

DISSERTATION

COSMOGENIC BACKGROUND REJECTION FOR THE STERILE NEUTRINO
SEARCH WITH THE SHORT-BASELINE NEUTRINO PROGRAM FAR
DETECTOR

Submitted by
Christopher M. Hilgenberg
Department of Physics

In partial fulfillment of the requirements
For the Degree of Doctor of Philosophy
Colorado State University
Fort Collins, CO
Fall 2020

Doctoral Committee:

Advisor: Robert J. Wilson

John Harton
Walter Toki
Wen Zhu

Copyright by Christopher M. Hilgenberg 2020
All Rights Reserved

ABSTRACT

COSMOGENIC BACKGROUND REJECTION FOR THE STERILE NEUTRINO SEARCH WITH THE SHORT-BASELINE NEUTRINO PROGRAM FAR DETECTOR

As the Short-Baseline Neutrino Program Far Detector, the ICARUS liquid argon time-projection chamber will operate at shallow depth and therefore be exposed to the full surface flux of cosmic rays. This poses a problematic background to the sterile neutrino search. A direct way to reject this background is to surround the cryostat with a detector capable of tagging incident cosmic muons with high efficiency, the Cosmic Ray Tagging System (CRT). I have worked to develop the CRT detector hardware and the simulation and reconstruction software. This system, currently undergoing installation, is approximately 25% commissioned and taking data. I have developed a powerful method for cosmogenic background rejection utilizing the CRT and Photon Detection System. Results from a simulation based study, informed by CRT data, demonstrate the power of the technique in rejecting cosmogenic events with little adverse impact on the neutrino sample.

ACKNOWLEDGEMENTS

I would not have reached this point without tremendous support from funding agencies, colleagues, family, and friends.

I gratefully acknowledge all of the funding that I received throughout my PhD program. Colorado State University provided stipend support my Masters Program, supporting me when I was just starting out in the CSU HEP laboratory. The most significant financial support was provided by multiple grants through the Department of Energy, Office of Science. Finally, I received significant stipend support over two years during my research at Fermilab through the Universities Research Association's Visiting Scholars Program.

Throughout my Masters and PhD programs, Professor Robert J. Wilson has been my academic and research adviser. Outside of these formal capacities, he has also been a mentor, helping me to navigate often complex situations and practical aspects that I encountered throughout my research. I want to thank him for his time, patience, and input throughout the writing process.

I want to thank the members of my Masters/PhD committee: Professor John Harton, Professor Walter Toki, and Professor Wen Zhu. I very much appreciate their time and input that made this dissertation the best it could be.

Performing research at Fermilab has been one of the most rewarding and challenging things I have done. I have met and learned from so many kind and talented researchers there.

Anne Schukraft, a research scientist working in Fermilab's Neutrino Division and former co-convener of the ICARUS Cosmic Ray Tagger technical working group, was an excellent mentor to me when I first moved to Fermilab. She taught me how to navigate the lab and find what I needed to get the job done. I learned a lot from her including how to set up a test stand and how to carry out large scale measurement campaigns.

Simone Marocci was a postdoctoral research associate who worked at Fermilab. He

started there shortly after I moved to Fermilab. We spent many hours in the lab together understanding complicated electronics and optimizing the Cosmic Ray Tagging System. He always had a positive attitude and practical wisdom to share that made the working day both fun and productive. Outside of work he was also a good friend. We shared many coffee breaks, lunches, and even the occasional barbeque at the Fermilab Village. It pains me to write that he passed away last Fall. He will be sorely missed.

Tyler Boone is fellow PhD student who has worked with me since 2016. He has helped me significantly in carrying out research, including taking and analyzing data, setting up experiments, and carrying my backpack at the Neutrino Summer School after I had fractured my clavicle.

Biswaranjan Behera, a postdoctoral researcher from CSU based at Fermilab, has been a very helpful in getting lab work and analysis done. He has been an integral member of the CSU ICARUS group and is always happy to jump in and assist with whatever was needed.

Minerba Betancourt is a co-convener of the ICARUS Cosmic Ray Tagger technical working group. She has provided leadership and support to get measurements done at the Side CRT test stand and has made the Side CRT installation happen.

Antoni Aduszkiewicz is a postdoctoral research associate at the University of Houston. He has played a significant role in getting the DAQ system working for the Cosmic Ray Tagger, without which there would be no data. He has also been a kind friend, lending an ear when I need to talk through research problems. He also helped me find the best Polish food in the Chicagoland area.

I want to thank David Warner, our HEP engineer at CSU, and Jay Jablonski, our HEP technician at CSU, who both provided crucial technical support that made the hardware I studied and worked with possible.

I would also like to thank all of my SBN collaborators. Experimental particle physics is a monumental undertaking, and these complex experiments are only possible with a lot

of hard work from dedicated and talented people. I could not have completed my work without the tools developed by my colleagues.

Finally, I would be remiss if I did not mention all of the support and patience provided by my family. My wonderful fiancée and soon-to-be wife, Danielle, has been there for me since my first year of graduate school. I appreciate all of her patience through the years, from weathering my complain..calculating through multi-page Lagrangians or working long nights to complete analyses. I want to thank her for leaving beautiful Colorado for the flat lands of Illinois. I also want to thank my loyal dog, Gozer. He has not always gotten the play time he deserves.

TABLE OF CONTENTS

ABSTRACT	ii
ACKNOWLEDGEMENTS	iii
LIST OF TABLES	ix
LIST OF FIGURES	xi
1 Introduction	1
1.1 The Standard Model	2
1.2 Neutrino History	8
1.3 Neutrino Oscillation Phenomenology	12
1.4 Landscape of Neutrino Experiments	17
1.5 Short-Baseline Anomalies	21
2 The Short-Baseline Neutrino Program	28
2.1 Liquid Argon Time-Projection Chambers	28
2.2 ICARUS: The World's First Large LAr TPC	32
2.3 Overview of the Short-Baseline Neutrino Program	35
2.4 Sterile Neutrino Search	43
2.5 Cosmogenic Background	51
3 The Cosmic Ray Tagging System	59
3.1 Conceptual Design	60
3.2 Early R&D at CSU	62
3.2.1 Scintillator	64
3.2.2 Silicon Photomultipliers (SiPMs)	66
3.2.3 Front-End Electronics	68

3.2.4	Lab Setup	71
3.2.5	Calibration	74
3.2.6	Prototype Development	76
3.3	Change of Plans	83
3.4	A New Optical Readout for the MINOS Scintillator Modules	89
3.5	MINOS Module Testing, Characterization, and Readout Development	95
3.5.1	Optimizing the Readout Design	100
3.6	Optical Readout Production and Testing	105
3.7	Installation	116
4	Simulation and Reconstruction	125
4.1	Overview of LArSoft and icaruscode	125
4.1.1	Geometry	126
4.1.2	Generators	127
4.1.3	Detector Simulation	128
4.1.4	TPC Readout Simulation and Reconstruction	132
4.1.5	PDS Readout Simulation and Reconstruction	133
4.1.6	CRT Readout Simulation and Reconstruction	135
4.2	CRT Detector Simulation	136
4.2.1	Detector Response Model Development	137
4.2.2	Validation	146
4.3	CRT Reconstruction	149
4.3.1	CRT Hits and Tracks	149
4.3.2	High-Level Reconstruction with the CRT	153
5	First Data from the Cosmic Ray Tagging System	155

5.1	Data Acquisition System and Data Taking	155
5.2	Calibration	157
5.2.1	Pedestal	158
5.2.2	Gain	159
5.2.3	Stability	160
5.3	Low-Level Analysis	161
5.3.1	Charge Sharing	162
5.3.2	Charge Spectra	164
5.3.3	Trigger Rates	168
5.4	Data-Monte Carlo Comparison	173
6	Cosmogenic Background Rejection	179
6.1	Cosmogenic Background Rates and Efficiencies	179
6.2	Auto-Veto	184
6.3	Time-of-Flight Veto	190
6.4	Impact on Sterile Neutrino Search	197
7	Conclusion and Outlook	199
7.1	Summary	199
7.2	Next Steps	201
	References	203

LIST OF TABLES

1.1	3ν mixing parameters from global analysis, shown for different assumptions of normal (NH) or inverted (IH) hierarchy. 1σ errors shown. Reproduced from [7].	15
1.2	(3ν) Mixing parameters probed by different neutrino sources at different baselines (MBL=medium-baseline, LBL=long-baseline). Reproduced from [7]. . . .	19
2.1	Selected Fermilab neutrino beam parameters.	38
2.2	SBN detector values of interest.	40
2.3	Cosmogenic photon rates in the SBN Far Detector (ICARUS) without (all events) and with dE/dx and distance from the muon cuts (events w/TPC cuts) with 6.6×10^{20} POT. Reproduced from [34].	56
3.1	Selected SiPM specifications reproduced from data sheets at an ambient temperature of 25 °C and recommended bias voltage [58], [59].	68
3.2	Summary of measured light yields for all CSU CRT strip prototypes. All strips are 1-cm thick except for the Mu2e strips, which are 2-cm thick. All strips are 30-cm long and all measurements were made at 50 cm from the SiPM (except where noted).	78
3.3	Scintillator parameters are summarized for the different ICARUS CRT subsystems. Note that due to scheduling constraints, the Top modules used different scintillator for each layer.	86
3.4	Optical fiber parameters for the different ICARUS CRT subsystems are summarized. Note that all subsystems use Kuraray Y11 WLS fibers and Hamamatsu photodetectors.	86
4.1	Parameter values used in the CRT detector simulation are summarized below (see Equation 4.1, 4.2, 4.4)	140

4.2	The spatial and temporal resolution of the CRTHit reconstruction are summarized using local coordinates in the CRT region frame with x as the lateral and y as the longitudinal coordinates.	153
6.1	For each true cosmic ray muon crossing different volumes, the true rate and efficiency are summarized, in the presence of 3-m concrete overburden, for different topologies.	180
6.2	The true fraction [%] of cosmogenic muons crossing the CRT including the first hit only, separated by CRT subsystem, in the presence of 3-m concrete overburden, is given for different topologies.	183
6.3	The true rates [kHz] of cosmogenic muons crossing the CRT including the first hit only, separated by CRT subsystem, in the presence of 3-m concrete overburden, is given for different topologies.	183
6.4	The true rate [kHz] of cosmogenic muons crossing the CRT including the first hit only, separated by CRT subsystem and region, in the presence of 3-m concrete overburden, is given for different topologies.	184
6.5	The true fraction [%] of cosmogenic muons crossing the CRT including the first hit only, separated by CRT subsystem and region, in the presence of 3-m concrete overburden, is given for different topologies.	184
6.6	For BNB neutrinos, auto-veto fractions [%] are summarized for different samples integrated over the full range in energy. The fractions are broken down by the volume containing the true neutrino vertex and whether the interaction is charged current or neutral current.	185

LIST OF FIGURES

1.1	Table of Standard Model particles. (Image credit: https://en.wikipedia.org) .	4
1.2	Feynman diagrams for neutrino-nucleon charged-current scattering (a) and neutrino-nucleon or neutrino-lepton neutral-current scattering (b) where $\ell = e, \mu, \tau$, and $f = n, p$ or e	5
1.3	Charged current inclusive muon neutrino scattering cross-section per nucleon measurements [1].	6
1.4	Masses of all of the Standard Model fermions showing the departure of the neutrino mass scale from the 0.5 MeV - 200 GeV range populated by the other fermions [3].	7
1.5	Two orderings of the neutrino mass eigenvalues are possible, the normal or inverted hierarchies. The mass-squared differences shown are not to scale. The colors indicate the “flavor content” of each mass eigenstate. (Image credit: http://ignatz.phys.rpi.edu)	16
1.6	Ratio of observed to expected ν_e events from four measurements using MCi ^{51}Cr and ^{37}Ar calibration sources in the SAGE and GALLEX experiments. The 1σ error band is shown [26][27][28][29].	22
1.7	Ratio of observed to expected $\bar{\nu}_e$ events by reactor experiments at different baselines. The 1σ error band is shown [30].	23
1.8	LSND(a) and MiniBooNE (b) observe an excess of electron neutrinos at short-baselines.	24
1.9	Exclusion limits and allowed regions are shown for the short-baseline anomalies [33].	25
1.10	Exclusion limits from neutrino disappearance data and allowed regions from $\nu_e/\bar{\nu}_e$ appearance data in $\theta_{\mu e}, \Delta m_{41}^2$ space [33].	27

2.1 An illustration of the basic working principles of a liquid argon time-projection chamber (LAr TPC). On the left, a neutrino traveling into the page interacted inside the volume of liquid argon. The charged particles produced in the interaction locally ionize the argon, generating free electrons and scintillation light. The light propagates to an array of photomultiplier tubes with the flash providing the time of the event. Shown on the right, the ionization electrons drift, over millisecond time scales, left to right in a uniform electric field to a set of three anode wire planes. There, the charge is sampled over time, generating different views of the same event. 29

2.2 An overview of the ICARUS T600 LAr TPC shows the important functional features [42]. 33

2.3 Inside one of T300’s TPC volumes the field cage, central cathode, PMT array, and anode wires are visible [42]. 34

2.4 An aerial view of the Short-Baseline Neutrino Program hosted at Fermilab. The three detectors are shown along with the BNB beam axis. In addition, the intermediate and far detectors will be exposed to an off-axis component of the NuMI beam. 37

2.5 Illustration of neutrino beam production for Fermilab’s Booster Neutrino Beam. Note that the subdominant kaon production and resulting neutrino beam contaminants are not shown. 38

2.6 SBN Booster Neutrino Beam flux predictions at the three detector Top row, left to right: SBND, MicroBooNE, ICARUS. Bottom row: ICARUS-to-SBND flux ratio (left) and MicroBooNE-to-ICARUS flux ratio (right). Reproduced from [34]. 39

2.7 Inside one of the ICARUS T300’s while being refurbished at CERN. 40

2.8 ICARUS cold vessels being installed into the warm vessel at Fermilab. 41

2.9 Installing feedthrough flanges on the T600. 42

2.10 ICARUS just before the start of cold commissioning. 43

2.11 ν CC candidates from data are shown. The color indicates the ionization density with blue to red being low to high. In both panels, the beam is going from left to right. 44

2.12 Two representations of the SBN flux covariance matrix are shown with the fractional error matrix on the left and the correlation matrix on the right. There are three rows and columns, one for each SBN detector. Within each row and column, there are sub rows or columns for the individual ν_μ and ν_e components. The 11 bins span 0.2 to 3 GeV in neutrino energy. Reproduced from [34]. 47

2.13 Top: absolute ν_e (left) and ν_μ (right) event rate uncertainties for the three SBN detectors due to cross-section uncertainties alone. Note that SBND was previously known as LAr1-ND. Bottom: cross-section covariance matrix for ν_e interactions spanning 11 bins over 0.2-3 GeV shown in two representations: fractional error matrix (left) and correlation matrix (right). Reproduced from [34]. 48

2.14 SBN sterile neutrino (3+1 model) search sensitivity in appearance (left) and disappearance (right) modes. Sensitivities shown here are reproduced from [53]. 51

2.15 Cosmogenic energy spectrum just above the SBN experimental buildings for different particle species [34]. 52

2.16 Event from the ICARUS technical surface run at Pavia without overburden. The greyscale indicates the level of ionization density. Several cosmogenic muon tracks are visible. Also shown are few-GeV CNGS ν_e CC events for scale. 53

2.17 A comparison of energies of cosmogenic gammas that convert in the active LAr volume (left) and the electrons produced in BNB ν_e CC interactions show that most gammas are below the analysis energy threshold. However, there are a significant number of cases where the gamma energy can be problematic. . . . 53

2.18	One way to cut out secondary photons generated by muons passing through the active LAr volume is to exclude the volume contained in a cylinder centered on the muon track. The plot on the right is a histogram of photon conversion distance from the parent muon track. For cylinders with a 15-cm radius, 99% of secondary photons are excluded with a loss of about 1% in fiducial volume. Reproduced from [34].	54
2.19	Impact of muon tagging on the SBN sterile neutrino search in the ν_e appearance channel shown as the projected SBN significance along the LSND 99% C.L. contour. Muon tagging enhances the sensitivity by about 1σ significance below sterile mass splittings of about 1 eV^2 . Reproduced from [34].	57
3.1	An illustration of the CRT conceptual design from MINOS [61] shows the basic elements described here.	60
3.2	The radiogenic background (energy spectrum measured at CSU by J. Mann, shown on the left), can generate CRT triggers when using just a single layer. The trigger rate versus threshold was measured showing the transition to begin dominated by cosmogenic muons at about six photoelectrons (credit: R.J. Wilson).	62
3.3	Our first setup at CSU was simple, but it was an effective first exercise in preparing for a full R&D campaign.	63
3.4	A sample of scintillator from the Fermilab Extrusion Facility sits in the CSU HEP Lab, ready for R&D.	65
3.5	A simplified circuit diagram of an SiPM with five microcells (left). An illustration of photons impinging on an SiPM (right) [59].	66
3.6	The CAEN A1702/DT5550 32-channel front-end board [56] is used for all SBN CRTs.	69
3.7	The FEB provides full analog readout and digitization (a) with configurable trigger logic (b) as shown in the block diagrams [56].	70

3.8	Setup at CSU using a CAEN VME crate for trigger logic and digitization.	71
3.9	Setup at CSU using a FEB for light yield measurements and a second one for operating a new hodoscope.	73
3.10	The steps from the calibration procedure are illustrated showing how the pedestal, gain, and photoelectron threshold equivalence were determined.	75
3.11	A MINOS-style strip is shown stacked on top of our SBND-like strip.	76
3.12	The 17-cm wide P2H prototype is pictured in the darkbox between tests.	78
3.13	Our final prototype, pictured here ready for light yield measurements, using a coil of fiber to simulate our target 5.4-m length.	81
3.14	SiPMs connected in parallel on a single FEB channel provide the analog sum of two optical fibers as shown here for the P4SG design.	82
3.15	A diagram shows the Double Chooz veto module components on the left, and one of the modules poses for a photo on the right.	83
3.16	The composition of a MINOS scintillator module is shown on top with a de- tailed view of the scintillator strip design on bottom [61].	87
3.17	A concept drawing of the Top CRT module (left). A photo of a single scintil- lator strip equipped with two SiPMs (right). Image credit: U. Kose.	88
3.18	On the right, The most up-to-date CRT layout is shown, including the cold vessels (magenta) and the argon active volumes (cyan). The Top and Side CRT subsystems are visible. On the right, a top-down view of the area below the Warm Vessel shows the Bottom CRT subsystem. Image credit: J. Tillman, C. James, and A. Schukraft.	88
3.19	My first concept drawing (right) for the SiPM board for the new MINOS optical readout used 1-mm SiPMs to read out each scintillator strip and a spacer to set the SiPM-fiber distance where the board would be mounted directly to the MINOS snout (left).	89

3.20	Our first measurement of light output from a MINOS module with SiPMs was crude, using a makeshift stand to position the SiPMs against the optical interface (a), but it was successful and demonstrated that our new test stand was ready (b).	90
3.21	Our first self-triggered run (a) was dominated by radiogenic photons with a rate consistent with measurements as CSU. Our first hodoscope-trigger run (b) showed a peak at roughly 27 PE with the hodoscope less than 1 m from the readout.	91
3.22	Our first SiPM board with six 1-mm SiPMs (d) was tested at CSU with a pulsed LED (a) and with light generated by muons tagged with CRT prototypes run (b),(c),(d).	92
3.23	A stack of MINOS modules awaits testing at Wideband.	94
3.24	The first light yield measurement with our partially populated prototype from a full-sized MINOS module using the CDF hodoscope at 30 cm from the readout end yielded 12 PE, a 30% improvement over the MINOS result.	94
3.25	The summer interns were hard at work testing one of the MINOS modules. . .	97
3.26	The light yield is shown for each scintillator strip at the three different hodoscope positions with the runs with known SiPM mounting problems removed: 1.5 m (a), 4.0 m (b), and 6.5 m (c). Distributions for both readout ends are included separately with the north end in blue and the south end in red.	99
3.27	A comparison of typical attenuation curves from MINOS (a) and from our tests at Wideband (b). The individual channels are shown as well as the sum at both ends.	100
3.28	Optical fiber-SiPM distance and SiPM size dependencies demonstrated the need to adopt a readout based on larger SiPMs. Here, results are shown for 1 x 1 mm ² (yellow) and 2 x 2 mm ² (blue) SiPMs.	101

3.29	To test the 3 x 3 mm ² SiPM board concept (top), two prototypes were fabricated: one with three SiPMs, skipping every other channel (middle), and one with four consecutive SiPMs (bottom).	102
3.30	Data from three runs are shown: two with all but one fiber covered (blue, red) and one run with no fibers covered (black). Each histogram is run-time normalized, and the taped runs are summed (magenta).	103
3.31	The light yield as a function of longitudinal distance from the readout end (a) and the resulting efficiency with respect to the hodoscope (b) demonstrated the success of our prototype based on 3-mm SiPMs.	104
3.32	An exploded drawing of the ORM shows all of the components (Image credit: David Warner).	106
3.33	D. Warner test fits a machined housing prototype on a MINOS module observed by B. Behera (left). S. Marocci prepares an ORM for first test measurements (center). The ORM fits perfectly and is ready for testing (right). . .	107
3.34	It took a lot of work from a good team to get the test stand ready for its next task. Pictured from left to right: A. Aduszkiewicz, A. Wood and D. Munford (all UH) and A. Heggstuen (CSU).	108
3.35	A direct ground connection to the FEB common was important to avoid electrical noise.	110
3.36	Pedestal values for each channel (one per color) over time for both the south and north side readouts (top and bottom) comparing tested to reference ORMs (left and right) show stability within 10 ADC. Shifts in pedestal value are coherent across all channels within but not across FEBs.	112

3.37 Gain values for each channel (one per color) over time for both the south and north side readouts (top and bottom respectively) comparing tested to reference ORMs (left and right respectively) show stability in most channels to about 2 ADC/PE and uniformity to about 10 ADC/PE. The gain fit occasionally fails, usually causing an overestimation. 113

3.38 Light yield values for each channel (one per color) over time for both the south and north side readouts (top and bottom respectively) comparing tested to reference ORMs (left and right respectively) show stability to 10 ADC. . . . 114

3.39 Test ORM efficiency relative to the reference ORMs is shown for both the south (green) and north (blue) side readouts plotted against light yield showing a strong correlation. A few ORMs had significantly lower efficiencies than the rest caused by a single bad SiPM. 114

3.40 The channel-by-channel light yield relative to the channel mean over all tests was compared to the SiPM radial deviation measured at CSU during fabrication, and it showed that performance was uniform within our tolerance range. Image credit: Tyler Boone. 115

3.41 We have eight different Side CRT regions with three different views shown here from a CAD model. Image credit: J. Tillman and C. James. 116

3.42 The Side CRT West Utility Rack is pictured just before the rack’s enclosure panels were added. 117

3.43 It took three people to carefully install the FEBs in the proximity cryogenics area at pit level. T. Coan (right) and I (left) are shown practicing excellent team work. Not pictured(and photo credit): T. Boone. 119

3.44 Stabilizer plates for the Side CRT West Rolling Wall can be seen on either side of the gap between the mezzanine and the Warm Vessel. 120

3.45 From left to right: T. Boone, J. Judd, and K. Hardin work to install the next MINOS module into the West Rolling Wall. 122

3.46	A panoramic view of the West Rolling Wall shortly after the MINOS modules were installed attempts to show the scale of the detector.	123
3.47	Two views of me laboring away to ensure ICARUS is protected against cosmogenic backgrounds showing me securing the second row of MINOS modules (right) and me installing the first ORMs on the West Rolling Wall (left). Photo credit: Tyler Boone.	123
4.1	The simulation workflow for SBN showing the simulated data produced for the three detector subsystems: TPC, PDS, and CRT.	126
4.2	Two different views of the CRT system are shown with the individual CRT modules (green), top CRT support I-beams (grey), and warm vessel (red) visible.	127
4.3	The photon arrival probability is shown for two different slices along the drift coordinate: immediately adjacent to the anode wire planes (a) or immediately adjacent to the cathode plane (b).	129
4.4	An example of a photon arrival time parameterization fit is shown on the left. On the right, the photon propagation versus distance to the PMT is shown with error bars representing the RMS.	131
4.5	The reconstruction chain is shown for the three detector subsystems: TPC, PDS, and CRT.	131
4.6	A TPC waveform is shown after noise filtering with three hits identified. Image credit: F. Varanini.	132
4.7	Individual clusters can be seen within PANDORA-reconstructed tracks from a ν_μ CC event. Credit: B. Howard.	132
4.8	134

4.9	Optical hit finding identifies PMT signals and extracts the time of the signal pulse as well as the number of photoelectrons associated with it (left). Time resolution depends on the pulse amplitude and the level of scattering as shown by protons with two different energies spanning a full drift distance (right). Image credit: K. Terao.	134
4.10	The barycenter of photoelectron distribution traces the approximate neutrino vertex position as shown for ν_μ CC (top) and ν_e CC (bottom) events. Image credit: A. Menegolli.	135
4.11	The light yield for four different SiPM channels for cosmogenic muons at different distances from the SiPM, as measured at our test stand, is fit with a quadratic.	137
4.12	The coordinates obtained from Geant4 are used as input to the light yield model.	138
4.13	I plotted the fractional change in light output from polystyrene based scintillator vs. the energy deposited in 1 cm. The full range in ionization density up to the saturation point is shown on the left while the range relevant to cosmogenic muons is shown on the right.	141
4.14	A transverse scan with a collimated β source shows the impact of bulk attenuation on the light yield of a Top CRT strip.	142
4.15	The simulated light yield (blue) compared to the input light yield model (red) as a function of the distance from the SiPM (a) and the trigger efficiency vs threshold from the simulation (blue) compared to test stand measurements (red) (b) show reasonable agreement.	147
4.16	The time resolution at different distances from the SiPM from the detector simulation (blue) compared to test stand measurements (pink) without time-walk corrections (a) show good agreement. The light propagation time as a function of distance (b) was then extracted from the simulation and later used in the reconstruction.	148

4.17	The relative difference between true and reconstructed hit coordinates in the local CRT module frame and hit time are shown for the top CRT subsystem regions. The left column is for the local longitudinal coordinate, the center column is for the local lateral coordinate, and the right column is for time. The top row and bottom row are for the roof and rim regions respectively.	150
4.18	The relative difference between true and reconstructed hit coordinates in the local CRT region frame and hit time are shown for the side CRT subsystem regions. The left column is for the local longitudinal coordinate, the center column is for the local lateral coordinate, and the right column is for time. The top, middle, and bottom rows are for the East/West walls, North Wall, and South Wall respectively.	151
4.19	The relative difference between true and reconstructed hit coordinates in the local CRT module frame and hit time are shown for the bottom CRT subsystem. The left panel is for the local longitudinal coordinate, the center panel is for the local lateral coordinate, and the right panel is for time.	151
5.1	Pedestal measurements from a single run show the values for specific front-end boards (a), the one-dimensional projection of (a) onto the pedestal axis (b), and the RMS for the distribution of pedestals on each front-end board (c).	158
5.2	Gain measurements from a single run show the values for specific front-end boards (a), the one-dimensional projection of (a) onto the gain axis (b), and the RMS for the distribution of gains on each front-end board (c).	159
5.3	Covering a period of two months, the pedestal value for each front-end board channel is shown at left. The gains over time for one FEB is shown at right (For all FEB gains over time, see Fig. 5.4 below).	160
5.4	The gains for each SiPM channel is plotted over a period of two months. Each panel is one front-end board.	161
5.5	Charge sharing	162

5.6	Calibrated charge spectra are shown for each SiPM channel from a North Wall front-end board (mac5 005 / west pit level). Channels 22-31 read out a 3.1-m module while the remaining channels read out 2.6-m modules.	165
5.7	Calibrated charge spectra are shown for each SiPM channel from a West Rolling Wall front-end board (mac5 024 / south mezzanine level).	166
5.8	The most probable values for the light yield over the full module lengths are given for the North Wall (left) and the West Rolling Wall (right). For the North Wall, the distributions for the mezzanine-level (5.1-m modules) and the pit-level(3.1-m and 2.6-m modules) show the dependence on module length, especially when compared to the full 8-m length modules in the West Rolling Wall.	167
5.9	The fraction of all triggers coming the North Wall inner layer due to each front-end board is shown by front-end board position.	170
5.10	The fraction of all triggers coming the West Rolling Wall outer layer due to each front-end board is shown by front-end board position.	170
5.11	North Wall trigger rates by channel show the inner layer in the top two rows and the outer layer in the bottom two. From left to right, the columns are organized by mezzanine level and pit level.	171
5.12	West Rolling Wall trigger rates by channel show the inner layer in the top two rows and the outer layer in the bottom two. From left to right, the columns are organized by mezzanine level, mezzpit level, and pit level.	172
5.13	Shape comparisons for North Wall show qualitative agreement between data and simulation.	175
5.14	Shape comparisons for West Rolling Wall.	176
5.15	Fractional differences in charge spectral shapes are shown for the North Wall (a) and the West Rolling Wall (b).	177

5.16	A comparison between the mean FEB trigger rates measured in simulation and in data show agreement to within 5-10% for both the North Wall (a) and the West Rolling Wall (b).	177
6.1	Truth-level kinematic information for cosmogenic muons, in the presence of 3-m concrete overburden, with the origin at the T600 center.	181
6.2	Truth-level CRT cosmogenic muon tagging efficiencies for different kinematic variables describing the muons at the point they enter the active volume. Here, 3-m concrete overburden is present and the origin is at the T600 center.	182
6.3	The auto-veto fraction is dependent on the location of the neutrino vertex, the neutrino energy and whether the interaction is charged-current or neutral-current as shown for the three different BNB neutrino samples: ν_μ (top row), intrinsic ν_e (middle row), and oscillated ν_e . Charged-current and neutral-current interactions are in the left and right columns respectively.	187
6.4	The neutrino energy dependence of the auto-veto fraction leads to distortion of the neutrino energy spectrum as shown by comparing the true spectrum before (blue) and after auto-veto (red) for the three different BNB neutrino samples: ν_μ (top row), intrinsic ν_e (middle row), and oscillated ν_e . Charged-current and neutral-current interactions are in the left and right columns respectively. The spectra are shown in the upper subpanels while the bin-by-bin fractional differences are shown in the lower subpanels.	188
6.5	An illustration of the time-of-flight veto concept shows a muon entering from the top CRT before entering into the active liquid argon volume (AV). The time-of-flight is given by the delay between the CRT hit and the first PMT signal, approximated here by the shortest path to the PMT closest to the AV entry point	190

6.6	The TOF distribution for cosmogenic muons from the CRT to the earliest possible light to arrive at a PMT in truth (a) or using the full simulation and reconstruction chain for the CRT and applying a 2-ns smearing to the PMT arrival time (b) demonstrate the feasibility of the TOF veto method.	191
6.7	A simple test of the combined CRT-PDS reconstruction chain using vertical muons starting either downgoing from above CRT or upgoing from the TPC center show the absolute TOF provides a clean separation with a 19 ns mean and a 3 ns RMS.	192
6.8	Shown here are the TOF distributions for cosmogenic particles (a) and BNB neutrino samples (b,c,d) with $\sin^2 2\theta=0.013$, $\Delta m^2=0.43 \text{ eV}^2$	193
6.9	The performance of the TOF veto method is shown in terms of separating cosmogenic muons from neutrinos (a) and in terms of the purity and efficiency of the ν_e CC appearance sample in terms of cosmogenic backgrounds and auto-veto respectively. Oscillation parameters used here are $\sin^2 2\theta=0.013$, $\Delta m^2=0.43 \text{ eV}^2$	196

Chapter 1

Introduction

The field of elementary particle physics seeks to understand the most fundamental building blocks of nature, which are, in our current understanding, quantum fields. Elementary particles are excitations of these fields, and, in turn, the fields govern the interactions between these particles. This picture is formally described by the Standard Model (SM) of particle physics, which is our current best picture of the subatomic world. The SM has been tested to remarkable precision and increasingly stringent tests are underway. Because of the importance of the SM, I will discuss it in more detail in Section 1.1 with emphasis on the neutrino sector.

Neutrinos are some of the oddest and least understood constituents of the SM. Though neutrinos were proposed 80 years ago, we still do not understand many of their basic properties. As we will see in Section 1.2, the effort to understand neutrinos better has led to some interesting findings.

Over the last several decades, few attempts to find physics beyond the Standard Model (BSM) have been successful. The first significant extension came from the discovery of the phenomenon of neutrino flavor oscillations, which I will introduce in Section 1.3. This phenomenon provides direct evidence that neutrinos are not massless as assumed by the SM. There are several approaches for adding neutrino mass to the SM, some of which will be briefly introduced in Section 1.1.

Many neutrino oscillation experiments have been carried out, and almost all of the results match the SM prediction with the ad hoc addition of nonzero neutrino masses. However, there have been some anomalous results that may be hinting at more new physics from the neutrino sector. I will discuss these anomalies in detail in Section 1.5 and the

pressing need to resolve them, underscoring the importance for the multi-experiment Short-Baseline Neutrino Program (SBN) discussed in Chapter 2. Particular emphasis will be on the SBN experiment that I work on, ICARUS. I will discuss how challenging resolving neutrino anomalies can be. Great care must be taken to understand backgrounds and systematics that can generate false positives or, put another way, a new neutrino anomaly.

In particular, my work addresses the challenge introduced by cosmogenic backgrounds that are of concern for SBN as all of its detectors will be operating at a shallow depth. My first step in this effort was to develop an auxiliary detector capable of tagging cosmogenic muons passing in close proximity to the primary neutrino detector, the Cosmic Ray Tagging system (CRT). I will discuss the CRT at length in Chapter 3 from its first conceptual design to installation. The remainder of my work has been to develop simulation and reconstruction techniques utilizing the CRT for cosmogenic background rejection. An overview of the simulation and reconstruction code will be given in Chapter 4. With the hardware and software in place, I will demonstrate the validity of the simulation and basic reconstruction tools in Chapter 5 using the first CRT data. Finally, in Chapter 6, I will tie everything together and present novel, high-level reconstruction techniques that I have developed for cosmogenic background rejection, including results from a simulation based study.

1.1 The Standard Model

In this section, I will give a brief overview of the Standard Model (SM), our current best picture of the subatomic realm. This will be important for understanding how neutrinos fit into the bigger picture and how recent discoveries in the neutrino sector may already be pointing the way towards beyond the SM (BSM) physics.

The SM is a Lorentz invariant quantum field theory. The SM Lagrangian, where all of

the fields and their symmetries are encoded, is renormalizable, meaning it is naively valid up to arbitrarily high energies, and is uniquely determined once its internal symmetries and particle content are specified. The symmetries are gauged $SU(3)\times SU(2)\times U(1)$ invariance demanded by the conservation of color charge, electric charge, weak isospin, and weak hypercharge. Accidental symmetries that arise from $U(1)$ invariance are baryon number, electron number, muon number, and tau number. The particle content is specified by the quantum fields listed below.

- Fermion fields ψ giving rise to the tri-generational doublet structure $(u,d),(s,c),(t,b)$ quarks and $(e,\nu_e),(\mu,\nu_\mu),(\tau,\nu_\tau)$ leptons.
- Electroweak fields W_1, W_2, W_3, B giving rise to the W and Z bosons and the photon.
- Gluon field G_a giving rise to the gluons.
- Higgs field ϕ giving rise to the Higgs boson.

The collection of matter particles, force carriers, and the Higgs is shown in Figure 1.1.

In addition to the fields and symmetries, the most general SM Lagrangian with massless neutrinos requires 19 free parameters that must be determined by experiments. Among these are the quark, lepton, and Higgs masses, the quark mixing angles, and gauge couplings.

Central to my work are neutrinos and their interactions with other matter. Neutrinos are the only particles in the SM that interact exclusively via the weak force and gravity. The weak force is mediated by the electrically charged W^+ and W^- bosons as well as the electrically neutral Z boson. Thus, neutrinos interact in two modes: charged current (CC) and neutral current (NC). These interactions are demonstrated in the tree-level Feynman diagrams in Figure 1.2. The CC interactions involve the transmutation of neutrinos from or into a charged lepton partner through the flow of charge via a W boson from or to a quark confined in a composite particle (hadron), a neutron or proton in this case. The

Standard Model of Elementary Particles

		three generations of matter (fermions)			interactions / force carriers (bosons)	
		I	II	III		
mass		$\approx 2.2 \text{ MeV}/c^2$	$\approx 1.28 \text{ GeV}/c^2$	$\approx 173.1 \text{ GeV}/c^2$	0	$\approx 124.97 \text{ GeV}/c^2$
charge		$\frac{2}{3}$	$\frac{2}{3}$	$\frac{2}{3}$	0	0
spin		$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1	0
		u up	c charm	t top	g gluon	H higgs
	QUARKS	$\approx 4.7 \text{ MeV}/c^2$	$\approx 96 \text{ MeV}/c^2$	$\approx 4.18 \text{ GeV}/c^2$	0	
		$-\frac{1}{3}$	$-\frac{1}{3}$	$-\frac{1}{3}$	0	
		$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1	
		d down	s strange	b bottom	γ photon	
	LEPTONS	$\approx 0.511 \text{ MeV}/c^2$	$\approx 105.66 \text{ MeV}/c^2$	$\approx 1.7768 \text{ GeV}/c^2$	$\approx 91.19 \text{ GeV}/c^2$	
		-1	-1	-1	0	
		$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1	
		e electron	μ muon	τ tau	Z Z boson	
		$< 2.2 \text{ eV}/c^2$	$< 0.17 \text{ MeV}/c^2$	$< 18.2 \text{ MeV}/c^2$	$\approx 80.39 \text{ GeV}/c^2$	
		0	0	0	± 1	
		$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1	
		ν_e electron neutrino	ν_μ muon neutrino	ν_τ tau neutrino	W W boson	
						GAUGE BOSONS VECTOR BOSONS
						SCALAR BOSONS

Figure 1.1: Table of Standard Model particles. (Image credit: <https://en.wikipedia.org>)

three generations of neutrinos are categorized by flavors, defined by the charged lepton to which it couples in CC interactions. Neutrino flavors cannot be determined in NC interactions, which are flavor agnostic.

In Section 1.2, I will show that neutrinos are notoriously difficult to detect. Due to the weak force being about four orders of magnitude weaker than the electromagnetic force and having a range of about 10^{-3} fm, about 0.1% the width of a proton, neutrinos have a relatively small probability of interacting with other matter. This probability is described by the cross-section for the scattering process. Charged current inclusive cross-section ($\nu_\mu N \rightarrow \mu^- X$) measurements are displayed in Figure 1.3 scaled by the inverse of the neutrino energy. You can see that the cross-section increases with neutrino energy,




Figure 1.2: Feynman diagrams for neutrino-nucleon charged-current scattering (a) and neutrino-nucleon or neutrino-lepton neutral-current scattering (b) where $\ell = e, \mu, \tau$, and $f = n, p$ or e .

transitioning between logarithmic and linear scaling at 100 GeV. Note the size of the error bars, a reflection of the difficulty of these measurements. Neutrino production processes are complex, making neutrino fluxes difficult to predict. Neutrino interactions themselves are difficult to model, especially when nuclear effects and final state interactions are considered.

The SM has been remarkably successful in describing most observed phenomena. However, we know it is not a complete picture with some [2] describing it as an effective field theory. One obvious missing piece is gravity, currently described at large distances by General Relativity. The problem of quantum gravity is notoriously difficult to solve. From cosmology, we know that the SM is missing something in that it cannot account for dark energy or dark matter. Searches for particle based dark matter have so far come up empty handed. Astronomical surveys are beginning to give us some understanding of the dynamics of dark energy, but a deeper understanding is likely far away. Finally, there have been surprises in the neutrino sector that are the focus here. The most well known of these surprises is that of nonzero neutrino masses.

In the SM, neutrinos are assumed to be massless. This is due to the fact that neutrinos only couple to the weak fields, having no electrical charge, and only come in “left-handed”




Figure 1.3: Charged current inclusive muon neutrino scattering cross-section per nucleon measurements [1].

chiral states. However, experiments have shown that neutrinos oscillate between weak flavor states as they propagate. As I will show in Section 1.3, this requires that neutrinos are massive and nondegenerate. This requires 3 additional neutrino mass parameters and 4 neutrino mixing parameters be added to some extension of the SM. Furthermore, constraints from cosmology and β -decay set limits on the absolute neutrino mass scale to below 1 eV, several orders of magnitude below all of the other SM particles (see Figure 1.4). This surprise from the neutrino sector raises some interesting questions.

First, what is the mechanism behind the neutrino mass? Broadly speaking, there are two possibilities. One is that neutrinos are Dirac fermions, the other is that neutrinos are Majorana fermions. Both require the existence of additional neutrinos. I will briefly explain the implications of each case below [3].

If neutrinos are Dirac fermions, right-handed neutrinos would need to exist. Further-




Figure 1.4: Masses of all of the Standard Model fermions showing the departure of the neutrino mass scale from the 0.5 MeV - 200 GeV range populated by the other fermions [3].

more, these right-handed neutrinos could not interact via any SM force. These “sterile” neutrinos could be the right-handed neutrino fields missing from the SM. In addition, this extended SM would remain renormalizable. These features make this mechanism appealing; however, this option comes at a cost. It does not provide any explanation for why the neutrino masses are so much smaller than the other SM fermions. Said another way, the Yukawa couplings giving rise to the other SM fermion masses via the Higgs mechanism would be much larger than the new set of Yukawa couplings of the neutrinos for no apparent reason. Other SM extensions exist that could explain this effect, with the addition of extra spatial dimensions and a new gauge symmetry for example, but this topic is beyond the scope of my work.

On the other hand, if neutrinos are Majorana fermions, they would be the first such fermion found in nature. Majorana fermions are their own antiparticles. Thus, they violate the hitherto observed conservation of lepton number by two units. In addition, this option requires us to give up on the renormalizability of the SM. This would be the first direct evidence that the SM is an effective field theory. However, this case provides

a natural explanation for the smallness of the neutrino mass, via the seesaw mechanism for example. In this case too, we obtain right-handed neutrinos, except they would be required to be very heavy, decaying early in the history of the universe.

In addition to a new mass mechanism, neutrino flavor oscillation requires six new parameters that uniquely determine the mixing physics: three mixing angles and two mass splittings (outlined in Section 1.3) in addition to a CP-violating phase. This last parameter is especially interesting for cosmology since, if the phase is near maximal, it could explain why we live in a universe dominated by matter rather than radiation.

Before getting into the phenomenology of neutrino flavor oscillation (Section 1.3), I will first provide a brief history of the neutrino, beginning with the original postulation of its existence and ending with the surprise of neutrino flavor oscillation. This will illustrate some of the challenges in neutrino physics with many parallels today.

1.2 Neutrino History

A dissertation on neutrino physics is not complete without some neutrino history. The year was 1911 and physicists studying β -decay, which was assumed to be a two-body process at the time, found something interesting: the energy spectrum of the emitted electron was continuous as opposed to discrete, apparently violating the law of conservation of energy. In 1931, not ready to give up on energy conservation, Wolfgang Pauli proposed the existence of a new, electrically neutral particle that he dubbed the “neutron.” This “neutron,” as part of a three-body β -decay process, explained the continuous energy spectrum. To be consistent with observations, Pauli’s “neutron” had to be very light, of the same order of magnitude of the electron mass or less. The following year, Enrico Fermi had devised a theory of β -decay incorporating Pauli’s “neutron.” Fermi renamed Pauli’s “neutron” to neutrino, an Italian play on words roughly translating to “little neutral one,” as a new, electrically neutral particle with approximately the same

mass as a proton had just been discovered and was a better match for the neutrino's former moniker. However, Pauli had his doubts that his "neutron" would ever be experimentally observed as Fermi's theory predicted that the cross-section of the inverse of the β -decay process, neutrino absorption, was impractically small. The story goes that he once bet a case of champagne that it would never be detected, declaring, "I have done a terrible thing, I have postulated a particle that cannot be detected."

