feb. 2021). ISSN: 1079-7114. DOI: [10.1103/physrevlett.126.051801](https://doi.org/10.1103/physrevlett.126.051801).
- [50] J. A. Formaggio and G. P. Zeller. “From eV to EeV: Neutrino cross sections across energy scales”. In: *Rev. Mod. Phys.* 84 (3 Sept. 2012), pp. 1307–1341. DOI: [10.1103/RevModPhys.84.1307](https://doi.org/10.1103/RevModPhys.84.1307).
- [51] H. Gallagher et al. “Neutrino-Nucleus Interactions”. In: *Annu. Rev. Nucl. Part. Sci.* 61 (2011), pp. 355–378. DOI: [10.1146/annurev-nucl-102010-130255](https://doi.org/10.1146/annurev-nucl-102010-130255).
- [52] C. Llewellyn Smith. “Neutrino reactions at accelerator energies”. In: *Physics Reports* 3.5 (1972), pp. 261–379. ISSN: 0370-1573. DOI: [https://doi.org/10.1016/0370-1573\(72\)90010-5](https://doi.org/10.1016/0370-1573(72)90010-5).
- [53] R. Smith and E. Moniz. “Neutrino reactions on nuclear targets”. In: *Nuclear Physics B* 43 (1972), pp. 605–622. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(72\)90040-5](https://doi.org/10.1016/0550-3213(72)90040-5).
- [54] A. A. Aguilar-Arevalo et al. “First measurement of the muon neutrino charged current quasielastic double differential cross section”. In: *Phys. Rev. D* 81 (9 May 2010), p. 092005. DOI: [10.1103/PhysRevD.81.092005](https://doi.org/10.1103/PhysRevD.81.092005).
- [55] A. A. Aguilar-Arevalo et al. “First measurement of the muon antineutrino double-differential charged-current quasielastic cross section”. In: *Phys. Rev. D* 88 (3 Aug. 2013), p. 032001. DOI: [10.1103/PhysRevD.88.032001](https://doi.org/10.1103/PhysRevD.88.032001).
- [56] O. Lalakulich et al. “Many-Body Interactions of Neutrinos with Nuclei - Observables”. In: (2014). eprint: [1203.2935](https://arxiv.org/abs/1203.2935). URL: <https://arxiv.org/pdf/1203.2935>.
- [57] S. Singh and E. Oset. “Quasielastic neutrino (antineutrino) reactions in nuclei and the axial-vector form factor of the nucleon”. In: *Nuclear Physics A* 542.4 (1992), pp. 587–615. ISSN: 0375-9474. DOI: [https://doi.org/10.1016/0375-9474\(92\)90259-M](https://doi.org/10.1016/0375-9474(92)90259-M).
- [58] I. Esteban et al. “The fate of hints: updated global analysis of three-flavor neutrino oscillations”. In: *Journal of High Energy Physics* 2020.9 (Sept. 2020). ISSN: 1029-8479. DOI: [10.1007/jhep09\(2020\)178](https://doi.org/10.1007/jhep09(2020)178).
- [59] “Precision electroweak measurements on the Z resonance”. In: *Physics Reports* 427.5 (2006), pp. 257–454. ISSN: 0370-1573. DOI: <https://doi.org/10.1016/j.physrep.2005.12.006>.
- [60] C. Athanassopoulos et al. “The liquid scintillator neutrino detector and LAMPF neutrino source”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 388.1–2 (Mar. 1997), pp. 149–172. ISSN: 0168-9002. DOI: [10.1016/S0168-9002\(96\)01155-2](https://doi.org/10.1016/S0168-9002(96)01155-2). URL: [http://dx.doi.org/10.1016/S0168-9002\(96\)01155-2](http://dx.doi.org/10.1016/S0168-9002(96)01155-2).