Fortunately for us, Pauli lost the bet. Though it took 26 years from his proposal, Clyde Cowan and Frederick Reines successfully made the first definitive observation of neutrinos produced by a nuclear reactor in 1956. In particular, they had observed the anti-electron neutrino ($\bar{\nu}_e$) via the inverse β -decay (IBD) process ($\bar{\nu}_e p \rightarrow n e^+$). This discovery earned Cowan and Reines the 1995 Nobel Prize.

After the existence of the neutrino was confirmed, physicists were wondering if there existed multiple generations of neutrinos in a similar manner as had been observed for the charged leptons (at the time, this was the electron (e) and the muon (μ)). To this end, Leon Lederman, Melvin Schwartz and Jack Steinberger devised the first accelerator based neutrino experiment. Their method was successful, and in 1962, they announced the discovery of a second flavor of neutrino, the muon neutrino (ν_μ). Their discovery demonstrated the weak doublet structure of the leptons. For these achievements, they were awarded the Nobel Prize in 1988.

Then, in 1975, a third generation of charged lepton was discovered, the tau (τ). Immediately following this discovery, physicists set out to find the corresponding tau neutrino (ν_τ). While indirect evidence seemed to hint at its existence, direct evidence for the ν_τ would not be discovered until 2000 by the DONUT collaboration. This last neutrino was the second to last missing particle of the Standard Model (with the last being the Higgs boson).

As we have seen, neutrinos are fascinating particles that completed duality in the Standard Model between the quarks and leptons where both groups have a double struc-

ture with three generations. However, the richness of neutrino physics does not end there. As we will see next, neutrinos also gave us our first look at Beyond the Standard Model Physics.

We must back up a bit now to the period just after the discovery of the ν_μ to the year 1964. John Bahcall and Raymond Davis proposed the Homestake Experiment to verify the idea that nuclear fusion was the source of energy in stars by detecting neutrinos produced in the Sun [4], enshrined in the Standard Solar Model (SSM). The radiochemical experiment, based on inverse β -decay of ν_e on chlorine, was carried out by Davis, who announced the first results [5] in 1968. He had successfully detected solar neutrinos. However, there appeared to be a problem; Davis had detected significantly fewer neutrinos than Bahcall had predicted with the SSM.

During the twenty or so years following Davis' initial results, both the experiment and the theory were scrutinized. The result, which had become known as the solar neutrino problem, appeared to be robust, leaving many physicists puzzled.

Later, in 1989, a three-kiloton water Cherenkov detector in Japan, Kamiokande, confirmed Davis' result. Originally designed as a proton decay experiment (not observed to this day), it turned out the detector was also well suited to observe neutrinos with its large mass, low background and energy threshold. The experiment showed that only about one-third of the neutrino flux predicted by the SSM was observed. What was going on? The prime suspect was the SSM. The neutrinos observed in the Homestake and Kamiokande experiments come from the production of ${}^8\text{B}$ and are on the high-energy tail of the solar neutrino flux. Furthermore, the ${}^8\text{B}$ production rate is strongly dependent on the core temperature of the Sun. This spurred another set of experiments to sample the entire solar neutrino energy spectrum rather than just the high-energy tail.

Two new radiochemical experiments, SAGE and GALLEX, based on inverse β -decay of ν_e on gallium, instead of chlorine as in Davis' experiment, had observation thresholds of 233 keV. This enabled the observation of the low-energy neutrinos produced early in

the pp-fusion chain. In the early 1990's, these experiments both showed that a significant deficit was present across the entire solar neutrino flux. Thus, the plot thickened.

Returning to the Kamiokande experiment and another water Cherenkov experiment, Irvine-Michigan-Brookhaven (IMB), a new problem arose. Both experiments were sensitive to ν_μ 's in addition to ν_e 's, unlike the radiochemical experiments. They made the first observations of atmospheric neutrinos, which are produced primarily by the decays of pions and muons in atmospheric showers. While the absolute normalization on the primary cosmic ray flux had about a 20% uncertainty at the time, the branching fractions for the different decay modes of pions and muons were precisely known. In turn, the ratio of events from ν_μ 's to those from ν_e 's was well predicted. However, both experiments observed a statistically significant deficit of ν_μ 's with respect to ν_e 's compared to predictions. The first result came from Kamiokande in 1988. Improved results from them along with consistent results from IMB in the early 1990's bolstered this surprising result, which came to be known as the atmospheric neutrino anomaly.

Following the intriguing results from Kamiokande and IMB, Kamiokande was upgraded in order to cope with the limited neutrino statistics. The next iteration, Super-Kamiokande (SK), had ten times the active volume and photo-coverage of its predecessor. The improved sensitivity enabled SK to map the atmospheric neutrino rate as a function of zenith angle. They confirmed their previous result; there was indeed a deficit of atmospheric ν_μ 's. In 1998, SK published conclusive evidence of neutrino flavor oscillation in atmospheric neutrinos. This explained the atmospheric neutrino anomaly in that the deficit was produced by the transition of ν_μ 's into ν_e 's.

In parallel, another water Cherenkov experiment, using heavy water as a target, was coming online. The Sudbury Neutrino Observatory (SNO), designed to follow up on the solar neutrino problem, confirmed previous results that only about a third of solar ν_e 's were detected. Furthermore, SNO was sensitive to neutral current scattering, which is neutrino flavor agnostic. Through the neutral current channel, SNO observed that there

were actually no missing neutrinos; the measured flux agreed precisely with predictions from the SSM. This result, published in 2001, was conclusive evidence that ν_e 's had changed flavor, proving the presence of neutrino flavor oscillation in solar neutrinos, thus solving the solar neutrino problem.

For their contributions to these profound discoveries, Takaaki Kajita of SK and Arthur McDonald of SNO were awarded the Nobel Prize in 2015. The confirmation that neutrinos oscillate was highly unexpected and incompatible with the Standard Model. As we will see in the next section, neutrino flavor oscillation requires that neutrinos have mass. This was the first instance of BSM physics. We have seen that neutrinos are full of surprises, and as we will see over the last sections in this chapter, there is reason to expect that more surprises are on the way.

1.3 Neutrino Oscillation Phenomenology

As we saw in the previous section, neutrino flavor oscillation was the solution to the decades old problems of the missing solar and atmospheric neutrinos. This qualitatively describes the phenomenon. Let us now go deeper into the neutrino flavor oscillation phenomenology [6].

In the language of quantum mechanics, neutrino flavor oscillation is the creation of a neutrino in a weak eigenstate ν_α and the interaction of this neutrino, after propagating some macroscopic distance, in a different weak eigenstate ν_β . Here, the weak eigenstates are defined by the charged leptons associated with the charged current neutrino production or interaction processes, e , μ , or τ . As neutrinos propagate in a vacuum, they are free particles and are therefore eigenstates of the free particle Hamiltonian with eigenvalues equal to the neutrino masses, m_i .

$$\hat{H}_{free}|\nu_i\rangle = m_i|\nu_i\rangle, i = 1, 2, 3. \quad (1.1)$$

Generally speaking, the weak basis and the mass basis are not the same, but one can rotate from one into the other via some unitary matrix. In the neutrino case, this matrix is known as the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix, U_{PMNS} . The case for three neutrino flavors is shown in Equation 1.2, but the phenomenology can easily be generalized to an arbitrary number of flavors.

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \quad (1.2)$$

From this, we can calculate the flavor transition probability $P_{\alpha\beta}$ from flavor state α to flavor state β as a function of the neutrino energy E_ν and the propagation distance L .

$$P_{\alpha\beta}(E_\nu, L) = \left| \sum_j U_{\beta j} U_{\alpha j}^* e^{-i \frac{m_j^2 L}{2E_\nu}} \right|^2 \quad (1.3)$$

The conventional parameterization of the PMNS matrix is in terms of rotation matrices with three associated mixing angles, θ_{13} , θ_{23} , and θ_{12} . We also have one physical CP-violating phase, δ_{CP} . For convenience, define $s_{ij} \equiv \sin(\theta_{ij})$ and $c_{ij} \equiv \cos(\theta_{ij})$ with $i = 1, 2, 3$.

$$U_{PMNS} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13} e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13} e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (1.4)$$

To illustrate the important features of this phenomenology, it is sufficient to consider one typical experimental regime and approximate the flavor transition probabilities to leading order. Note that, while these approximations were sufficient for early oscillation experiments, modern experiments use the exact descriptions. Define the mass-squared difference $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$. From experiment, we know $\Delta m_{21}^2 \ll |\Delta m_{31}^2| \simeq |\Delta m_{32}^2|$. When

$\Delta m_{21}^2 L/E_\nu \ll 1$, we have, to leading order, the following probabilities for muon neutrino disappearance ($\nu_\mu \rightarrow \nu_x$, Equation 1.5), electron neutrino disappearance ($\nu_e \rightarrow \nu_x$, Equation 1.6), and electron neutrino appearance ($\nu_\mu \rightarrow \nu_e$, Equation 1.7). These approximate probabilities are the same for the anti-neutrino case). This regime is appropriate for atmospheric neutrinos, long-baseline accelerator neutrino experiments, and short-baseline reactor neutrino experiments. We will discuss the different oscillation experiments in Section 1.4.

$$P_{\mu\mu}(E_\nu, L) \approx 1 - \sin^2(2\theta_{23}) \sin^2\left(\frac{\Delta m_{31}^2 L}{4E_\nu}\right) \quad (1.5)$$

$$P_{ee}(E_\nu, L) \approx 1 - \sin^2(2\theta_{13}) \sin^2\left(\frac{\Delta m_{31}^2 L}{4E_\nu}\right) \quad (1.6)$$

$$P_{\mu e}(E_\nu, L) \approx \sin^2(2\theta_{13}) \sin^2(\theta_{23}) \sin^2\left(\frac{\Delta m_{31}^2 L}{4E_\nu}\right) \quad (1.7)$$

The defining feature of the oscillation probabilities is that the frequency of oscillation depends on the propagation distance divided by the neutrino energy L/E_ν (what the experimentalists can choose) scaled by the mass-squared difference Δm^2 . This illustrates how any oscillation implies non-zero neutrino mass. Note that oscillations are not sensitive to the absolute neutrino mass scale. For neutrino disappearance, the amplitude of the mixing is set by one mixing angle. In contrast, for neutrino appearance, the amplitude depends on two mixing angles with the dominant one associated with the disappearance probability for each flavor involved.

Note that the oscillation phenomenology becomes significantly more complex when we introduce effects from neutrinos propagating in matter, known as the MSW effect. This is the primary driver of solar neutrino oscillation, and it can have important effects in atmospheric and long-baseline oscillation experiments. As this effect is not important for ν_μ disappearance or ν_e appearance in a ν_μ neutrino beam at short baselines, the most important modes for this dissertation, I will not discuss matter effects further.

Table 1.1: 3ν mixing parameters from global analysis, shown for different assumptions of normal (NH) or inverted (IH) hierarchy. 1σ errors shown. Reproduced from [7].

	$\theta_{12}[\circ]$	$\theta_{23}[\circ]$	$\theta_{13}[\circ]$	$\Delta m_{21}^2 \times 10^5 [eV^2]$	$\Delta m_{32}^2 \times 10^3 [eV^2]$	$\delta_{CP}[\circ]$
NH	$34.5_{-1.0}^{+1.2}$	$47.7_{-1.7}^{+1.2}$	$8.45_{-0.14}^{+0.16}$	$7.55_{-0.16}^{+0.20}$	$2.424_{-0.03}^{+0.03}$	218_{-27}^{+38}
IH	$34.5_{-1.0}^{+1.2}$	$47.9_{-1.7}^{+1.0}$	$8.53_{-0.15}^{+0.14}$	$7.55_{-0.16}^{+0.20}$	$-2.50_{-0.03}^{+0.04}$	281_{-27}^{+23}

Oscillation experiments have successfully determined many of the parameters in the phenomenology presented here. Table 1.1 summarizes these values. For a graphical view, see Figure 1.5, which shows the relative size of the mass-squared differences as well as the flavor content of each mass eigenstate. Despite this tremendous progress, there are still some things we do not yet know.

- What is the sign of Δm_{31}^2 (i.e. normal or inverted hierarchy)?
- Is θ_{23} precisely maximal (i.e. $\pi/4$) and if not, in which octant is it?
- What is the precise value of δ_{CP} ?

It is useful to differentiate between two classes of oscillation experiments using the “baseline” set by the L/E_ν range. Since the baseline is scaled by the Δm^2 , whatever baseline is used in an experiment will determine how sensitive the experiment is to different oscillation parameters. Experiments will choose a neutrino source that sets the flavor and energy content (discussed in Section 1.4). Then, depending on what oscillation parameters are being probed, the propagation distance from the source is set. For example, neutrinos produced at particle accelerators typically have energies in the range of 1-10 GeV. A “short-baseline” experiment would operate at a distance of hundreds of meters from the source and be sensitive to large Δm^2 ($>0.1 \text{ eV}^2$). With the same source, a “long-baseline” experiment would operate at distances of hundreds of kilometers and be sensitive to Δm^2 in the range of 10^{-2} - 10^{-3} eV^2 .

In experimentally probing oscillation parameters, there are several challenges to overcome. First, the neutrino source is seldom monoenergetic. Instead, the source is described by a flux $\Phi(E_\nu)$ that is not perfectly known. As described in Section 1.1, the neutrino

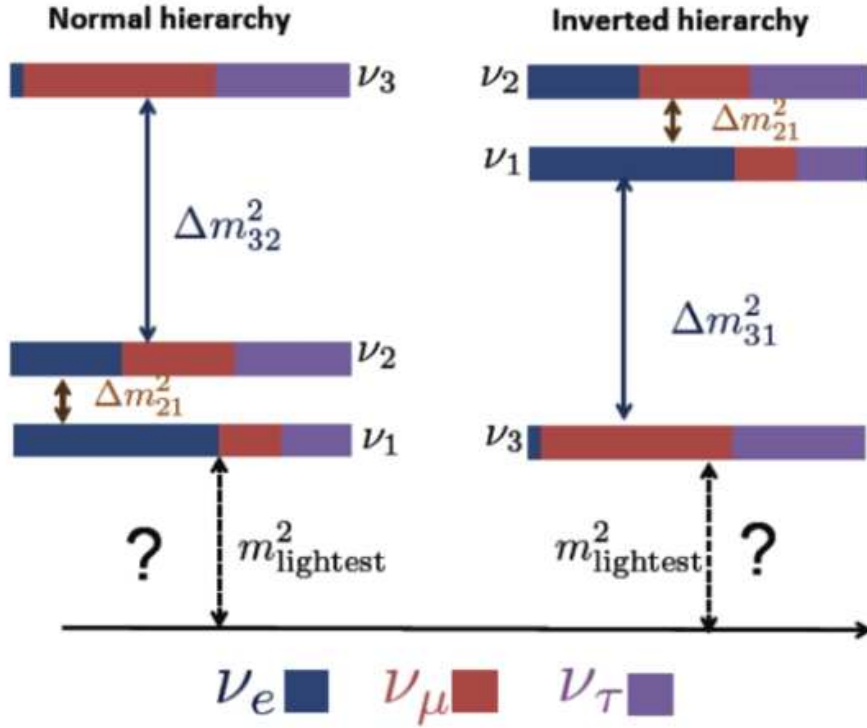


Figure 1.5: Two orderings of the neutrino mass eigenvalues are possible, the normal or inverted hierarchies. The mass-squared differences shown are not to scale. The colors indicate the “flavor content” of each mass eigenstate. (Image credit: <http://ignatz.phys.rpi.edu>)

scattering cross-section, $\sigma(E_\nu)$, is energy dependent and not precisely known. Then there are detector effects, including limited resolution, observation efficiency, etc., that depend on the interaction channel and energy, $\epsilon(E_\nu)$. What is measured by the experiment is the average oscillation probability in some bin of neutrino energy, E_ν . To account for all of these effects, the average oscillation probability is calculated as follows [7], integrating over the energy range of the bin.

$$\langle P_{\alpha\beta} \rangle = \frac{\int dE_\nu \frac{d\Phi(E_\nu)}{dE_\nu} \sigma(E_\nu) \epsilon(E_\nu) P_{\alpha\beta}(E_\nu)}{\int dE_\nu \frac{d\Phi(E_\nu)}{dE_\nu} \sigma(E_\nu) \epsilon(E_\nu)} \quad (1.8)$$

To further complicate matters, oscillation experiments also must contend with backgrounds. The types of backgrounds depend on the experimental techniques among other

things. Neutrino experiments tend to use high- Z materials as targets in order to enhance the neutrino interaction rate. The drawback to this strategy is that the heavy nuclei introduce complex nuclear effects. In addition, the nuclear environment can affect observed final state interactions. These effects are difficult to model and can cause misidentification of the neutrino interaction, thereby introducing systematic uncertainties into the analysis. In general, neutrino experiments contend with significant cosmogenic backgrounds. This is why many neutrino experiments opt for operation below ground where the significant overburden strongly suppresses most of the cosmogenic backgrounds, especially the hadronic and electromagnetic components. Even below ground, there is still some component remaining in the form of muons and secondary production.

One strategy for isolating the signal from the background and reducing systematic uncertainties is to introduce a near detector close to the neutrino source that can observe the neutrino flux before the neutrino oscillation probability becomes nontrivial. The near detector can be used to predict the background at the far detector. In doing so, correlated uncertainties between the two detectors are suppressed, especially when the near detector is functionally identical to the far detector. We will see this technique in use at the Short-Baseline Neutrino Program discussed in Chapter 2.

In the next section, I will introduce the landscape of neutrino experiments. This will be important for understanding which types of experiments probe which oscillation parameters. This will also be important for understanding the short-baseline neutrino anomalies introduced in Section 1.5.

1.4 Landscape of Neutrino Experiments

The full landscape of neutrino experiments is vast, and the purpose of this section is not to provide an exhaustive review of past, present, and future experiments. Rather, I will focus on those experiments that are important for constraining the short-baseline

neutrino oscillation anomalies that will be discussed in the next section. Of course, not all neutrino experiments study neutrino oscillation. We also have neutrino telescopes, using neutrinos to do astrophysics; cross-section programs, which seek to better understand neutrino interactions; neutrinoless $\beta\beta$ -decay searches, which aim to determine whether the neutrino is Dirac or Majorana in nature; and direct neutrino mass measurements, which utilize extremely precise kinematic measurements of the electron emitted in ordinary β -decay. Several results from these experimental programs have direct implications for our understanding of neutrino oscillation.

As discussed in the previous section, experimentalists studying neutrino oscillation need to identify neutrino flavors, neutrino energies, and the distance between the neutrino source and where the neutrino interacts. It will be useful to categorize experiments by the neutrino source: astrophysical sources, accelerators, reactors, and radioactive sources. Table 1.2 summarizes the oscillation parameters probed by different neutrino sources and baselines. I will not discuss neutrinos produced by radioactive decays in Earth's interior known as geoneutrinos.

Astrophysical sources include atmospheric, solar, and supernovae in addition to as of yet unknown sources, both galactic and extra-galactic in origin. As a whole, this category spans keV through EeV energies and all neutrino flavors. Experiments using these sources take advantage of the fact that neutrinos have a relatively small cross-section, thus having the unique ability to point back directly to the source while also maintaining their initial state (modulo flavor oscillation). For example, we saw in Section 1.2 how neutrinos validated the proposed mechanism for stellar energy production incorporated into the Standard Solar Model. As another example, neutrinos carry away about 99% of the energy released in core-collapse supernovae and reveal much about the underlying dynamics [8].

Aside from deepening our understanding of astrophysical phenomena, these sources can be used to study flavor oscillation. We already saw examples of this in Section 1.2. so I will not dwell on this point here. It is important to mention the IceCube experiment [9],

Table 1.2: (3ν) Mixing parameters probed by different neutrino sources at different baselines (MBL=medium-baseline, LBL=long-baseline). Reproduced from [7].

Experiment	Dominant	Important
Solar (GALEX, SAGE, Borexino)	θ_{12}	$\Delta m_{21}^2, \theta_{13}$
Atmospheric (Super-K, IceCube(Deep Core))		$\theta_{23}, \Delta m_{31,32}^2 , \theta_{13}, \delta_{CP}$
Reactor LBL (KamLAND)	Δm_{21}^2	θ_{12}, θ_{13}
Reactor MBL (Daya Bay, RENO, Double Chooz)	$ \Delta m_{31,32}^2 , \theta_{13}$	
Accel. LBL disapp. (K2K, MINOS, T2K, NO ν A)	$ \Delta m_{31,32}^2 , \theta_{23}$	
Accel. LBL appear. (MINOS, T2K, NO ν A)	δ_{CP}	θ_{13}, θ_{23}

which consists of 1 km³ of ultrapure ice at the South Pole instrumented with an array of photomultiplier tubes. While its primary objective is to use neutrinos for astrophysics, the experiment has conducted studies of atmospheric neutrino oscillation (muon neutrino disappearance). Their results are in agreement with previous results. In addition, IceCube has set a unique constraint on anomalous oscillation results with its ability to contain especially high energy neutrino events. As a particularly relevant example for the next section, they measured atmospheric mixing parameters in the energy range 320 GeV through 20 TeV.

Reactor sources produce exclusively anti-electron neutrinos ($\bar{\nu}_e$) from fission of heavy nuclei, usually ²³⁵U and/or ²³⁹Pu, and have energies in the range of about 1-10 MeV. This source is often used for $\bar{\nu}_e$ disappearance experiments and probes the solar mass splitting (Δm_{21}^2) and mixing angles θ_{12} and θ_{13} . Some current and past examples include Daya Bay [10], Double Chooz [11], RENO [12], and Bugey [13]. Some intriguing results have come out that highlight the challenges in modeling the reactor neutrino flux that I discuss in the next section. This has spurred a new generation of movable neutrino detectors, DANSS [14], PROSPECT [15], STEREO [16], and NEOS [17] among others.

Radioactive decay sources produce electron or anti-electron neutrinos in β -decays with typical energies of tens to hundreds of keV. Experiments use these sources for different reasons. Some are using them for calibration, the solar neutrino experiments SAGE [18] and GALEX [19] for example. Tritium sources are used in direct neutrino mass mea-

surements. One such experiment, KATRIN, recently published their first upper limit on the neutrino mass at about 2 eV [20]. Some experiments, such as EXO [21], CUORE [22] and GERDA [23] to name only a few, use particular isotopes, e.g. ^{116}Cd , ^{130}Te , ^{136}Xe and ^{76}Ge , known to produce $\beta\beta$ -decays, to search for the hypothetical neutrinoless $\beta\beta$ -decay process. If observed, this would confirm the Majorana nature of neutrinos. Finally, as we will see in the next section, radioactive sources can also be used to look for new oscillation physics.

Accelerator sources produce primarily muon and anti-muon neutrinos, with a small admixture of electron and anti-electron neutrinos, with typical energies of hundreds of MeV to tens of GeV. Experiments using these sources study neutrino interactions and neutrino oscillations. These sources offer several advantages including the ability to model the flavor and energy content of the beam as well as having precisely known timing structure, providing an effective method for suppressing backgrounds. We will see these techniques applied in the Short-Baseline Neutrino Program in Chapter 2. Some important experiments that used or are using accelerator sources include LSND, MiniBooNE, MINOS, T2K, and *NOvA*. These experiments have probed the atmospheric mass splitting and mixing angle. The first two experiments observed anomalous oscillations, as we will see next. Though the other experiments, MINOS in particular, have provided strong constraints on an intriguing interpretation of these anomalous results. More recently, some hints at the mass hierarchy [24] (the sign of Δm_{13}^2) and a non-zero CP-violating phase [25] (δ_{CP}) have been observed.

Taken as a whole, nearly all of these experiments provide a self-consistent picture of neutrino flavor oscillation involving the three known neutrinos from the Standard Model. However, as mentioned throughout this chapter, there have been some anomalous findings, discussed in the next section, that may be hinting at something beyond this three-neutrino picture.

1.5 Short-Baseline Anomalies

Neutrino flavor oscillation physics is entering an age of precision measurements, enabled by cutting-edge detector technology and intense neutrino sources. We have seen the substantial success of the three-flavor picture in describing observations. With the mixing angles being better measured with each passing generation of experiments, and the future DUNE and Hyper-Kamiokande experiments promising to make a precision measurement of the CP-violating phase and an unambiguous determination of the mass hierarchy, it would appear that our picture of neutrino oscillation is on track to completion. However, over the last twenty years or so, three classes of anomalous results have emerged that question the completeness of the three-flavor picture or, at the least, hint at some form of BSM physics. All three of these have been observed at short baselines. Collectively, these are referred to these as the short-baseline anomalies (SBAs).

The first SBA is the so-called Gallium Anomaly. This is the apparent deficit of electron neutrinos from MCi β -sources used for calibration by the SAGE and GALLEX experiments. The experiments used ^{51}Cr and ^{37}Ar and found a ratio of 0.84 between the observed and predicted counts with a significance of about 3σ [29], shown in Figure 1.6.

The second anomaly is the so-called Reactor Anomaly. This is the apparent deficit of anti-electron neutrinos from nuclear reactors compared to the predicted flux. A recent improvement in theoretical models used to predicted these reactor fluxes has only exacerbated the discrepancy. Figure 1.7 shows different reactor neutrino measurements in the form of the ratio of measured to expected counts as a function of distance to the reactor [30]. The ratio is about 0.93 with a significance of about 3σ . This motivates an approach at reactors where the absolute normalization of the observed anti-electron spectrum is left as a free parameter and shape-only analyses are given more weight. This method marginally reduces the significance of the deficit, but is the more conservative approach so is what is considered here. This approach has spurred a new generation




Figure 1.6: Ratio of observed to expected ν_e events from four measurements using MCi ^{51}Cr and ^{37}Ar calibration sources in the SAGE and GALLEX experiments. The 1σ error band is shown [26][27][28][29].

of reactor neutrino experiments with movable detectors and is complimentary to the approach that will be taken by the Short-Baseline Neutrino Program, discussed in Chapter 2.

The final class of SBAs consists of the LSND and MiniBooNE anomalies. This is the apparent excess of electron and anti-electron neutrinos compared to the three-neutrino prediction from accelerator sources. LSND was a liquid scintillator detector that operated at Los Alamos using a stopped pion source (decaying to muons and muon anti-neutrinos with a branching fraction of more than 99.99%) with a mean neutrino energy $\langle E_\nu \rangle$ of about 30 MeV at an L/E_ν of about 1 m/MeV. The significance of this excess (Figure 1.8a) is 3.8σ [31]. This result spurred a followup experiment, MiniBooNE, using a different neutrino source, this time at Fermilab, with a higher energy $\langle E_\nu \rangle \approx 700$ MeV but a similar L/E_ν . MiniBooNE observed a similar excess (Figure 1.8b) with a significance of 4.7σ [32]. A combined analysis of the two data sets increased the significance of the excess to 6σ .

Taken at face value, each class of SBA can individually be explained by adding N sterile neutrinos to the Standard Model, known as $3+N$ models with $N \geq 1$. As discussed





Figure 1.7: Ratio of observed to expected $\bar{\nu}_e$ events by reactor experiments at different baselines. The 1σ error band is shown [30].

in Section 1.1, sterile neutrinos are theoretically motivated. There are other possibilities, but those are beyond the scope of this work.


To get a sense of the situation, we will consider here the simplest possibility of adding one stable, sterile neutrino to the SM, the 3+1 case. We can use the mixing phenomenology introduced in Section 1.3 to describe 3+1 mixing. We now have a 4x4 PMNS matrix, adding 7 new matrix elements, one new mass eigenstate, and one new flavor eigenstate in addition to our standard three-flavor picture.

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \\ \nu_s \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} & U_{e4} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} & U_{\mu4} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} & U_{\tau4} \\ U_{s1} & U_{s2} & U_{s3} & U_{s4} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \\ \nu_4 \end{pmatrix} \quad (1.9)$$

Let us assume the new sterile neutrino mass scale, characterized by a new mass-squared difference Δm_{41}^2 , is much larger than that for the three known neutrinos, i.e. $\Delta m_{41}^2 \gg |\Delta m_{31}^2|, \Delta m_{21}^2$. Two new effective mixing angles, $\theta_{\mu\mu}$ and $\theta_{\mu e}$, can be defined



(a) LSND [31]



(b) MiniBooNE [32]

Figure 1.8: LSND(a) and MiniBooNE (b) observe an excess of electron neutrinos at short-baselines.

in terms of new PMNS matrix elements as shown in Equations 1.10 and 1.11. In the regime of short baselines where $\Delta m_{31}^2 L/E_\nu \ll 1$ and $\Delta m_{21}^2 L/E_\nu \ll 1$, we have, to leading order, the following oscillation probabilities (Equations 1.12 and 1.13). The amplitude of these new oscillations are set by effective mixing angles for muon neutrino disappearance ($\nu_\mu \rightarrow \nu_x$) and electron neutrino appearance ($\nu_\mu \rightarrow \nu_e$) respectively (including their anti-neutrino counterparts). These three new parameters will describe the new sterile neutrino phase space.

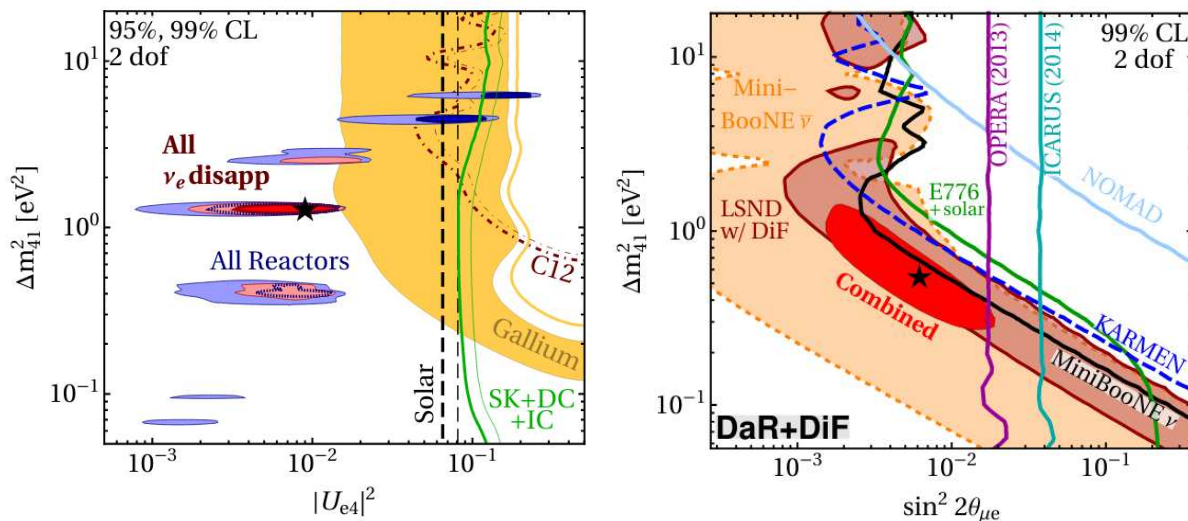
$$\sin^2(2\theta_{\mu\mu}) \equiv 4|U_{\mu 4}|^2(1 - |U_{\mu 4}|^2) \quad (1.10)$$

$$\sin^2(2\theta_{\mu e}) \equiv 4|U_{\mu 4}|^2|U_{e 4}|^2 \quad (1.11)$$

$$P_{\mu\mu}(L, E_\nu) \approx 1 - \sin^2(2\theta_{\mu\mu}) \sin^2\left(\frac{\Delta m_{41}^2 L}{4E_\nu}\right) \quad (1.12)$$

$$P_{\mu e}(L, E_\nu) \approx \sin^2(2\theta_{\mu e}) \sin^2\left(\frac{\Delta m_{41}^2 L}{4E_\nu}\right) \quad (1.13)$$

With this, I now discuss the limits and allowed regions in the 3+1 phase space. I start by examining the results from the electron neutrino disappearance experiments, including the Gallium and Reactor Anomalies. Figure 1.9a shows exclusion limits from electron- ^{12}C scattering, solar neutrino experiments and IceCube (excluding the region to the right). Also shown are the allowed regions from the Gallium (yellow band) and Reactor Anomalies (blue islands). The best fit value is shown by the black star. The 95% and 99% confidence level (CL) regions are shown. There is a slight tension between the gallium and reactor data with a p-value of 9%. However, the 3+1 case is still preferred over the null hypothesis. The reactor data alone (with flux normalization free) prefer the 3+1 scenario at a significance of 2.9σ . Addition of the gallium data increases the significance to 3.2σ . The best fit value for Δm_{41}^2 is about 1 eV^2 .



(a) $|U_{e4}|^2, \Delta m_{41}^2$ space for ν_e disappearance data. (b) $\theta_{\mu e}, \Delta m_{41}^2$ space for ν_e appearance data.

Figure 1.9: Exclusion limits and allowed regions are shown for the short-baseline anomalies [33].

Next, I discuss the results from electron neutrino appearance anomalies observed by LSND and MiniBooNE. The excluded and allowed regions in the $(\theta_{\mu e}, \Delta m_{41}^2)$ plane are

shown in Figure 1.9b. Results from the NOMAD, KARMEN, OPERA, ICARUS, E776, and solar neutrino experiments exclude the region to the right of their respective curves. The solid black curve is an exclusion limit from MiniBooNE neutrino data. The tan region is the MiniBooNE allowed region from anti-neutrino data. The brown region is the LSND allowed region using the full LSND data set (includes decay in flight contaminants rendering a more conservative result). The combined allowed region is the red island with the best fit point shown by the black star. All limits are 99% confidence level. Here, as in the electron neutrino disappearance case, the best fit value for Δm_{41}^2 is about 1 eV^2 . Reiterating from earlier, the significance of this result is 6.1σ .

Finally, I discuss the combination of oscillation results, including electron neutrino disappearance ($\nu_e \rightarrow \nu_x, \bar{\nu}_e \rightarrow \nu_x$), electron neutrino appearance ($\nu_\mu \rightarrow \nu_e, \bar{\nu}_\mu \rightarrow \bar{\nu}_e$) and muon neutrino disappearance ($\nu_\mu \rightarrow \nu_x, \bar{\nu}_\mu \rightarrow \bar{\nu}_x$). This acts to overconstrain the phase space. The results are shown in the $(\theta_{\mu e}, \Delta m_{41}^2)$ plane in Figure 1.10 [33]. The blue contours are exclusion limits from the combined disappearance data sets, excluding the space to the right. The red region is the same allowed region from the previous plot. All contours shown here are at the 3σ confidence level. For the disappearance data, the case for both free and fixed flux normalization is shown. For the appearance data, results for both the decay-at-rest-only LSND data as well as with the decay-in-flight contaminants are shown. There is tension between the disappearance and appearance results with a significance of 4.7σ .

It is important to note that additional stable sterile states do not alleviate the appearance-disappearance tension. It is quite possible that some other BSM physics may be required. In any case, it is clear that a precision short-baseline oscillation experiment is urgently needed to make sense of the current situation. Furthermore, this experiment should investigate both muon neutrino disappearance *and* electron neutrino appearance. Fortunately, such an experiment is about to get underway, the Short-Baseline Neutrino Program, the topic of the next chapter.




Figure 1.10: Exclusion limits from neutrino disappearance data and allowed regions from $\nu_e/\bar{\nu}_e$ appearance data in $\theta_{\mu e}, \Delta m_{41}^2$ space [33].

Chapter 2

The Short-Baseline Neutrino

Program

To address the pressing need for a resolution to the short-baseline neutrino oscillation anomalies discussed in the previous chapter, the Short-Baseline Neutrino Program (SBN) was proposed [34] in 2015. Hosted at Fermilab, SBN will utilize the world’s most intense neutrino beams. It will utilize a precision detection technology that reached maturity over the last decade and continues to advance through dedicated R&D projects around the world, the liquid argon time-projection chamber (LAr TPC), which will be described in Section 2.1. The demonstration of the technology at the hundred-ton scale and the analysis techniques developed by the ICARUS collaboration will be presented in Section 2.2.

A broad overview of SBN and its science goals will be provided in Section 2.3. There, I will introduce the neutrino beams and the three SBN LAr TPC detectors, each an experiment in their own right. Considerations and sensitivities for the sterile neutrino search will be presented in Section 2.4. I will end the chapter with a detailed discussion on cosmogenic backgrounds to the sterile neutrino search as these are of particular relevance for my work.

2.1 Liquid Argon Time-Projection Chambers

The liquid argon time-projection chamber (LAr TPC) was first proposed in 1977 by Carlo Rubbia [35] as an extension of Nygren’s gaseous TPC [36]. The basic working principles

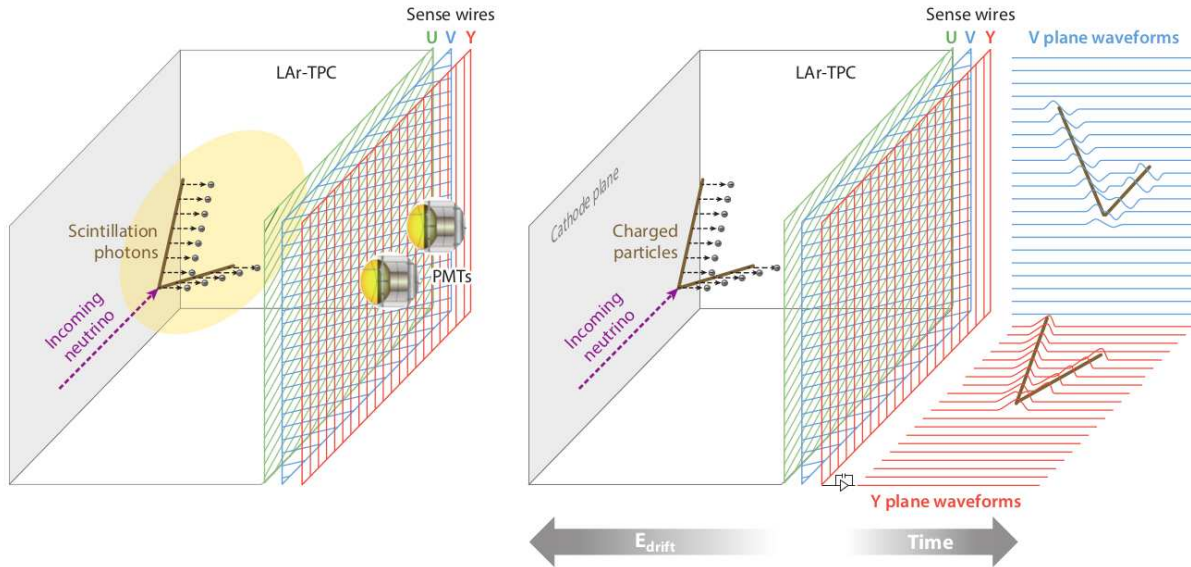


Figure 2.1: An illustration of the basic working principles of a liquid argon time-projection chamber (LAr TPC). On the left, a neutrino traveling into the page interacted inside the volume of liquid argon. The charged particles produced in the interaction locally ionize the argon, generating free electrons and scintillation light. The light propagates to an array of photomultiplier tubes with the flash providing the time of the event. Shown on the right, the ionization electrons drift, over millisecond time scales, left to right in a uniform electric field to a set of three anode wire planes. There, the charge is sampled over time, generating different views of the same event.

of the technology are illustrated in Figure 2.1.

LAr is an ideal medium for neutrino detection as it is relatively dense (1.4 g/mL), is an excellent scintillator ($\sim 50,000$ photons/MeV for a minimum ionizing particle), and ionizes readily with a low recombination probability, allowing it to be used as both a target and detection material. In addition, LAr also has dielectric properties compatible for use in an electric field, comfortably accommodating 500 V/cm. Important for practicality and scalability, argon is relatively cheap as it makes up about 1% of our atmosphere, is nontoxic, and chemically inert.

A hollow box is formed by parallel planes of a cathode and three, closely spaced anode wire planes with the remaining faces covered by field shaping electrodes. This box is filled with LAr, defining the active volume. With appropriate voltage biasing, a static, uniform electric field is generated normal to and between the anode and cathode planes. The

active volume is contained within a cryostat to maintain the liquid argon temperature and pressure. LAr also exists outside of the active volume, contained within the cryostat, forming an inactive volume. While no charge measurements are made in the inactive volume, scintillation light can still be produced and be detected at the photodetectors.

When a sufficiently energetic, electrically charged particle traverses the active volume, it leaves behind a trail of ionization electrons and argon ions which drift in opposite directions along the electric field lines. With a nominal field strength of 500 V/cm, ionization electrons drift with a velocity of 0.16 cm/ μ s [37] toward the anode wire planes. For example, for a 1.5-m drift length, the maximum drift time is 0.91 ms. The first two wire planes encountered by the ionization electrons, the induction wire planes, perform a nondestructive charge measurement while the last plane, the collection plane, collects the electrons for a direct charge measurement. The different wire planes provide complimentary, two-dimensional views of the same event, giving spatial, temporal, and calorimetric information.

In addition to ionization, a throughgoing, charged particle also excites argon atoms, creating argon excimers (leading to a “fast” and a “slow” component with lifetimes of 6 ns and 1.6 μ s respectively) that, when they decay, produce light peaking at 128 nm (in the VUV). LAr is highly transparent to its scintillation light with a Rayleigh scattering length of about 1 m and an absorption length of over 2 m at the peak emission wavelength [37].

Behind the virtually transparent anode wire planes, a plane of photodetectors, in this case photomultiplier tubes (PMTs), captures the scintillation light. With currently available photodetectors, a VUV-to-visible wavelength shifter is required. The most commonly used wavelength shifter for this application is tetraphenyl butadiene (TPB). The propagation time of the scintillation light in LAr is 7.5 ns/m [38]; thus, the fast component of this light provides nanosecond-level time information which enables reconstruction of the particle position along the drift coordinate by providing the start time of the event (t_0).

The number of ionization electrons and the number of excimers produced are anti-

correlated, a fact that allows the optical measurements to improve the calorimetric measurements made with charge alone. The fraction of deposited energy allocated to each process depends on the applied electric field. Larger electric fields yield more ionization electrons and fewer argon excimers.

These features, taken together, facilitate millimeter-resolution three-dimensional imaging and full sampling calorimetry. LAr TPCs are basically electronic bubble chambers. I will soon discuss the powerful techniques these features facilitate that make LAr TPCs a prime candidate for precision neutrino experiments.

The first demonstration of the LAr TPC technology was made by the ICARUS collaboration led by C. Rubbia. Decades of R&D resulted in the first hundred-ton-scale LAr TPC was finally realized in 2001 by the ICARUS collaboration. ICARUS paved the way for future LAr TPC experiments not only by demonstrating the maturity of the technology but also by developing many of the basic simulation and reconstruction tools later utilized and built upon by these experiments. This will be the focus of the next section.

Shortly after the demonstration by ICARUS, the first LAr TPC R&D efforts were launched in the US. From this work came the first LAr TPC in the US to operate in a neutrino beam, the hundred-litre-scale ArgoNeuT [39] detector. The experiment was a success and spurred the use of the technology in a proposed followup experiment to investigate the MiniBooNE anomaly, MicroBooNE [40]. This would be the first ton-scale LAr TPC built and operated in the US. I will discuss MicroBooNE further in Section 2.3.

With the success of the technology and its promise in furthering the field, LAr TPCs were chosen as the basis of the future DUNE [41] long-baseline oscillation experiment. DUNE will use multi-kiloton-scale LAr TPCs for its far detector. Two of these will be advanced versions of the basic LAr TPC design I have discussed here. One of the other two will be a new style of LAr TPC known as a dual-phase LAr TPC; here, an additional gaseous region of Ar between the liquid volume and the anode plane acts to significantly enhance the ionization signal from the LAr. There are pros and cons of both approaches,

but this topic is beyond the scope of this work. The final module, dubbed “the module of opportunity,” may use a novel technology.

The Short-Baseline Neutrino Program will utilize three hundred-ton-scale, single-phase LAr TPCs. The experiment will build on much of the original techniques pioneered by ICARUS.

2.2 ICARUS: The World’s First Large LAr TPC

The ICARUS collaboration, originally based in Italy, designed, built, and operated the world’s first hundred-ton-scale LAr TPC. I will briefly summarize some relevant details of the experiment and the detector here. For more information, see [42]. In order to achieve an active mass of 476 tons of LAr for their detector, the T600, the collaboration opted to divide the volume into two identical “cold vessels”. Each of these T300’s contained two identical drift volumes with a common central cathode. The dimensions of the active volume of each T300 were 18.0 m (L) x 3.0 m (W) x 3.2 m (H). Each of the cold vessels was housed, side by side, in a “warm vessel”. This vessel consisted of a steel structure lined with passive insulation. Between the passive insulation and the T300’s, there was a “cold shield”, an array of aluminum honeycomb planes connected to pipes full of circulating liquid nitrogen, to intercept the heat flux from the outside world and keep the LAr below its boiling point. Figure 2.2 provides an overview of the T600’s major components. Figure 2.3 shows the inside of one of the T300’s TPC volumes with the PMT array, anode wire planes, central cathode, and field cage all visible.

Following construction in Pavia, Italy, a technical run was executed there in 2001 to test the T600 design. Only one of the T300 vessels was filled with high-purity LAr with the other one empty. Following cryogenic commissioning and stabilization, a period of time was required for the LAr to continuously circulate through a purification system that removed any remaining electronegative impurities, primarily water and oxygen. The free




Figure 2.2: An overview of the ICARUS T600 LAr TPC shows the important functional features [42].

electron lifetime increased from about $50 \mu\text{s}$ to 1.8 ms after several days of purification. With this, the detector was able to observe cosmogenic ionization events. During 68 days of data taking, about 30,000 events were collected; the total rate of cosmogenic ionization events in the active volume is at the level of several kHz. The technical run at Pavia was a success and demonstrated the experiment was ready to move into the next phase.

The T600 was moved into an experimental hall at Laboratori Nazionali del Gran Sasso (LNGS) located in L'Aquila, Italy. Here, the detector operated at full capacity from 2010 to 2013 while exposed to the CERN Neutrinos to Gran Sasso (CNGS) high-energy neutrino beam with a mean energy of about 17 GeV . ICARUS observed thousands of ν_μ CC events. In addition, a few atmospheric neutrino events and some ground penetrating cosmogenic muons were observed.




Figure 2.3: Inside one of T300's TPC volumes the field cage, central cathode, PMT array, and anode wires are visible [42].

The LNGS run was a success, and ICARUS had demonstrated that the LAr TPC technology had reached maturity. ICARUS was able to set a limit in the sterile neutrino oscillation phase space (see Section 1.5), and a set of simulation and reconstruction tools were developed for the novel LAr TPC technology. I summarize some of these key results below:

- three-dimensional track [43] and shower reconstruction;
- demonstration of precision calorimetry by reconstructing the π^0 and η^0 invariant masses[44];
- particle identification by measuring ionization density (dE/dx) vs. residual range for fully contained particles [43];
- muon momentum measurement via multiple coulomb scattering [46];
- electron/gamma separation via calorimetric and topological selection [45].

These tools have been the foundation of LAr TPC experiments that followed including ArgoNeuT and MicroBooNE. The technology continues to improve with the growing interest in the technology for neutrino experiments as will be shown in the next section.

2.3 Overview of the Short-Baseline Neutrino Program

The Short-Baseline Neutrino Program (SBN), hosted at Fermilab, will build on the ICARUS experience and the growing LAr TPC community to carry out an exciting science program. The science goals are fourfold:

- test the sterile neutrino hypothesis by covering the LSND allowed region in $(\theta_{\mu e}, \Delta m_{41}^2)$ space at 5σ significance in addition to the globally allowed regions in both $(\theta_{\mu e}, \Delta m_{41}^2)$ and $(\theta_{\mu\mu}, \Delta m_{41}^2)$ spaces at better than 3σ significance;

- explore other BSM physics possibilities for the short baseline anomalies;
- carry out a set of world leading ν -Ar cross-section measurements;
- develop and test new hardware, simulation, reconstruction, and analysis techniques for the next generation of LAr TPC based experiments.

Besides having a significant impact on the broader neutrino field, SBN’s results will be of particular importance for the future DUNE experiment.

Originally, SBN consisted of three LAr TPCs located along the axis of the Booster Neutrino Beam (BNB): a near detector (SBND), an intermediate detector (MicroBooNE), and a far detector (ICARUS). All three detectors, discussed in more detail later in this section, are functionally identical, which is a crucial feature for the suppression of systematic uncertainties in the sterile neutrino search. In addition, the intermediate and far detectors will also be exposed to an off-axis component of the Neutrinos from the Main Injector (NuMI) beam. Due to some scheduling delays, only SBND and ICARUS will operate concurrently for oscillation data taking. For an aerial view of SBN, see Figure 2.4. The combination of precision LAr TPC technology with the world’s highest intensity neutrino sources will enable SBN’s world-leading measurements.

Fermilab hosts the world’s most intense neutrino sources. These sources’ significance to SBN’s science goals cannot be overstated. I will provide an overview here of the basic concepts behind generating a neutrino beam before going into more detail about the two different neutrino beams at Fermilab.

The conventional neutrino beam begins as a source of protons, typically accelerated to a few to tens of GeV energies. The protons impinge on a target, usually graphite or beryllium, generating a shower of mesons. The mesons, mostly pions and kaons, pass through a magnetic lens or “horn” that deflects the charged mesons toward or away from the beam axis, depending on the charge sign. This sign can be chosen by the direction of current in the magnetic horn, an important feature as it facilitates the production of a

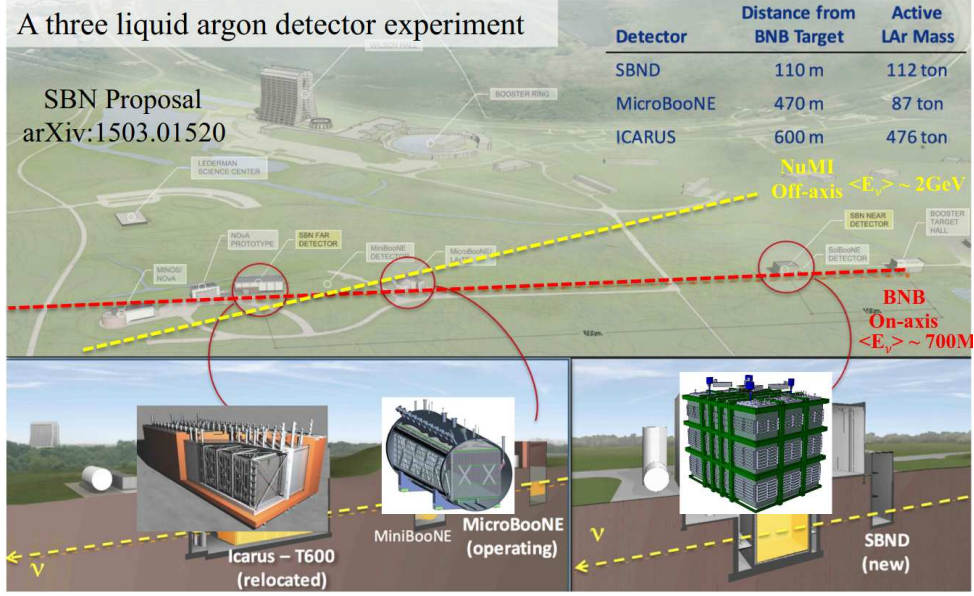


Figure 2.4: An aerial view of the Short-Baseline Neutrino Program hosted at Fermilab. The three detectors are shown along with the BNB beam axis. In addition, the intermediate and far detectors will be exposed to an off-axis component of the NuMI beam.

neutrino or anti-neutrino beam. The mesons proceed down a decay pipe, filled with air at atmospheric pressure, where they decay in flight into charged lepton-neutrino pairs. The branching ratio for charged pions decaying into neutrinos is $\text{BR}(\pi^+ \rightarrow \mu^+ \nu_\mu) = \text{BR}(\pi^- \rightarrow \mu^- \bar{\nu}_\mu) > 99.99\%$ while for kaons it is $\text{BR}(K^+ \rightarrow \mu^+ \nu_\mu) = \text{BR}(K^- \rightarrow \mu^- \bar{\nu}_\mu) \approx 64\%$ with the remaining modes being semi-leptonic or hadronic, dominated by pions. These non-leptonic kaon decay modes introduce wrong-sign muon neutrino and electron/anti-electron neutrino contamination in the neutrino beam. At the end of the decay pipe is an absorber, usually concrete and/or steel. The absorber, aptly named, absorbs the charged particles leaving only the neutrinos to propagate down the beam line. This process is illustrated in Figure 2.5 for the case of the BNB.

For the sterile neutrino search, SBN will use the BNB. For the cross-sections program, both the BNB and NuMI beams will be used. Only MicroBooNE and ICARUS will be exposed to a significant flux from the NuMI beam however. I summarize relevant beam parameters in Table 2.1. The remainder of this chapter will be focused on the BNB as it pertains to the sterile neutrino search.




Figure 2.5: Illustration of neutrino beam production for Fermilab’s Booster Neutrino Beam. Note that the subdominant kaon production and resulting neutrino beam contaminants are not shown.

Table 2.1: Selected Fermilab neutrino beam parameters.

Beam	Mean Energy [GeV]	Spill Rate [Hz]	Spill Width [μ s]	Bunch Width [ns]	Bunches/Spill
BNB	0.7	5.00	1.6	2	81
NuMI	2.0	0.75	9.5	2	486

One of the more difficult parameters to calculate in neutrino experiments is the neutrino flux, reflected in a large associated systematic uncertainty. Due to significant efforts by the MiniBooNE collaboration [47], hadron production measurements from the HARP experiment [48], and the stable running configuration of the BNB over the last twelve years or so, the systematics uncertainties of the BNB flux prediction at each SBN detector are relatively low. The flux prediction at each detector is presented in the SBN proposal [34] but reproduced in Figure 2.6 for convenience. The BNB flux prediction is shown for “forward horn current” running, which produces a predominantly ν_μ beam. Also shown are the wrong-sign ($\bar{\nu}_\mu$) and electron neutrino ($\nu_e/\bar{\nu}_e$) contamination. The




Figure 2.6: SBN Booster Neutrino Beam flux predictions at the three detector Top row, left to right: SBND, MicroBooNE, ICARUS. Bottom row: ICARUS-to-SBND flux ratio (left) and MicroBooNE-to-ICARUS flux ratio (right). Reproduced from [34].

systematic uncertainty at the peak is about 9% and increases away from the peak. The flux is slightly harder at MicroBooNE and ICARUS due to a narrower acceptance. The ratio of total integrated flux at ICARUS to SBND is roughly 0.25, shown in the bottom left plot in Figure 2.6, where the acceptance mismatch is evident. In the bottom right of the same figure, the ratio of fluxes at MicroBooNE to ICARUS, about 0.6, shows that the two detectors are in the r^{-2} regime.

I now provide a brief overview of the detectors including relevant parameters, summarized in Table 2.2, and schedule information.

The near detector, located 110 m from the BNB target, is called the Short-Baseline Near Detector (SBND). SBND consists of an 112-ton active LAr mass contained in a single cryostat housing two TPC volumes with a common central cathode. It is new construction featuring several important R&D components relevant for the future DUNE experiment. The cryostat is the same design used by ProtoDUNE, building on synergies between Fermilab and the CERN Neutrino Platform. The light detection system consists of conventional PMTs as well as new photodetectors based on silicon photomultipliers

Table 2.2: SBN detector values of interest.

Detector	Dist. [m]	Total LAr [ton]	Active LAr [ton]	No. TPCs	Drift Dist. [m]
SBND	110	260	112	2	1.5
μ BooNE	470	170	87	1	2.5
ICARUS	600	760	476	4	1.5




Figure 2.7: Inside one of the ICARUS T300's while being refurbished at CERN.

known as the ARAPUCA and X-ARAPUCA. It also features a VUV-to-visible wavelength shifting cathode mesh that renders the effective response of the photon detection system significantly more uniform. The detector is in the assembly phase with installation at the SBND building expected to start in 2021.

The intermediate-distance detector, located 470 m from the BNB target, is Micro-BooNE. The detector is the first hundred-ton-scale LAr TPC built in the US. It consists of a single TPC volume containing 87 tons of LAr. It was under construction at the time of writing of the SBN proposal and began taking data in Fall of 2015. It completed data taking in 2020. The experiment has made significant contributions in the development of simulation and reconstruction tools for LAr TPCs in addition to several important cross-section measurements.

The far detector, located 600 m from the BNB target is the ICARUS T600, adding

significant expertise to SBN. ICARUS will be the focus for the remainder of this dissertation. Following the successful run at LNGS, the T600 was transported to CERN for refurbishing, which took place over 2016-2017. Figure 2.7 shows some of our collaborators in action. The improvements are listed below:

- new Photon Detection System (PDS) with improved photo-coverage [49];
- new, compact anode wire biasing and readout electronics [50];
- improved flatness of the cathode plane;
- new cryostats;
- new warm vessel;
- upgraded cryogenics system.



Figure 2.8: ICARUS cold vessels being installed into the warm vessel at Fermilab.

After refurbishing, the T600 was transported to Fermilab, arriving in 2018. In 2019, the cold vessels (filled with air only) were installed into the warm vessel at the experimental hall (Figure 2.8). A flurry of installation activities commenced shortly thereafter. I was

fortunate to be present at Fermilab while these activities were ongoing and was one of the first people to be able to go into the cold vessels. Specifically, I reached my arm through the feedthrough ports on top of the cold vessels in order to dismount the anode plane analog signal cables as well as the PMT signal and high voltage cables from the inner frame and pull the cables through a second feedthrough flange that we installed as shown in Figure 2.9 (to bridge the gap from the cold vessel to the outside of the warm vessel). Following this, I was part of a major effort to test all of these cables for continuity. After the warm vessel roof was installed, the last part of the feedthrough was installed and connected to feedthrough flanges. Following this step, the cables all had to be tested again, except this time it was from the outside world. I helped with this effort as well.




Figure 2.9: Installing feedthrough flanges on the T600.

With a lot of work from a lot of people, the detector was fully equipped with cryogenics, electronics racks, power supplies, front-end electronics, and more. Immediately before the start of cold commissioning, the detector appeared as in Figure 2.10. Cold commissioning,

including cooldown of the cryostats, filling of the T300's with LAr, stabilization, and start of LAr recirculation and purification, was successfully completed in May 2020. The first ionization events in the active LAr volume are expected to be observed in early Fall 2020.

We have seen that SBN is equipped with cutting edge, precision neutrino detectors ready to observe neutrino interactions from the world's most intense neutrino sources. This dissertation is being written at an exciting but uncertain time with full physics running about to commence but with inevitable delays due to the COVID-19 pandemic. With the introduction of these basic components, I turn my attention to the sterile neutrino search at SBN.

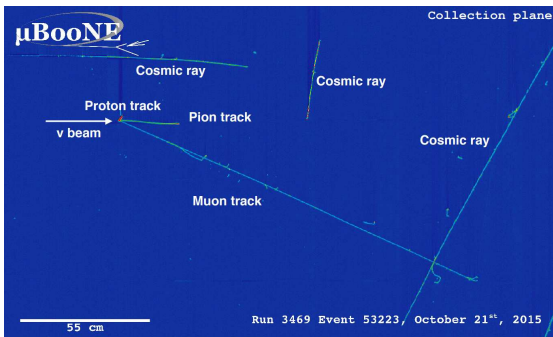


Figure 2.10: ICARUS just before the start of cold commissioning.


2.4 Sterile Neutrino Search

The flagship measurement of SBN is the confirmation or the refutation of the sterile neutrino hypothesis, bringing clarity to the short-baseline anomalies discussed in Section

1.5. To achieve this, the SBN program must search for short-baseline oscillations covering the entire allowed region of the sterile neutrino oscillation phase space with better than 3σ significance. With the most significant contribution to the allowed phase space coming from the LSND ν_e appearance anomaly, SBN should cover the entire LSND region with 5σ significance. Furthermore, in order to resolve the disappearance/appearance tension, SBN should perform the search in both ν_e appearance ($\nu_\mu \rightarrow \nu_e$) and ν_μ disappearance ($\nu_\mu \rightarrow \nu_x$) channels. We must be careful in this measurement to control for systematics so that we avoid generating yet another anomaly. In this section, I will discuss the systematics under consideration for the sterile neutrino search as well as backgrounds that must be mitigated. I will then show the predicted spectra and sensitivities.



(a) ν_μ CC event from MicroBooNE with a proton, charged pion, and muon in the final state [51]. Cosmic rays are also visible.



(b) ν_e CC event from ArgoNeuT with a charged pion and electron in the final state. The gaps in the electron track are due to dead wires [52].

Figure 2.11: ν CC candidates from data are shown. The color indicates the ionization density with blue to red being low to high. In both panels, the beam is going from left to right.

For the purposes of measuring event rates, each detector will reconstruct a neutrino energy distribution for each inclusive search channel. ν_μ CC events ($\nu_\mu n \rightarrow \mu^- X$; $\bar{\nu}_\mu p \rightarrow \mu^+ X$) are identified by the presence of a muon in the final state, visible as a tens-of-cm- to few-meters-long, minimum ionizing track (Figure 2.11a). ν_e CC events ($\nu_e n \rightarrow e^- X$; $\bar{\nu}_e p \rightarrow e^+ X$) are identified by an electron in the final state that will generally generate an electromagnetic (EM) shower (Figure 2.11b).

As discussed in Section 1.3, the usual quantities an oscillation experiment must understand in order to measure oscillation probabilities include the neutrino flux, neutrino cross-section, and signal selection efficiencies. A convolution of these parameters determines the signal event rates in the detector. In addition to signal events, there are also backgrounds. Taking all of these details into account, the following approach is taken for calculating the sensitivity for different values of the new sterile neutrino mass-squared difference, Δm_{41}^2 , and the effective sterile neutrino mixing angles, $\theta_{\mu e}$ and $\theta_{\mu\mu}$. A distribution of reconstructed neutrino candidate energies over 11 bins spanning 200 MeV through 3 GeV is considered for the analysis. In the i^{th} energy bin, there is the number of events from the sterile neutrino induced oscillations for a particular choice of oscillation parameters, N_i^{osc} , and the number of events in the absence of oscillations (the null hypothesis), N_i^{null} . All statistical and systematic uncertainties are accounted for in the total covariance matrix, E^{total} . With these parameters, we can compute a χ^2 surface in the $(\sin^2 2\theta, \Delta m_{41}^2)$ plane as shown in Equation 2.1 for each effective sterile mixing angle. From this surface, we can determine our oscillation search sensitivity contours in terms of confidence levels defined by the deviation of the χ^2 relative to the global minimum value.

$$\chi^2(\sin^2 2\theta, \Delta m_{41}^2) = \sum_{i,j} [N_i^{\text{null}} - N_i^{\text{osc}}(\sin^2 2\theta, \Delta m_{41}^2)] (E_{ij}^{\text{total}})^{-1} [N_j^{\text{null}} - N_j^{\text{osc}}(\sin^2 2\theta, \Delta m_{41}^2)] \quad (2.1)$$

The total covariance matrix can be separated into statistical and systematic parts, E^{stat} and E^{syst} respectively, so that we have $E^{\text{total}} = E^{\text{stat}} + E^{\text{syst}}$. E^{stat} is the completely uncorrelated (all off-diagonal elements are zero-valued) statistical error matrix with $E_{ii}^{\text{stat}} = N_i^{\text{null}}$. The elements of the systematic covariance matrix is calculated using reweighting techniques using \mathcal{N} universes as shown in Equation 2.2. N_{CV}^i is the number of entries in the i^{th} energy bin of the nominal event distribution, and N_m^i is the number of entries in the i^{th} energy bin of the m^{th} universe. i and j correspond to the neutrino

energy bins across all three SBN detectors.

$$E_{ij} = \sum_{m=1}^{\mathcal{N}} [N_{CV}^i - N_m^i] \times [N_{CV}^j - N_m^j] \quad (2.2)$$

More useful forms of the covariance matrix are the fractional error matrix F , defined in Equation 2.3, and the correlation matrix ρ , defined in Equation 2.4.

$$F_{ij} = \frac{E_{ij}}{N_{CV}^i N_{CV}^j} \quad (2.3)$$

$$\rho_{ij} = \frac{E_{ij}}{\sqrt{E_{ii}} \sqrt{E_{jj}}} \quad (2.4)$$

The total systematic covariance matrix can be broken down into several independent components that account for the signal event rate mentioned above as well as different sources of background. These are listed, in no particular order, in Equation 2.5. I discuss each briefly below.

$$E^{total} = E^{flux} + E^{cross-section} + E^{detector} + E^{dirt} + E^{cosmogenic} \quad (2.5)$$

Important for suppression of systematic errors arising from the flux prediction are the correlations in the flux bins between detectors. The flux covariance matrix is shown in Figure 2.12, reproduced from the SBN proposal, including contributions from both the ν_μ and ν_e beam components. There are two representations shown. The fractional error matrix is shown on the left. The flux correlation matrix is shown on the right. For both, there are three rows and columns one for each SBN detector. Within each row or column, eleven bins span 0.2 to 3 GeV in neutrino energy. There is strong correlation between the flux energy bins between the detectors. Along the diagonal boxes, you can see the correlations within each detector. In the off-diagonal cells, the correlations between the flux bins between detectors can be seen. Notice that the correlation is stronger for

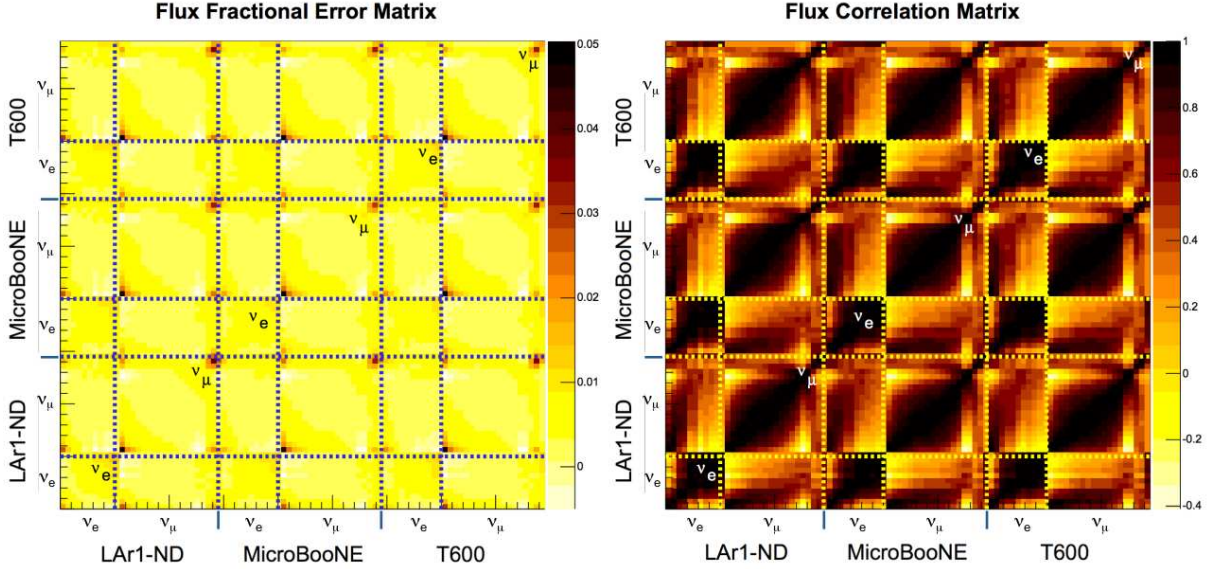


Figure 2.12: Two representations of the SBN flux covariance matrix are shown with the fractional error matrix on the left and the correlation matrix on the right. There are three rows and columns, one for each SBN detector. Within each row and column, there are sub rows or columns for the individual ν_μ and ν_e components. The 11 bins span 0.2 to 3 GeV in neutrino energy. Reproduced from [34].

the intermediate and far detectors than for either of these with the near detector. This is simply due to different acceptance between the two groups. Due to this high level of correlation, there is a suppression of the fractional error, mostly sub-percent level, compared to the approximate 6% minimum (at the flux peak) in the prediction at each detector individually. This is crucial for the overall SBN sensitivity.

The event rate at each detector is determined by the product of the flux and the cross-section. Simulation results presented in the SBN proposal indicate a 10-15% absolute uncertainty (Figure 2.13, top) in the event rate due to cross-section uncertainties alone in each individual detector. The power of multiple, functionally identical detectors can be seen in Figure 2.13 (bottom). We see two representations of the cross-section covariance matrix with the fractional error matrix on the left and the correlation matrix on the right. In the error matrix, we see the significant suppression in uncertainties due to the high level of correlation (nearly 1.0) shown in the correlation matrix. As with the flux correlations, the interaction model correlations are a crucial feature for the overall SBN sensitivity.

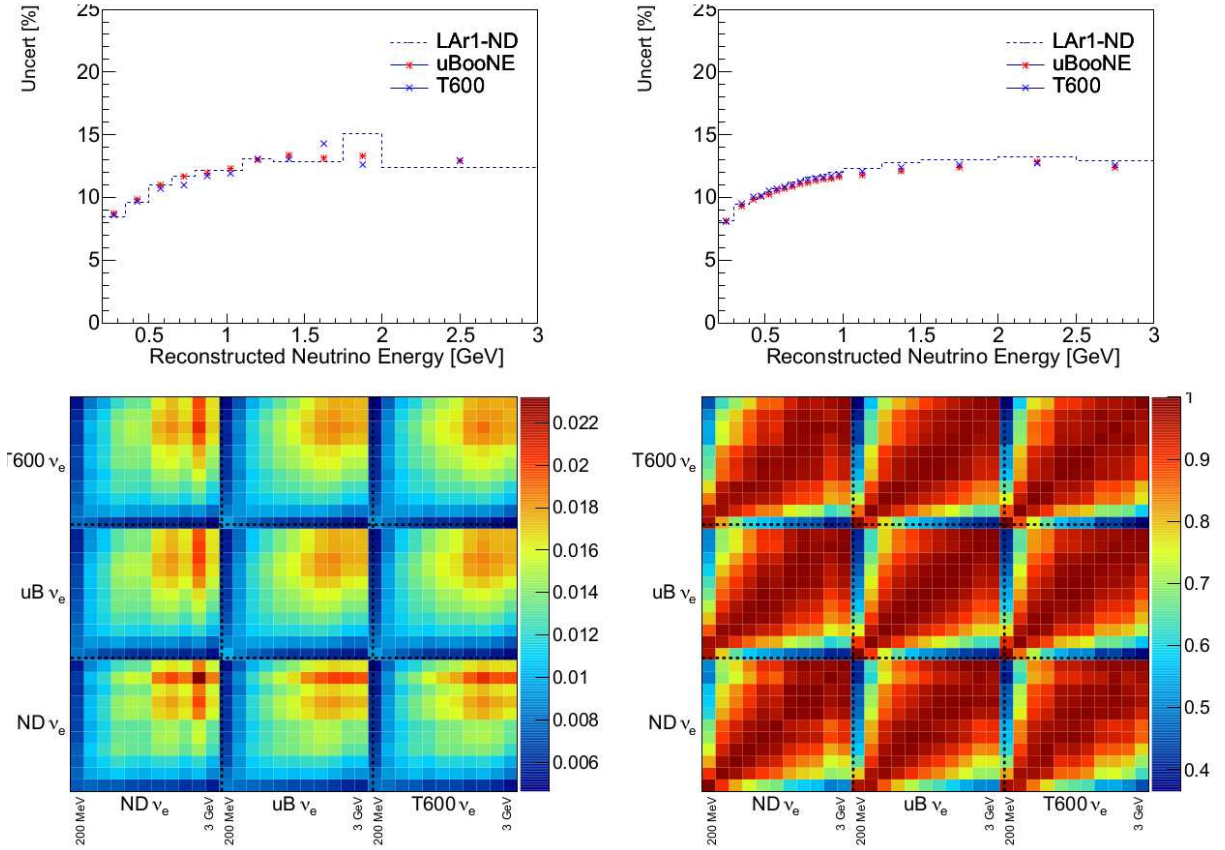


Figure 2.13: Top: absolute ν_e (left) and ν_μ (right) event rate uncertainties for the three SBN detectors due to cross-section uncertainties alone. Note that SBND was previously known as LAr1-ND. Bottom: cross-section covariance matrix for ν_e interactions spanning 11 bins over 0.2-3 GeV shown in two representations: fractional error matrix (left) and correlation matrix (right). Reproduced from [34].

With functionally identical detectors, the detector response modeling essentially cancels in the far-to-near and intermediate-to-near rate ratios. However, there are small differences between the detectors, e.g. anode wire orientation, photon detection system, etc. These effects can be well measured with data and corrected. Preliminary studies suggested that a detector systematic uncertainty in the range of 2-3% would preserve the required sterile search sensitivity and was used as a design requirement.

There are three sources of background: intrinsic wrong-sign and electron neutrino beam contamination; secondaries produced by beam induced neutrino interactions in materials surrounding the LAr volume; cosmogenic interactions in and around the LAr

volume.

The beam intrinsic background (Figure 2.6) is irreducible. However, the contribution to the overall event rate can be well constrained by the near detector.

“Dirt” events are any secondaries produced by the neutrino beam interacting with materials outside of the active LAr that go on to interact in the active LAr volume and mimic a ν_μ CC or a ν_e CC interaction. Preliminary studies have shown that this background is negligible in the case of ν_μ CC interactions. However, for the ν_e CC case, dirt photons, from π^0 production and subsequent decays for example, can propagate into the active volume and induce an EM shower. This background is challenging to simulate as it requires simulating a large number of neutrino interactions over a large volume. Any study performed for SBND cannot be extrapolated to the far detector due to different acceptance effects. However, a large Monte Carlo production was carried out for MicroBooNE, the intermediate detector, which already had a detailed geometric description available. Since both detectors are in the r^{-2} regime, the results for MicroBooNE can be extrapolated to ICARUS. Of course, such a study will need to be done for ICARUS as the geometries are not identical. As of this writing, this study has not yet been carried out.

There are a couple of straightforward methods for reducing the dirt background. First, we can apply the gamma/electron shower discrimination via dE/dx (Section 2.2). In the analysis for the SBN proposal, the implementation of this method rejected 94% of pair-producing gammas. Second, thanks to the relatively short radiation length ($X_0=14$ cm) in LAr, the outer shell of LAr is effective in shielding the inner LAr volume. For example, defining a fiducial volume with faces about two radiation lengths from the outer active volume faces rejects an additional 80% of the dirt background.

The cosmogenic background poses a significant challenge for all SBN detectors since they operate at shallow depths. This is especially true for ν_e CC events. As this background is at the core of my work, I will devote the next section to a detailed discussion

on this topic. For the purposes of this section, the same 94% rejection factor from the dE/dx cut as used for the dirt events has been applied. In addition, an external tracking system is assumed to be present around the LAr volume that is capable of rejecting 95% cosmogenic backgrounds. I discuss this external tracking system in Chapter 3.

With this, I can now present the projected sterile neutrino search sensitivity for SBN in both appearance and disappearance modes. The projected sensitivities include all of the statistical and systematic uncertainties as well as all of the backgrounds discussed above. The planned exposure is 6.6×10^{20} protons on target (POT). You will find projected sensitivities in the SBN proposal that show the LSND 99% C.L. and LSND/MiniBooNE combined 3σ allowed regions for ν_e appearance are covered by SBN at 5σ significance. The entire 90% C.L. exclusion limit set by ν_μ disappearance in MiniBooNE combined with SciBooNE are also covered at 5σ . In the years since the SBN proposal was written, new data has become available and the 3+1 sterile phase space has become tightly constrained.

The sensitivities presented in the SBN proposal have been reproduced by [53] with comparisons to the most recent global analyses (presented in Section 1.5), shown in Figure 2.14. In the plot on the left, for ν_e appearance, the 5σ (red dashed) and 3σ (red solid) limits are shown. The solid yellow regions are the global allowed region with all global oscillation data included. It is covered by SBN at 3σ significance with a large fraction of the region covered at 5σ significance. In the plot on the right, for ν_μ disappearance, the 3σ exclusion limit set by ν_μ disappearance is shown. This limit has become stronger since the SBN proposal, however SBN covers nearly this entire limit at 5σ significance. As in the appearance case, the entire globally allowed region, with all oscillation data included, is covered by SBN at 3σ significance with most of this region covered at 5σ significance.

As shown by the projected sensitivities, SBN is well positioned to make a definitive determination on the 3+1 sterile neutrino hypothesis, thereby bringing some needed clarity to the short-baseline anomalies. While I have focused exclusively on the 3+1 case here, SBN will have similar sensitivities for the 3+N case. Of course, other BSM physics could

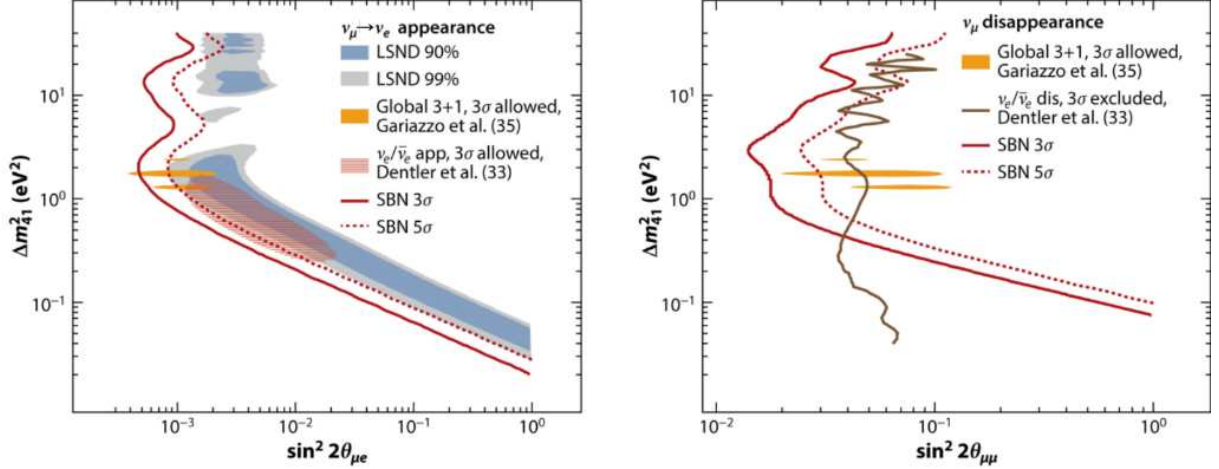


Figure 2.14: SBN sterile neutrino (3+1 model) search sensitivity in appearance (left) and disappearance (right) modes. Sensitivities shown here are reproduced from [53].

be the cause the short-baseline anomalies. SBN will be well positioned to study these as well. See [53] for further discussion on this point. To make these sensitivities possible, I will now turn my attention to mitigating the cosmogenic background.

2.5 Cosmogenic Background

Primary cosmogenic backgrounds originate from the cosmogenic surface flux, before the particles have a chance to interact with something besides air (i.e. the building, the TPC, etc.). The primary components can be categorized as follows: charged hadrons, neutrons, gammas (high-energy photons), electrons, and muons. Figure 2.15 shows the primary cosmogenic flux distribution by particle species scaled by energy. Figure 2.16 shows a data sample of cosmogenic events in one T300 detector during the surface run in Pavia where no overburden was present. A preliminary study performed in preparation for the SBN proposal found that 3 m of concrete overburden would provide a factor of about 400 in suppression for charged hadrons, primary gammas, and primary electrons. The dominant background remaining comes from muons and associated production. Unfortunately, MicroBooNE was not able to obtain overburden. However, both SBND and




Figure 2.15: Cosmogenic energy spectrum just above the SBN experimental buildings for different particle species [34].

ICARUS will install the 3 m of concrete overburden as proposed. Thus, for the remainder of this (ICARUS-centric) dissertation, I will consider only this operating condition.

With overburden, the focus of cosmogenic background mitigation efforts will be on muons and associated production. For the ν_μ disappearance measurements, most cosmogenic muons can be unambiguously identified and removed from the analysis. The rate of ν_μ CC interactions in all SBN detectors is sufficiently large that there is some tolerance in the harshness of event selection cuts without appreciable loss in signal sample size. This is not the case for the golden channel of ν_e appearance, however, where the expected number of signal events ranges from zero to a few hundred depending on what nature has in store.

The most important cosmogenic background for the ν_e appearance channel comes from secondary gammas produced by muons passing through the LAr volume itself or through




Figure 2.16: Event from the ICARUS technical surface run at Pavia without overburden. The greyscale indicates the level of ionization density. Several cosmogenic muon tracks are visible. Also shown are few-GeV CNGS ν_e CC events for scale.




Figure 2.17: A comparison of energies of cosmogenic gammas that convert in the active LAr volume (left) and the electrons produced in BNB ν_e CC interactions show that most gammas are below the analysis energy threshold. However, there are a significant number of cases where the gamma energy can be problematic.

nearby materials. In Figure 2.17, I show a comparison of the energies of these secondary gammas compared to the energy spectrum of electrons produced by BNB intrinsic ν_e CC interactions in the active volume. At these energies, the dominant photon interactions are pair production, leading to an electromagnetic (EM) shower, and Compton scattering, which can produce an electron with sufficient energy to generate an EM shower.

It is useful to define two timing categories and two event topologies:

- timing category A: cosmogenic photon interacts anywhere in the active volume during the beam spill;

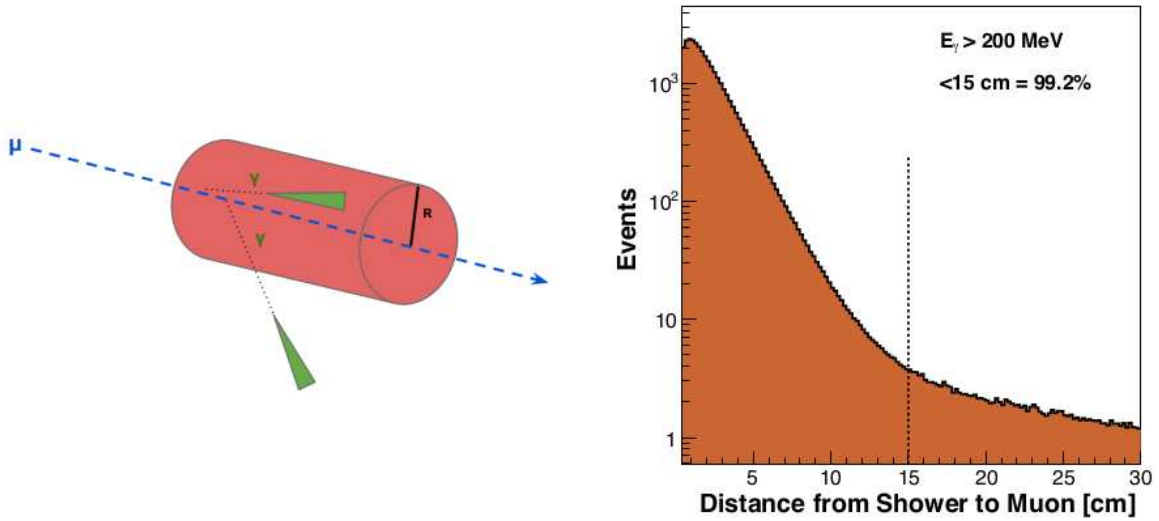


Figure 2.18: One way to cut out secondary photons generated by muons passing through the active LAr volume is to exclude the volume contained in a cylinder centered on the muon track. The plot on the right is a histogram of photon conversion distance from the parent muon track. For cylinders with a 15-cm radius, 99% of secondary photons are excluded with a loss of about 1% in fiducial volume. Reproduced from [34].

- timing category B: cosmogenic photon interacts anywhere in the active volume during the drift time, and a different cosmogenic particle interacts in the active volume during the beam spill (the photon could be mistaken for the in-spill event);
- topology category I: secondary cosmogenic photon interacts in the fiducial volume, and the parent muon enters the active volume;
- topology category II: cosmogenic photon interacts in the fiducial volume and is a primary, the parent is not visible (e.g. a neutron), or the parent does not enter the active volume.

Given the importance of the ν_e appearance channel, my primary focus will be on the cosmogenic background directly relevant to it. There are several methods outlined in the SBN proposal for mitigating this background. The most important of these will be reproduced here with some added commentary.

- Beam timing: Using nanosecond-level timing information provided by the Photon

Detection System and the early-warning signal from the accelerator, any activity outside of some trigger window centered on the beam spill can be rejected. Within the beam spill, the spill structure, presented in Section 2.3, can be exploited to reject up to 90% of cosmogenic events.

- **dE/dx:** Using precise spatial and calorimetric information, gamma- and electron-induced EM showers can be distinguished. A demonstration of this technique was presented in Section 2.2 where ICARUS (@LNGS) was able to reject π^0 backgrounds from the ν_e CC signal. In one particular application of this technique used in the simulation for the SBN proposal, about 94% of pair-producing gammas were rejected. Of course, this is ineffective against Compton electrons.
- **Distance to the muon track:** In cases where a muon track is present in the active LAr volume, simulation studies show that 99% of all gammas produced by the muon convert within 15 cm of the track (Figure 2.18). This motivates a strategy where a cylinder with a 15-cm radius centered on the muon track is used to exclude a portion of the active volume. This results in a loss of about 1% of the fiducial volume.
- **Gap cut with hadronic activity at the vertex:** If hadronic activity is visible at the neutrino vertex, the mm-level spatial resolution can be exploited to check for a gap between the vertex and the beginning of the candidate EM shower. This is highly effective in rejecting gammas as they usually propagate a few cm before interacting while electrons begin to ionize the LAr immediately. However, preliminary studies suggest this cut would discard about 25% of ν_e events, so it is not considered here.
- **Muon tagging:** Gammas produced by muons that miss the active LAr volume can propagate into the active LAr volume, effectively rendering selection cuts based solely on topological information ineffective. A secondary detector that surrounds the cryostat capable of tagging throughgoing muons with high efficiency (e.g. 95%)

Table 2.3: Cosmogenic photon rates in the SBN Far Detector (ICARUS) without (all events) and with dE/dx and distance from the muon cuts (events w/TPC cuts) with 6.6×10^{20} POT. Reproduced from [34].

Interaction Description	Timing Cat.	Topology Cat.	All Events	Events w/TPC Cuts
γ Compt. in spill, primary μ enters AV	A	I	599	<4
γ pair prod. in spill, primary μ enters AV	A	I	32,000	21
γ Compt in spill, primary μ misses AV	A	II	<4	<4
γ pair prod. in spill, primary μ misses AV	A	II	11	<1
γ Compt. in drift, primary μ enters AV	B	I	3300	30
γ pair prod. in drift, primary μ enters AV	B	I	176,000	113
γ Compt in drift, primary μ misses AV	B	II	<4	<4
γ pair prod. in drift, primary μ misses AV	B	II	60	<4
Total cosmogenic backgrounds			211,970	164

can address this problem. In addition, the supplemental spatial and temporal information provided by such a system would enhance the TPC based methods.

The TPC based methods are powerful. Using the dE/dx and the distance-to-the-muon cuts, rejection power shown in Table 2.3 is achieved. There, the event rates for the different timing and topology categories listed above are shown separated by photon interaction type (i.e. pair production or Compton scattering) assuming the nominal 6.6×10^{20} protons on target delivered in 1.32×10^8 beam spills amounting to 211 seconds of beam time. These rates come from a truth-level analysis based on FLUKA. The TPC based methods together reject 99.9% of cosmogenic photons.