- [61] A. Aguilar et al. “Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam”. In: *Phys. Rev. D* 64 (11 Nov. 2001), p. 112007. DOI: [10.1103/PhysRevD.64.112007](https://doi.org/10.1103/PhysRevD.64.112007).
- [62] A. D. Sakharov. “Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe”. In: *Pisma Zh. Eksp. Teor. Fiz.* 5 (1967), pp. 32–35. DOI: [10.1070/PU1991v034n05ABEH002497](https://doi.org/10.1070/PU1991v034n05ABEH002497).
- [63] E. Alexeyev et al. “Detection of the neutrino signal from SN 1987A in the LMC using the INR Baksan underground scintillation telescope”. In: *Physics Letters B* 205.2 (1988), pp. 209–214. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(88\)91651-6](https://doi.org/10.1016/0370-2693(88)91651-6).
- [64] K. Hirata et al. “Observation of a neutrino burst from the supernova SN1987A”. In: *Phys. Rev. Lett.* 58 (14 Apr. 1987), pp. 1490–1493. DOI: [10.1103/PhysRevLett.58.1490](https://doi.org/10.1103/PhysRevLett.58.1490).
- [65] A. Back et al. “Measurement of beam-correlated background neutrons from the Fermilab Booster Neutrino Beam in ANNIE Phase-I”. In: *Journal of Instrumentation* 15.03 (2020), p. 3011. ISSN: 1748-0221. DOI: [10.1088/1748-0221/15/03/p03011](https://doi.org/10.1088/1748-0221/15/03/p03011).
- [66] A. R. Back et al. *Accelerator Neutrino Neutron Interaction Experiment (ANNIE): Preliminary Results and Physics Phase Proposal*. 2017. DOI: [10.48550/ARXIV.1707.08222](https://doi.org/10.48550/ARXIV.1707.08222). eprint: [1707.08222](https://arxiv.org/abs/1707.08222). URL: <https://arxiv.org/abs/1707.08222>.
- [67] A. R. Back et al. “Design and Construction of the ANNIE Detector at Fermilab”. In progress.
- [68] *Kentain Products Ltd.* URL: <https://www.kentain.com/>.
- [69] R. B. Firestone and V. S. Shirley. *Table of Isotopes*. John Wiley, New York, 1996.
- [70] J. F. Beacom and M. R. Vagins. “Antineutrino Spectroscopy with Large Water Čerenkov Detectors”. In: *Physical Review Letters* 93.17 (Oct. 2004). ISSN: 1079-7114. DOI: [10.1103/physrevlett.93.171101](https://doi.org/10.1103/physrevlett.93.171101).
- [71] S. Dazeley et al. “Observation of neutrons with a Gadolinium doped water Cherenkov detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 607.3 (2009), pp. 616–619. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2009.03.256>.
- [72] S. Dazeley et al. “A search for cosmogenic production of β -neutron emitting radionuclides in water”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 821 (2016), pp. 151–159. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2016.03.014>.

- [73] P. Fernández. “Status of GADZOOKS!: Neutron Tagging in Super-Kamiokande”. In: *Nuclear and Particle Physics Proceedings* 273-275 (2016). 37th International Conference on High Energy Physics (ICHEP), pp. 353–360. ISSN: 2405-6014. DOI: <https://doi.org/10.1016/j.nuclphysbps.2015.09.050>.
- [74] J. He. *UC Davis Material Compatibility Studies*. Date accessed: 2024-05-16. URL: <http://svoboda.ucdavis.edu/experiments/watchman/materials-testing/>.
- [75] L. Marti et al. “Evaluation of Gadolinium’s Action on Water Cherenkov Detector Systems with EGADS”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 959 (2020), p. 163549. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2020.163549>.
- [76] V. Fischer et al. “Development of an ion exchange resin for gadolinium-loaded water”. In: *Journal of Instrumentation* 15.07 (July 2020), P07004–P07004. ISSN: 1748-0221. DOI: [10.1088/1748-0221/15/07/p07004](https://doi.org/10.1088/1748-0221/15/07/p07004).