In the time since the SBN proposal, the TPC based methods discussed above have been studied with more sophisticated simulations. So far, results are consistent with prior projections. A brief summary of recent results on this topic will be given in Chapter 4. Concerning the muon tagging option, some limited investigations were made with simple simulations for and shortly after the proposal that demonstrated the potential of the muon tagging scheme. Assuming a full coverage (4π) Cosmic Ray Tagging system (CRT) with 95% tagging efficiency, a 99% rejection on all cosmogenic triggers can be obtained, defined

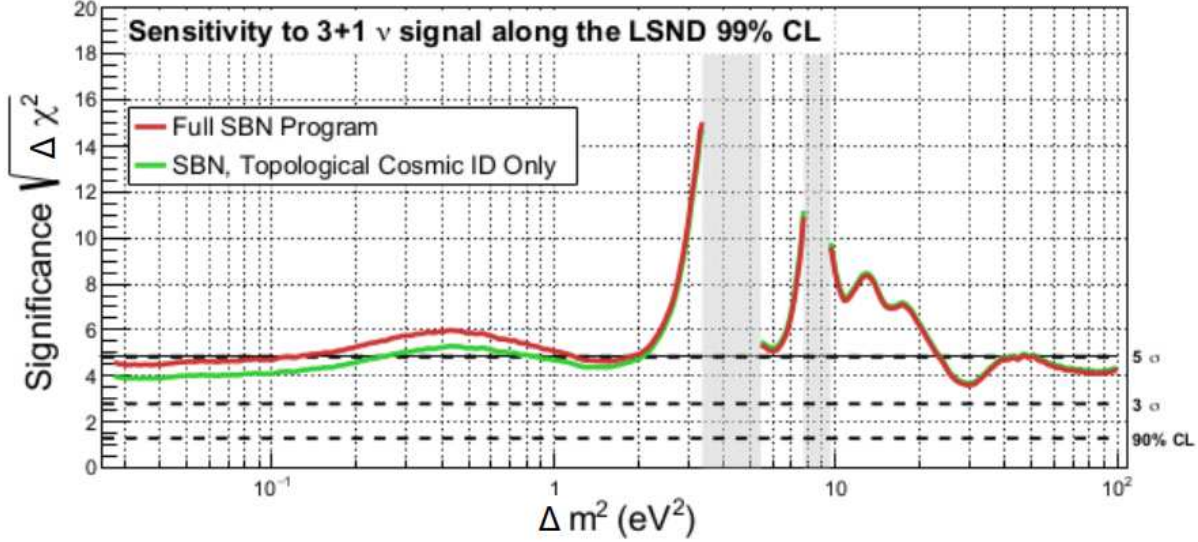


Figure 2.19: Impact of muon tagging on the SBN sterile neutrino search in the ν_e appearance channel shown as the projected SBN significance along the LSND 99% C.L. contour. Muon tagging enhances the sensitivity by about 1σ significance blow sterile mass splittings of about 1 eV^2 . Reproduced from [34].

by any event with more than 200 MeV deposited in the active volume in coincidence with the beam spill triggering a readout of the TPC. With our nominal beam exposure, we expect 2.5×10^6 cosmogenic triggers. It was shown that the CRT can reduce this number to 2.4×10^4 cosmogenic triggers.

For the sterile neutrino search sensitivities (Figure 2.14), all candidate neutrino events with a CRT tag in time with the beam spill were rejected. This, combined with the dE/dx and distance-from-the-muon cuts, brought the number of cosmogenic background events during the nominal beam exposure to less than one. This is to be compared to the ~ 150 events left over using TPC based methods alone. The impact of this CRT enhanced rejection power is shown in Figure 2.19 where the sensitivity for the SBN sterile neutrino search, projected along the LSND 99% C.L. contour, is compared in the cases of TPC based methods only and the addition of the CRT. For Δm_{41}^2 below about 1 eV^2 , the CRT provides an additional $\sim 1\sigma$ significance.

These results were sufficiently promising to motivate R&D studies on possible hardware implementations. Eventually, Cosmic Ray Tagger (CRT) systems were constructed

for each SBN detector. A thorough discussion on the CRT hardware development, installation, and commissioning will be given in the next chapter.

There are different possibilities for the application of the additional information from the CRT. The one used in the SBN proposal is the most straightforward. However, this method comes at a cost. Even in beam spills with no cosmogenic activity, if particles from the neutrino interaction produces a tag in the CRT, the event will be rejected. I revisit this point later in Chapter 6 and provide alternative approaches considering background rejection power as well as the impact on the neutrino signal.

Chapter 3

The Cosmic Ray Tagging System

To cope with the cosmogenic background discussed in Section 2.5, all three of the SBN detectors will be equipped with a Cosmic Ray Tagging System (CRT). The primary function of the CRT is to tag muons passing through or near the cryostats. To be worthwhile, the muon tagging efficiency must be about 95% or better in order to exceed the rejection power possible in the absence of a CRT (see Section 2.5). To enhance the existing rejection power of the TPC, the CRT should be able to unambiguously match tagged muon tracks to those reconstructed in the TPC active volume. The poorest spatial resolution in the TPC is along the drift direction, limited by the time resolution of the Photon Detection System. The drift coordinate, x , is reconstructed as $x_{reco} = x_{true} + v_{drift}(t_0 - t_{true})$, depending on the electron drift speed, v_{drift} and the absolute trigger time, t_0 . With an expected PMT time resolution of about 1 ns, the maximum useful spatial resolution for the CRT is about 30 cm. In addition, the CRT time resolution should be at the nanosecond-level in order to enhance beam spill structure exploitation. These features must be obtained without introducing false positives, from ambient, low-energy radiogenic activity for example. While important, the CRT cannot be allocated a large fraction of the budget as it is not central to the SBN science goals. Thus, cost was a significant driver of the CRT design. This constraint makes the CRT design challenging especially when the surface area to be covered is considered. In the case of SBND and MicroBooNE, the CRT surface area is of the order of hundreds of square meters. In the case of ICARUS, it is about 1000 m²

Whatever design is chosen, it should be functionally identical for all three SBN detectors in order to minimize systematics in the oscillation analysis. In Section 3.1, I will discuss

the conceptual design for the CRT systems that meet the requirements outlined above. For the remainder of this chapter, I will focus only on the ICARUS CRT as this was the bulk of my PhD work. I will discuss the full history from our early design research and development (Section 3.2), to “bumps in the road” (Section 3.3), to production and testing (Sections 3.4 through 3.6), and ending with installation of the system at ICARUS (Section 3.7).

3.1 Conceptual Design

The most promising base detector technology satisfying the requirements outlined above is solid organic scintillator. The material is cheap, easy to manufacture, and easy to machine. Using this material, we can design a two-dimensional tracker. To that end, we also need wavelength-shifting (WLS) fiber and photosensors, either photomultiplier tubes (PMTs) or silicon photomultipliers (SiPMs).




Figure 3.1: An illustration of the CRT conceptual design from MINOS [61] shows the basic elements described here.

The WLS fiber shifts the blue scintillation light toward green wavelengths and guide the shifted light to the photosensor. One of the primary reasons to use WLS fibers is

that it significantly increases the minimum segment size, defined by optically isolated scintillator volumes, when considering light output in relation to muon tagging efficiency. The bulk attenuation length in typical solid organic scintillators is of the order of 10 cm while the attenuation length in WLS fibers is typically of the order of a few meters. This concept is illustrated in Figure 3.1.

For the choice of photosensor, SiPMs have become the preferred choice for detecting photons from optical fibers. Compared to PMTs, SiPMs are less fragile, have higher photon detection efficiencies, are more compact, and require only tens of volts as opposed to the hundreds of volts typically required for PMTs. The primary downside to SiPMs compared to PMTs is the significantly higher rate of spontaneous discharge (“dark noise”) when operating above cryogenic temperatures.

From the point of view of cost, the number of electronics channels must be minimized. From the point of view of tagging efficiency, enough light must be collected to ensure that the efficiency remains at or above about 95% at every point in the active CRT system. Minimizing the number of segments per unit area tends to reduce the number of channels while also reducing the tagging efficiency due to attenuation losses. Fewer gaps in coverage will modestly improve the total tagging efficiency. Even if light yield was not a factor, the size of the segments in the system is limited by the requirement of avoiding multiple muon occupancy. This requirement avoids introducing additional ambiguity in removing cosmogenic backgrounds from candidate neutrino events.

In terms of geometry, at least two optically isolated layers of scintillator are required to reject false muon tags, which primarily originate from ambient, low-energy, radiogenic activity, by applying interlayer coincidence. See Figure 3.2 for a radiogenic background spectrum measured at CSU. This background is similar across the world as evidenced by measurements at five sites across the US and Europe being within about 30% of the mean [54].

A two-layer design is advantageous in designing a two-dimensional tracking system.

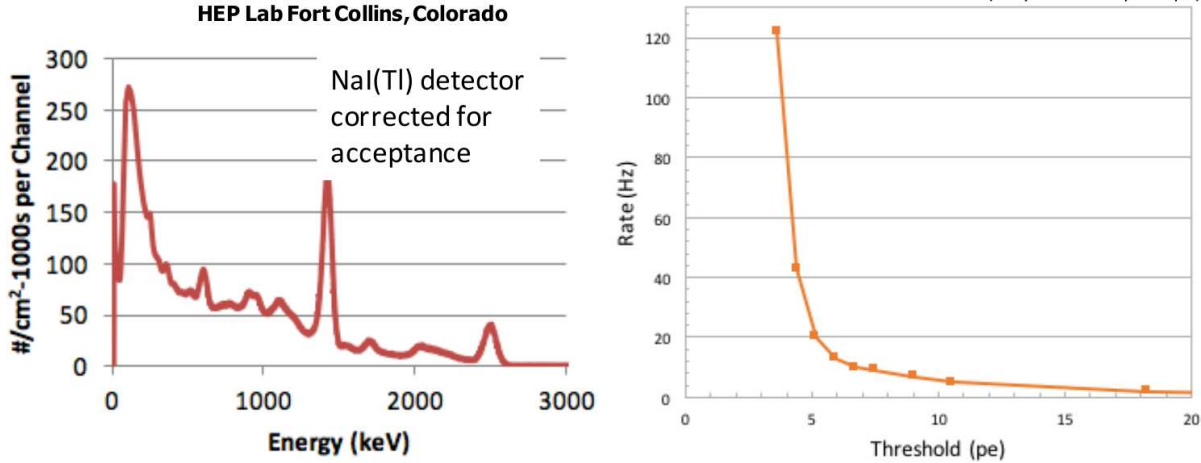


Figure 3.2: The radiogenic background (energy spectrum measured at CSU by J. Mann, shown on the left), can generate CRT triggers when using just a single layer. The trigger rate versus threshold was measured showing the transition to begin dominated by cosmogenic muons at about six photoelectrons (credit: R.J. Wilson).

Two planes of long, narrow scintillator strips with WLS fibers embedded inside, running the length of the strips can be constructed. The two planes are then stacked on top of each other with the orientation of the strips orthogonal between the two layers, forming a two-dimensional grid or “X-Y configuration”. As an added bonus, Cartesian geometries work well in a modular design, making it scalable and flexible in covering a large surface.

3.2 Early R&D at CSU

In 2015, my advisor, Professor Robert J. Wilson, and I joined the ICARUS experiment as part of SBN. Our initial contribution would be the detailed design of a CRT system. This was shortly after the SBN proposal was accepted and granted funding by DOE. At this point, the conceptual design discussed in the previous section was mostly worked out. We then had two objectives. First, we needed to demonstrate to ourselves and the collaboration that the conceptual design could achieve the tagging efficiency outlined in the SBN proposal. Second, we needed to optimize the design, driven by cost and schedule, while maintaining sufficient tagging efficiency. We would compete with another group,

based at CERN, to develop and propose a design for a new, 4π coverage CRT.




Figure 3.3: Our first setup at CSU was simple, but it was an effective first exercise in preparing for a full R&D campaign.

We built a test stand in the High Energy Physics Lab in the basement of the Physics Department at CSU. There was already some equipment available for us to get started: spare wavelength-shifting fiber and plastic scintillator; a $3 \times 3 \text{ mm}^2$ SensL 30035C-SMT SiPM; a power supply; a 100x amplifier; and a Tektronix 2024C oscilloscope that was able to run with a laptop. This setup allowed us to practice the basic techniques we would need for a full R&D campaign. The SiPM and a first CRT strip test setup are shown in Figure 3.3. With the basic tools in place, we were ready to proceed.

Starting with the conceptual design, we determined the design parameters to be studied. The guiding principle was to make the scintillator strips as large as possible, up to the maximum width where multiple muon occupancy becomes likely, while minimizing the number of electronics channels required. A secondary consideration was the amount of labor required to fabricate the particular prototype. The primary quantity that drives the tagging efficiency is the light yield of the design. This is determined by the intrinsic yield of the scintillator, the attenuation of the scintillator and optical fibers, and the efficiency with which that light is collected, which has a strong dependence on geometry.

We needed to determine the following items, guided by their impact on the total light yield of each design:

- maximum strip width and length that satisfy the tagging efficiency requirement;

- number of optical fibers per strip;
- position(s) of optical fiber(s) in each strip;
- optical fiber diameter;
- optical fiber - scintillator coupling method;
- optical fiber - SiPM coupling method;
- scintillator composition;
- SiPM model.

In this section, I will discuss in detail the different design components, the experimental setup, prototype development, and our final proposed design.

3.2.1 Scintillator

Plastic scintillator consists of a base polymer with one or several added fluors and wavelength shifters. The cheapest and most widely available base material is polystyrene. One of the advantages of plastic scintillators is their ability to be cast or extruded into virtually any shape. Casting produces pure, high quality polymers with excellent scintillation performance. However, this method renders the product relatively expensive, especially if additional machining is required to achieve a particular geometry. Also, individual scintillator volumes need to be coated with a surface layer of diffuse reflector in order to protect the scintillator from moisture and oxidation, optically isolate the volume from adjacent scintillator volumes, and promote locally isotropic optical propagation. This added labor increases the cost of cast scintillator. These considerations motivate the extrusion production method.

Extrusion significantly reduces cost in several ways: commercial polystyrene pellets are the starting point, avoiding monomer purification and bulk polymerization steps; extrusion templates make it possible to manufacture complicated geometries without added




Figure 3.4: A sample of scintillator from the Fermilab Extrusion Facility sits in the CSU HEP Lab, ready for R&D.

machining; holes for optical fibers or a diffuse reflector layer can be coextruded; and the process minimizes the required labor and production time. One drawback is the scintillator performance is marginally worse than cast scintillator. However, the cost savings make large scintillator detectors, such as MINOS [61], possible. Since cost was a significant design driver, extruded polystyrene scintillator was an obvious choice.

Uniplast, the manufacturer for the extruded plastic scintillator used by SBND, had indicated that they would not be able to produce a sufficient amount of scintillator for our needs on our required timescale. We turned to the Fermilab Extrusion Facility, where the scintillator for MINOS and MINER ν A was produced. They confirmed that they would be able to meet our requirements. The Uniplast and Fermilab scintillator were comparable in terms of performance.

The extrusion templates at the Fermilab Extrusion Center set a limit on the geometry of our design. Without modifying the production line, which would add tens of thousands of dollars to the total cost, the maximum width and thickness of scintillator bars that could be produced was 17 cm x 1.3 cm. By piggybacking on another order, we were able to obtain a sample (Figure 3.4) of 137 newly manufactured scintillator bars, all 5-cm wide,

1-cm thick, and several meters long. Each bar had a coextruded TiO_2 diffusive layer. 37 of the bars had a 3-mm diameter coextruded hole 1.2-cm from the side wall. We could machine these bars and glue them together to achieve different effective bar sizes.

3.2.2 Silicon Photomultipliers (SiPMs)

SiPMs are solid state, silicon based photodiodes arranged in parallel and operating in Geiger mode. An equivalent circuit is shown in Figure 3.5. These devices are typically produced as a square array with mm-scale side length containing thousands of micron-sized microcells, each a photodiode. When a flash of light impinges on an SiPM, several microcells discharge simultaneously, each generating one photoelectron (PE) worth of current. The total output current from the SiPM discharge is the coherent sum of its individual microcell discharges. This functionality yields a linear relation between the number of detected photons and the amount of charge produced, so long as the number of photons per microcell per tens of nanoseconds is below unity.

SiPM performance is characterized by several quantities: the ability to detect a single photon, photon detection efficiency (PDE); the amount of charge generated by a single photoelectron, gain; the electrical isolation of a microcell from adjacent microcells, crosstalk; and the rate of spontaneous discharge in the absence of light, dark count rate (DCR).

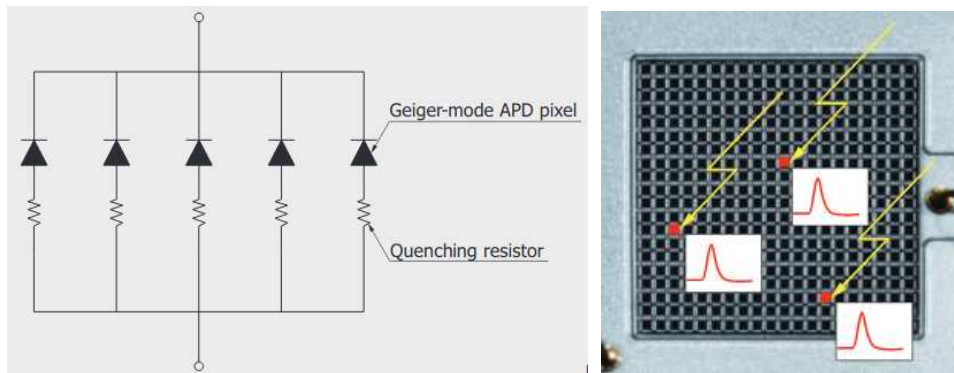


Figure 3.5: A simplified circuit diagram of an SiPM with five microcells (left). An illustration of photons impinging on an SiPM (right) [59].

The CSU HEP group had recent experience working with SensL C-Series SiPMs [58] as part of development for the DUNE photon detection system. To set a baseline of SiPM performance with our setup, we made our first measurements with SensL 30035 SMT SiPMs. Relevant specifications can be found in Table 3.1. To understand how our measurements compared to the SBND CRT group, we compared the light yield from one of our CRT prototype strips that was similar to the SBND design. There was a discrepancy between two groups' measured light yields. As previously mentioned, we used scintillator produced at Fermilab while SBND was using Uniplast Scintillator. Both products were expected to be similar in terms of performance.

We later obtained a sample of Hamamatsu [59] S12825-050Ps from the Bern group, the same model used in the SBND CRT. Relevant specifications are summarized in Table 3.1. We found that our measured light yield, with this model of SiPM, now agreed with that measured by the Bern group. The difference could be accounted for by noting the difference of around 20-30% in crosstalk probabilities between the two models. This was verified by checking for compatibility with Poisson statistics for low-level signals. The SensL SiPMs were found to agree reasonably well while the Hamamatsu SiPMs required a correlation correction compatible with the expected rate of crosstalk.

A higher performing SiPM product line from Hamamatsu, considered by the SBND group but later rejected due to cost, was the S13360 series. We made measurements of the dark rate and light yield with the S13360-2050VE, -3050VE, -3050PE, and -3070PE. I will discuss this further when I summarize prototype development. Most of our measurements with this series were on the S13360-2050VE. Relevant specifications are summarized in Table 3.1. As this model had recently come down in price, along with the excellent performance specifications, these SiPMs were a great candidate for our final design.

Table 3.1: Selected SiPM specifications reproduced from data sheets at an ambient temperature of 25 °C and recommended bias voltage [58], [59].

Model	Dark Rate [kHz/mm ²]	Gain [x10 ⁶]	Breakdown [V]	PDE [%]	X-Talk [%]	Temp. Coeff. [mV/°C]
SensL C-series	33	3	24.5	31	7	21.5
Hama. S12825	70	1.3	N/A	37	40	N/A
Hama. S13360	60	1.7	53	40	4	54

3.2.3 Front-End Electronics

The front-end boards (FEBs) used by SBND were designed by Igro Kreslo and others at the University of Bern specifically for their particular CRT design, providing the minimum cost per electronics channel. Even with this optimization, the cost per FEB channel remained a significant design driver. A full discussion of the requirements and the functionality of the Bern FEB design can be found in [56]. For convenience, I reproduce a summary below. The FEB

1. provides bias voltage in the range of 20-90 V individually adjustable for each of 32 SiPMs;
2. amplifies and shapes the SiPM output pulse on each of 32 channels (Fig. 3.7a);
3. performs discrimination of shaped signals at configurable level equivalent to 0 to 50 photoelectrons (PEs);
4. allows for requiring coincidence of signals from each pair of adjacent channels (optional);
5. allows for triggering only on events that are validated by external signals, such as by a signal in a group of other FEBS (Fig. 3.7b);
6. generates a trigger for digitization of all 32 signal amplitudes;
7. generates a time stamp with respect to the input reference pulse with an accuracy of 1.3 ns RMS;

8. performs digitization of the signal amplitude of each of the 32 channels;
9. provides on-board data buffering;
10. provides efficient back-end communication based on the 100 Mbps Ethernet standard;
11. allows firmware upgrade via back end Ethernet link;
12. requires input power 5 V with the consumption ranging from 450 mA to 550 mA depending on channel configuration.


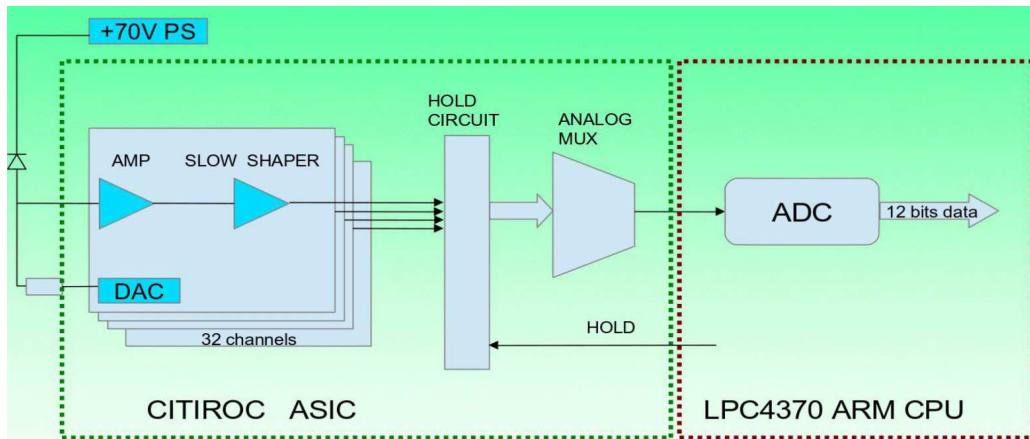


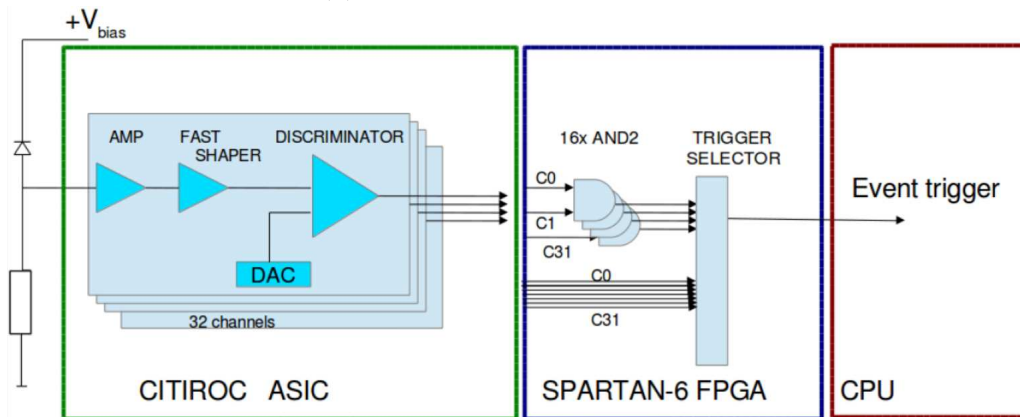
Figure 3.6: The CAEN A1702/DT5550 32-channel front-end board [56] is used for all SBN CRTs.

The FEB was originally designed to readout a single CRT module consisting of 16 scintillator strips, each with two WLS fibers. SiPMs are known to be noisy devices with typical dark rates in the range of kHz–MHz. To reduce the rate of false triggers due to dark noise, an approximate 30-ns coincidence gate is applied to channels x and $x+1$ with $x = 0, 2, 4, \dots, 30$. This coincidence window is optimized for the SBND geometry. Furthermore, the FEB can be configured to require an external validation signal in order to trigger a readout. See Figure 3.7b. The intended use of this feature is to generate coincidence gates between FEBs reading out scintillator modules in adjacent layers. The SBND CRT has two layers of modules, where each module is read out by a single FEB. The two layers are

arranged perpendicular to each other in a X-Y layout, providing the approximate point of entry of a cosmogenic muon. When any of the single strips in a layer generates a signal above threshold, a validation pulse is sent out from that FEB to the FEBs in the adjacent layer. The event is validated if coincidence across two layers is observed within 150 ns, the maximum delay incurred for a validation signal due to SBND’s cable layout. If a readout occurs, the triggered FEBs incur a 22- μ s deadtime. Otherwise, the deadtime is reduced to the coincidence gate width of 150 ns.



(a) FEB analog readout circuit.



(b) FEB trigger circuit.

Figure 3.7: The FEB provides full analog readout and digitization (a) with configurable trigger logic (b) as shown in the block diagrams [56].

Individual FEBs can be daisy-chained through the Ethernet interface, up to 256 in total, on a single network for easy interfacing with a single PC. Slow control of the FEB via the Ethernet interface facilitates real-time configuration of the preamplifier gain,

SiPM bias voltage, and discriminator threshold values. For small-scale applications, this functionality is provided through a standalone DAQ software package equipped with a dedicated GUI. The software is written in a combination of C and ROOT6, and data is written in root tree format for quick analysis. In addition, the software provides real-time, low-level analysis results including charge spectra for each FEB channel, trigger rate, and period measurements for the timing signals.

3.2.4 Lab Setup

All of our measurements were performed in the HEP lab at CSU (elevation 5000 ft or 1.52 km). According to [60], we should expect a vertical muon flux of $\sim 100 \text{ m}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$. Our group has prior experience working with PMTs and SiPMs. Working from that, we built a “dark box” to optically isolate the prototype from ambient light. Electrical connections between components in the dark box and the outside world were made through dedicated feedthroughs.

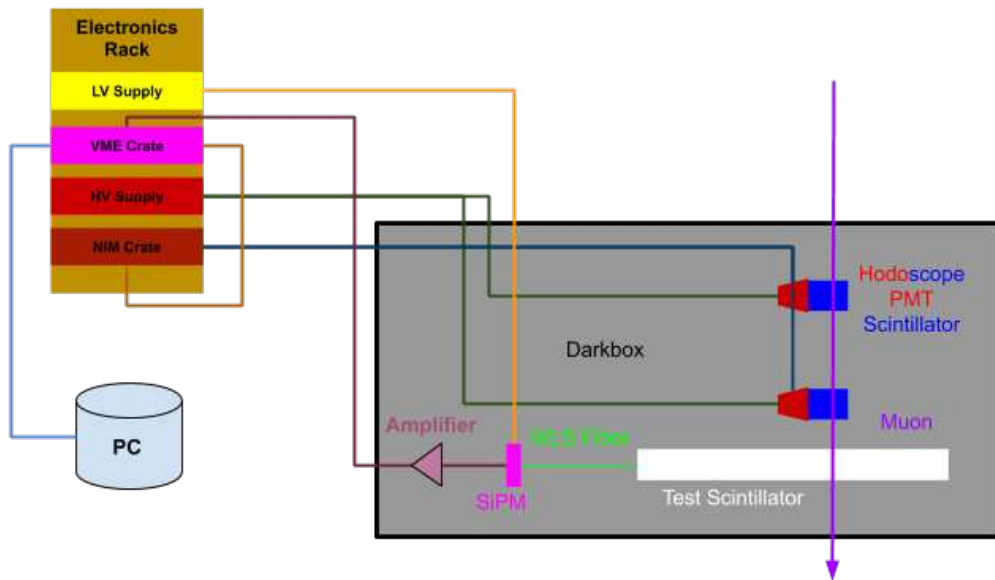


Figure 3.8: Setup at CSU using a CAEN VME crate for trigger logic and digitization.

In our initial setup, we had a single electronics rack housing two high-voltage power supplies, a single low-voltage power supply, a NIM crate operating a discriminator and

coincidence module, a CAEN VME crate operating a CAEN digitizer board and providing a control and data transmission connection to a single PC. With this DAQ system, facilitated by the CAEN WaveDump software, we could record full SiPM waveforms.

In order to study the response of our CRT strips to throughgoing cosmogenic muons at well defined positions, we used a homemade muon telescope, also known as a hodoscope. This hodoscope consisted of two 3 cm x 2 cm segments of scintillator coupled to a waveguide coupled to a PMT. The whole unit was optically isolated with a generous application of black electrical tape. Each PMT was connected to a high-voltage power supply and a discriminator. The output from the discriminator for each PMT was sent, as an input, to a coincidence unit. If a coincidence occurred between the PMTs within a window of several nanoseconds, we interpreted the coincident signal as having originated from a throughgoing muon. The coincidence signal was used as an external trigger signal to the digitizer board. I will refer to this run configuration as hodoscope-triggered. The board could be configured to disregard the external trigger. In this case, the system is self-triggered. The setup is illustrated in Figure 3.8. Analyzing the full waveform allowed us to check for anomalous behavior with our SiPM readout, unstable baselines or after-pulsing for example. This would serve as an important cross-check when we adopted new front-end electronics.

After making several measurements with the CAEN VME system, the SBN Cosmic Ray Mitigation Task Force was charged with finding a common technical solution for the three different CRT systems. It was natural and economical for the far detector to utilize the FEB designed by Bern which, while making up a large fraction of the total cost, had been optimized for this task rendering it the cheaper and more robust option. We received two loaner front-end boards (FEBs) designed by the Bern group for SBND.

The FEBs would replace the CAEN VME crate. However, we would no longer have access to the full SiPM waveforms as the Bern FEBs only write the integrated charge per pulse to disk. The benefit of this was simplified analysis and reduced disk usage. Some

simple tests showed that the two DAQ systems produced consistent results.

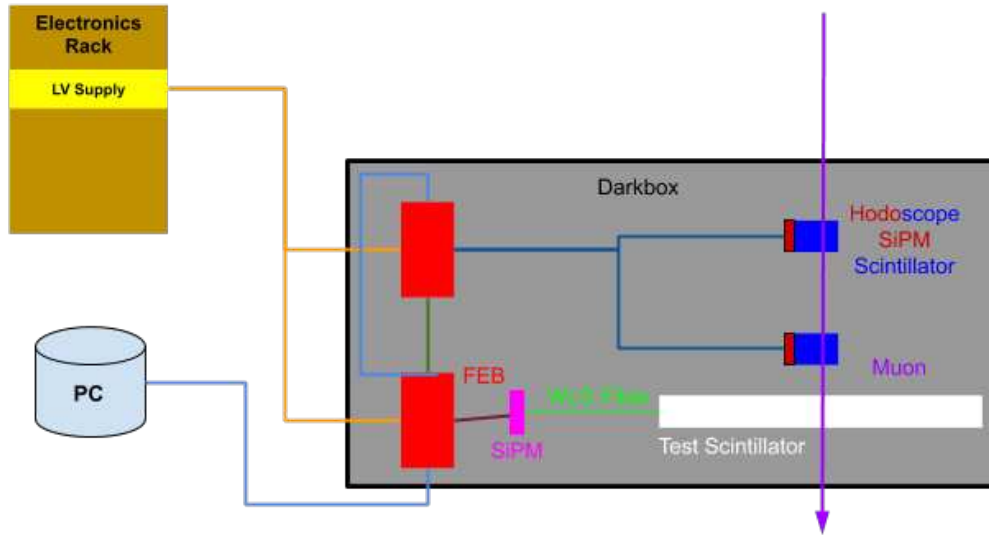


Figure 3.9: Setup at CSU using a FEB for light yield measurements and a second one for operating a new hodoscope.

The rate from the PMT based hodoscope was painfully low, producing a few hundred triggers over several hours. This motivated us to fabricate a new hodoscope. We could use one of the FEBs to read out the SiPMs reading out our prototypes. The second FEB could read out SiPMs connected to multiple small scintillator bars. We could use the trigger output signal from the latter FEB as an external validation signal to the former FEB. By making the new hodoscope scintillator bars larger than the PMT based hodoscope, we could improve our data acquisition rate.

We produced several 15 cm x 2 cm x 1.3 cm scintillator bars with two 3x3-mm² SensL C-series SiPMs glued directly to the scintillator with optical-grade epoxy. The whole unit was then wrapped in black electrical tape for light-tightness. Our new hodoscope consisted of two of these bars mounted in parallel to each other and to the prototype. One SiPM per bar was assigned to a single coincidence pair (two pairs in all) on the dedicated FEB. Using the pairwise channel coincidence feature, the pair of bars would act as a hodoscope.

This setup is illustrated in Figure 3.9. We verified that our larger hodoscope produced


results consistent with the previous hodoscope. All results presented below were obtained using the setup shown in Figure 3.9.

3.2.5 Calibration


We took the pedestal to be the inherent electronics noise present on each FEB channel with all active components connected, including SiPMs. The pedestal distribution was obtained using the same configuration settings as those used in normal operation except with a low threshold (~ 0.5 PE). The central value was obtained by assuming the noise is normally distributed. An example of a Gaussian-fit pedestal is shown in Figure 3.10c.

To measure the effective gain, determined by the intrinsic SiPM gain as well as the FEB preamplifier, we took advantage of the discrete nature of the SiPMs that produced distinct photopeaks resolvable with our system (Figure 3.10a). A low-level signal was required as it was only in this regime where the photopeaks could be clearly resolved. We used an uncollimated ^{60}Co source (γ emitter) to this end. A collimated β emitter would have worked better since it allows for localized measurements on the prototype; however, the ambient radiogenic background rate on our $30 \times 17\text{-cm}^2$ prototypes was too large for our β sources. Several of the low-valued peaks in the signal distribution were each fitted with a Gaussian, and the central values were used to define the ADC peak values. A linear fit to these values, plotted against the photopeak number (Figure 3.10b), produced a slope equivalent to the effective gain with units of ADC/PE.


Calibration runs measuring the pedestal and gain were performed immediately before and after each light yield measurement. The average of these values was used for the run calibration. Equation 3.1 was used to convert the raw ADC signal amplitude into PEs. Additionally, this data set provided information on the stability of the pedestal and the SiPM gains, which are temperature dependent. To account for temperature fluctuations, we took daily measurements of the ambient lab temperature. The temperature changed slowly enough that no drastic changes in calibration were observed during our runs.




(a) Fit photopeaks.



(b) Linear fit to ADC versus photopeak number.



(c) Gaussian fit to pedestal.



(d) Trigger rate versus threshold.

Figure 3.10: The steps from the calibration procedure are illustrated showing how the pedestal, gain, and photoelectron threshold equivalence were determined.

$$\text{Amplitude[PE]} = \frac{\text{Amplitude[ADC]} - \text{Pedestal[PE]}}{\text{Gain[ADC/PE]}} \quad (3.1)$$

To determine the discriminator threshold-to-photoelectrons correspondence, we measured the dark rate while scanning across threshold values. Due to the discrete nature of SiPMs, one obtains the characteristic shape shown in Figure 3.10d. The plateaus correspond to valleys between photopeaks. The central value on these plateaus was used to define discriminator thresholds in terms of half-integral PE values. Note that the correspondence is not linear.

3.2.6 Prototype Development

In mid-2015, we received the shipment of sample scintillator from the Fermilab Extrusion Facility. We decided that all prototypes would be 30-cm long, based on attenuation lengths found in the literature for this type of scintillator, 5-10 cm [55]. Thus, our prototypes would be 3-6 attenuation lengths long. We established a baseline of performance by taking light yield measurements on these 5-cm wide bars before constructing wider ones.

To determine the importance of optical coupling of the fiber to the scintillator, we compared measurements taken with the hodoscope centered and along each lateral edge. We compared the light yield with a single 1-mm fiber inside a hole to one inside a machined slot, glued with optical-grade epoxy, and covered with Mylar film. The glued fiber configuration (Figure 3.11) produced a factor of two more light than the configuration with a single fiber in a hole.

Next, we investigated how light yield scales with the number of optical fibers used. We verified the 3-mm diameter holes could accommodate up to three 1-mm diameter fibers. Use of extra fibers can act to offset the loss in performance incurred by the suboptimal optical coupling. We found that two optical fibers inside of a hole can double the light yield relative to a single fiber in the same hole. For the slotted, glued case with two fibers in the slot, we also observed a factor of two increase in light yield. We also found that this doubling in light yield can be achieved by using a larger diameter fiber, 1.4 mm in this case.




Figure 3.11: A MINOS-style strip is shown stacked on top of our SBND-like strip.

The last measurement we needed to make before constructing wider bars was to determine the dependence of light yield due to the proximity of the throughgoing muons to the diffuse reflector. We triggered on muons passing through the center of a 5-cm wide bar with a single fiber in a coextruded hole. Using an end mill, we then removed 3 mm of material off each side to completely remove the TiO_2 layer (the top and bottom layers were still intact). This step was followed by polishing of the exposed scintillator. By measuring the light yield in the same configuration, we determined that the difference in light yield was a factor of two. This result would help us to understand the validity of results derived from prototypes constructed from several bars machined and glued together.

To construct bars with a larger effective width, we machined off interior, edges from the 5-cm bars with an end-mill, polished the exposed scintillator, then glued the bars together with optical-grade epoxy. We verified the transparency of the glue joints by triggering on muons passing through the far side of the joint relative to an optical fiber that was read out. Its signal was compared to the signal obtained with the original 5-cm bars. The results were found to be satisfactory after adjusting for the loss of the edge reflector.

The width of the SBND bars is 10.8 cm and the SBND CRT covers a factor of three less area. The SBND CRT group completed a R&D campaign similar to ours. Using their design as a reference, we could set 10.8 cm as our minimum practical width limited by cost. To provide a benchmark of the SBND design performance, we fabricated a 9.4-cm wide bar (Figure 3.11) from two 5-cm wide sections. This width was the most practical width similar to the SBND width we could build without wasting a significant amount of scintillator. As in the SBND design, we cut slots into the side faces of the bar and glued 1-mm diameter fibers into them with optical-grade epoxy. The slots were then covered with Mylar film. We found that our 9.4-cm prototype produced similar results to those obtained by the SBND group.

¹Light yield from all fibers was summed.

²Distance to SiPM was 5.4 m.

Table 3.2: Summary of measured light yields for all CSU CRT strip prototypes. All strips are 1-cm thick except for the Mu2e strips, which are 2-cm thick. All strips are 30-cm long and all measurements were made at 50 cm from the SiPM (except where noted).

Name	Width [cm]	Description	Max. Light Yield [PE]
P1G (old)	5.0	MINOS-like with 1 fiber glued in center slot	15
P1H	5.0	Coextruded hole 12-mm from the edge	13
P1R	5.0	Same as P1G with Mylar film covering slot	22
SBND-like	9.4	One fiber glued into a slot on each strip side	17 ¹
Mu2e	5.0	Two coextruded holes; different fluors	28
P2H01	17	Two coextruded holes 6.2-cm from the edge	7
P4SG	17	Four slots with fibers glued	38 ^{1 2}

The next logical step was to measure the light yield from the widest bars that the Fermilab Extrusion Center could extrude, 17 cm. The significant electronics cost incurred per FEB channel motivated a strategy of extracting the most light from each scintillator bar while keeping the number of electronics channels to read out the SiPMs at a minimum. This could be achieved by making the scintillator strips as large as possible up to the limit set by double occupancy and spatial resolution for matching to TPC tracks, about 30 cm..

The simplest version of this design, P2H01, used two coextruded holes for the fibers (Figure 3.12). The fiber positions in the prototype were motivated by the hole position in the 5-cm wide bars. We measured the base performance of this width using just a single 1-mm fiber per hole and used the hodoscope to scan laterally across the bar. We found that the signals obtained from muons at the far edges were too low to meet our minimum tagging efficiency requirement.




Figure 3.12: The 17-cm wide P2H prototype is pictured in the darkbox between tests.

Motivated by cost, we sought to optimize a 17-cm wide design. Attenuation and

geometric effects motivated the use of multiple fibers distributed throughout the bulk scintillator. Effective attenuation measurements (Figure 3.13c) were used to optimize the fiber positioning. Fibers were glued into four slots to optimize the optical collection efficiency of the fibers, motivating the name for this prototype, P4SG. CRT layout considerations motivated lengths of 5.4 m for the individual bars. We simulated a full length bar by using 5.4 m of WLS fiber with our 30-cm long scintillator bar with the bulk of the excess fiber contained in a coil.


The fibers had attenuation lengths of about 5 m. To obtain sufficient tagging efficiency at the far end of the bar, we found it necessary to use either two 1-mm fibers stacked in a slot or one 1.4-mm fiber per slot. We fabricated a prototype using one or the other fiber diameter and found that they produced comparable results. Based on cost and ease of fabrication, we opted to use the 1.4-mm fiber in our final prototype. Furthermore, we needed to add a Mylar reflector at the end of the bar to mirror the far end of the fibers and boost the light output. A fine grained transverse light yield scan with normally incident muons (Figure 3.13d) was used for a final evaluation of the prototype.

The end result of the 17-cm wide prototype optimization is shown in Figure 3.13. Using the Hamamatsu S13360 series SiPMs, we were able to achieve a normal-incidence muon tagging efficiency of 98% at the far end of the bar, 5.4 m from the SiPM. For two layers in coincidence, the equivalent efficiency becomes 96%, exceeding our target efficiency.


To further reduce the cost of the design, we investigated the possibility of reducing the number of electronics channels required by electrically ganging two SiPMs together (encouraged by the success of the CSU DUNE photon detector group ganging SensL SiPMs together). By connecting several SiPMs together in parallel, creating one effective SiPM, we could read out multiple SiPMs from the same FEB channel using the same applied bias voltage (Fig. 3.14). We were successful in ganging two Hamamatsu S13360 2x2 mm² SiPMs without significant performance degradation. With two-fiber ganging,

the number of required channels is reduced by a factor of two.


Ultimately, a full coverage, costed design based on this last prototype was developed by the CSU ICARUS group. Unfortunately, we would not get to move it into production as I will discuss next.




(a) P4SG-1.4 in the darkbox from the front.



(b) P4SG-1.4 during a transverse scan.



(c) Effective bulk attenuation length.



(d) Light yield versus transverse position.

Figure 3.13: Our final prototype, pictured here ready for light yield measurements, using a coil of fiber to simulate our target 5.4-m length.




Figure 3.14: SiPMs connected in parallel on a single FEB channel provide the analog sum of two optical fibers as shown here for the P4SG design.

3.3 Change of Plans

After the CSU and CERN designs were proposed, the collaboration decided that there was inadequate funding available to build the full, 4π coverage system from scratch. This decision was partially motivated by convenient timing of an alternative solution that had become available. In this section, I will introduce the adopted detector components from experiments that had recently been decommissioned that are based on the same basic design concept introduced in Section 3.1. At this stage, our primary task had evolved from building a CRT from scratch into incorporating existing components into a hybrid system that could accomplish the same task, tagging cosmogenic muons with better than 95% efficiency.

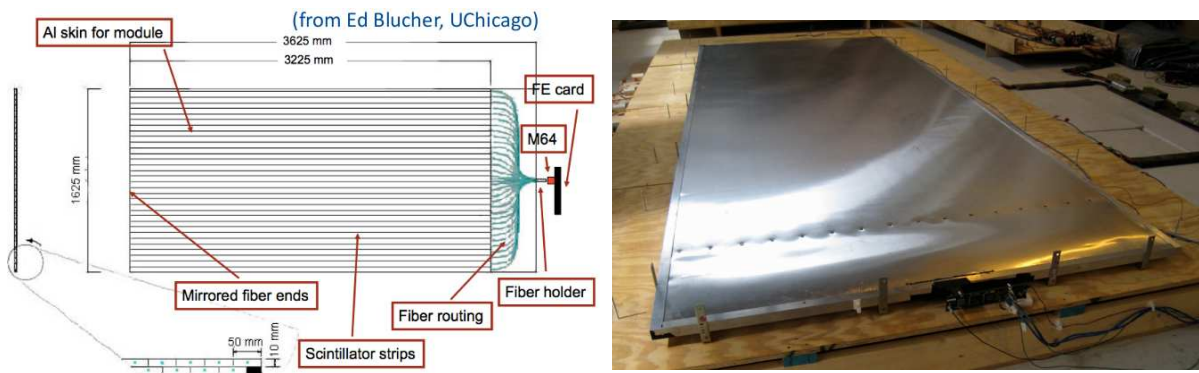


Figure 3.15: A diagram shows the Double Chooz veto module components on the left, and one of the modules poses for a photo on the right.

Prof. Edward Blucher, from University of Chicago, was a member of Double Chooz, which had just stopped taking data. Double Chooz had produced cosmic ray veto modules (Figure 3.15) for their experiment. A subset of these modules were designated for use in ProtoDUNE while the SBN leadership was able to convince Blucher to provide us about 20 modules for use as part of our CRT system. In addition, we would be able to use the full readout from the original system. The DAQ system would be reworked by ProtoDUNE to be compatible with their artDAQ based system. As we are using a similar, artDAQ based system as ProtoDUNE, little extra work would be needed on our end. The system

would essentially be plug-and-play. Due to installation constraints, any CRT coverage existing below the cryostat would need to be installed before or during cryostat assembly. This made the Double Chooz modules prime candidates for the bottom portion of the CRT as we only had a few months before assembly was scheduled to begin.

As I was already familiar with analyzing the simulation used for the SBN proposal, I was asked to help Anna Schukraft, our new co-convener of the CRT technical working group at the time, quantify the impact of different layouts of the Double Chooz modules below the cryostat. Choosing a layout was a difficult task as the modules would need to be installed into place before the warm vessel construction began. Cabling of the detectors would come after warm vessel installation was complete. We settled on a layout that facilitated later access of the modules should the need arise. Engineers at Fermilab had come up with the idea to install ball bearings underneath the modules so they could be rolled. Rails could be installed below the warm vessel that could guide the modules as they were pushed from one side or the other. While this solution all but guaranteed access to all of the modules up until the start of cold commissioning, it did not maximize coverage (Fig. 3.18).

In parallel, the MINOS far detector had just been decommissioned. While the primary detector was essentially destroyed in the process, the cosmic ray veto shield was salvaged. In total, 173 MINOS scintillator modules (Figure 3.16) were extracted from the Soudan Mine and shipped to Fermilab. The modules were gifted to ICARUS for use in the CRT.

The bulk of the work to incorporate the MINOS modules into the CRT would be to develop a new front-end for the system. The previous PMTs and front-end electronics could be provided; however, the electronics were not designed to operate at the kHz rates associated with surface running (the MINOS far detector operated over 700 m underground and only had to cope with a few Hz of cosmogenic muons). Furthermore, the scintillator had aged enough to cause a meaningful drop in light yield that could impact their performance. The muon tagging efficiency had been observed by MINOS

to degrade linearly at a rate of about 0.15% per year with the last measurement, taking place on March 7th, 2011, producing a two-layer-coincident efficiency of 96.1%. This implied the efficiency would have dropped to 95.1% by the time ICARUS was scheduled to begin commissioning in 2018. At the end of planned data taking, this efficiency would be about 94.5%, below the tagging efficiency goal outlined in the SBN proposal. With the requirement of new front-end electronics, the loss in light yield, and the desire to use the same front-end electronics as the other SBN detectors, it was clear that we needed to develop a new readout. It would be worthwhile, however, as the MINOS modules could provide enough coverage, about 450 m², with spares for the sides of the cryostat.

There was still some 400 m² to be covered above the cryostat. This surface also intercepts about 80% of the cosmogenic muon flux that enters into the cryostat. Our CERN colleagues decided to move into production with their proposed design in partnership with INFN groups. Thus, the top portion of the CRT would be all new construction, shown in Figure 3.17.

Relevant parameters for the different CRT subsystems are summarized in Tables 3.3 and 3.4. The final CRT layout, updated from the basic conceptual layout presented in the SBN proposal, is shown in Figure 3.18. According to my simulation studies, this layout achieves a 97% geometric efficiency for intercepting muons that enter the cryostat. For more discussion on this, see Section 6.1.

The remainder of this chapter will focus on the Side CRT subsystem since this is where I made large contributions. Over the remaining sections, I will present my work in the context of the larger CRT technical working group. This will cover development of a new readout for the MINOS modules, testing and characterization of the MINOS modules, readout prototyping, production and testing, and, finally, installation of the first 25% of the Side CRT subsystem.

Table 3.3: Scintillator parameters are summarized for the different ICARUS CRT subsystems. Note that due to scheduling constraints, the Top modules used different scintillator for each layer.

System	Scintillator	Composition	Reflector	LxWxH [cm]	Strip Conf.
Top	cast	polystyrene 2% pTP 0.05(0.03)% POPOP	paint	184x23x1(1.5)	8 per layer 2 layers X-Y
Side	extruded	polystyrene 1% pTP 0.03% POPOP	coextruded TiO ₂	800x4x1	20 1 layer N/A
Bottom	extruded	polystyrene 1% pTP 0.03% POPOP	coextruded TiO ₂	322x5 x1	32 per layer 2 layers X-X

Table 3.4: Optical fiber parameters for the different ICARUS CRT subsystems are summarized. Note that all subsystems use Kuraray Y11 WLS fibers and Hamamatsu photodetectors.

System	No./Strip	Position	Diam. [mm]	Photodetector
Top	2	6-cm from side, slot, glue	1.0	S13360-1350CS (SiPM)
Side	1	center, slot, glue	1.2	S14160-3050HS (SiPM)
Bottom	1	center (hole)	1.5	H8804 (MAPMT)

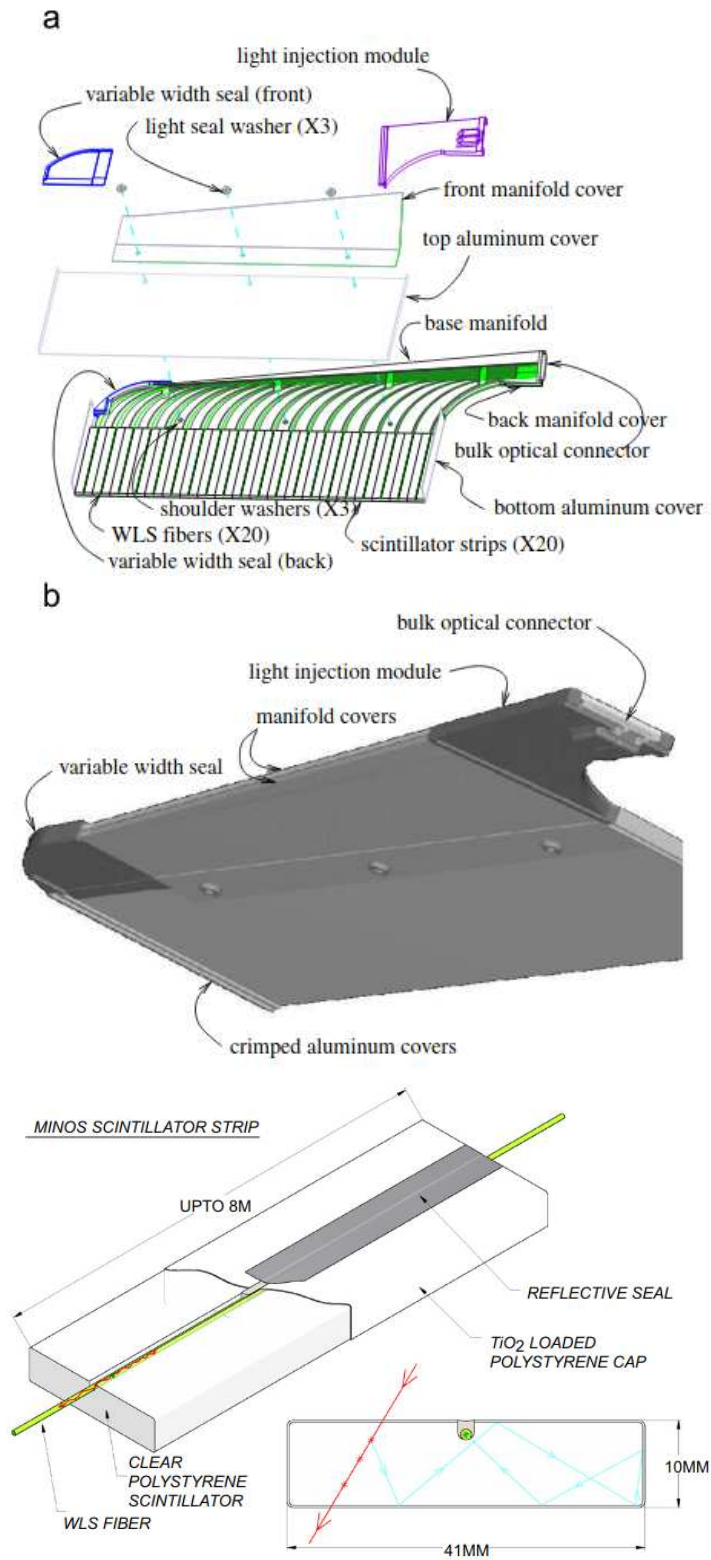


Figure 3.16: The composition of a MINOS scintillator module is shown on top with a detailed view of the scintillator strip design on bottom [61].

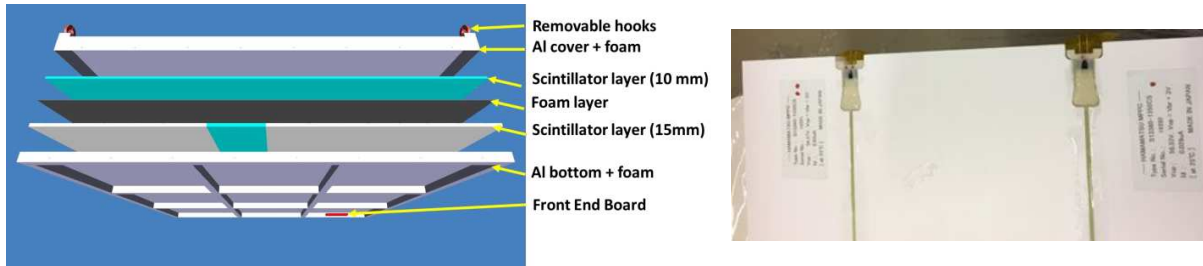


Figure 3.17: A concept drawing of the Top CRT module (left). A photo of a single scintillator strip equipped with two SiPMs (right). Image credit: U. Kose.




Figure 3.18: On the left, The most up-to-date CRT layout is shown, including the cold vessels (magenta) and the argon active volumes (cyan). The Top and Side CRT subsystems are visible. On the right, a top-down view of the area below the Warm Vessel shows the Bottom CRT subsystem. Image credit: J. Tillman, C. James, and A. Schukraft.

3.4 A New Optical Readout for the MINOS Scintillator Modules




Figure 3.19: My first concept drawing (right) for the SiPM board for the new MINOS optical readout used 1-mm SiPMs to read out each scintillator strip and a spacer to set the SiPM-fiber distance where the board would be mounted directly to the MINOS snout (left).

The first task in repurposing the MINOS scintillator modules for the ICARUS CRT was to develop a new optical readout. Having just spent two years on R&D for the CRT design, I was well positioned to lead this effort. It was clearly beneficial to use the same front-end electronics as the other SBN CRT systems as well as the Top CRT subsystem for ICARUS. These were the same front-end electronics that were part of our proposed design in Section 3.2. With these considerations, we adopted a solution based on SiPMs rather than the PMT based version used in MINOS. Front-end electronics aside, the SiPM based solution offers several advantages over the PMT based solution. First, in MINOS, clear fiber cables guided the light from the scintillator modules to PMT multiplexing boxes. We could offset the impact from scintillator aging by avoiding optical interfaces, which produce scattering losses, as well as attenuation from the fiber cable using SiPMs for direct readout from the scintillator module. I calculated that this could avoid 5-10% in losses. In addition, PMTs have lower photon detection efficiencies (about 10-20%) than their SiPM counterparts (about 40-50%). Based on these considerations, I estimated that we could, in principle, not only compensate for scintillator aging but also exceed the original light output obtained in MINOS by 250%.



Figure 3.20: Our first measurement of light output from a MINOS module with SiPMs was crude, using a makeshift stand to position the SiPMs against the optical interface (a), but it was successful and demonstrated that our new test stand was ready (b).

I spent three weeks at Fermilab in November 2016 with Anne Schukraft, a Fermilab scientist and co-convener with Professor Wilson for the CRT technical working group. From CSU, I brought two of our single SiPM mounts equipped with Hamamatsu S13360-2050VEs that we had used for our R&D studies and two FEBs. Working with Anne, we succeeded to build a test stand from scratch in some available space in building PB7, also known as “Wideband.” This is where the crates of MINOS scintillator modules were stored. I guided our efforts with my experience from the test stand at CSU with a goal of obtaining an estimate of the light output from one of the MINOS modules and the trigger rate due to cosmogenic muons and radiogenic photons. We needed to demonstrate that we could find an operating condition where dark noise would not compete with our signal.

In our initial setup, we had a single electronics rack housing one server, a NIM crate, and a control unit that operated a stack of counters, originally used in the CDF experiment. For first testing, we used a short, spare MINOS module that could easily be moved by hand. We procured a small dark box from one of Anne’s colleagues with a hole that was large enough to accommodate the module snout and a few cables. We inserted the module snout with cables and made a light tight seal at the dark box feedthrough using a generous amount of black electrical tape. I came up with a rather crude way of “mounting” the SiPMs on the optical connector. I tacked the SiPM stack onto a foam block that




Figure 3.21: Our first self-triggered run (a) was dominated by radiogenic photons with a rate consistent with measurements as CSU. Our first hodoscope-trigger run (b) showed a peak at roughly 27 PE with the hodoscope less than 1 m from the readout.

was held up by a slotted wooden block. This assembly was used to align the SiPMs with two fibers using the classic guess and check method (Fig. 3.20a). With this all in place, we were ready to test our readout.

We took two different data runs with the FEB operating in a self-triggered or hodoscope-triggered mode. Given the crudeness of the setup, we were pleasantly surprised to find that we had successfully obtained a first measurement of the light yield due to cosmogenic muons close to the readout end. Figure 3.21 shows a self-triggered spectrum from a 20-minute run and a hodoscope-triggered spectrum from a 2.5-hour run. There appeared to be a muon peak at around 20-30 PE in the hodoscope-triggered spectrum. This result proved that the basic concept was sound.

After returning to CSU, the next challenge was to understand how to mount SiPMs directly on the module optical interface. As the MINOS modules were not originally designed for direct readout, the geometry did not lend itself to such a use case. There did not exist any commercially available SiPMs that could completely cover the optical fibers of every channel without interference. The optical fibers had an outer diameter of 1.2 mm with a 2-mm pitch. The smallest SiPM active area available was $1.3 \times 1.3 \text{ mm}^2$, the Hamamatsu S13360-1375PE. Including the frame, the minimum outer SiPM dimension




Figure 3.22: Our first SiPM board with six 1-mm SiPMs (d) was tested at CSU with a pulsed LED (a) and with light generated by muons tagged with CRT prototypes run (b),(c),(d).

was 2.1 mm, too large to fit without interference. In order to collect light from each strip, the only possibility was to accept some collection inefficiency and only cover most of the fiber. The best SiPM for this was the 1 x 1 mm² SensL MICROFC-10035-SMT. If the SiPM was in direct contact with the fiber, the photon flux exiting the fiber was radially and azimuthally uniform, and the fiber axis was perfectly aligned with the SiPM center, our geometric efficiency would have been 84%.

One of our undergraduate electrical engineering students, Blake Troksa, designed a PCB on which we could mount the 20 SensL SiPMs. This PCB was also designed to mate to the interface, using the existing alignment pins from the original system. The electrical connections were made on the backside of the PCB using through vias connecting to male micro-coaxial connectors, compatible with the FEB header board we had used in R&D.


After fabricating two partially populated prototypes, each with six 1-mm SiPMs, I tested them for basic functionality in our dark box with a pulsed LED followed by a readout test from two R&D prototypes with a hodoscope trigger (Fig. 3.22). The SiPMs performed as expected. I observed uniform gain and dark rate in each channel. With these in hand, we were ready to test them on MINOS scintillator modules.

I moved to Illinois in January 2017 to continue this work at Fermilab. I brought the two readout prototypes with me. Anne worked with some Fermilab technicians to get a stackable shelving system (Fig. 3.23) that would allow us to access several of the full-length MINOS modules at once. We were ready to make our first measurements with the actual modules from which the Side CRT would be built. Figure 3.24a shows one of the prototypes mounted on a MINOS module.


Our initial tests were a success. Each SiPM channel had an evident signal from the scintillator module. In fact, we found our light yield was 30% larger than the original MINOS result (Figure 3.24b). This last confirmation put our design on firm footing and justified the next step of fabricating a few fully populated SiPM boards, which we could use to test all of the MINOS modules. The only change from this base design was to switch from using microcoaxial cables to using a single 20 pin connector compatible with one twisted-pair cable equipped with a global shield. The shield was grounded to the FEB common. This change was made to make disconnection and connection cycles easier with a large testing campaign in mind. The MINOS module testing campaign is the topic of the next section.



Figure 3.23: A stack of MINOS modules awaits testing at Wideband.



(a)



(b)

Figure 3.24: The first light yield measurement with our partially populated prototype from a full-sized MINOS module using the CDF hodoscope at 30 cm from the readout end yielded 12 PE, a 30% improvement over the MINOS result.

3.5 MINOS Module Testing, Characterization, and Readout Development

With our optical readout prototypes validated, the next task was to understand what fraction of the salvaged scintillator modules would be suitable for the CRT. We had no way of knowing if all or none of them were in working order. This necessitated a testing campaign to assess the health of each of the 173 modules. These tests also served as an opportunity to characterize the detector response of each module.

Fortunately for me, Anne had just hired a new postdoc, Simone Marcocci, to help with the testing on a half-time basis. Working with Simone, I developed a testing procedure, building on my R&D experience at CSU. I determined we would need to test three things:

1. light tightness across the entire module body;
2. damaged scintillator or optical fibers;
3. total light output in each scintillator strip.

With this information we could determine the uniformity of the response across all of the scintillator strips. In addition to testing and characterizing the MINOS modules, this would also test the robustness of our optical readout design.

The testing procedure for a single MINOS module required two people and proceeded as follows:

1. using pure ethyl alcohol and a cotton swab, carefully clean the optical interfaces;
2. use a handheld air blower, clear away any cotton fibers which may remain;
3. visually inspect the optical interface from different angles to ensure it is free of debris;

4. verify optical fiber continuity by injecting light in one end of the fiber and ensuring a clear image at the opposite end;
5. apply a thin layer of optical coupling compound to each optical interface;
6. mount the SiPM readout board to each snout;
7. place a light tight bag over each snout, and seal it with black electrical tape;
8. connect the SiPMs to the FEBs;
9. in the DAQ interface, enable to the SiPM high voltage and enable data acquisition;
10. with one person monitoring the trigger rate, the other person slowly scans along the optical readout then along the full module body using a bright LED flashlight;
11. if there were no sudden rate increases, the module is deemed light tight so move on to the next step 7;
12. if there was a sudden increase in rate, a light leak has been found → seal it with tape and verify rate is stable;
13. using an uncollimated ^{60}Co button source and enabling triggering on a single channel at a time, verify the trigger rate increases when the source is over the scintillator strip in the longitudinal center and 1.5 m from each readout end;
14. record the ambient temperature;
15. using the muon telescope to trigger the FEBs, take 15 minutes of data at the same positions as in step 7;
16. record the ambient temperature;
17. disable the SiPM high voltage;
18. carefully disconnect and dismount the SiPM readouts;

19. using pure ethyl alcohol and a cotton swab, carefully clean the optical interfaces, and wrap them with plastic film.

Simone and I tested the procedure. It took approximately 1-2 hours to test a single module depending on how difficult any light leaks were to find and repair. Fortunately for us, Anne hired two summer interns to help with the testing. While the interns performed the tests (Figure 3.25), I developed an analysis software package.




Figure 3.25: The summer interns were hard at work testing one of the MINOS modules.

The analysis software built on code that I developed previously for our R&D studies at CSU. The ^{60}Co data was written to disk; however, it was not used in the analysis. The hodoscope-triggered runs provided sufficiently resolvable photopeaks for the calibration step. Thus, calibration and light yield measurements were made from the same data set. After calibration, the spectrum for each SiPM channel was fitted to extract the most probable light yield. With the light yields extracted from the three different hodoscope

positions, I fit a simple exponential to a light yield versus distance plot to extract an effective attenuation length.

In total, only a single MINOS module was unusable. This was due to severe corrosion of the aluminum casing that rendered the module unusable due to the level of ambient light leaking into the module. The damage was significant enough that it could not be repaired with tape or something like RTV. Nearly all of the corrosion was confined to a single half of the module. There was a possibility that some of the MINOS modules would need to be cut to a short length to be used in certain areas of the Side CRT, so we decided to test cutting and sealing MINOS modules with this one. The corroded half was cut off and discarded, while the usable half was set aside for later testing of a procedure for cutting and sealing modules.

The remaining 172 modules were tested in the same way using the above procedure. This amounted to 516 hodoscope-triggered runs with a total run time of 129 hours. We tested a total of 3440 scintillator strips and, with each strip being read out by two SiPMs, generated 20,640 light yield spectra from cosmogenic muons. The sheer volume of data to analyze made this task much more challenging than the analysis of our R&D data set.

Unfortunately, our prototype optical readout was not perfect. After hundreds of mount-dismount cycles, the alignment pins had become loose. This led to suboptimal alignment of the SiPMs with respect to the optical fibers leading to anomalously low light yield. The decline in alignment proceeded gradually at first before several consecutive tests suggested that all scintillator strips in several modules were performing significantly worse than observed up to that point in other modules. The modules with some questionable results from when the alignment began to worsen had already been packed back into the shipping crates and buried beneath other crates. As we were already using a large amount of technician time to move modules for regular testing, we were reluctant to track down and retest these modules with questionable data. We found a workaround for the alignment issue and proceeded with testing.

The modules with data that may have been affected by the alignment issue were excluded from the light yield analysis. The remaining modules produced the distributions in Figure 3.26. I found that our mean light yield at each of the three hodoscope positions, shown in Figure 3.26, is better than the MINOS result. However, the level of improvement depends on the position. Compared to the light yield values extracted from Figure 3.27a, the fractional increase in the light yield over the MINOS result is 5%, 23%, and 56% at 1.5 m, 4.0 m, and 6.5 m respectively. This scaling with distance can be understood by a difference in attenuation length. The effective attenuation length was measured to be about 5-7 m across all strips (Fig. 3.27b). 5 m is the attenuation length extracted from Figure 3.27a.




Figure 3.26: The light yield is shown for each scintillator strip at the three different hodoscope positions with the runs with known SiPM mounting problems removed: 1.5 m (a), 4.0 m (b), and 6.5 m (c). Distributions for both readout ends are included separately with the north end in blue and the south end in red.

Based on the uniformity we observed in modules with quality data, we concluded that nearly all of the modules would be suitable for use in the Side CRT. We found six modules with 1-2 damaged fibers. Thus, we had identified 166 good modules.

With this light yield distribution, our total cosmogenic muon tagging efficiency, integrated along the full module length, would have exceeded our goal of 95%. However, we were concerned with the tagging efficiency at the position with the lowest light yield. With the MINOS modules read out at both ends, this point is the center of the module where the light yield was 4.3 PE on average. Based on these results, we estimated that,

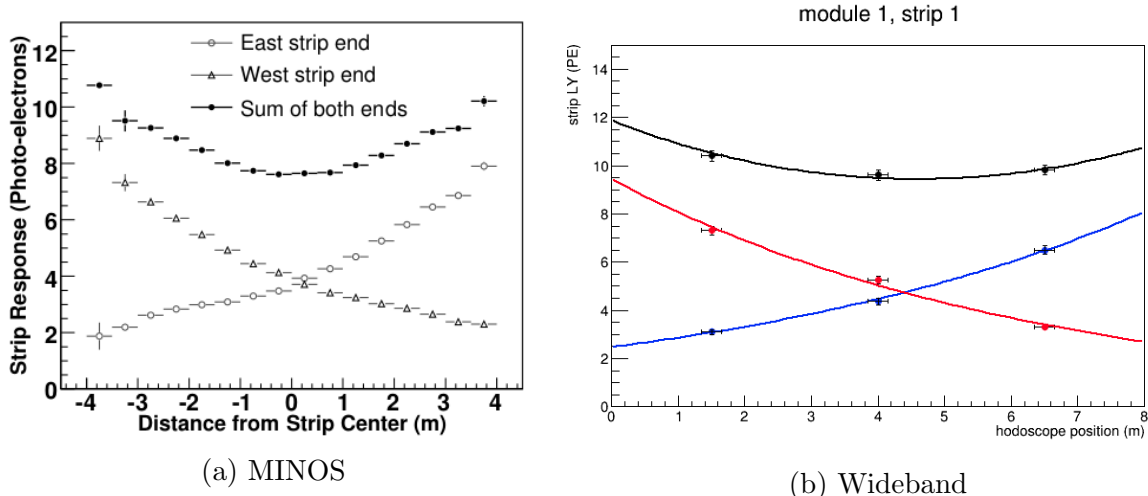


Figure 3.27: A comparison of typical attenuation curves from MINOS (a) and from our tests at Wideband (b). The individual channels are shown as well as the sum at both ends.

with two layers in coincidence, the lowest tagging efficiency would be 64% with a 2.5 PE threshold. This led us to reconsider our optical readout design.

3.5.1 Optimizing the Readout Design

We found a few shortcomings of our design that was used for MINOS module testing. One problem that we already knew about was that the SiPMs did not fully cover the optical fiber causing us to throw away some of the light. By considering the numerical aperture of the optical fiber, we realized that there was strong dependence on the fiber-to-SiPM distance. To make matters worse, the light output from these optical fibers was not radially uniform but rather biased by 9% at the outer edge with respect to the fiber center [62].

To understand how important providing full SiPM coverage of the fiber was, we fabricated a new SiPM board using ten $2 \times 2 \text{ mm}^2$ SiPMs where we skipped every other fiber. Then, to understand the effect of the distance between the SiPM and optical fiber, we set up a test using precision washers to set a few different distances for both our original 1-mm SiPM boards and our new 2-mm SiPM board. We used a hodoscope-triggered run




Figure 3.28: Optical fiber-SiPM distance and SiPM size dependencies demonstrated the need to adopt a readout based on larger SiPMs. Here, results are shown for $1 \times 1 \text{ mm}^2$ (yellow) and $2 \times 2 \text{ mm}^2$ (blue) SiPMs.

configuration with the hodoscope positioned 1.5 m away from the readout end, one of the distances used for our summer testing. We mapped out the light yield measured as a function of the SiPM-fiber distance. The results, shown in Figure 3.28, demonstrate that the light yield gets better for the 1-mm board as the distance is reduced. For the 2-mm board, the light yield is constant up to a distance of about 0.3 mm where it then begins to decline. To interpret this effect, we used a simple, geometric model based on the numerical aperture of the fiber and assumed the index of refraction of the optical grease for the gap between the SiPM and the optical fiber. The 1-mm SiPM board is reasonably described by the model, matching the predicted exit angle of 25° , consistent with an index of refraction of 1.7. For the 2-mm board, the exit angle from the fit of 50° is more consistent with an index of refraction of 1.1. It could be that air bubbles were introduced in the optical coupling compound or that an insufficient amount was applied for the larger gap sizes. Regardless of the explanation for the larger exit angle, we observed a factor of about two increase in the light yield using the 2-mm SiPMs. This result confirmed that we needed to adopt a design with a larger SiPM that could fully intercept the light cone from the optical fiber. Furthermore, it is important to minimize the fiber-SiPM distance while avoiding the risk of the making physical contact with the SiPM.

For testing purposes, it was useful to be able to read out each individual scintillator




Figure 3.29: To test the $3 \times 3 \text{ mm}^2$ SiPM board concept (top), two prototypes were fabricated: one with three SiPMs, skipping every other channel (middle), and one with four consecutive SiPMs (bottom).

strip. However, for the purposes of the CRT, this granularity was unnecessary given the 4-cm width of the strips was finer than the amount of multiple Coulomb scattering a muon would undergo as it propagated from the CRT into the cryostat. Thus, we considered the possibility of optically combining two fibers onto a single SiPM while providing complete active SiPM coverage of the fibers. There was only one way to make this work, and the idea came from Prof. Wilson. We would have to use $3 \times 3 \text{ mm}^2$ SiPMs, and they needed to be rotated on the PCB plane by an angle of 28° from vertical (Figure 3.29, top).

A downside to this design was that approximately 75% of the SiPM active area would go unused. Apart from not contributing to our light collection, the unused SiPM area would contribute to the overall dark rate. As dark rate scales with active area, we would need to contend with a factor of about 9 increase in the dark rate. The larger dark rate required that we increase the discriminator threshold on the FEB to maintain an acceptable trigger rate. From a tagging efficiency point of view, a higher threshold would only be acceptable if the new design provided a significant improvement in light collection.

A couple of prototypes, shown in Figure 3.29, with Hamamatsu S13360-3050VE $3 \times 3 \text{ mm}^2$




Figure 3.30: Data from three runs are shown: two with all but one fiber covered (blue, red) and one run with no fibers covered (black). Each histogram is run-time normalized, and the taped runs are summed (magenta).

SiPMs were fabricated at CSU and sent to us at Fermilab. We measured the light yield for these prototypes using a hodoscope-triggered configuration and the same MINOS module that was used for the 1-mm versus 2-mm SiPM comparison. With the hodoscope at 4.0 m from the readout end, we observed an average of 15.5 PE, consistent with expectations from the results obtained with the 2-mm SiPMs.

To ensure our results were not inflated by optical crosstalk from an adjacent fiber not intended to illuminate the SiPM in question, Simone and I performed some dedicated tests to search for such an effect. Comparing the two SiPM prototypes, one with four consecutive SiPMs and one with three SiPMs skipping every other channel, we observed no difference in light yield. Next, we searched for an effect from both optical fibers combined on a single SiPM. As shaping time for a single SiPM channel was about 30 ns, we did not expect multiple muons crossing the same scintillator strips. As only a small fraction of muons crossing the scintillator strips would cross the boundary between two strips readout by the same SiPM, we expected to see little effect in comparing the spectrum for a single SiPM with all but one fiber covered with tape to a spectrum obtained with neither fiber covered. The result of the latter test, shown in Figure 3.30, demonstrated no discernible level of crosstalk. The slight difference in peak position for the summed single-fiber spectra and the spectrum with neither fiber covered was attributed to the

small fraction of muons that do cross both strips.

Lastly, before we could be confident that our new prototype would succeed, we needed to measure the cosmogenic muon tagging efficiency, at normal incidence, along the module. We used the four-SiPM prototype since the read out strips spanned nearly the whole hodoscope width. Anna and I carried out a hodoscope-triggered, longitudinal scan of the light yield at six positions spanning the full MINOS module length. The resulting light yield versus distance curve for each SiPM channel is shown in Figure 3.31a.

We used a scaler module in the NIM crate to count hodoscope triggers. The muon tagging efficiency was defined as ratio of the number of FEB triggers to the total number of hodoscope triggers, corrected for accidental coincidences due to muons or other particles interacting outside of the hodoscope footprint that artificially inflate the measured tagging efficiency. The obtained efficiency curve is shown in Figure 3.31b.

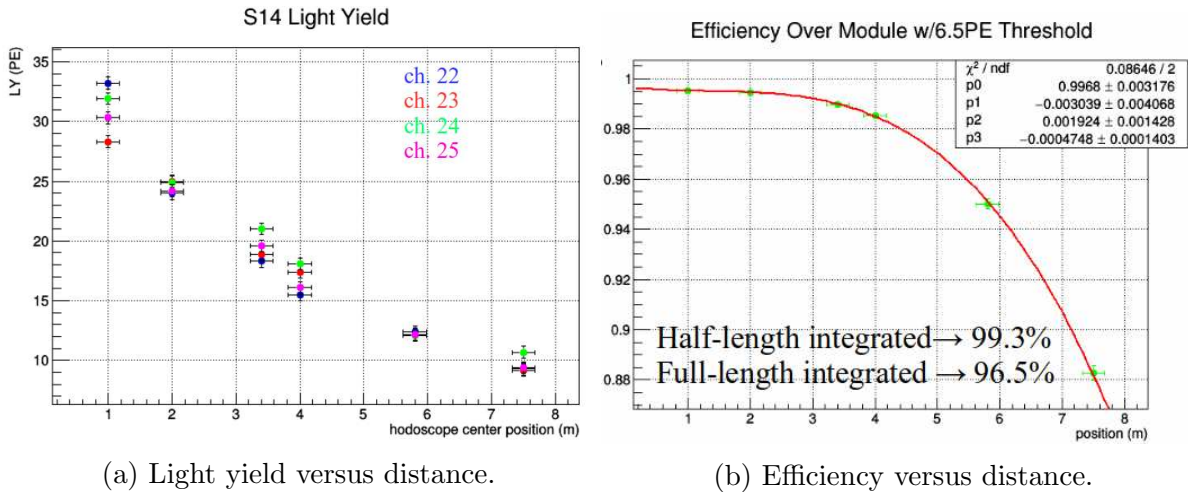


Figure 3.31: The light yield as a function of longitudinal distance from the readout end (a) and the resulting efficiency with respect to the hodoscope (b) demonstrated the success of our prototype based on 3-mm SiPMs.

At the longitudinal center, the measured efficiency was 98.5%, or 97% for two layers in coincidence, exceeding our goal by 2%. The half-length integrated efficiency was 99.3%, or 98.6% for two layers in coincidence. From the point of view of reconstruction, it is useful to match FEB triggers at both ends of a module. The full-length integrated efficiency was

96.5%, or 93% for two layers in coincidence. While this is too low for applying dual-ended coincidence as a trigger condition, it is still large enough to enhance hit reconstruction along the length of the module.

While all of our measurements were performed with the MINOS module in a horizontal orientation, the modules will all be vertically oriented in the CRT. For normally incident muons, the path length through the scintillator is 1 cm. With the module on its side, the path length 8 cm. Of course, not all muons will be normally incident. Nevertheless, moving from a horizontal to a vertical orientation should increase the average path length of throughgoing muons in the scintillator, thereby increasing the average energy deposited and the most probable light yield. Thus, we should expect that the tagging efficiency that we measured in the horizontal configuration should be maintained in moving to the vertical orientation, if not improved.

With the validity of the 3-mm board design demonstrated, we proceeded to develop a complete prototype.

3.6 Optical Readout Production and Testing

Simone created a Geant4 optical simulation including a truncated MINOS geometry (two scintillator strips, with one optical fiber per strip) and one 3 x 3-mm² SiPM. He simulated a sample of thoroughgoing muons in the scintillator strips and measured the photon collection efficiency at the SiPM, varying different position and orientation related parameters. From his simulation, he was able to estimate the required tolerances for the design: within 3° of the nominal 28° tilt angle, less than 0.15 mm of lateral displacement from the nominal position, and within 0.25 mm of a nominal 0.5 mm fiber-SiPM distance.

With a lot of cleverness and perseverance from our HEP engineer at CSU, David Warner, we were able to achieve the required tolerances and produce a prototype. A few fully populated SiPM boards were fabricated using a newer module from Hamamatsu,

the S14160-3050HS, which advertised 25% larger PDE, 47% larger gain, and 75% larger crosstalk. The level of crosstalk was still acceptable while the increase in gain and PDE could only improve the performance of our design.




Figure 3.32: An exploded drawing of the ORM shows all of the components (Image credit: David Warner).

In addition, to the SiPM board, David designed a housing assembly that would facilitate the precise alignment that was required while also providing light tightness and strain relief of the SiPM cable. A few prototypes were machined by hand for testing. The full SiPM board with housing assembly was named the Optical Readout Module (ORM). An exploded CAD view is shown in Figure 3.32.

David brought the prototypes to Fermilab for a test run with one of our MINOS modules. The housing worked well, but we found it somewhat difficult to mount and seal (light tightness was achieved with the usual generous application of black tape). Some of these moments are shown in Figure 3.33. With the ORM successfully mounted and a round of light yield measurements, I found that the light yield had improved by 10-20%




Figure 3.33: D. Warner test fits a machined housing prototype on a MINOS module observed by B. Behera (left). S. Marcocci prepares an ORM for first test measurements (center). The ORM fits perfectly and is ready for testing (right).

compared to our first prototypes. With this successful demonstration, we were ready to propose the design.

The final design passed a production readiness review at Fermilab in January 2019 and production commenced shortly thereafter. I devised a set of quality control tests that could be performed at CSU using the dark box and the pulsed LED setup used previously to supplement metrics for position and SiPM connectivity. These tests would be performed in parallel with production to ensure that any issues would be flagged and addressed promptly. The production would be overseen by David and by CSU graduate student, Tyler Boone, who would also conduct the quality control tests.

While the tests at CSU could assess the basic functionality of the SiPMs, these were no substitute for testing the ORMs with real MINOS modules. We needed to understand that the SiPM-optical fiber alignment was correct and that the housing could mate correctly with the optical interface. In addition, there was a possibility that there were SiPM performance dependencies on the pulse shape and pulse rate that would be difficult to study with the pulsed LED. I devised a second set of tests that could be performed at the Wideband test stand at Fermilab.

We would use a stack of three MINOS modules with the outer two serving as a control



Figure 3.34: It took a lot of work from a good team to get the test stand ready for its next task. Pictured from left to right: A. Aduszkiewicz, A. Wood and D. Munford (all UH) and A. Heggestuen (CSU).

and the inner module serving as the test module. All three modules would be equipped to be read out at both ends, and each ORM would be driven by its own FEB. Two runs would be performed. One would serve as a sanity check by measuring cosmogenic muon tagging efficiency without the use of a hodoscope. This would ensure that any strange behavior induced by a subtly defective SiPM could be observed. The second would use the hodoscope to measure light yield in a similar fashion as what was done for the MINOS module testing. This would serve as a cross-check on the efficiency obtained in the first test and further quantify the alignment and efficiency of individual SiPM channels. With the control modules kept in a constant configuration throughout the testing, we could also monitor system stability. This experience would go a long way in preparing to scale up to the full CRT system.

The self-triggered efficiency test would measure cosmogenic muon tagging efficiency integrated over the full 8-m module length. A NIM crate was configured to count coincidences between the reference modules alone and all three modules together using the trigger signal output from each FEB. The efficiency was calculated as the ratio of the

latter to the former.

The second test measured the light yield at the longitudinal center, 4 m from all readouts. This was to understand the worst-case efficiency that each SiPM channel might achieve. The setup was similar to that used for the MINOS module testing except the hodoscope footprint was doubled in order to increase the trigger rate by a factor of two.

I led the effort to repurpose the test stand for our tests. I had a lot of help from a postdoc, Antoni Aduszkiewicz, and two graduate students, Anthony Wood and Donald Munford, all from the University of Houston. Additional assistance was provided by CSU postdoc, Biswaranjan Behera, an incoming CSU graduate student, Anna Heggestuen, and an ICARUS collaborator from INFN-Catania, Francesco Tortorici. Figure 3.34 shows some of the team after a hard day of work. Together, we pulled and dressed lots of cables, cleaned up the test stand, modified the existing standalone DAQ code for our specific needs, and wrote the documentation needed to pass the operational readiness review.

Unfortunately, we encountered a significant problem early on while exercising the test procedure. In moving to a more complex test setup involving many more FEBs and connections between them, we observed a significant effect due to electrical noise. We could not access the raw SiPM waveforms directly, but the effect of the noise was observed as a non-zero FEB trigger rate when no bias voltage was applied to the SiPMs. While our trigger rate was not saturated, I was worried that the noise would have an adverse impact on the ORM test measurements. Thus, we spent the next month attempting to identify the cause.

In the end, there were three significant findings. We were using switched power supplies provided with the FEBs by CAEN. These are known to be a source of electromagnetic interference. We changed over to bench-top power supplies and observed a reduction in the noise rate, but this was not a complete solution. The second finding was that some of the ORM cables were sensitive to touch. Grasping a cable in one's hand, changing the capacitive coupling, reduced the noise rate. We fixed this by replacing the cables.



Figure 3.35: A direct ground connection to the FEB common was important to avoid electrical noise.

After some discussion with one of the electrical engineers at Fermilab who was familiar with the FEB, I learned that the FEB was not grounded anywhere. In order to fix this, I grounded the electronics rack to building ground via one of the electrical conduits. Then, I chained a ground braid from the electronics rack to all FEBs at either end of the MINOS modules. I determined that the best place to make the ground connection was through one of the bolts that secure the outer box. This bolt was in electrical contact with one of the annular rings on the FEB PCB as verified with a multimeter. This was enough to provide a common reference to all FEBs and solved the problem. Daily testing shifts commenced shortly thereafter.

Antoni, Biswaranjan, Francesco and I divided two daily shifts amongst ourselves. About halfway through ORM testing at Fermilab, production at CSU had completed. Shortly thereafter, Tyler moved to Illinois to assist long-term with CRT efforts. At around

the same time, another INFN-Catania collaborator, Catia Petta, arrived to assist with shifts for a few weeks. With this large team effort, we tested a total of 360 ORMs.

To analyze the data, I built upon and improved the analysis code I had developed for the MINOS module testing. To start, I analyzed the Wideband test data while Tyler further analyzed the CSU test data. The analysis procedure proceeded in the same way as for the MINOS testing.

The raw data was calibrated by measuring the pedestal central value and the resolvable photopeaks were used to extract the gain. The pedestal fitting procedure worked well and produced values stable to within 10 ADC over all of the runs (Fig. 3.36). This is the expected result as the pedestals should be fairly independent of the connected SiPMs as the terminal capacitance should be uniform. Any shifts that did occur in the pedestal were coherent across all SiPM channels on a given FEB and were independent across FEBs.

The gain fitting algorithm was less reliable than that for fitting the pedestal as it was inherently more complex. In Figure 3.37, a sparse population of higher gain values can be seen. This is due to failed gain fits, as understood through hand scanning, where a photopeak would be missed leading to a larger apparent increase in ADC per change in photopeak number. When these instances are removed from the analysis, the gain is stable to within about 2 ADC/PE. This was the expected result from the known breakdown voltages supplied by Hamamatsu for each SiPM in our production batch.

After calibration, I applied the same light yield fitting algorithm used for the MINOS testing to extract the most probable light yield. The results are shown for each channel in Figure 3.38. The impact from bad gain fits can be clearly seen. Artificially larger gains will produce artificially smaller light yields.

Finally, I combined the light yield and efficiency measurements. Shown in Figure 3.39, I found a strong correlation between efficiency and light yield where there is one point per channel. Vertical bands are due to association with a single test ORM. South test ORMs

were routinely observed to produce lower relative efficiencies than north test ORMs by about 1-2%. This can be understood from lower light yields compared to the north ORMs. The interpretation of this difference is that the light yield at the south end of the test MINOS module was intrinsically lower than the north end. A similar light yield imbalance was observed in several MINOS modules during testing. A few ORMs produced relatively low efficiencies. Upon inspection, this was due to a single SiPM producing little to no signal.