- [77] M. O’Flaherty. “The Accelerator Neutrino Neutron Interaction Experiment: Design, Construction and Preparation for First Physics”. PhD thesis. University of Sheffield, 2020.
- [78] A. Meyer. “The CDF Experiment at the Tevatron: The First Two Years of Run II”. In: *Modern Physics Letters A* 18.24 (Aug. 2003), pp. 1643–1660. ISSN: 1793-6632. DOI: [10.1142/s0217732303011599](https://doi.org/10.1142/s0217732303011599).
- [79] Y. Kurimoto. “The SciBooNE Experiment at FNAL”. In: *Nuclear Physics B - Proceedings Supplements* 188 (2009). Proceedings of the Neutrino Oscillation Workshop, pp. 182–184. ISSN: 0920-5632. DOI: <https://doi.org/10.1016/j.nuclphysbps.2009.02.041>.
- [80] Y. Sugiyama et al. “The Data Acquisition System for the KOTO Experiment”. In: *IEEE Trans. Nucl. Sci.* 62.3 (2015), pp. 1115–1121. DOI: [10.1109/RTC.2014.7097532](https://doi.org/10.1109/RTC.2014.7097532). arXiv: [1406.3907 \[physics.ins-det\]](https://arxiv.org/abs/1406.3907).
- [81] B. Richards. *ToolDAQ*. URL: <https://github.com/ToolDAQ/ToolDAQFramework>.
- [82] Richards, Benjamin. “The ToolDAQ DAQ Software Framework & Its Use In The Hyper-K & ANNIE Detectors”. In: *EPJ Web Conf.* 214 (2019), p. 01022. DOI: [10.1051/epjconf/201921401022](https://doi.org/10.1051/epjconf/201921401022). URL: <https://doi.org/10.1051/epjconf/201921401022>.
- [83] A. A. Aguilar-Arevalo et al. “Neutrino flux prediction at MiniBooNE”. In: *Phys. Rev. D* 79 (7 Apr. 2009), p. 072002. DOI: [10.1103/PhysRevD.79.072002](https://doi.org/10.1103/PhysRevD.79.072002).
- [84] J. Crawford et al. *Booster Rookie Book*. 2009. URL: https://operations.fnal.gov/rookie_books/Booster_V4.1.pdf.
- [85] Fermilab Accelerator Division Operations Department. *Concepts Rookie Book*. 2020. URL: https://operations.fnal.gov/rookie_books/concepts.pdf.

- [86] H. Zhang. “Neutron tagging and its application to physics in Super-Kamiokande IV”. In: *32nd International Cosmic Ray Conference* (2011). DOI: [10.7529/ICRC2011/V04/0353](https://doi.org/10.7529/ICRC2011/V04/0353).
- [87] B. Aharmim et al. “Measurement of neutron production in atmospheric neutrino interactions at the Sudbury Neutrino Observatory”. In: *Phys. Rev. D* 99 (11 June 2019), p. 112007. DOI: [10.1103/PhysRevD.99.112007](https://doi.org/10.1103/PhysRevD.99.112007).
- [88] A. Collaboration et al. *Deployment of Water-based Liquid Scintillator in the Accelerator Neutrino Neutron Interaction Experiment*. 2024. arXiv: [2312.09335 \[hep-ex\]](https://arxiv.org/abs/2312.09335).
- [89] E. Oberla and H. J. Frisch. “The design and performance of a prototype water Cherenkov optical time-projection chamber”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 814 (Apr. 2016), pp. 19–32. DOI: [10.1016/j.nima.2016.01.030](https://doi.org/10.1016/j.nima.2016.01.030).
- [90] J. Dalmasson et al. “Distributed imaging for liquid scintillation detectors”. In: *Physical Review D* 97.5 (Mar. 2018). DOI: [10.1103/physrevd.97.052006](https://doi.org/10.1103/physrevd.97.052006).