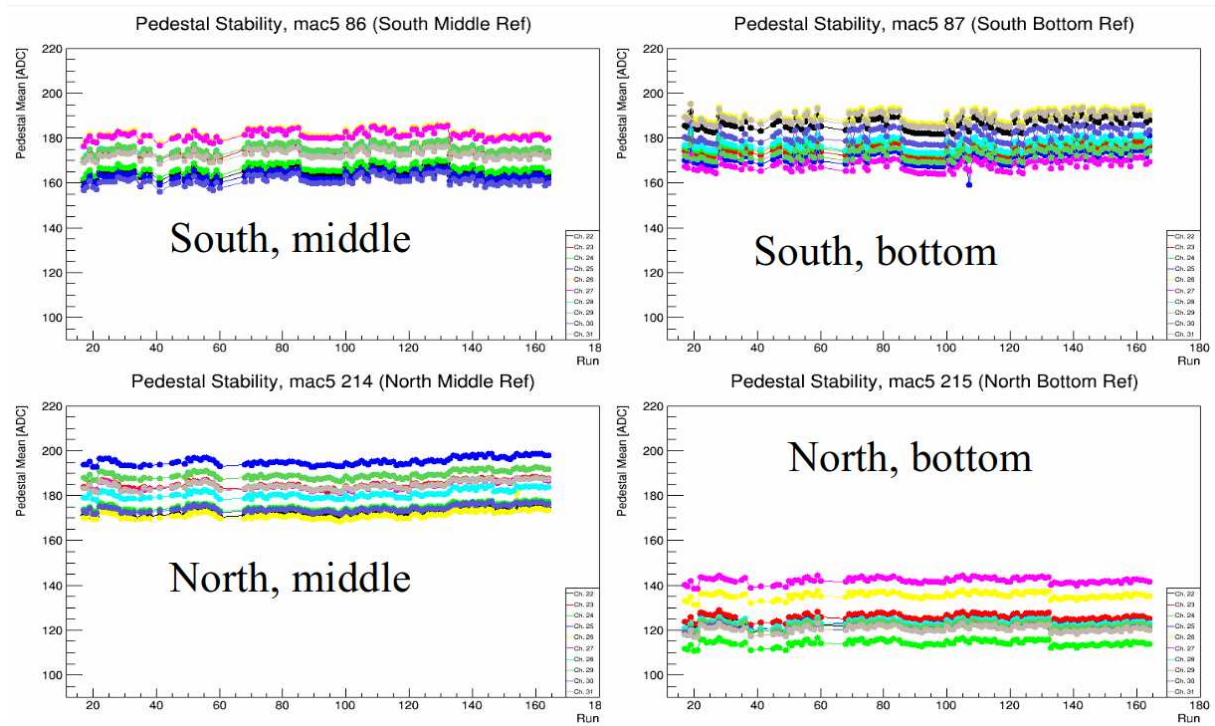


Figure 3.36: Pedestal values for each channel (one per color) over time for both the south and north side readouts (top and bottom) comparing tested to reference ORMs (left and right) show stability within 10 ADC. Shifts in pedestal value are coherent across all channels within but not across FEBs.

Tyler and I worked together to prepare a common analysis framework that would facilitate cross-analysis of the CSU and Fermilab test data sets. I had the idea to extract light yields for each pair of scintillator strips and readout end relative to the mean for all tested SiPMs. This would put channel-to-channel performance comparisons on equal footing so long as the all SiPM gains were similar. Tyler implemented this approach to

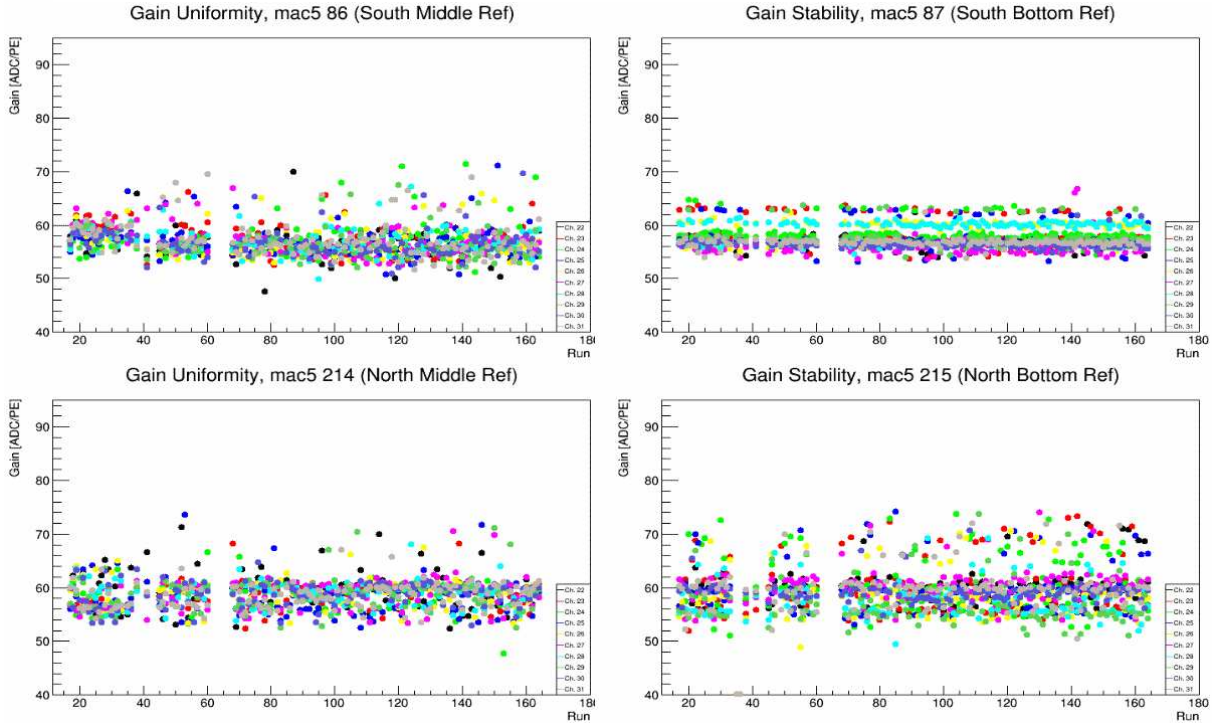


Figure 3.37: Gain values for each channel (one per color) over time for both the south and north side readouts (top and bottom respectively) comparing tested to reference ORMs (left and right respectively) show stability in most channels to about 2 ADC/PE and uniformity to about 10 ADC/PE. The gain fit occasionally fails, usually causing an overestimation.

generate a plot of the relative light yield for each SiPM yield versus the SiPM lateral displacement, also referred to as radial deviation, that was measured at CSU. The result, shown in Figure 3.40, indicates that the tolerances set by Simone’s simulation based study had produced the intended result with all SiPM relative light yields being uniform within our tolerance range.

The ORM production and testing campaign was a success. With 360 ORMs produced and tested, only 13 (4%) failed quality control checks. The data was used to select ORMs for the final CRT installation. Many lessons were learned along the way including how to monitor and fix grounding problems, how to most efficiently mount and seal ORMs on the MINOS modules, and how to automate and improve the calibration procedure. We gained some confidence in the system after observing good stability over several weeks.

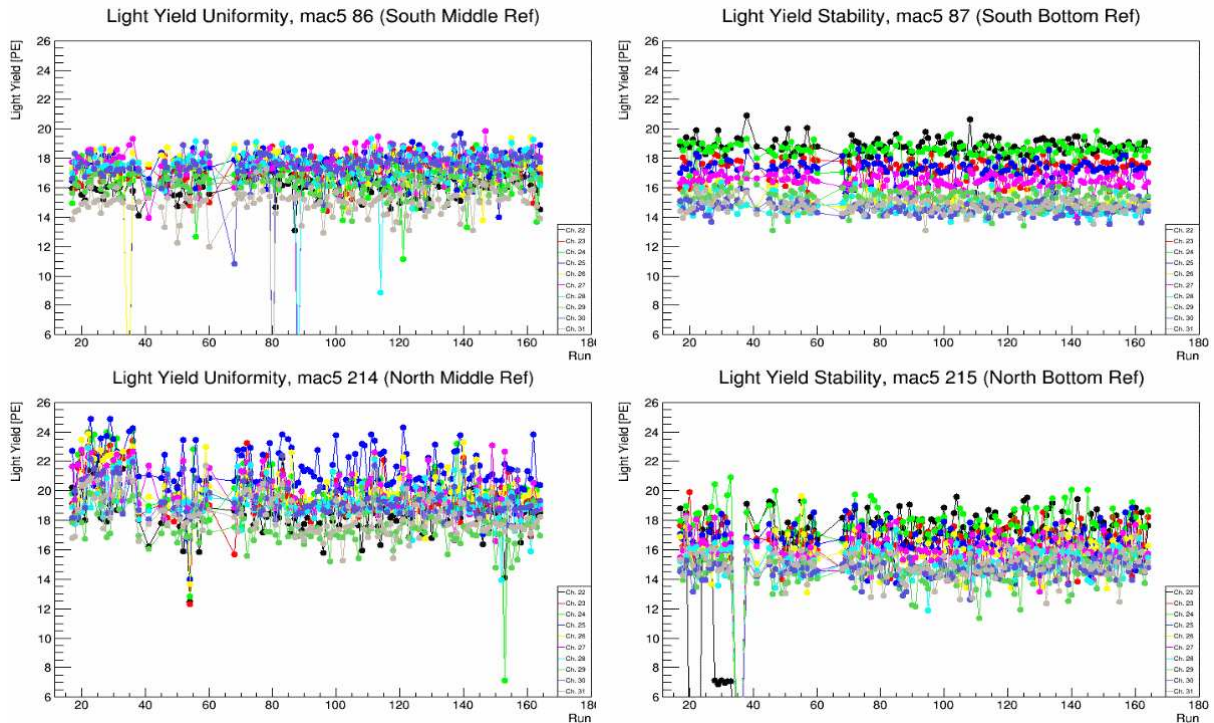


Figure 3.38: Light yield values for each channel (one per color) over time for both the south and north side readouts (top and bottom respectively) comparing tested to reference ORMs (left and right respectively) show stability to 10 ADC.




Figure 3.39: Test ORM efficiency relative to the reference ORMs is shown for both the south (green) and north (blue) side readouts plotted against light yield showing a strong correlation. A few ORMs had significantly lower efficiencies than the rest caused by a single bad SiPM.






Figure 3.40: The channel-by-channel light yield relative to the channel mean over all tests was compared to the SiPM radial deviation measured at CSU during fabrication, and it showed that performance was uniform within our tolerance range. Image credit: Tyler Boone.


3.7 Installation



(a) East/West CRT Walls.



(b) North Wall.



(c) South Wall.

Figure 3.41: We have eight different Side CRT regions with three different views shown here from a CAD model. Image credit: J. Tillman and C. James.

The Side CRT subsystem can be divided into eight regions, each with two layers of MINOS modules: north, south, west-north, west-center, west-south, east-north, east-center, and east-south. The South Wall (Figure 3.41c), upstream of the active volume and parallel to the drift direction, is deemed the most important for the Side CRT. This is the only Side CRT region with an X-Y configuration planned. To achieve this, we had to cut nine modules in half. The cut end would be the load bearing side, enabling the module to be oriented vertically. With the exception of the North Wall (Figure 3.41b), all other Side CRT regions (Figure 3.41a) consist of full-length modules with strips oriented

horizontally and parallel to the cryostat wall. Since three module lengths is longer than the cryostat, both center walls must be offset laterally in order to allow the modules to overlap. To guarantee access to the ORMs and FEBs that are covered by the overlap, these walls are also required to be movable. Thus, the mounting structure is fixed to a rolling platform, earning these walls the alternate names, East Rolling Wall and West Rolling Wall. These walls must also clear cable trays that run between the top of the cryostat and racks on the mezzanine level, which sets the number of rows of modules at eight. The stationary east and west walls use nine rows of modules.



Figure 3.42: The Side CRT West Utility Rack is pictured just before the rack’s enclosure panels were added.

The first portion of the Side CRT to be installed was the North Wall. Due to a lack of available space, the North Wall CRT modules had to be installed prior to the proximity cryogenics. ORM and FEB installation and cabling would take place at a later date. A total of 24 cut modules of three different lengths needed to be installed (Figure 3.41b),

twelve at 508 cm, four at 309 cm, and eight at 256 cm. The lengths were chosen to provide clearance for cryogenic piping while maintaining as much coverage as possible. These modules and the South Wall vertical modules were cut and prepared at Wideband following the procedure developed in Section 3.5.

The CRT has a total of four utility racks, providing power and timing signals, and two server racks, together housing 12 DAQ servers. The Top CRT will have one utility rack that has yet to be built. I built the east and west utility racks for the Side CRT. The test rack for the bottom CRT that I built with Simone was translated into a production utility rack for the Bottom CRT by Ryan Howell from the University of Rochester.

Continuing to focus on the Side CRT, I will provide a brief overview of the Side CRT utility racks. Each rack houses a single Wiener PL512 low-voltage power supply that provides a constant, 5-V DC output on up to 12 channels. The ground connection is made through the AC connection to the building power panel. This ground is propagated to all connected FEBs via the DC power return line and should be a significant improvement of what was implemented at Wideband for ORM testing.

In addition to the power supply, each rack houses a custom timing distribution box built and provided to us by our collaborator from the University of Pittsburgh, Vittorio Paolone. These boxes take two timing signals, one from the experiment wide GPS-disciplined PPS node, and another from the TPC trigger system, and fan them out to the CRT system.

Both racks have been completed and passed an operational readiness review. So far, only the bottom and west utility racks are in use. See Figure 3.42 for a photo of the rack in operation.

After cryogenic activities around the North Wall had slowed, we were approved to instrument the Side CRT North Wall. Significant assistance was provided by Tyler Boone and Prof. Thomas Coan of Southern Methodist University. To start, Tyler and I installed 24 ORMs and tested them for light tightness using a spare FEB and the standalone DAQ

running on a laptop. With Prof. Coan's help, we installed eight FEBs and routed the timing and network cables. This was a difficult task due to the proximity to the cryogenic equipment creating limited access. The pit-level FEBs required all three of us (Figure 3.43) with one person suspending the FEB mount from a rope from the mezzanine level, one person to stabilize the assembly, and one person to fasten the assembly to the Unistrut support post. Shortly after, the power distribution boxes arrived. Tyler and I mounted two power distribution boxes, added the individual twisted pair power cables to the bus bar, tested that all eight cables had the correct output voltage polarity, then powered up the FEBs. We were happy to observe that all channels were operational. Assistance with initial debugging work on the North Wall was provided by Francesco Tortorici of U. Catania.



Figure 3.43: It took three people to carefully install the FEBs in the proximity cryogenics area at pit level. T. Coan (right) and I (left) are shown practicing excellent team work. Not pictured (and photo credit): T. Boone.



Figure 3.44: Stabilizer plates for the Side CRT West Rolling Wall can be seen on either side of the gap between the mezzanine and the Warm Vessel.

Due to a delay in the start of cryogenic commissioning, we were afforded the opportunity to install a second portion of the Side CRT, the West Rolling Wall. There was a further possibility that the East Rolling Wall could also be installed. Due to a shortage in technicians at Fermilab, Tyler Boone and I were tasked with installing a large fraction of the hardware required for the mounting structure. In particular, we installed nearly all of the fiberglass Unistrut standoffs which electrically isolate the Unistrut support posts from the Warm Vessel. This was followed by stabilizer brackets that tied the columns of three posts together. All of this was a prerequisite for the next step, mounting stabilizer plates for the rolling walls. This was a difficult job as many places where plates had to be attached were difficult to reach. These plates were custom pieces, and ease of installation was not a consideration. Nevertheless, we prevailed. These plates can be seen, after installation, in Figure 3.44.

With the utility rack already in place and the experience from installing and instrumenting the North Wall, this second round went much faster. I led the effort as I had for the North Wall with significant assistance from Tyler and additional support from recent CSU PhD, Matthew Hogan. We had a strict deadline to motivate us set by the start of cold commissioning. We installed and instrumented the entire West Rolling Wall in three short weeks.

Figure 3.45 shows one of the full-length MINOS modules being lowered into the pit to be installed in the West Rolling Wall. Installing the MINOS modules was the most difficult step in the West Rolling Wall installation. The typical lifting fixtures that would be used to handle these modules could not be used due to limited crane access at the destination. Getting the modules out of the shipping crates and into hoisting slings, braced by Unistrut channel, required about five people with a change in slings on top of the cryostat.

A panoramic photo, Figure 3.46, shows the West Rolling Wall after all modules were installed from the mezzanine. This photo fails to capture the scale of the detector, covering a surface area of 54 m². Photos taken at intermediate stages of the installation and the beginning of instrumentation are shown in Figure 3.47.

Following the installation and instrumentation of the North Wall and the West Rolling Wall, Tyler, Matthew, and I set out to run some preliminary tests in order to debug the system. We found that there was a non-zero contribution to the FEB trigger rate due to electrical noise, checking in the same way we as we did during ORM testing (Section 3.6). We identified some poor connections at the ORM cables. Most of these could be fixed with a simple disconnect-reconnect cycle, while two cables appeared to be defective. The noise rate improved following this intervention, but the problem was not completely resolved. With support from Linda Bagby, the SBN electrical coordinator, we determined that the entire Side CRT support structure was missing a reference to ground. A technician installed a daisychain of ground braids that connected all of the Unistrut



Figure 3.45: From left to right: T. Boone, J. Judd, and K. Hardin work to install the next MINOS module into the West Rolling Wall.

support posts. The daisychain was connected to building ground via a nearby I-beam. With the mounting structure grounded, additional ground connections were made between the the MINOS modules and the Unistrut support posts. Following this intervention, the noise rate disappeared.

With this, we turned our attention to additional, likely, future sources of externally induced electrical noise. Our primary concern was the proximity of the pit-level FEBs to the cryogenic pumps. We arranged a time with Linda to operate the variable frequency drives (VFDs) that power the pump motors. We found that the liquid nitrogen system VFDs induced several kilohertz of noise rate while the liquid argon system VFDs caused our FEBs' trigger rate to become saturated at about 35 kHz. We worked with a member of the cryogenics group, Trevor Nichols, to perform several cross-checks of the power distribution with regard to the VFDs. Unfortunately, no solution was found prior to the start of cold commissioning, but we monitored the situation during the CRT commissioning phase, discussed in Chapter 5.

Work on resolving the VFD-induced noise issue will continue pending access to Fermi-



Figure 3.46: A panoramic view of the West Rolling Wall shortly after the MINOS modules were installed attempts to show the scale of the detector.




Figure 3.47: Two views of me laboring away to ensure ICARUS is protected against cosmogenic backgrounds showing me securing the second row of MINOS modules (right) and me installing the first ORMs on the West Rolling Wall (left). Photo credit: Tyler Boone.

lab during the COVID-19 pandemic. Overall, we successfully installed and instrumented about 25% of the Side CRT. Along the way, we learned some valuable lessons that should make the remaining Side CRT installation process smoother.

Chapter 4

Simulation and Reconstruction

In this chapter, I start with a broad overview of the simulation, reconstruction, and analysis tools currently available. Next, I will focus on realistic detector simulation and reconstruction tools for the CRT that I have developed. I will end the chapter with a discussion about combining the different detector subsystems together in higher level reconstruction.

4.1 Overview of LArSoft and icaruscode

In an effort to promote synergy and to facilitate direct performance comparisons between the different experiments, the SBN program as a whole has opted to use the LArSoft [63] software base, which is built upon the HEP art event processing framework. LArSoft is a C++ based infrastructure and algorithm set for reconstruction, simulation, and analysis of data for and from LAr TPCs. The LArSoft collaboration maintains the code and implements improvements or new features based on input from LAr TPC experiments. Notable contributions have come from ArgoNeuT [39] and MicroBooNE [40]. LArSoft utilizes abstract classes for experiment specific code. This means that a significant portion of the work is done for us.

Any experiment that uses LArSoft is responsible for implementing the experiment specific code including geometry and the detector readout simulation. These implementations reside in experiment specific repositories. In ICARUS, this repository is called icaruscode.

ROOT libraries are imported to LArSoft. Analysis is facilitated using abstract ana-

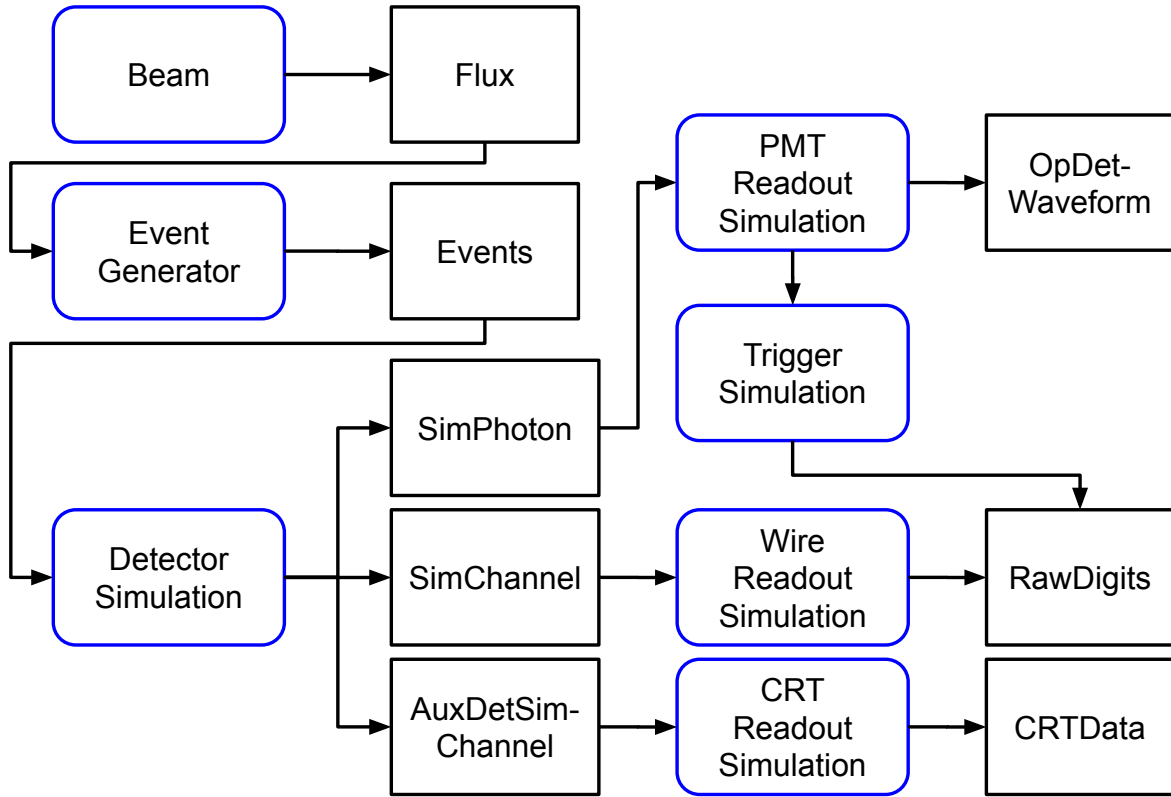


Figure 4.1: The simulation workflow for SBN showing the simulated data produced for the three detector subsystems: TPC, PDS, and CRT.

lyzer modules that convert the native artROOT formatted files into ROOT histograms, graphs, trees, or ntuples.

Below, I will briefly describe the different elements of the simulation and reconstruction chains. I begin with the simulation chain, shown as a flow diagram in Figure 4.1.

4.1.1 Geometry

A realistic geometric description has been developed from a combination of survey measurements, engineering drawings and CAD models. It has been written in the gdm language and includes the following:

- surface building, concrete pit, dirt surrounding the pit, 3-m concrete overburden;
- steel outer cryostat approximated as a hollow box with the same mass as the real




Figure 4.2: Two different views of the CRT system are shown with the individual CRT modules (green), top CRT support I-beams (grey), and warm vessel (red) visible.

cryostat;

- foam insulation between outer and inner cryostats;
- aluminum inner cryostats;
- inner stainless steel structure for TPC (beams and columns);
- each individual TPC wire;
- PMTs approximated by hemispheres;
- CRT modules approximated by hollow aluminum boxes containing polystyrene strips;
- Top CRT steel support I-beams.

With my intimate knowledge of the CRT system, I wrote a Python generator script that is responsible for generating all of the CRT gdml fragments that eventually get merged into the master gdml file that is input to the simulation chain.

4.1.2 Generators

We rely on external, open source software packages interfaced with LArSoft to handle often complex particle interactions that go on to induce activity in our detectors. We can

group these into three categories: single particles, neutrinos, and cosmic rays. I will briefly summarize them here.

Single particles are the simplest of the generators. We use Geant4 [67] to simulate this initial particle. I will discuss more on this in the next subsection.

There are several neutrino event generators available. GENIE [64] has been growing in popularity within the experimental neutrino physics community. The generator features several appealing features, including

- flexible, experimentalist-centric design;
- consistent, validated, and efficient implementation of standard and alternative physics models;
- integrated comparison, tuning, and reweighting tools;

and it is the generator of choice for SBN. Support for GENIE is provided by active GENIE collaborators that are members of our collaboration.

LArSoft supports two different cosmic ray generators, CRY [65] and CORSIKA [66]. Each has advantages and disadvantages over the other. CRY uses precompiled input tables that account for primary cosmic rays with energies in the range of 1 GeV through 100 TeV. CORSIKA simulates the full airshower starting from the primary cosmic ray with energies in the range 1 GeV through 100 EeV. It performs four-dimensional simulations and includes the interaction and decay of nuclei, hadrons, muons, electrons and photons in the atmosphere. These particles are tracked through the atmosphere until they undergo interactions or decay into secondary particles. CRY is preferred in terms of computational speed while CORSIKA is generally considered to produce a more realistic simulation.

4.1.3 Detector Simulation

The event generators provide particles with initial state position and momentum four-vectors. The next step is to take these initial states and propagate them through our

detectors and surrounding materials. For this task, we use the Geant4 toolkit. This provides support for user defined geometry, particle tracking, and a comprehensive library of physics models covering electromagnetic, hadronic, and optical processes. The optical simulation is of particular importance for my work, presented in Chapter 6.




Figure 4.3: The photon arrival probability is shown for two different slices along the drift coordinate: immediately adjacent to the anode wire planes (a) or immediately adjacent to the cathode plane (b).

Scintillation light is produced in the same location as the ionization charge when a charged particle passes through the liquid argon. The light has two components, dubbed “fast” and “slow”, with characteristic times of 6 ns and 1.6 μ s respectively. The Rayleigh scattering length in liquid argon at the scintillation peak wavelength is about 95 cm, preserving directionality from the point of view of the PMTs. Measuring both components of the light can be used to more accurately measure the energy deposited in the argon, improving calorimetric measurements as ionization and scintillation are anti-correlated.

Optical simulation is computationally challenging. In liquid argon, ionizing particles generate on the order of 40,000 photons/cm with the precise value depending on the

particle species, the deposited energy per unit length, and the magnitude of the electric field. For example, protons, charged pions, and muons generate 19,200 photons per MeV of deposited energy in an 500 V/cm drift field. With such large numbers, tracking each generated optical photon for every interaction is simply not practical. There are several approaches to address this issue. The most common approach is to use a lookup table, sometimes called a photon library. This is the default approach taken by ICARUS; however other approaches, such as a semi-analytic approach being adopted in SBND, are under consideration. I will only consider the photon library approach here.

To start, the entire liquid argon volume is divided into cubic volumes called voxels. In ICARUS, 5-cm voxels are used amounting to about two million for a single cryostat. In each voxel, 10^5 photons are generated isotropically with energies matching the scintillation spectrum for liquid argon. A table is filled for each PMT with the probability that a photon originating from a particular voxel with a given initial momentum will arrive at the PMT. The photon library, queried for two different slices along the drift coordinate in one TPC, shows the visible fraction of photons arriving at each PMT in Figure 4.3. Close to the PMT array, the photon visibility strongly depends on the proximity to the PMT. As the distance to the PMT increases, the visibility becomes uniform. A clear boundary can be seen between the active and inactive volumes due to the field cage.

Since individual photons are not tracked, we must use a parameterization for the photon propagation time. ICARUS has adopted a model developed by Prof. Diego Garcia Gamez for SBND. The photon sample generated for filling the photon library is used to generate time-of-arrival distributions for each voxel. The distributions are fit with a combined Landau-exponential function. The parameters are then stored in the library. An example of such a fit is shown in Figure 4.4 along with the distance dependence of the arrival time and the spread in arrival time. As the distance to the PMT increases, Rayleigh scattering becomes a significant effect. Note that the arrival time is larger than the expectation from the group velocity alone due to scattering and reflections.




Figure 4.4: An example of a photon arrival time parameterization fit is shown on the left. On the right, the photon propagation versus distance to the PMT is shown with error bars representing the RMS.

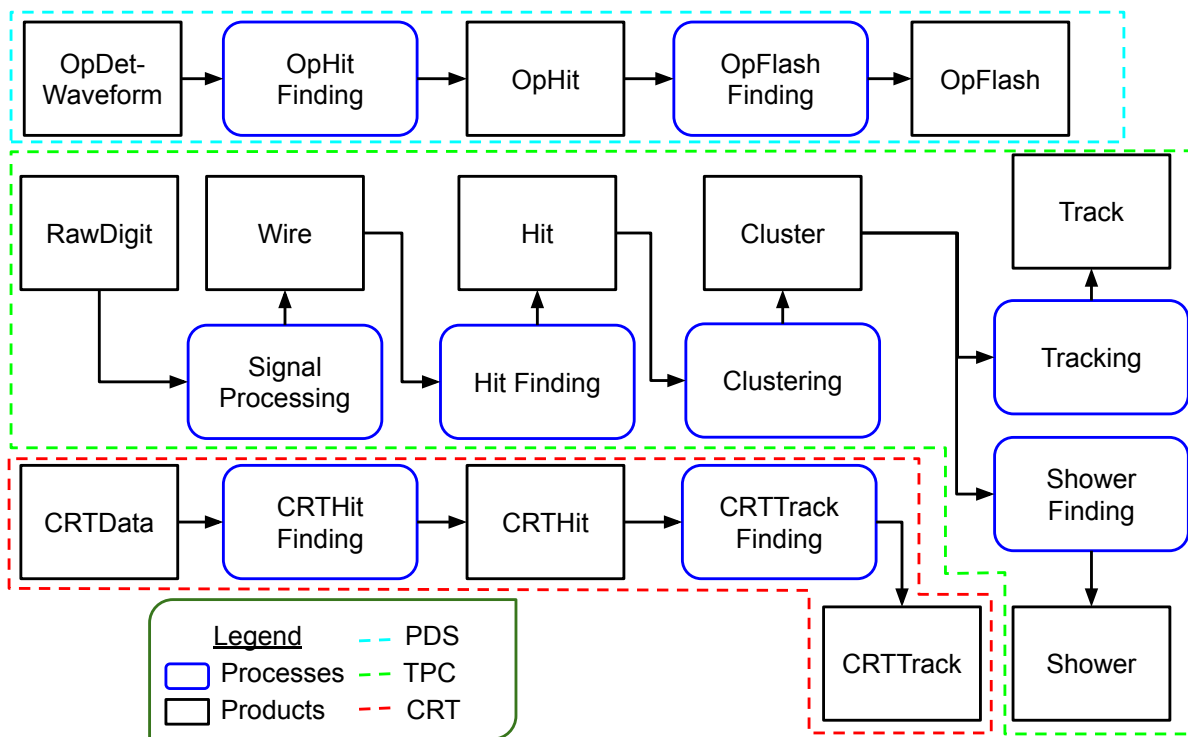


Figure 4.5: The reconstruction chain is shown for the three detector subsystems: TPC, PDS, and CRT.

4.1.4 TPC Readout Simulation and Reconstruction




Figure 4.6: A TPC waveform is shown after noise filtering with three hits identified. Image credit: F. Varanini.

The drift field response to ionization charges is modeled using a two-dimensional Garfield simulation [68]. Starting from the ionization charge generated in the detector simulation, the charge is “drifted” to the wire planes. This also includes the diffusion of ionization electrons along the drift path. The field response of the wires due to the drifting charge is also described here. As the charge passes the induction planes or is collected on the collection plane, voltage pulses are induced on the wires. This simulation chain is illustrated in Figure 4.1.




Figure 4.7: Individual clusters can be seen within PANDORA-reconstructed tracks from a ν_μ CC event. Credit: B. Howard.

Identifying the charge signals and characterizing their width and amplitude is the task of TPC hit finding. There are two approaches currently under study. One approach was developed by the original ICARUS collaboration using the raw waveforms and is known as raw hit finding. The other approach attempts to recover the intrinsic charge versus

time, expected to have a Gaussian profile, through a deconvolution process [68]. This approach is known as Gaussian hit finding. An example of this, performed after signal processing, is shown in Figure 4.6 above.

The signals provide information about the spatial charge distribution and the charge density of the ionization charge that generated them. As all of the charge is ultimately collected and measured on the last wire plane, TPC hits can be used for calorimetry in addition to reconstructing the event topology.

Correlating the TPC hits in space and time, a process known as clustering (illustrated in Fig. 4.7), allows one to reconstruct a track or shower. There are several ways to do this. The current approaches used in SBN are PANDORA [69] and TrajCluster [68]. The TPC reconstruction chain is summarized diagrammatically in Figure 4.5.

I will not make use of the TPC readout or reconstruction in my analysis, so I will not provide any more detailed discussion on this topic. However, I will say more about future work that could incorporate TPC reconstruction into an extension of my work on background rejection (Section 6.4).

4.1.5 PDS Readout Simulation and Reconstruction

The PMT signal formation and digitization modeling (Figure 4.1) builds on experience from the CERN PMT test stand. In situ calibration of the PMT response using well controlled, laser sourced optical pulses delivered by optical fibers to each PMT will provide additional data for tuning the PMT readout simulation. An example of a simulated PMT waveform is shown for a single photoelectron in Figure 4.8.

PMT hit finding (OpHits) and flash reconstruction (OpFlash) is modeled after the MicroBooNE approach [70]. Similar to TPC hit finding, PMT hit finding identifies PMT signals above the baseline, determines the amplitude of the signal, and at what time the signal occurred. Note that we can have multiple OpHits per PMT waveform due to early and late light for example; a single ionization event can generate many OpHits. See




Figure 4.8

Figure 4.9. OpFlashes are built from coincidence-gated OpHits. Based on the total light observed in each OpHit, we can reconstruct the two-dimensional projection of the event onto the PMT array. This reconstruction chain is illustrated in Figure 4.5.




Figure 4.9: Optical hit finding identifies PMT signals and extracts the time of the signal pulse as well as the number of photoelectrons associated with it (left). Time resolution depends on the pulse amplitude and the level of scattering as shown by protons with two different energies spanning a full drift distance (right). Image credit: K. Terao.

The topology of the event is loosely correlated with the light distribution of the OpFlash. See Figure 4.10 for an example using simulated BNB neutrino events. This can be used to match different clusters of ionization activity in the TPC readout to OpFlashes, a process known as flash matching. By associating a precise time stamp to each

cluster, this process yields the drift coordinate, thereby disambiguating three-dimensional reconstructions of the clusters occurring at different times in the TPC readout window. This step is crucial since, due to the long drift time, interactions outside of the beam spill will be overlaid onto the in-spill event, from cosmogenic muons for example. The time measured from the wire signals cannot be used to this end as this is only the time of arrival of the ionization charge at the wire planes.

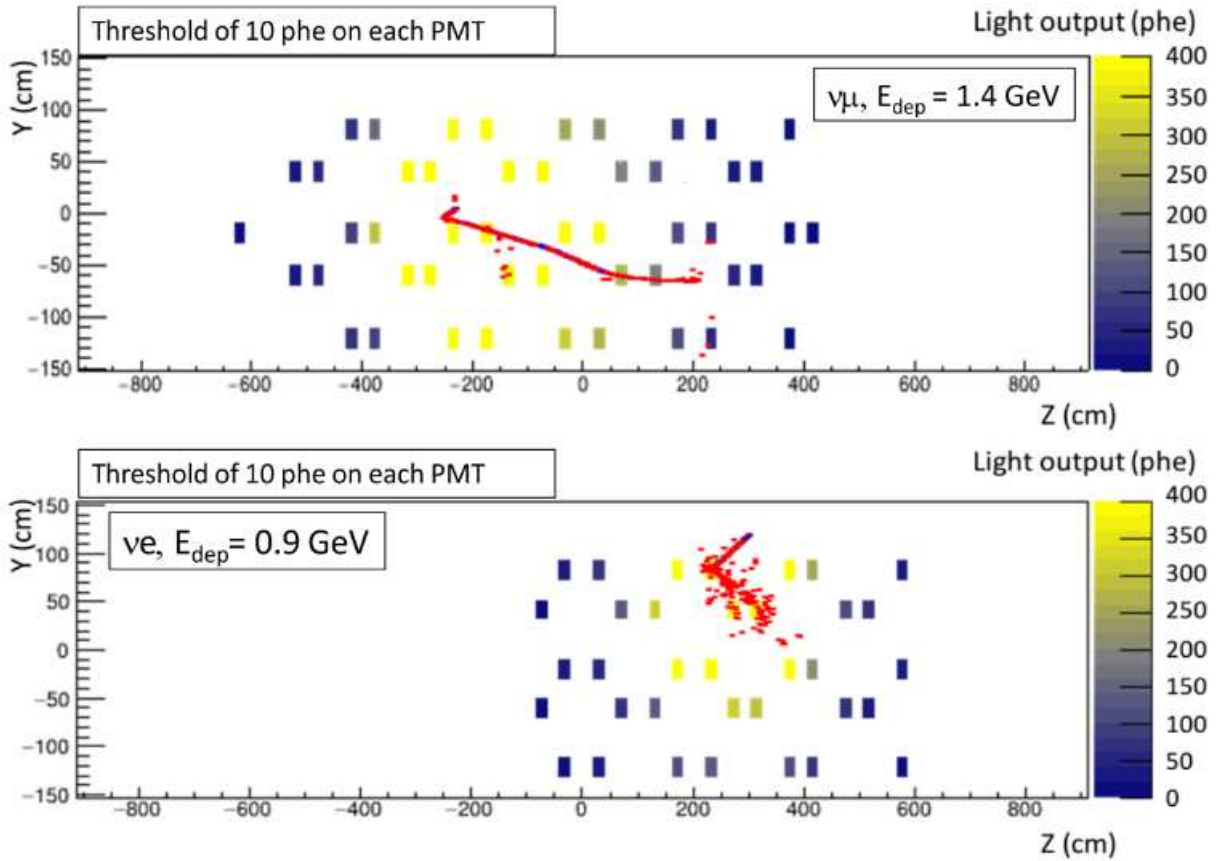


Figure 4.10: The barycenter of photoelectron distribution traces the approximate neutrino vertex position as shown for ν_μ CC (top) and ν_e CC (bottom) events. Image credit: A. Menegolli.

4.1.6 CRT Readout Simulation and Reconstruction

The simulation and reconstruction for the CRT had to be implemented from scratch as LArSoft does not yet have this feature. What LArSoft does support is the presence

of auxiliary detectors present outside of the cryostat. This allows the basic detector simulation (i.e. Geant4) to be extended to the CRT volumes as defined in the geometry. Thus, the CRT detector simulation starts from the deposited energy, track IDs, entry and exit positions and times. The rest of the detector simulation, e.g. how much light is produced in the scintillator, the readout simulation, and any reconstruction must be implemented from scratch. See Figure 4.1 for an illustration of the simulation workflow.

With LAr TPCs being inherently slow and several instances of these detectors operating on the surface, it is foreseen that the CRT related code developed for the SBN detectors will eventually be introduced into LArSoft. As is the case for the other subsystems, SBN collaborators working on the CRT simulation and reconstruction are making an effort to make the code as flexible and detector agnostic as possible. Unfortunately, with the ICARUS CRT being comprised of three different subsystems with some significant differences from the other SBN CRTs, this is a difficult task. Most of the difficulty comes in the detector readout simulation as well as the first step in the reconstruction. However, the reconstruction products (Fig. 4.5) can all be defined in the same way for all SBN CRTs. There are some ideas about how to abstract all of the CRT related classes and algorithms, but I will not discuss them here. It suffices to mention that the work on this topic in SBN could find uses in other current or future LAr TPC experiments operating on the surface.

I devote the next two sections to my work on the CRT detector simulation and low-level reconstruction respectively. The results presented here will be crucial for the remaining chapters in this dissertation.

4.2 CRT Detector Simulation

The task of the detector simulation is to convert energy deposited in the scintillator strips into an analog SiPM signal. The signal is then injected into the detector readout simu-

lation, simulating the front-end electronics, including gain, charge resolution, threshold, time stamp generation, and trigger logic. In short, the detector simulation should reproduce the data obtained from measurement. I used data obtained at the Wideband test stand (Chapter 3) to develop models for the detector response that I will discuss below.

4.2.1 Detector Response Model Development

There are two approaches to modeling the light production in the scintillator. First, if all quantities related to scintillation, attenuation, and reflection are well known, a full optical simulation can be carried out, facilitated by Geant4. While the CRT volume is not nearly as large as the argon volume, the number of photons to be tracked per energy deposition is nonetheless impractically large. We encounter the same issue with insufficient computing power as described in Section 4.1.3. We could employ a similar solution as used for the optical simulation in the argon, computing a photon library, but this approach has not yet been pursued. The second method is to use an empirical model developed from measurements of the detector response. This approach has been adopted and is described below.




Figure 4.11: The light yield for four different SiPM channels for cosmogenic muons at different distances from the SiPM, as measured at our test stand, is fit with a quadratic.

From the test stand measurements, I know the light yield from cosmogenic muons

as a function of longitudinal distance for the MINOS modules. These measurements are a convolution of the scintillator light yield, the scintillator bulk attenuation length, the reflection coefficients of the diffuse reflector at the scintillator surface, the trapping efficiency of the optical fiber, the attenuation length of the optical fiber, the optical coupling efficiency between the optical fiber and the SiPM, the SiPM quantum efficiency, and Poisson statistics of the observed number of photons. From the point of view of the detector simulation, I am only concerned with accurately modeling the amount of charge measured by the front-end board for a particular energy deposit at a particular position in the scintillator strip and the resulting time stamp. I do not need to concern myself with what effect each parameter has on the resulting spectrum so long as the net effect is captured.




Figure 4.12: The coordinates obtained from Geant4 are used as input to the light yield model.

The model I developed for calculating the number of optical photons that arrive at the photodetector is given in Equation 4.1, and it depends on the average longitudinal and transverse positions of a throughgoing track in the frame of the scintillator strip, z and x respectively, and the amount of energy deposited, E_{dep} . The domain of this function is set by the module geometry where the strip length is L and the strip width is W . See Figure 4.12 for an illustration. The index of Y refers to the detector subsystem being modeled: Top (c), Side(m), Bottom(d). The model is factorized into three parts: the intrinsic light yield, $Y_{0,i}$, currently the same for all subsystems; the longitudinal attenuation (Eq. 4.2) primarily due to the optical fiber, $a_{l,i}$, currently the same for all subsystems; and the transverse attenuation primarily due to the bulk scintillator, $a_{t,i}$, currently ignored for

the Side and Bottom subsystems (Eq. 4.3) but modeled for the Top subsystem (Eq. 4.4). The transverse attenuation depends on the distance from the optical fiber at $x = x_f$.

To obtain the parameter values for Equation 4.2, I used the measurements shown in Section 3.5. Recall that, in order to obtain these light yield values, I fit the resulting hodoscope-triggered spectrum at each longitudinal position with a convolution of a Landau with a Gaussian. The most probable value of the fit is taken as the light yield. I performed a quadratic fit to each of the four channels (Fig. 4.11) and took the average of the fit parameters for my model. This fit, evaluated at $z = 0$, yields the intrinsic light yield parameter, Y_0 . Dividing the fit parameters by this value yields the parameter values shown in Table 4.1. Note that Equation 4.2 includes not only the attenuation effects from the optical fiber but also also corrections for the photodetector quantum and optical coupling efficiencies.