- [91] T. Kaptanoglu et al. “Cherenkov and scintillation light separation using wavelength in LAB based liquid scintillator”. In: *Journal of Instrumentation* 14.05 (May 2019), T05001–T05001. DOI: [10.1088/1748-0221/14/05/t05001](https://doi.org/10.1088/1748-0221/14/05/t05001).
- [92] H. Iams and B. Salzberg. “The Secondary Emission Phototube”. In: *Proceedings of the Institute of Radio Engineers* 23.1 (Jan. 1935), pp. 55–64. DOI: [10.1109/JRPROC.1935.227243](https://doi.org/10.1109/JRPROC.1935.227243).
- [93] Hamamatsu. URL: <https://www.hamamatsu.com>.
- [94] Wikipedia. *Photomultiplier Tube and Scintillator*. Date accessed: August 2022. URL: https://en.wikipedia.org/wiki/Photomultiplier_tube.
- [95] Photonis. *Photomultiplier tube basics*. URL: <https://www.photonis.com>.
- [96] T. Gys and C. Joram. “Position-sensitive vacuum photon detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 970 (2020), p. 163373. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2019.163373>.
- [97] Wikipedia. *Microchannel plate detector*. Date accessed: August 2022. URL: https://en.wikipedia.org/wiki/Microchannel_plate_detector.
- [98] J. Ladislav Wiza. “Microchannel plate detectors”. In: *Nuclear Instruments and Methods* 162.1 (1979), pp. 587–601. ISSN: 0029-554X. DOI: [https://doi.org/10.1016/0029-554X\(79\)90734-1](https://doi.org/10.1016/0029-554X(79)90734-1).
- [99] LAPPD Collaboration. URL: <https://psec.uchicago.edu/>.
- [100] M. J. Minot et al. “Large Area Picosecond Photodetector (LAPPDTM)-Pilot Production and Development Status”. In: *NIMA Proceedings D* 18.00204 (Nov. 28, 2018).

- [101] B. W. Adams et al. “Timing characteristics of Large Area Picosecond Photodetectors”. In: *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 795 (June 2015), pp. 1–11. ISSN: 01689002. DOI: [10.1016/j.nima.2015.05.027](https://doi.org/10.1016/j.nima.2015.05.027).
- [102] T. Kaptanoglu et al. “Cherenkov and scintillation separation in water-based liquid scintillator using an LAPPD™”. In: *The European Physical Journal C* 82.169 (2022). DOI: [10.1140/epjc/s10052-022-10087-5](https://doi.org/10.1140/epjc/s10052-022-10087-5).
- [103] E. Oberla et al. “A 15 GSa/s, 1.5 GHz Bandwidth Waveform Digitizing ASIC”. In: (Sept. 2013). DOI: [10.1016/j.nima.2013.09.042](https://doi.org/10.1016/j.nima.2013.09.042). URL: <http://arxiv.org/abs/1309.4397>.
- [104] E. Oberla. “Charged particle tracking in a water Cherenkov optical time-projection chamber”. PhD thesis, University of Chicago. PhD thesis. 2015.
- [105] M. Bogdan et al. “A Modular Data Acquisition System using the 10 GSa/s PSEC4 Waveform Recording Chip”. In: (July 2016). DOI: [10.1109/RTC.2016.7543167](https://doi.org/10.1109/RTC.2016.7543167). URL: <http://arxiv.org/abs/1607.02395>.
- [106] B. Adams. “Power-Supply Options for the psec-4 Electronics”. In: (2021). ANNIE internal report (unpublished).
- [107] E. Angelico. “Development of Large-Area MCP-PMT photo-detectors for a Precision Time-of-Flight System at the Fermilab Test Beam Facility”. PhD thesis. 2020. DOI: [10.2172/1637600](https://doi.org/10.2172/1637600).