While only the most probable value of the light yield is included in the model, the Landau portion of the light yield spectrum is included via the dependence on the deposited energy, which comes from a sampled Landau distribution. The constant that appears in Equation 4.1, E_{MIP} , is set to the most probable value of a minimum ionizing muon depositing energy over 1 cm.

$$Y_i(x, z, E_{dep}) = Y_{0,i} a_{l,i}(z) a_{t,i}(x) \frac{E_{dep}}{E_{MIP}} \quad 0 < z < L, |x| < W/2, i = c, m, d \quad (4.1)$$

$$a_{l,i}(z) = a_l^2 z^2 + a_l^1 z + a_l^0 \quad (4.2)$$

$$a_{t,m}(x) \equiv a_{t,d}(x) \equiv 1 \quad (4.3)$$

$$a_{t,c}(x) = \begin{cases} a_t^{c5} x^5 + a_t^{c4} x^4 + a_t^{c3} x^3 + a_t^{c2} x^2 + a_t^{c1} x + a_t^{c0} & |x| \leq x_f \\ a_t^{n4} x^4 + a_t^{n3} x^3 + a_t^{n2} x^2 + a_t^{n1} x + a_t^{n0} & x > x_f \\ a_t^{f3} x^3 + a_t^{f2} x^2 + a_t^{f1} x + a_t^{f0} & x < -x_f \end{cases} \quad (4.4)$$

Simply scaling the light yield with the deposited energy is not completely valid. It

Table 4.1: Parameter values used in the CRT detector simulation are summarized below (see Equation 4.1, 4.2, 4.4)

Parameter	Value	Parameter	Value	Parameter	Value
Y_0 [PE]	36.54	E_{MIP} [MeV]	1.891	a_t^2 [cm ⁻²]	0.01024
a_t^1 [cm ⁻¹]	-0.1749	a_t^0	1.0	a_t^{c5} [cm ⁻⁵]	-3.891 x 10 ⁻⁵
a_t^{c4} [cm ⁻⁴]	1.480 x 10 ⁻⁴	a_t^{c3} [cm ⁻³]	6.366 x 10 ⁻⁴	a_t^{c2} [cm ⁻²]	-7.076 x 10 ⁻⁴
a_t^{c1} [cm ⁻¹]	-0.02045	a_t^{c0}	0.6830	a_t^{n4} [cm ⁻⁴]	0.001279
a_t^{n3} [cm ⁻³]	0.04492	a_t^{n2} [cm ⁻²]	0.5956	a_t^{n1} [cm ⁻¹]	3.546
a_t^{n0}	8.789	a_t^{f3} [cm ⁻³]	7.818 x 10 ⁻⁴	a_t^{f2} [cm ⁻²]	-0.01982
a_t^{f1} [cm ⁻¹]	0.1682	a_t^{n0}	0.1399		

is well known that scintillation is subject to quenching effects. The variables used in the quenching model that follows are distinct from the light yield model shown above. The larger the ionization density is (dE/dx), the less scintillation light (L) is produced. A commonly used empirical model describing this phenomenon, Birks' Law, is given in Equation 4.5. k_B is known as Birks' constant and is a property of the scintillator. For polystyrene based scintillator like we use in the CRT, $k_B=12.6$ cm/GeV. This parameter can be thought of as the saturation point of the scintillator, corresponding to a dE/dx of about 80 MeV/cm. For comparison, the most probable ionization density for a MIP muon is about 1.9 MeV/cm in polystyrene.

$$dL/dx = L_0 \frac{dE/dx}{1 + k_B dE/dx} \quad (4.5)$$

In Figure 4.13, I have plotted the fractional loss in light yield due to Birks' quenching. For MIP particles, the loss is at the level of a few percent. This correction should already be implicitly present in our measurements of the most probable light yield used to develop the model presented above. However, the correction should be made particles with a dE/dx greater than about 10 MeV/cm. This is primarily expected for heavier particles, i.e. protons. With 3 m of concrete overburden, the only source of such particles that could produce CRT tags are neutrino events. I will discuss this case at length in Chapter 6, including the impact of quenching. For now, I will not include quenching since I have

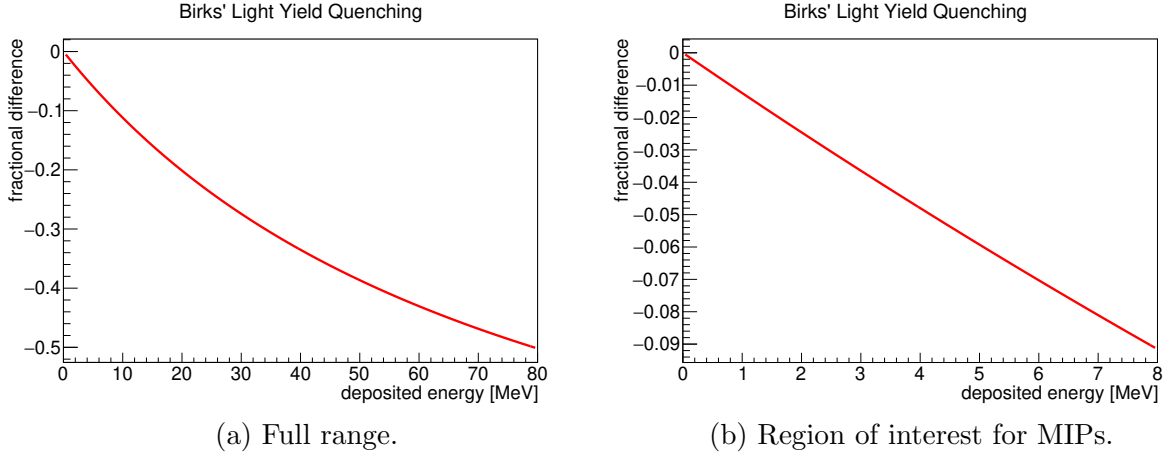


Figure 4.13: I plotted the fractional change in light output from polystyrene based scintillator vs. the energy deposited in 1 cm. The full range in ionization density up to saturation point is shown on the left while the range relevant to cosmogenic muons is shown on the right.

only recently implemented the model and it needs to be tested.

The number of photons detected by the photodetector obeys Poisson statistics. This is implemented by sampling from a Poisson distribution with a mean equal to the light yield produced by Equation 4.1. Without loss of generality, I will focus on the Top and Side subsystems using SiPMs. The charge generated by the SiPM is determined by the gain and the number of microcell discharges. The total charge produced is the coherent sum over the charge produced per microcell, equivalent to the product of the gain and number of microcell discharges. Since the number of photons impinging on the SiPM is much less than the number of microcells, I can safely ignore nonlinear effects.

Before continuing, I will comment on some caveats regarding the validity of this implementation of the CRT detector simulation. First, I should justify why we can reasonably ignore transverse attenuation for the Side and Bottom subsystem. The Bottom CRT strips are nearly identical to the MINOS strips except that they are 5-cm wide instead of 4-cm wide. Focusing on the Side subsystem, effects from attenuation in the bulk scintillator are expected to be negligible for two reasons. First, for horizontally oriented modules, a throughgoing track will be at most 2-cm from the optical fiber. From my previous R&D

studies at CSU where our 17-cm wide prototype had the same thickness as the MINOS strips, we measured an effective bulk attenuation length of 5 cm. [55] measured a bulk attenuation length of 38 cm in strips made from nearly identical scintillator that were 0.5 cm thick and 2 cm wide. The effective bulk attenuation length depends on the aspect ratio. From this, I estimate a decrease in the light yield at the strip edge relative to the strip center of 5-33%. Approximately 60% of throughgoing tracks will pass within 1.2 cm of the fiber leading to a maximum bulk attenuation loss of 3-21%. Now, all of the MINOS modules will be vertically oriented. Considering that there are very few horizontal cosmogenic muons, the distances from the optical fiber given above set the upper limit on these losses. However for the Bottom CRT modules are horizontally oriented. The second reason the transverse attenuation can be neglected in calculating light yield is that these effects are averaged out in the light yield measurements. They are effectively included in the parameters extracted for Equation 4.2.




Figure 4.14: A transverse scan with a collimated β source shows the impact of bulk attenuation on the light yield of a Top CRT strip.

The argument above for neglecting the transverse attenuation within the bulk scintillator does not apply to the Top CRT modules where each strip is 23-cm wide. In this

case, bulk attenuation is a significant effect. At the test stand at CERN, during the R&D phase, a collimated ^{90}Sr source (β emitter) was used to measure the light yield as a function of transverse position for a prototype strip design resembling what has been used for production. The result, shown in Figure 4.14, demonstrates up to a 60% decrease in the light yield relative to the peak yield at the fibers. Using this data, I developed a transverse attenuation model for the Top CRT modules only, shown in Equation 4.4. One caveat with this data set is that the fibers were not glued into the slots. This is expected to have a meaningful impact on the effective attenuation.

The most significant caveat of my implementation of the detector simulation is that I have applied the same light yield and longitudinal attenuation model, which also includes effects from SiPM quantum efficiency and optical coupling efficiency, to all CRT subsystems. This is done out of necessity due to a lack of well controlled light yield measurements for the Bottom and Top CRT modules. When these samples become available, the detector simulation can easily be updated by replacing the parameter values in Equation 4.2 with the subsystem specific values.

While each subsystem utilizes the same conceptual design, there are some important differences. For the Bottom CRT modules, the optical readout and front-end electronics are significantly different from that of the side CRT; however, the light yield, attenuation, and internal reflection characteristics should be similar to the side CRT since the scintillator material and aspect ratio as well as the optical fibers are nearly identical to those for the side CRT. One important difference is that the optical fibers for the Side CRT modules are glued into a slot, providing optimal coupling to the scintillator. The strips used in the Bottom CRT modules house the fibers in a coextruded hole. According to my studies in Section 3.2, this should reduce the coupling efficiency of the fiber by a factor of two. However, the Bottom CRT optical fibers are 25% larger than the Side CRT fibers. This should offset the coupling loss. Another significant difference will be in the coupling efficiency of the PMT to the optical fibers and the quantum efficiency of the

PMT. Based on these considerations, I would predict that my current model estimates the light yield of the Bottom CRT modules withing a factor of approximately two. Even if the light yield is reasonably modeled, there is no reason to expect the attenuation profile to be accurately modeled as the modules are several meters shorter than the MINOS modules, and the fibers are mirrored at the far end relative to the PMT.

For the Top CRT modules, the scintillator is similar and the same optical fibers are used albeit with a 17% smaller diameter. The SiPM-fiber coupling efficiency should be similar to the Side CRT, and the front-end electronics are the same. The primary difference is the significantly larger width of the scintillator strips where we cannot ignore the impact of transverse attenuation, but this has been accounted for. As with the Bottom CRT modules, the optical fibers are mirrored at the far end from the readout. With the modules being several meters shorter than the Side CRT modules, this is expected to impact the attenuation profile.

Having discussed the light yield model, we can move on to the other half of the detector simulation, modeling the front-end board (FEB) response. The raw SiPM signal passes through a low-noise preamplifier with a configurable charge multiplier. The amplified signal is then shaped by a RC-CR shaping circuit before being integrated with the total charge passed to a discriminator. If the charge is above threshold, the FEB trigger candidate proceeds to the application of coincidence logic. See Section 3.2.3 for more information.

For the charge measurement, an ASIC converts the analog charge measurement to ADC counts. The correspondence between analog charge value and ADC counts is determined through dedicated calibration measurements. See Section 3.2.5 for more details on the calibration procedure. The intrinsic electrical noise of the front-end board manifests itself as a pedestal corresponding to 0 photoelectrons. For a low-photon flux, like we have here, the relation between the total charge and the number of SiPM microcell discharges is linear and measured as an effective gain in units of ADC counts per photoelectron (PE).

The total ADC counts is given by the sum of the mean pedestal value and the product of the effective gain and the number of PE produced. In the detector simulation, I apply the same typical values for the pedestal and effective gain to all SiPM channels. The produced charge is compared to the discriminator threshold. If the total charge is above threshold, the charge in all 32 FEB channels produced within the 30-ns shaping time is latched into a sample-and-hold circuit while the coincidence logic is assessed.

The coincidence logic will differ for each CRT subsystem. Generally, a coincidence between adjacent layers of scintillator strips is applied. This is to reject radiogenic backgrounds and SiPM dark noise. However, differences in the implementation of each subsystem affect how this coincidence is applied.

For the top CRT, the coincidence is applied within a single FEB. I first require a coincidence between SiPMs within the same strip before the SiPM channels can be used in a coincidence between layers. The next condition is that there must be a coincidence between at least one strip from each of the adjacent layers within the same CRT module. In order to improve muon tagging efficiency while still suppressing the SiPM dark noise, the CERN group is investigating the possibility of making an OR between SiPMs in the same strip while still keeping the AND between layers; however this logic has not yet been implemented in the simulation.

For the side CRT, there are multiple FEBs within a layer. The coincidence is applied between FEBs reading out modules in adjacent layers of a particular Side CRT region. These regions were identified in Section 3.7. Both FEBs must have at least one SiPM channel above threshold. When this happens, a validation signal is distributed to the adjacent layer. If a FEB receives this validation signal and also has at least one channel above threshold, this will trigger a FEB readout in both FEBs.

For the bottom CRT, there is one fiber per strip. I apply a simple AND between each adjacent layer within a CRT module where at least one PMT channel from each layer must be above threshold.

For all CRT subsystems, when a FEB is triggered, a single time stamp is generated corresponding to the moment the triggering channel's SiPM pulse crosses the discriminator threshold. I have adopted the model developed for the SBND FEB clock as the FEBs used in SBND are identical to those used in ICARUS with the only exception being the firmware versions. This could have secondary impacts on timing, but the hardware aspects are expected to dominate. The model accounts for the 4-ns clock period, the signal amplitude dependence, and the delay due propagation in the optical fiber.

Lastly, for each FEB trigger candidate, some deadtime is incurred. If the coincidence condition is not met, a 150-ns deadtime is incurred. If a readout is triggered, the FEB incurs a 22- μ s deadtime.

4.2.2 Validation

As a cross-check that the side CRT model was implemented correctly, I generated samples of normally incident muons distributed over a single scintillator strip with a longitudinal footprint matching that of the muon telescope we used in our light yield measurements. There is one sample for each measurement position with the surface centered on the central measurement position. I then applied the same analysis code that was used to extract the most probable light yield from the data. I plot this most probable light yield versus distance from the SiPM and compare it to the input light yield model in Figure 4.15a. These results show that the light yield model is mostly correct but underestimates the light yield by about 8%, most likely due to an overestimation of the most probable energy deposited in 1 cm by a MIP muon. Some tuning of this parameter should bring the simulation output in line with the model. As I will explain below, this level of disagreement should not have a significant impact on the results presented in this dissertation. The simulation should be tuned eventually, but it is better to wait and compare to data from the production system after installation at the far detector. I will comment more on this in Chapter 5.

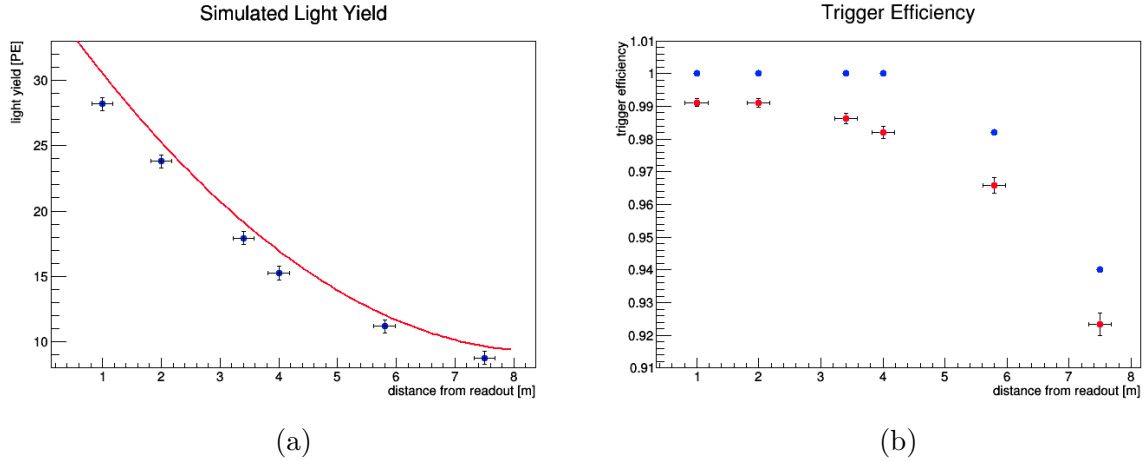


Figure 4.15: The simulated light yield (blue) compared to the input light yield model (red) as a function of the distance from the SiPM (a) and the trigger efficiency vs threshold from the simulation (blue) compared to test stand measurements (red) (b) show reasonable agreement.

The light yield is a primary driver of the muon tagging efficiency. I compare the tagging efficiencies for two layers in coincidence in simulation to data in Figure 4.15b. While the simulated light yield is somewhat lower than the input model, the tagging efficiency measured in this study was found to be 1-2% higher than measurements from the test stand. For the simulation all muons are perfectly centered on the scintillator strip. In reality, some muons cross the boundary between scintillator strips, not depositing enough energy in either strip to generate a trigger. Thus, this disagreement can be understood as a geometric effect. The 8% underestimation in light yield by the simulation is not expected to significantly impact the tagging efficiency as the most probable light yield from the simulation is 15 PE at the longitudinal center, producing a two-layer tagging efficiency of 98.5% at a threshold of 6.5 PE assuming Poisson statistics.

The other important quantity for cosmogenic background rejection is the time resolution. We currently have a limited number of measurements from the test stand available for comparison. Unfortunately, due to COVID-19 restrictions, we were unable to access the test stand to study time resolution in more detail. A comparison of two test stand measurements of the time resolution to the prediction from this simulation based study

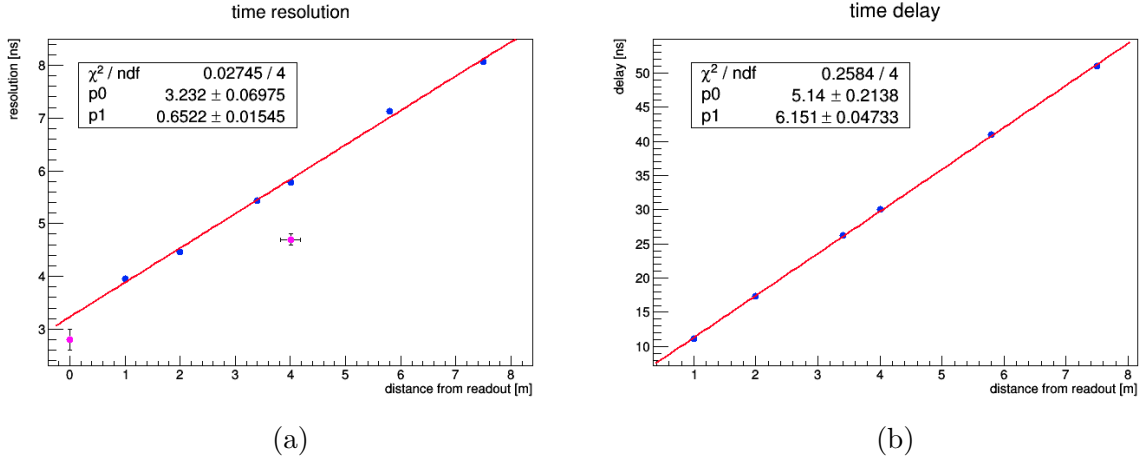


Figure 4.16: The time resolution at different distances from the SiPM from the detector simulation (blue) compared to test stand measurements (pink) without time-walk corrections (a) show good agreement. The light propagation time as a function of distance (b) was then extracted from the simulation and later used in the reconstruction.

is shown in Figure 4.16. A measurement of the intrinsic resolution of the FEB-SiPM system is shown at the 0-m readout distance. The point at 4.0 m was obtained using a full MINOS readout with a hodoscope. I find that the results are mostly consistent with the simulation producing a roughly 15% larger resolution. In terms of the impact on current studies using this simulation, simulating worse CRT time resolution should make any results based on CRT timing conservative. In addition, we can use this study to extract the light propagation time as a function of distance from the SiPM. These results are in agreement with the group velocity used in the simulation.

These test stand based validations of the detector simulation give us confidence in the modeling, but we should also compare to cosmogenic data taken with the production version of the CRT system following commissioning to make sure this result is robust. This work has been started, and first results will be presented in Chapter 5.

4.3 CRT Reconstruction

Once CRT data is generated, useful information can be extracted via the reconstruction chain. In this section, I focus on low-level reconstruction using the CRT only; incorporating reconstruction products from the other detector subsystems will be discussed in the next section. The first step in the reconstruction chain is to construct CRT hits defined as a point in space and time corresponding to a muon track crossing the CRT volume, discussed in the next subsection.

4.3.1 CRT Hits and Tracks

As a first step, all of the CRT data, in the form of FEB readouts, in a given event is time ordered and grouped by CRT region. The next step is CRT subsystem-specific. In the case of the Top or Bottom subsystems, each module is a self-contained coincidence unit; one FEB readout corresponds to a single ionization event with interlayer coincidence. Thus, there is one CRT hit per Top or Bottom CRT data object. The situation is more complicated for the side CRT where there is a coincidence between FEBs residing in adjacent layers. Furthermore, there can be multiple FEB readouts within the same layer. To identify a coincident grouping of CRT data objects, I apply a software based coincidence gate. The hardware based coincidence gate width is known and should be treated as the minimum width for the software gate. Making the gate too large will introduce false coincidences, from low-energy radiogenic gamma rays for example.

Once coincident groupings of CRT data are formed, the spatial information is extracted to reconstruct the position of the crossing track. The first step is to identify which scintillator strips should be considered. This is done by inspecting the charge amplitude in each channel within the data products, grouped by FEB. The channel with the largest amplitude is the channel that generated the FEB trigger signal. Secondary charge information should also be considered. When the FEB is triggered, the charge in all other

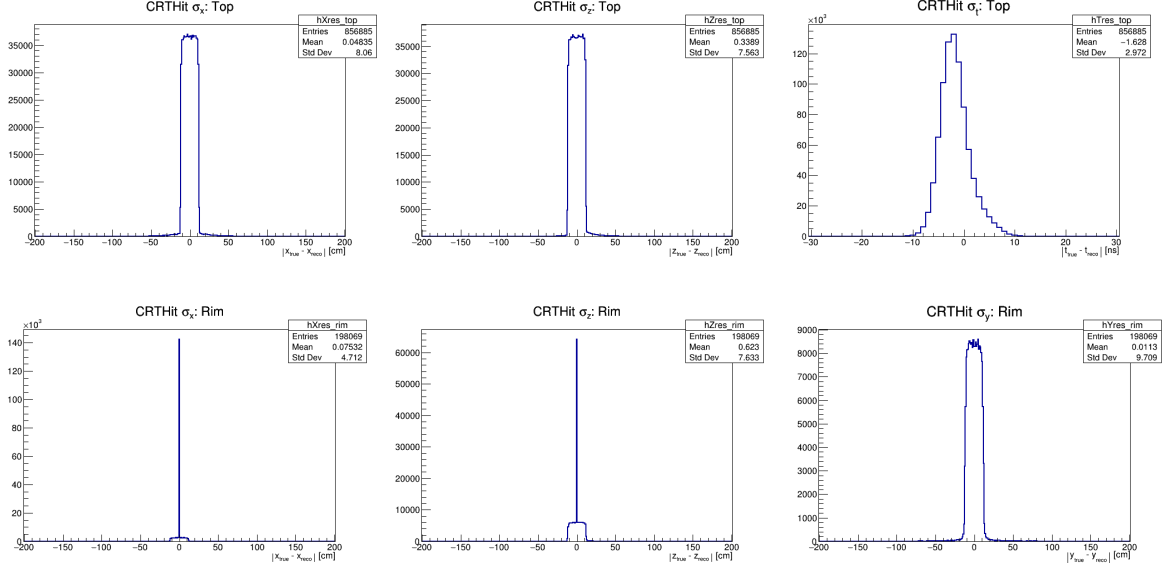


Figure 4.17: The relative difference between true and reconstructed hit coordinates in the local CRT module frame and hit time are shown for the top CRT subsystem regions. The left column is for the local longitudinal coordinate, the center column is for the local lateral coordinate, and the right column is for time. The top row and bottom row are for the roof and rim regions respectively.

FEB channels is sampled simultaneously. This is useful in cases where the track crosses multiple strips in each layer. With the crossed strips identified, the position is extracted. From the geometry, in the case of the simulation, or from a hardware database, in the case of data, the position of any given CRT scintillator strip in global coordinates can be extracted. The hit position is taken as the mean strip position weighted by the charge amplitude in each strip.

The way the individual strip position information enters into the average is subsystem dependent. If a given CRT region has an X-Y strip configuration (Top subsystem or Side subsystem, South Wall), the average is applied only within each layer and a single, complimentary spatial coordinate is extracted. If the CRT region has an X-X configuration, the average is applied over all strips and layers within the coincidence grouping.

The last step is to extract the hit time. The time stamp is generated by the FEB channel with the largest amplitude. There are a few effects that must be corrected in

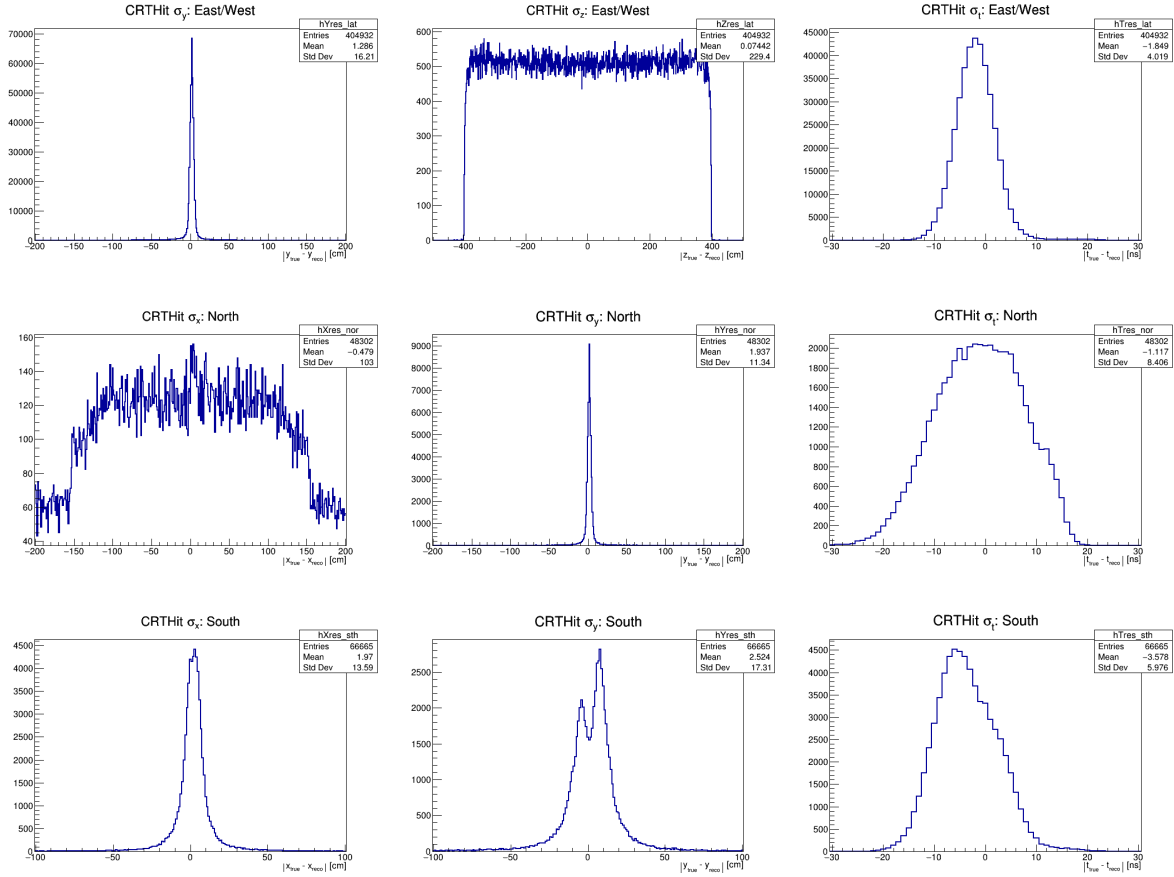


Figure 4.18: The relative difference between true and reconstructed hit coordinates in the local CRT region frame and hit time are shown for the side CRT subsystem regions. The left column is for the local longitudinal coordinate, the center column is for the local lateral coordinate, and the right column is for time. The top, middle, and bottom rows are for the East/West walls, North Wall, and South Wall respectively.




Figure 4.19: The relative difference between true and reconstructed hit coordinates in the local CRT module frame and hit time are shown for the bottom CRT subsystem. The left panel is for the local longitudinal coordinate, the center panel is for the local lateral coordinate, and the right panel is for time.

order to obtain the most accurate time stamp possible. First, there is a correction for the time walk with signal amplitude. The time stamp is generated when the charge amplitude exceeds the discriminator threshold. The larger the amplitude is, the shorter the rise time of the signal is. This relationship is established by measurement and can be corrected using a charge-dependent correction factor. Second, there is jitter on the front-end board clock. This is corrected by a simple scale factor obtained by assuming a perfect second supplied by the GPS-disciplined PPS signal and taking the ratio of the measured to the assumed PPS period. Third, the propagation time incurred by the scintillation light propagating from the point of creation to the photodetector must be subtracted from the time stamp after correcting for time walk. To do this, I assume the true track position in the scintillator strip is the same as the reconstructed point. Knowing the group velocity of the light propagating in the optical fiber, the propagation time can be calculated. Finally, there are all of the timing cable delays, different from FEB to FEB. These are corrected using the known cable lengths obtained from a hardware database at the DAQ level.

In simulation, only the time walk and propagation corrections are made as jitter and cable lengths are not included in simulating time stamp generation.

There are three metrics for quantifying the performance of the hit reconstruction algorithms: efficiency, spatial and temporal resolution. These are determined via a truth-matching exercise where I compare the data products from the detector simulation and subsequent reconstruction to the truth-level inputs. Efficiency is the ratio of the number of CRTData products incorporated into the CRTHits to the number of CRTData products available. The spatial and temporal resolution is the relative difference between true and reconstructed coordinates. The results are shown in Figure 4.17 for the Top subsystem, Figure 4.18 for the Side subsystem, and Figure 4.19 for the Bottom subsystem. The resolution is quoted as one standard deviation of the full distribution for each coordinate, and the results are summarized in Table 4.2.

Around 30% of cosmogenic muons that enter the cryostat will produce multiple CRT

Table 4.2: The spatial and temporal resolution of the CRTHit reconstruction are summarized using local coordinates in the CRT region frame with x as the lateral and y as the longitudinal coordinates.

Region	σ_x [cm]	σ_y [cm]	σ_t [ns]
Top, roof	8	8	3
Top, rim	10	8	3
Side, lateral	16	229	4
Side, north	11	103	8
Side, south	17	14	6
Bottom	79	50	6

hits. These can be combined to form a CRT track. The efficiency for these objects is insufficient for them to be used as a primary CRT reconstruction object. They are useful nonetheless, as I will discuss in the next section.

4.3.2 High-Level Reconstruction with the CRT

I refer to CRT/TPC hit or track and OpHit or OpFlash reconstruction as low-level reconstruction. These objects, on their own, provide a significant amount of information that is useful for calibration, background rejection, and event selection. Taken together, these different reconstruction objects provide complimentary information that can be combined to provide more information than the sum their parts. I will refer to this cross-subsystem reconstruction, incorporating the low-level objects, as high-level reconstruction. As an example, OpFlashes can be combined with TPC tracks and showers to provide precise time stamps and drift coordinates for each cluster of activity in the TPC. This is effective in rejecting out-of-time cosmogenic activity and greatly enhancing the full, three-dimensional reconstruction of tracks and showers.

Another powerful combination is CRT hits and tracks with TPC and PMT objects separately or even with OpFlash-matched TPC clusters. In terms of calibration, CRT hits can be used to construct t_0 -tagged muons tracks, independent of the PDS, to make a cross-check on the flash matching algorithm, including position and angular dependence. Cross-checks like these can help to identify channel mapping problems.

As flash matching is typically a computationally intensive process, the CRT can be used to select cosmogenic muons suitable for various studies. For example, with known time and spatial coordinates for a throughgoing muon coming from the CRT, space charge effects can be mapped throughout the TPC volume or the liquid argon purity can be measured. This is usually done by selecting tracks with the TPC that cross the anode and/or cathode planes, providing an unambiguous time for the event. These samples typically have low rates that could be remedied by inclusion of the CRT.

In Chapter 6, I will focus on the combination of CRT hits and tracks with OpHits for fast cosmogenic background rejection. I will show how this combination can improve on the trade off between cosmogenic background rejection power and neutrino signal efficiency. Incorporating all three subsystems should improve the background rejection even further. See Section 2.5 for discussion on this topic.

Chapter 5

First Data from the Cosmic Ray

Tagging System

As discussed at the end of Chapter 3, about 25% of the side CRT has been installed and is now in the commissioning phase. We have been actively taking data since February 2020. This affords us an opportunity to test the simulation. This chapter is devoted to describing the commissioning process, and the conditions and configurations used for the data taking. I will present the first analysis of the CRT data including calibration, cosmogenic muon light yield and rates, and overall system stability. The chapter will end with a data-Monte Carlo comparison, an important verification of the robustness of the simulation.

During the CRT commissioning phase, we took data on a regular basis in order to understand the behavior of the system and identify any problems. It was important for us to remedy any issues prior to the start of cold commissioning as access to the building would be significantly restricted during and immediately following cooldown, filling, and stabilization of the TPC.

5.1 Data Acquisition System and Data Taking

For all of the data that was presented in Chapter 3, we used a standalone DAQ package built on ROOT6, originally developed by Igor Kreslo and modified by us for our specific needs. This software was useful for debugging the system following installation, but it is inadequate for controlling and operating the full system. Both ICARUS and SBND

will use the artdaq [71] framework for all DAQ needs, The code resides in the sbndaq repository. Since both experiments' CRTs use the same front-end electronics (with the exception of the ICARUS bottom CRT), the DAQ software developed for ICARUS can easily be extended to SBND. This makes our first exercise in data taking and analysis all the more useful. The DAQ development effort is led by Antoni Aduszkiewicz from the University of Houston.

Each front-end board (FEB) daisychain is connected to one server Ethernet port, and each server services two ports. For each port, the most primitive DAQ process is a daemon which acts as the interface between the FEB and artdaq, known as febdrv. febdrv is largely based on a subset of the C code used in the standalone software. The most primitive artdaq process is artdaqDriver, and we have one instance of this for each port. artdaqDriver facilitates transmission of configuration bitstreams to the connected FEBs as well as management of the data packets. Due to some technical challenges in integrating the original drivers developed for the FEBs into sbndaq, the CRT DAQ is not yet in a production state. The CRT DAQ will eventually be controlled by a central manager program that can control all artdaqDriver instances across all the CRT DAQ servers, including simultaneously starting and stopping data taking runs, facilitating access to online monitoring tools, and forwarding the data streams to the central event-building server. In the meantime, the shifter must manage each instance of artdaqDriver manually. For the current system, this amounts to four instances.

In the configuration bitstream, the user can specify the following values:

- list of mac5 address to look for on the front-end board daisychain;
- choose which front-end boards will enable high voltage;
- set the preamplifier gain;
- set each front-end channel bias voltage;
- set the discriminator threshold value;

- choose which front-end channels to mask.

With this, we can define three run configurations used throughout our daily data taking. We wish to monitor the stability of our system. To this end, we monitor each channel's effective SiPM gain, pedestal, and trigger rate. In addition, we check for the presence of electrical noise. To access these quantities, we require three different configurations. In each, we use the same settings for the preamplifier gain, SiPM bias voltage, and discriminator mask.

1. Noise monitoring - set high voltage off on all FEBs, and set the discriminator threshold to the lowest value that prevents triggering on pedestal in the absence of electrical noise (~ 0.5 photoelectrons);
2. Calibration - enable the high voltage on all FEBs and set a low (~ 4.5 photoelectrons) threshold that leads to low-signal domination of the charge spectrum;
3. Cosmics - enable high voltage on all FEBs and set threshold below the edge of the expected charge spectrum for cosmogenic muons (~ 6.5 photoelectrons).

In total, we currently operate 20 FEBs with a total of 560 SiPM channels. The full side CRT system will use a total of 92 FEBs and 2680 SiPM channels.

We began daily noise monitoring at the start of cold commissioning as we expected a significant noise rate from the proximity cryogenics VFDs (see Section 3.7).

5.2 Calibration

The calibration procedure used here is the same as that described in Section 3.2. In this section, I will present the results of the calibration studies performed on the first CRT data taken at the far detector including pedestal uniformity, gain uniformity, and stability over time.

The run was taken on 12 February 2020, before the start of cold commissioning, and lasted for about one hour. 40 million events were acquired overall. The standalone DAQ was used to verify that the noise rate was 0 Hz. All FEBs from the North Wall and West Rolling Wall were included in the analysis.

5.2.1 Pedestal

It is expected that the central value of pedestals across 32 channels of a given FEB will vary from FEB to FEB. As shown in Chapter 3, these values should be stable in time and not depend on what SiPM is connected or the applied bias voltage.

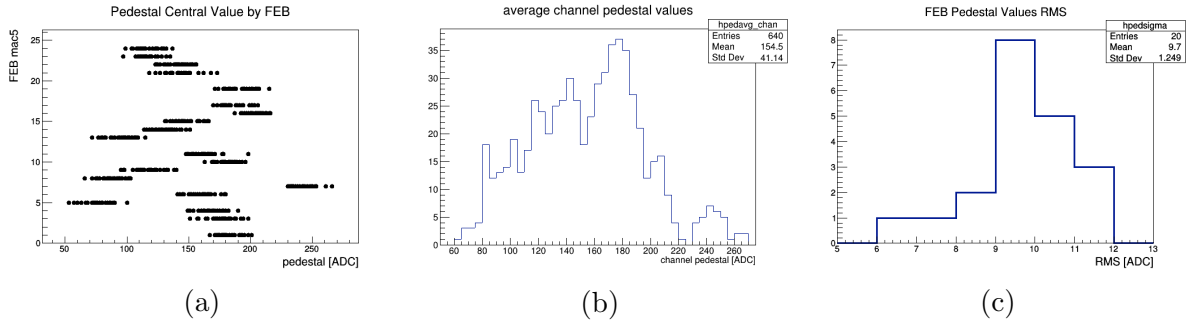


Figure 5.1: Pedestal measurements from a single run show the values for specific front-end boards (a), the one-dimensional projection of (a) onto the pedestal axis (b), and the RMS for the distribution of pedestals on each front-end board (c).

The central pedestal value is shown in two views. In Figure 5.1a, the pedestal central values for each channel is plotted against the FEB mac5 address. In Figure 5.1b, a one-dimensional version of the previous plot projected onto the pedestal value axis is shown. Figure 5.1c shows the RMS of the pedestal values across the individual FEBs. While the pedestal central value can range between about 50-260 ADC, the values are fairly uniform within a FEB with a typical RMS of 10 ADC, roughly 0.2 PE.

We are currently applying the same discriminator threshold to all FEBs. Provided the SiPM gains are sufficiently uniform, the mean pedestal value on a given FEB, averaged across all 32 channels, is the main driver of adjustments required to be made to the discriminator threshold, if any. For the gains we have measured (next subsection), the

full range in pedestal values spans about 3.9 PE. While this is the simplest way to set the initial configuration of the system, this spread is too large, and the pedestal will need to be set on an individual FEB basis.

5.2.2 Gain

For our first data taking, we used the same bias voltage for all SiPM channels. This is expected to result in some non-uniformity in the gain values as the breakdown voltage of each SiPM is somewhat different. Using the calibration data we have so far obtained, we can adjust the individual channels bias values in order to equalize all SiPM gains.

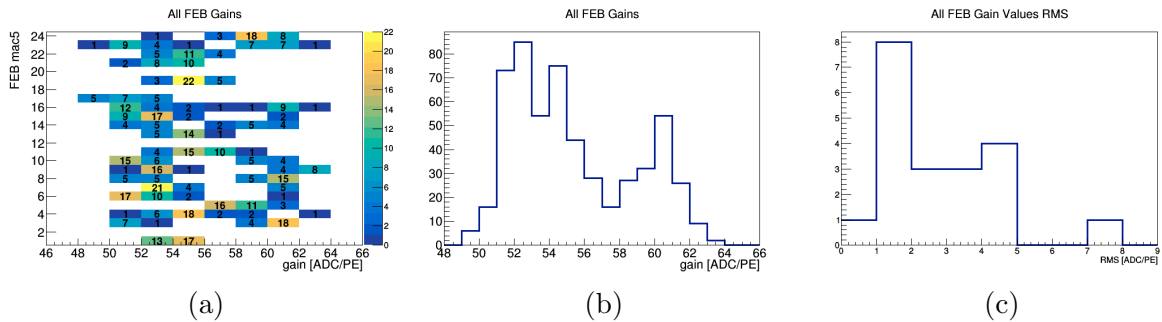


Figure 5.2: Gain measurements from a single run show the values for specific front-end boards (a), the one-dimensional projection of (a) onto the gain axis (b), and the RMS for the distribution of gains on each front-end board (c).

Using the same plot format as the previous subsection, I show two views of the distribution of SiPM gain values. In Figure 5.2a, the gain is plotted against the FEB mac5 address. This demonstrates that the reference bias voltage set on each FEB is nearly uniform. In the other view, you can see the previous plot projected onto the gain axis. This shows that there are two populations of breakdown values with some narrow spread. The RMS of the distribution as a whole is approximately 2 ADC/PE. This will have a marginal impact on the data-MC comparison presented in Section 5.5.

5.2.3 Stability

Calibration runs were taken almost daily, each lasting about one hour. This provided an opportunity to study the stability of the system. Our only other experience with monitoring system stability comes from the ORM testing campaign. The results shown here include data spanning about two months.

The pedestal values over time are shown in Figure 5.3 as a deviation from the mean value. There is some interesting behavior during the first week where there is a few ADC deviation from the long term trend. Aside from some occasional fluctuations, the pedestal values were stable, consistent with expectations.




Figure 5.3: Covering a period of two months, the pedestal value for each front-end board channel is shown at left. The gains over time for one FEB is shown at right (For all FEB gains over time, see Fig. 5.4 below).

As was the case for the ORM testing, the gain fitting algorithm occasionally fails to identify that a photopeak was missed or will incorporate a false peak into the fit. This can skew the gain fit by as much as 50%. We wish to exclude bad fits from the analysis. I used the reduced chi-squared and a requirement on the minimum number of photopeaks

identified to filter out the bad fits. For each SiPM channel, I plotted the gain over time. In Figure 5.4, there is one frame per FEB. Some FEBs are fairly stable while others have channels that seem to oscillate between two values. This effect could be due the few PE spread in threshold from FEB to FEB.

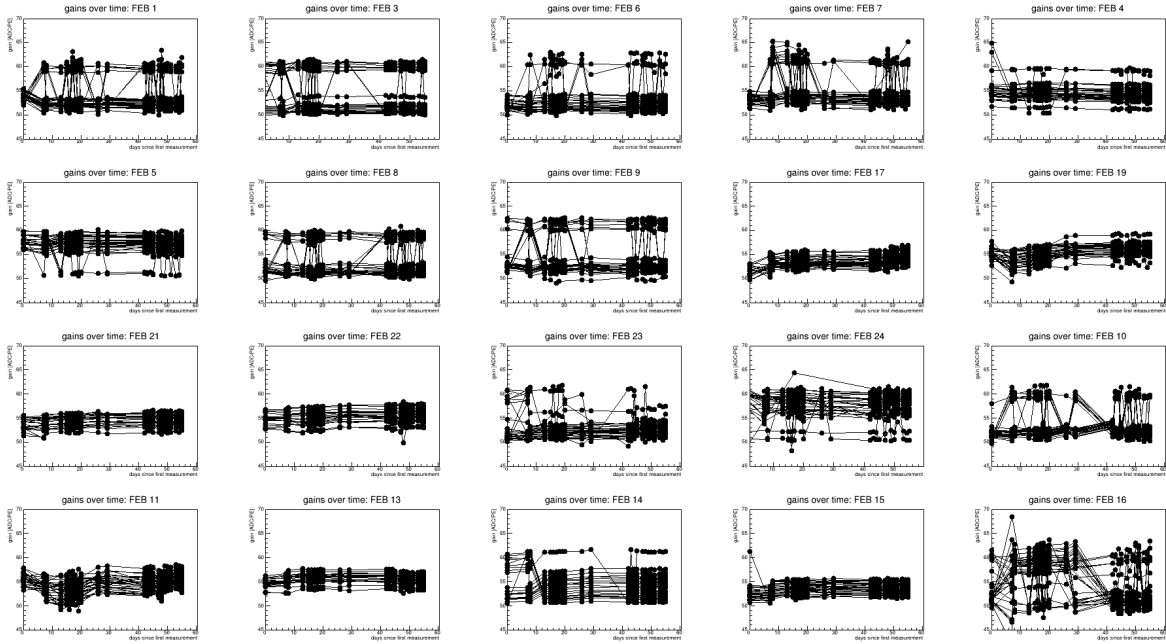


Figure 5.4: The gains for each SiPM channel is plotted over a period of two months. Each panel is one front-end board.