- [108] B. Adams. “LAPPD Trigger Board”. In: (2020). ANNIE internal report (unpublished).
- [109] J. He. *LVHV*. <https://github.com/jujuhub/lvhv.git>. GitHub. 2021.
- [110] *Incom, Inc.* URL: <https://www.incomusa.com>.
- [111] P.-A. Amaudruz et al. “In-situ characterization of the Hamamatsu R5912-HQE photomultiplier tubes used in the DEAP-3600 experiment”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 922 (2019), pp. 373–384. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2018.12.058>.
- [112] T. Pershing. “The Accelerator Neutrino-Neutron Interaction Experiment”. PhD thesis. University of California, Davis, 2020.
- [113] G. Caceres. *Gains status*. ANNIE Document Database, doc-id #5269.
- [114] M. Nieslony. “Towards a neutron multiplicity measurement with the Accelerator Neutrino Neutron Interaction Experiment”. PhD thesis. Johannes Gutenberg-Universität in Mainz, 2022.
- [115] S. Doran. *Laser Timing Analysis and Calibration*. ANNIE Document Database, doc-id #5404.

- [116] M. collaboration et al. *First demonstration of $\mathcal{O}(1\text{ ns})$ timing resolution in the MicroBooNE liquid argon time projection chamber*. 2023. arXiv: [2304.02076 \[hep-ex\]](https://arxiv.org/abs/2304.02076).
- [117] I. Murata et al. “Neutron and gamma-ray source-term characterization of AmBe sources in Osaka University”. In: *Progress in Nuclear Science and Technology* 4 (2014), pp. 345–348. DOI: [10.15669/pnst.4.345](https://doi.org/10.15669/pnst.4.345).
- [118] M. Nieslony. *Event Building Overview*. ANNIE Document Database, doc-id #4865.
- [119] D. Akimov et al. “The ZEPLIN-III dark matter detector: Instrument design, manufacture and commissioning”. In: *Astroparticle Physics* 27.1 (2007), pp. 46–60. ISSN: 0927-6505. DOI: <https://doi.org/10.1016/j.astropartphys.2006.09.005>.
- [120] F. Neves et al. “ZE3RA: the ZEPLIN-III Reduction and Analysis package”. In: *Journal of Instrumentation* 6.11 (Nov. 2011), P11004. DOI: [10.1088/1748-0221/6/11/P11004](https://doi.org/10.1088/1748-0221/6/11/P11004). URL: <https://dx.doi.org/10.1088/1748-0221/6/11/P11004>.
- [121] S. Dazeley et al. “The Watchboy Radionuclide Detector Deployment and Analysis”. In: (2014). eprint: [LLNL-TR-661658](https://arxiv.org/abs/1405.1658).
- [122] H. Maesaka. “Evidence For Muon Neutrino Oscillation In An Accelerator-based Experiment”. PhD thesis. Kyoto University, 2005.
- [123] T. Kikawa. “Measurement of Neutrino Interactions and Three Flavor Neutrino Oscillations in the T2K Experiment”. PhD thesis. Kyoto University, 2014.
- [124] D. E. Groom et al. “Muon Stopping Power and Range Tables 10 MeV–100 TeV”. In: *Atomic Data and Nuclear Data Tables* 78.2 (2001), pp. 183–356. ISSN: 0092-640X. DOI: <https://doi.org/10.1006/adnd.2001.0861>.
- [125] *Frustum*. Date accessed: May 2024. URL: <https://www.mathnasium.com/math-centers/hinsdale/news/frustum>.
- [126] T. Zhang. *Muon Vertex*. https://github.com/bond400apm/Muon_vertex. GitHub. 2023.
- [127] P. Sigmund. *Particle Penetration and Radiation Effects*. Vol. 151. Springer Series in Solid State Sciences. Springer, 2006. ISBN: 3-540-31713-9.