5.3 Low-Level Analysis

In this section, I will present low-level analysis results for “normal” data taking that should produce an almost pure sample of cosmogenic muons. After all channels have been calibrated, we can compare the charge spectra and trigger rates across all channels on equal footing. In this section, I will present the calibrated charge spectra, trigger rates, and charge sharing among front-end board (FEB) channels.

The data used in this analysis was acquired, immediately following the calibration data presented in the previous section, on 12 February 2020. The run lasted for approximately 2.5 hours with a total of 94.1 million events acquired.

5.3.1 Charge Sharing

To start, I will present a first charge sharing analysis. Measurements made at the Side CRT test stand showed that, most of the time, a single strip is hit. However, in the production system at the far detector, all of our modules are vertically oriented. It is expected that muons will be more likely to hit several adjacent strips.

There are a few metrics I have considered. The first is the fraction of the total charge held by the channel that generated the trigger. The second is the fraction of the total charge held by the two channels with the largest amplitudes. The third is to count the multiplicity of SiPM channels above a given analysis threshold. These results are shown in Figure 5.5.

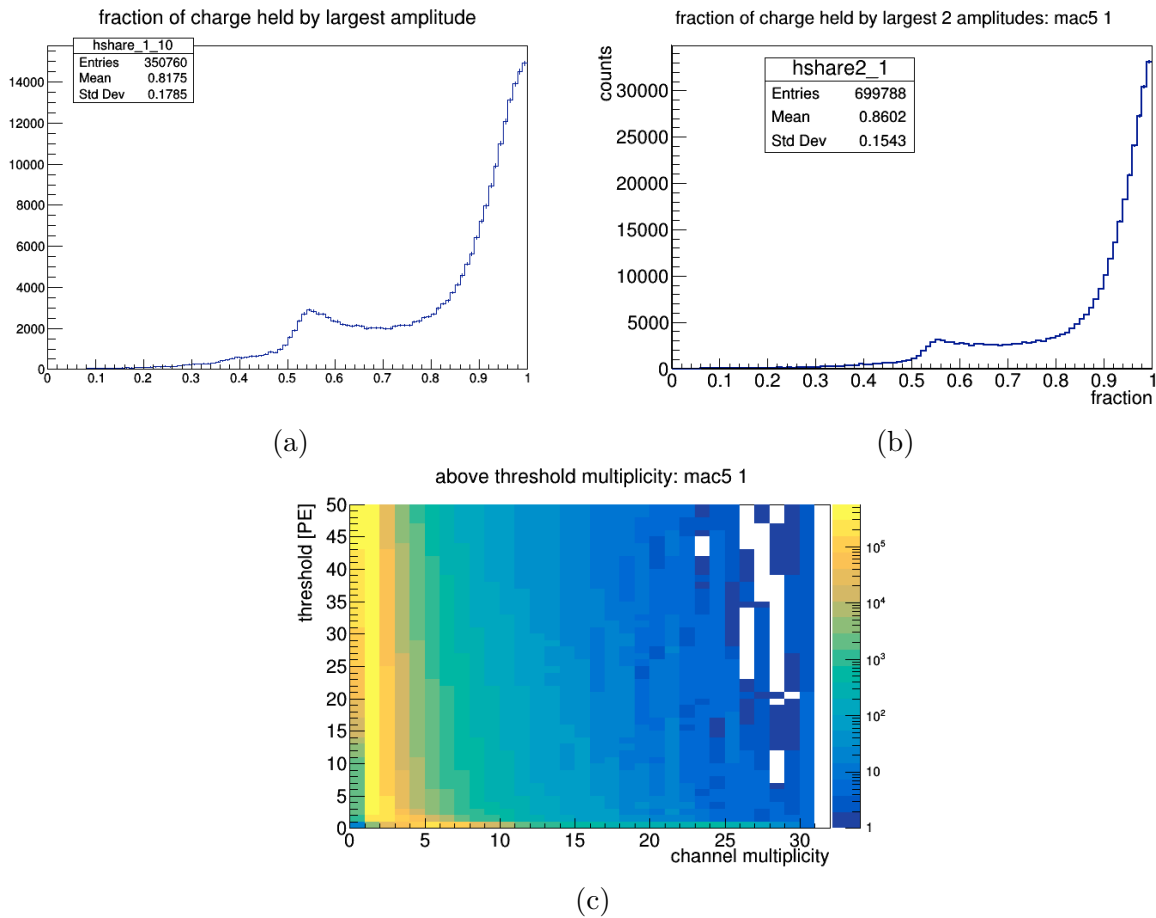


Figure 5.5: Charge sharing

Figure 5.5a shows that the trigger channel usually contains over 80% of the total charge measured by the FEB. The secondary peak shows that, sometimes, the charge is divided nearly equally among two channels. This can be interpreted as cases where the muon crosses two strips read out by different SiPMs. Then there are a few cases where the fraction of charge held by the trigger channel is in the range of 10-50%. These can be interpreted as coming from muons with nearly vertical orientations that deposit energy in several strips.

In Figure 5.5b, the features are nearly identical to Figure 5.5a except that the region between 50-80% is flatter. From the detector geometry with a coincidence between two layers where the gap between the two is about 8 cm, if the FEB trigger is due to a single muon, it is virtually impossible for the track to hit all strips in a given CRT module. In order to maintain the coincidence, it is only possible for the track to cross short portions of some strips and large portions of others.

The interpretations I have given above for the sharing histograms is supported by Figure 5.5c, which shows the number of channels on one FEB above a given analysis threshold with the multiplicity on the x-axis and the threshold in PE on the y-axis. The z-axis is log-scale. Zero-channel multiplicity only happens when the software threshold is lower than the hardware threshold. Also, in the zero-multiplicity bin, the drop in tagging efficiency is evident as the threshold approaches the most probable light yield. This plot shows that, most of the time, only a single channel is above threshold. There is a strong suppression of multiplicities above about three above a threshold of only a few PE. This is consistent with the interpretation of the charge sharing histograms that short path lengths in other strips, where only a small amount of energy is deposited, contribute to the tail below 50%.

5.3.2 Charge Spectra

Next, I will present the charge spectra. The North Wall contains exclusively cut modules at three different lengths: 2.6 m, 3.1 m, and 5.1 m. This changes the attenuation and light yield characteristics of the modules considerably. You can see the light yield distributions for the two different module lengths in Figure 5.6. Compared to full length modules, the cut modules have a significant shift in the light yield toward higher values. The distribution of muon entry points along the module is a uniform distribution. The lengths given above for the cut modules amount to a reduction in the distribution mean, and the variance, by 36-68%. Coupled with the exponential falloff of the light yield due to attenuation, the most probable value of the light yield is drastically increased.

In the West Rolling Wall, all modules are a full 8-m length and read out at both ends. The attenuation profile for these modules should match what was measured at the test stand. The light yield distribution profile should change however. We should expect the mean light to increase compared to what was observed at the test stand since the modules are now on their sides, which presents more scintillator to each cosmogenic muon track. The light yield distributions for each SiPM channel from one of the twelve West Rolling Wall front-end boards are shown in Figure 5.7.

The most probable values for the light yields, integrated over the full modules length, are shown in Figure 5.8. Comparing the light yield values between the two CRT regions, we observe a 3-PE and 8-PE RMS spread in the most probable light yields for the West Rolling Wall and North Wall respectively. From the test stand measurements, we know that there is approximately a 3-PE spread in the light yield across the different scintillator strips. An additional effect could come from possible damage to the modules during the installation process. I comment more on this below. As the larger than expected spread is confined to the North Wall modules, this could be hinting at an effect from the cutting process. It is worth noting that, with few exceptions, the the most probable light yields

being above 15 PE indicate that the two-layer tagging efficiency should be above 98% at the chosen threshold of 6.5 PE. In cases where the light yield is 12 PE, the two-layer efficiency drops to 91%. The most probable light yield at 4.0 m is likely above these values however as these values are integrated over the full 8 m length. This should bias the most probable value below light yields at positions closer than 4 m.

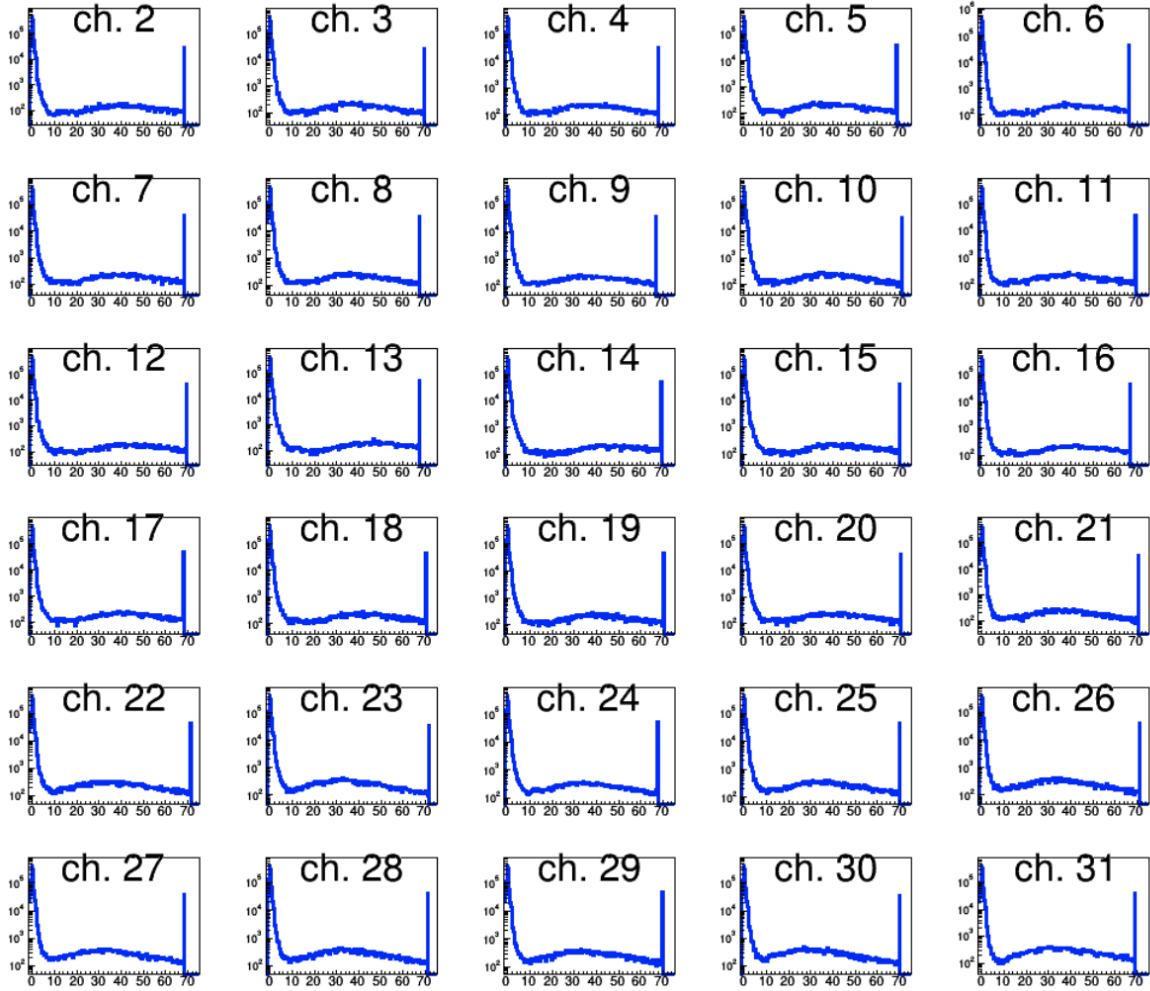


Figure 5.6: Calibrated charge spectra are shown for each SiPM channel from a North Wall front-end board (mac5 005 / west pit level). Channels 22-31 read out a 3.1-m module while the remaining channels read out 2.6-m modules.

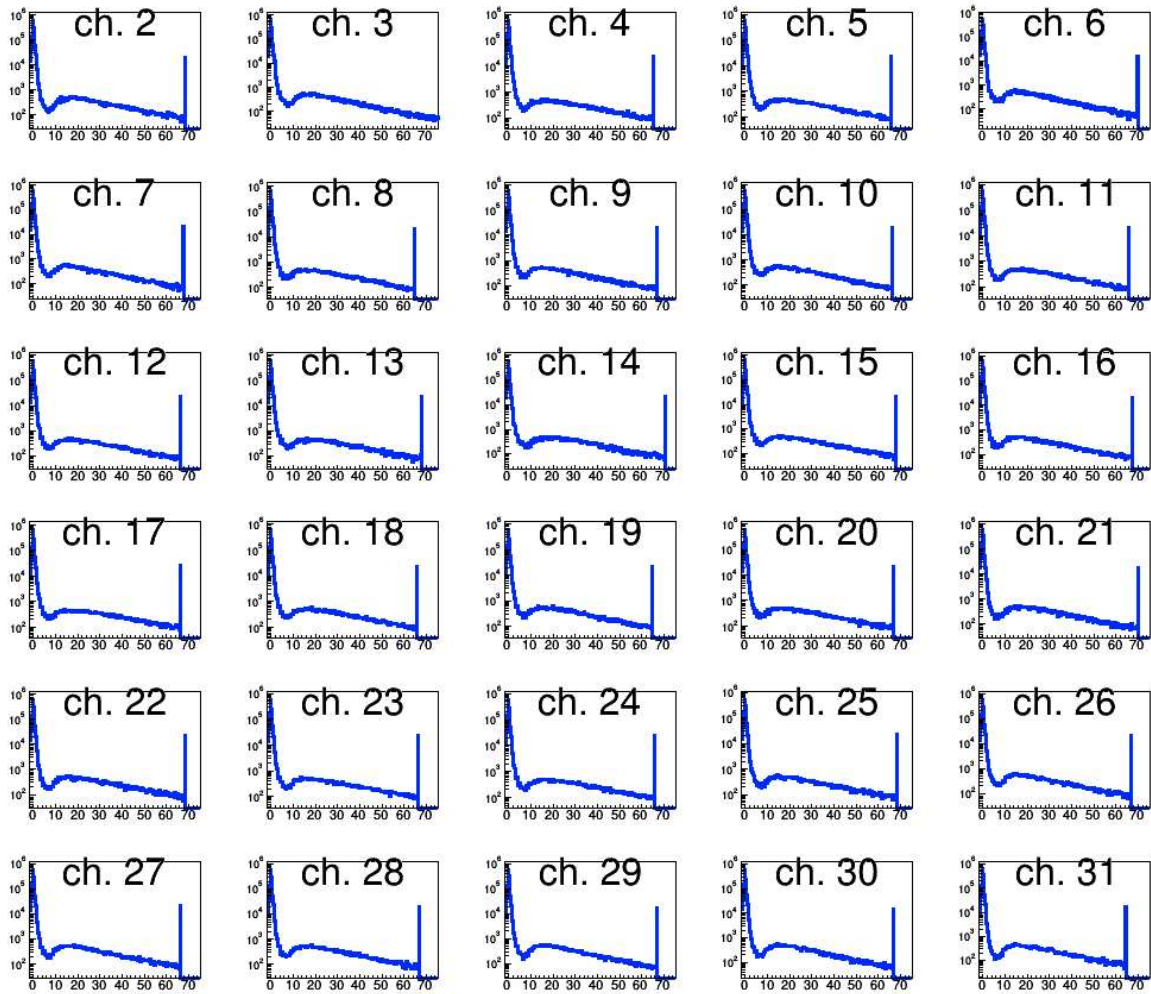


Figure 5.7: Calibrated charge spectra are shown for each SiPM channel from a West Rolling Wall front-end board (mac5 024 / south mezzanine level).

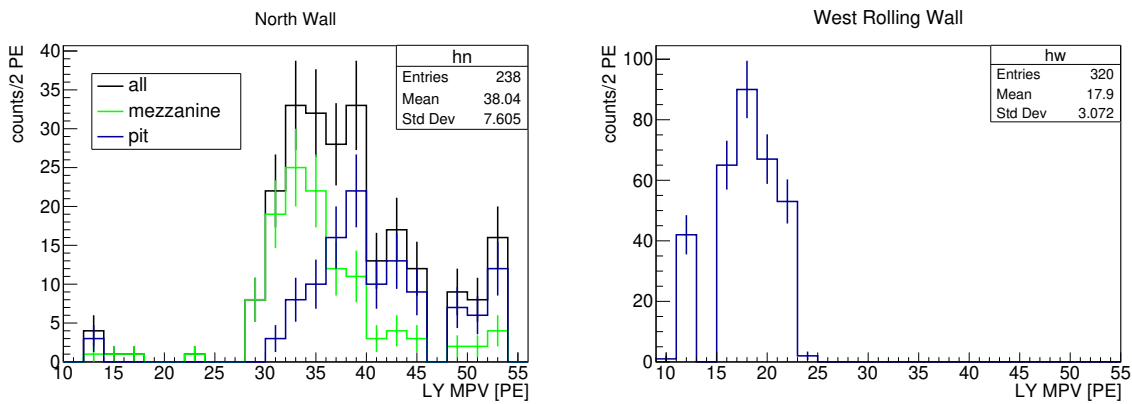


Figure 5.8: The most probable values for the light yield over the full module lengths are given for the North Wall (left) and the West Rolling Wall (right). For the North Wall, the distributions for the mezzanine-level (5.1-m modules) and the pit-level (3.1-m and 2.6-m modules) show the dependence on module length, especially when compared to the full 8-m length modules in the West Rolling Wall.

5.3.3 Trigger Rates

Finally, I show the trigger rate per FEB and per FEB channel. Figure 5.9 and Figure 5.10 show the fraction of all triggers generated within each region and layer generated by a given FEB grouped by FEB position.

In the North Wall, there are four FEBs per layer and two rows of FEBs, designated by the level of the building from where they are accessible, mezzanine-level and pit-level. About one third of all triggers are generated by each mezzanine-level FEB while about one sixth of all triggers are generated by pit-level FEBs. Each FEB reads out 30 SiPM channels. The difference in trigger fraction between mezzanine-level and pit-level FEBs can be explained by the lengths of the modules read out by each FEB. The mezzanine-level FEBs read out an area of 12.5 m^2 each while the pit-level FEBs read out an area of 6.7 m^2 . The pit-level to mezzanine-level area ratio is about 0.54 while the rate ratio is 0.43. An additional reduction in rate is expected for the pit-level FEBs as there is more shielding from incident muons with oblique angles, from the cryostat for example.

In the West Rolling Wall, there are six FEBs per layer reading out a total of eight Side CRT modules from both ends giving three rows of FEBs, also designated by their positions, mezzanine-level, mezzpit-level, and pit-level. The mezzanine-level and mezzpit-level FEBs each have 30 SiPM channels and generate about 20% of the triggers. The pit-level FEBs each have 20 SiPM channels and generate 10% of the triggers. As with the North Wall, the difference in the fraction of triggers can be understood by the area read out by each FEB. The pit-level to mezzanine-level and mezzpit-level area ratio is 0.67 while the rate ratio is 0.59. As was the case for the North Wall, there is more shielding for the pit-level FEBs than the mezzanine-level FEBs.

There is one anomaly in the trigger fraction plot for the West Rolling Wall, the north mezzanine FEB. After this analysis, it was discovered that the connector making the coincidence connection at the FEB was loose. This caused the FEB to operate in a self-

trigger mode. With a 6.5 PE threshold, the bulk of the triggers from the FEB were cosmogenic-muon related. However, some low-energy events, from radiogenic gammas for example, also contributed as evidenced by the higher fraction of total triggers. This issue has now been resolved.

We observe good uniformity in trigger rates across the FEBs once the associated CRT area is taken into account. Next, I analyzed the individual channels' trigger rates. The channel that generates the trigger is the one that has the largest signal amplitude (and fastest rise time). The trigger rates for each channel are shown for the North Wall in Figure 5.11 and the West Rolling Wall in Figure 5.12. Channels with no trigger rate are the non-connected channels. For the North Wall, the effect on the trigger rate due to different module lengths read out by the pit-level FEBs is evident.

One feature that stands out is the relatively low trigger rate in several edge channels, those connected to scintillator strips on outer edges of the scintillator modules. I believe that the most likely explanation for this feature is that some edge channels were damaged. The damage could have been inflicted by braces in the shipping containers that hold the modules in place, or it could have occurred during installation. In case the latter case is correct, we will make some minor adjustments to the installation to avoid damaging the edges. For example, adding some braces to distribute the load over the full edge during crane operations, could help. Overall though, aside from some edge channels, the trigger rates are uniform across all channels.




Figure 5.9: The fraction of all triggers coming from the North Wall inner layer due to each front-end board is shown by front-end board position.




Figure 5.10: The fraction of all triggers coming from the West Rolling Wall outer layer due to each front-end board is shown by front-end board position.

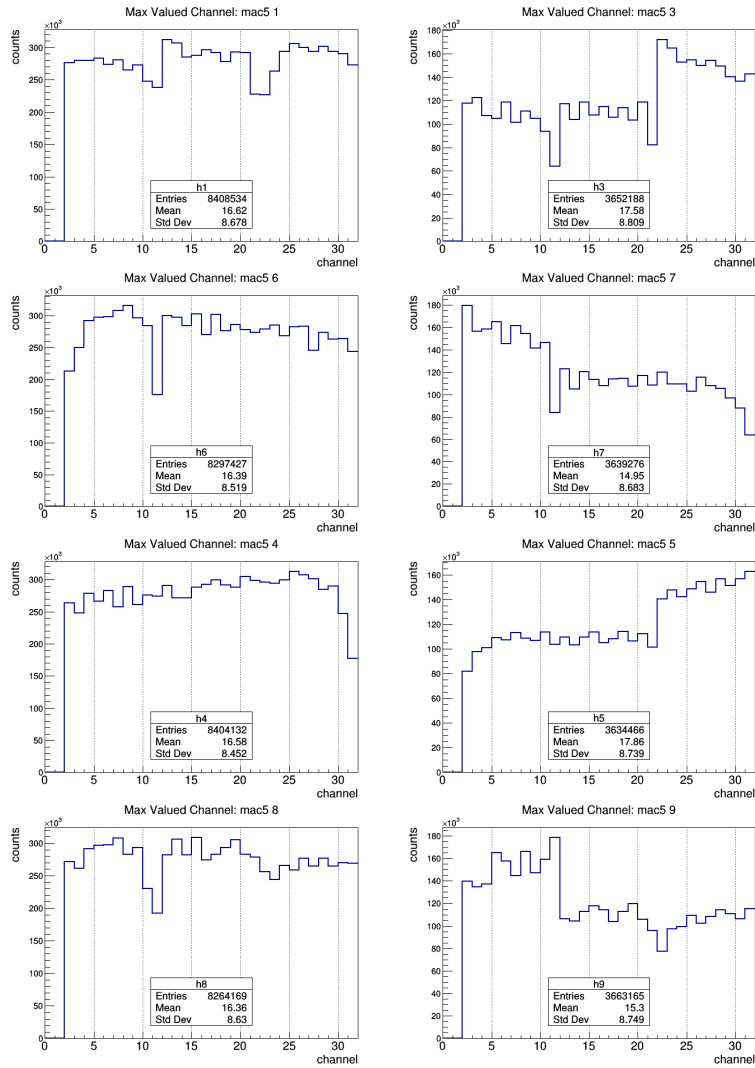


Figure 5.11: North Wall trigger rates by channel show the inner layer in the top two rows and the outer layer in the bottom two. From left to right, the columns are organized by mezzanine level and pit level.

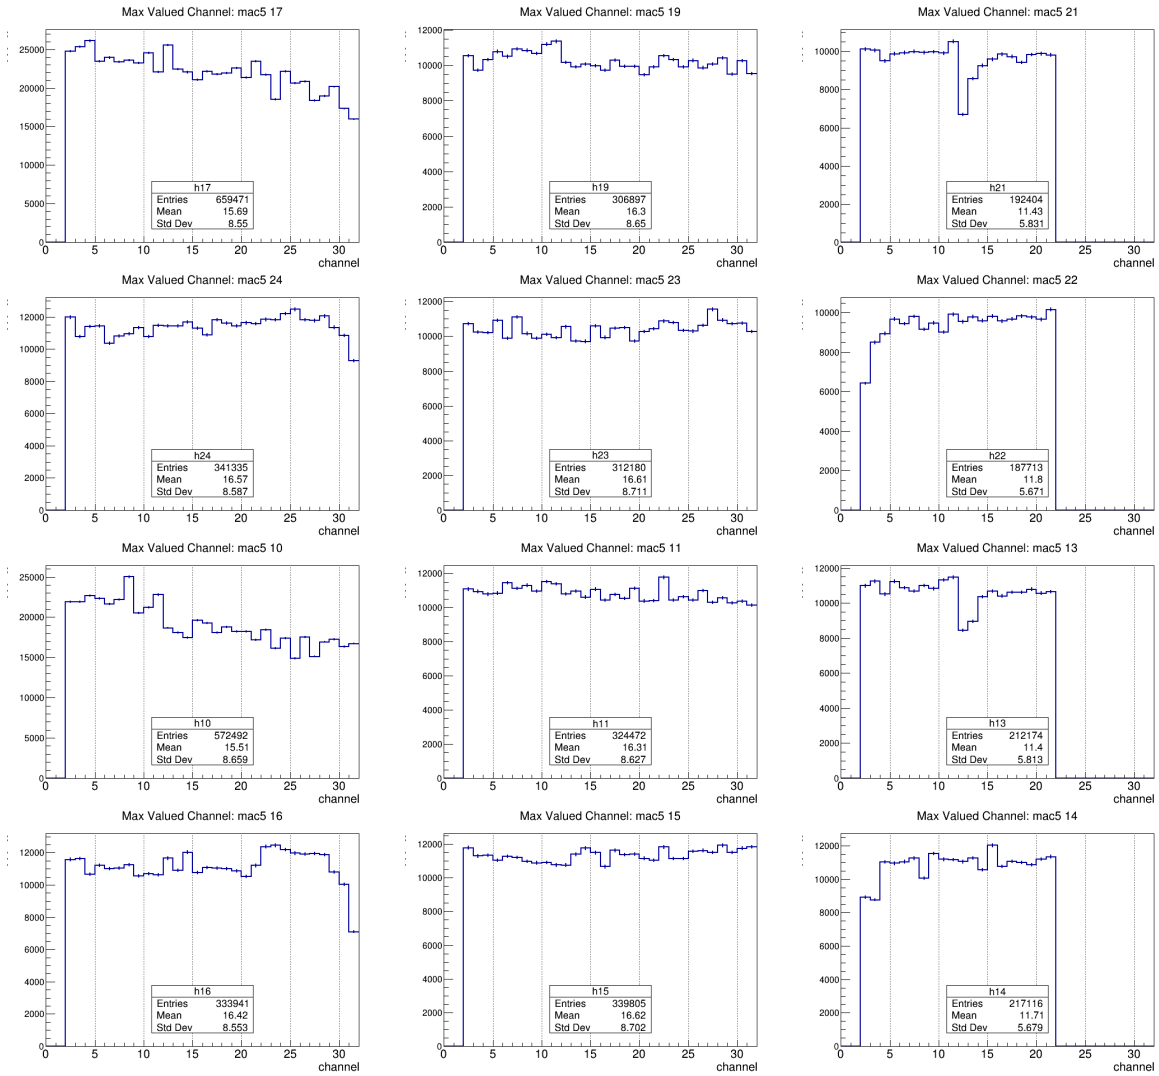


Figure 5.12: West Rolling Wall trigger rates by channel show the inner layer in the top two rows and the outer layer in the bottom two. From left to right, the columns are organized by mezzanine level, mezzpit level, and pit level.

5.4 Data-Monte Carlo Comparison

In Section 4.2, I showed a validation of the CRT detector simulation using test stand data and single muon samples for the simulation. This gives us some confidence in the detector simulation, but it is worthwhile to cross-check this validation using data from the production version of the side CRT. In addition to testing the quality of the detector simulation, this will also test the cosmic ray generator. All of this is important to demonstrate that we can trust projections coming from the simulation concerning cosmogenic muon tagging efficiency and rates. In this section, I will present the results of a data-Monte Carlo comparison study using the first data from the partially commissioned side CRT system.

For the simulation sample, I produced a CORSIKA sample without overburden representing about 30 seconds of exposure. This produced a combined 491,000 FEB triggers between the North Wall and the West Rolling Wall.

The data sample used for this study was acquired on 29 February 2020, still before the start of cold commissioning providing a low-noise condition. A total of 30 million FEB triggers were acquired per CRT wall over a period of 43 minutes for the West Rolling Wall and 112 minutes for the North Wall. The loose coincidence connection affecting one of the West Wall FEBs mentioned in the previous section was fixed by this point. All configurations used match those from the 12 February run.

I will show shape-only and rate comparisons that will test the cosmic ray generator and detector simulation. Most of this work was done, under my guidance, by CSU graduate students, Anna Heggestuen and Tyler Boone, using tools that I developed. Additional assistance was provided by CSU postdoc, Biswaranjan Behera.

First, I will present the shape comparison. We compare the North Wall and the West Rolling Wall separately as the North Wall consists of cut modules, ranging between 2.6 m and 5.1 m in length, while the West Rolling Wall consists of full-length modules. For simulation and data samples separately, the charge spectrum is drawn, calibrated to

photoelectrons. Next, each spectrum is normalized to the total number of entries. Finally, all the charge spectra are combined into a single spectrum and normalized to the number of spectra combined. The resulting spectrum is the average shape for the CRT region, and the contents of each bin are the fraction of all entries contained in that bin.


We fill the charge spectrum histograms two ways: select only the amplitude from the channel that generated the trigger (i.e. the one with the maximum amplitude) or include all the amplitudes above some threshold from each trigger. In the second case, this includes any coincident activity accompanying the FEB trigger, additional strips hit by the muon or delta rays for example.

One challenge in comparing the spectral shapes is the varying dynamic range from channel to channel. In the simulation, the pedestal and gain values are identical for all channels. Thus, the saturation bin is the same for all channels. This is not the case for the data. Even if all gains are identical, the few PE variation in pedestals across the front-end boards will change the PE value of the saturation bin. If we assume that the saturation bin contains physical events (i.e. not electrical noise), we should include the bin content in the overall normalization. Furthermore, there is potential for dark noise or radiogenic backgrounds to generate accidental coincidences in FEBs where the threshold is too low.

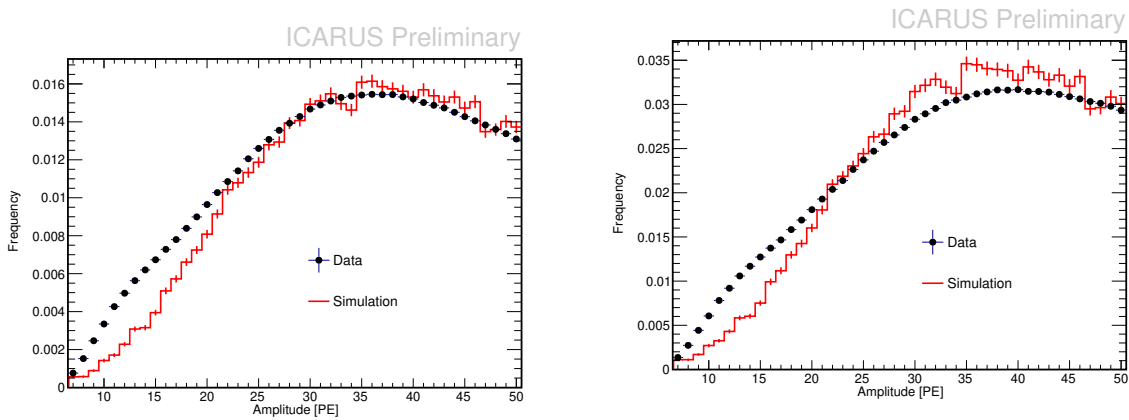
To account for threshold effects and differing interpretations of the saturation bin, we have used two different histogram ranges for this analysis: full range where all entries above the analysis threshold of 3.5 PE are included or a limited range set by the dynamic range used in the simulation, 6.5-50.5 PE, excluding the saturation bin. We will compare the shapes only in the simulation dynamic range. Using the full range or the limited range should only affect the relative normalization.

Figure 5.13 and Figure 5.14 show the shape comparisons for the North Wall and the West Rolling Wall respectively. Regardless of which range is included in the normalization or whether all channels or the trigger channel only are included in the shape, we observe

qualitative agreement between data and simulation in terms of the general spectral shape and the most probable light yield.



(a) All amplitudes and all channels above threshold. (b) Amplitudes in range 6.5-50.5 PE and all channels above threshold.

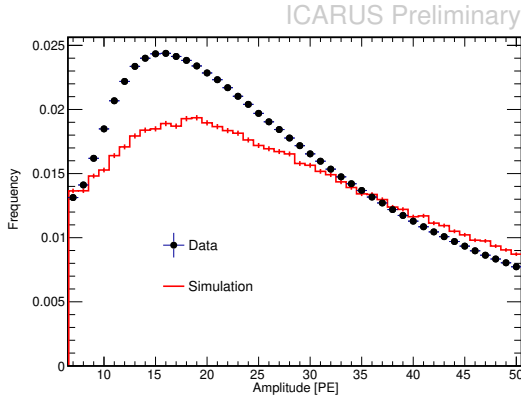


(c) All amplitudes and trigger channel only. (d) Amplitudes in range 6.5-50.5 PE and trigger channel only.


Figure 5.13: Shape comparisons for North Wall show qualitative agreement between data and simulation.

For the North Wall, the most probable light yield is predicted correctly. However, there is a 20-60% underestimation in the number of entries for bins in the range 6.5-20 PE for the all channel and trigger channel only cases respectively.


For the West Rolling Wall, the most probable light yield is overestimated by 25%. The number of entries per bin is underestimated by up to 40% in the range of 10-25 PE and overestimated by up to 20% at yields over 30-40 PE. This is not expected to have a significant impact on the overall robustness of the simulation based results. From the




(a) All amplitudes and all channels above threshold.



(b) Amplitudes in range 6.5-50.5 PE and all channels above threshold.



(c) All amplitudes and trigger channel only.



(d) Amplitudes in range 6.5-50.5 PE and trigger channel only.

Figure 5.14: Shape comparisons for West Rolling Wall.

point of view of tagging efficiency, a 15 PE light yield produces a two-layer efficiency of above 98% with a 6.5 PE threshold.

While this is good result considering that the simulation has not been tuned to the production CRT data, this data should be used to improve the level of agreement between data and simulation.

Next, I show a comparison between the mean FEB trigger rates within the North Wall and West Rolling Wall separately. Assuming the standard value for the cosmogenic muon flux of $1 \text{ cm}^{-2}\text{sr}^{-1}\text{min}^{-1}$, we would expect 1 kHz of cosmogenic muons cross a full-length, horizontally-oriented MINOS module. In a previous study, we found that the rate should decrease by about 59% when the module is vertically oriented. Thus, for the West Rolling

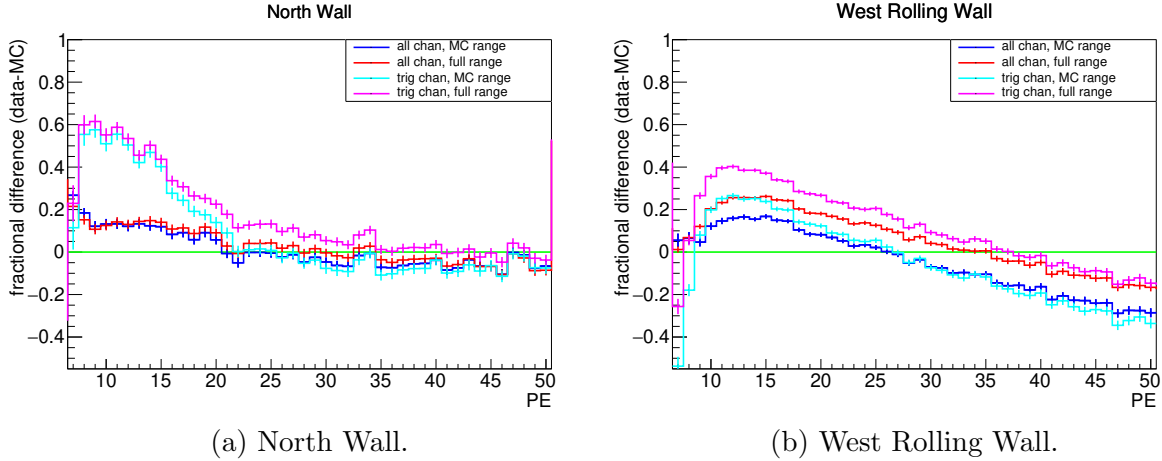


Figure 5.15: Fractional differences in charge spectral shapes are shown for the North Wall (a) and the West Rolling Wall (b).

Wall, we should expect a trigger rate per FEB of about 930 Hz.

Figure 5.16 shows a histogram of the mean FEB trigger rates with one entry per FEB. Comparing the mean of the simulation and data rate distributions, we found that the mean trigger rates are within 5-10% of each other. Furthermore, the trigger rates are in agreement with our rough prediction from above. Note that this prediction must be scaled by the area of the North Wall modules for the North Wall FEBs. Four of the West Rolling Walls readout two instead of three modules, hence the 33% lower rate.

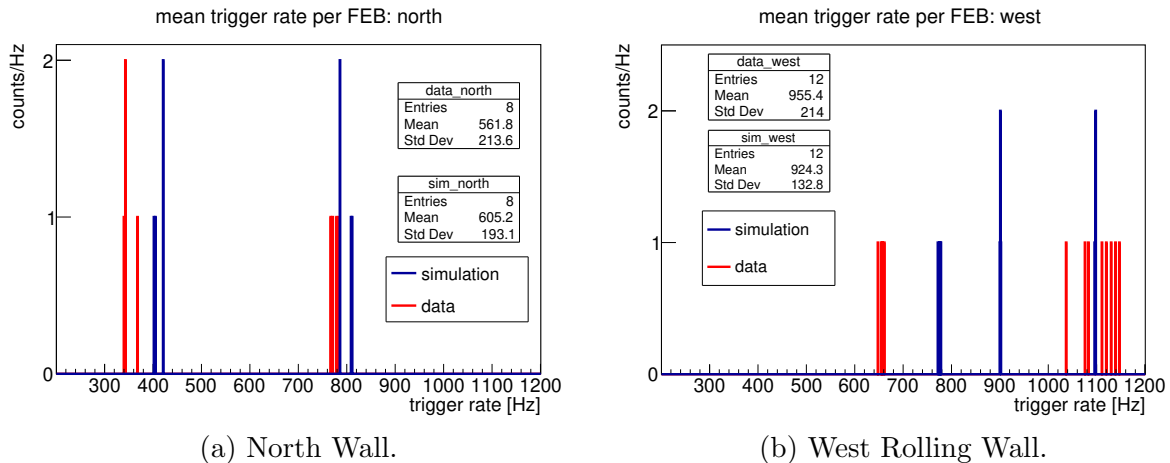


Figure 5.16: A comparison between the mean FEB trigger rates measured in simulation and in data show agreement to within 5-10% for both the North Wall (a) and the West Rolling Wall (b).

Given that this is the first comparison between data and simulation, these results are encouraging. This demonstrates that the simulation as a whole provides a reasonable approximation to reality for the purposes of estimating backgrounds and developing reconstruction and analysis tools. After the CRT system is fully installed and the nominal configurations are set, this study should be replicated. The results of this future study should then be used to tune the CRT detector simulation.

Chapter 6

Cosmogenic Background Rejection

In this chapter, I discuss how we can use the CRT to reject cosmogenic backgrounds while minimizing losses or the introduction of systematic uncertainties into the sterile neutrino search. The results shown will be purely simulation based. This is a major caveat as cosmogenic backgrounds are difficult to model. Even channels with low rates can become a concern when our possible sterile neutrino signal is comparably rare. However, simulation based results are the first step. Clearly, before SBN can adopt any of the following methods, they should be studied with large data sets. From the point of view of studying the cosmogenic background, we can obtain a large data set in a short time. I will comment on this in more detail in the concluding chapter.

In the results that follow, I use the default configuration of `icaruscode` presented in Chapter 4. I begin this chapter by presenting the different topological and timing cases, their rates, and related efficiencies from a background-centric point of view. I will follow this with a discussion of the simplest approach to cosmogenic background rejection using the CRT only. Next, I discuss a second method based on combining the Photon Detection System (PDS) with the low-level CRT reconstruction method. I finish the chapter with a discussion on the impact of these methods on the sterile neutrino search.

6.1 Cosmogenic Background Rates and Efficiencies

To start, I need to understand the true rates and tagging efficiencies for different topologies of cosmogenic muons. These will be calculated from my simulation. I need to know the rate of muons crossing the CRT boundary, entering the LAr inactive volume (IV), LAr

active volume (AV), and the LAr fiducial volume (FV). I adopt the same fiducial volume definition used in the SBN proposal: 25 cm from the lateral faces of the outer TPCs; 25 cm from the top and bottom faces of each TPC; 35 cm from the upstream face of each TPC; and 50 cm from the downstream face of each TPC.

In Table 6.1, the fraction of muons that hit the CRT and then go on to miss the argon volume completely, enter the IV but not the AV, enter the AV, or enter the FV are shown. About half of all CRT muon tags are for cases where the muon does not enter the argon volume so would not produce a scintillation flash. About 10% of muons go on to deposit energy in the IV but never enter the AV and would therefore not be reconstructed in the TPC. In this case, it is possible that a scintillation flash could be observed.

Table 6.1: For each true cosmic ray muon crossing different volumes, the true rate and efficiency are summarized, in the presence of 3-m concrete overburden, for different topologies.

Description	Fraction hitting CRT and [%]	Tagging Efficiency [%]	Rate [kHz]
CRT	1.0	100.0	43.8
CRT only	50.4	100.0	22.1
CRT and IV only	9.7	97.9	4.2
CRT and AV	39.9	97.9	17.8
CRT and FV	26.3	97.4	11.9

We should not only consider cosmogenic muon tagging efficiencies integrated over the entire CRT and momentum range but also position and angular dependence. Truth-level distributions for kinematic variables for all cosmogenic muons that enter the detector enclosure are shown in Figure 6.1. The angular distributions follow $\cos^2 \theta$ -like distributions as expected. The spectra fall off faster at oblique angles due the TPC being tens of meters below grade. The vast majority of muons enter the AV from the top face. The horizontal AV entry coordinate (x and z) distributions are flat as expected.

The CRT tagging efficiency as a function of these variables for the muons entering the AV are shown in Figure 6.2. For the position and angular distributions, we see the effect of gaps between the flat portion of the Top CRT and the rim. Oblique tracks can pass

through this gap and enter the AV while vertical muons that enter the TPC must have passed through the CRT. This leads to a spread of less than 4% in efficiency; however the efficiency exceeds 95% in all cases. In terms of momentum dependence, not surprisingly, the low-momentum muons, those below 2 GeV, are tagged with about 1% lower efficiency than those above 2 GeV where the efficiency becomes flat. This is also expected as the muon is MIP-like in this regime.

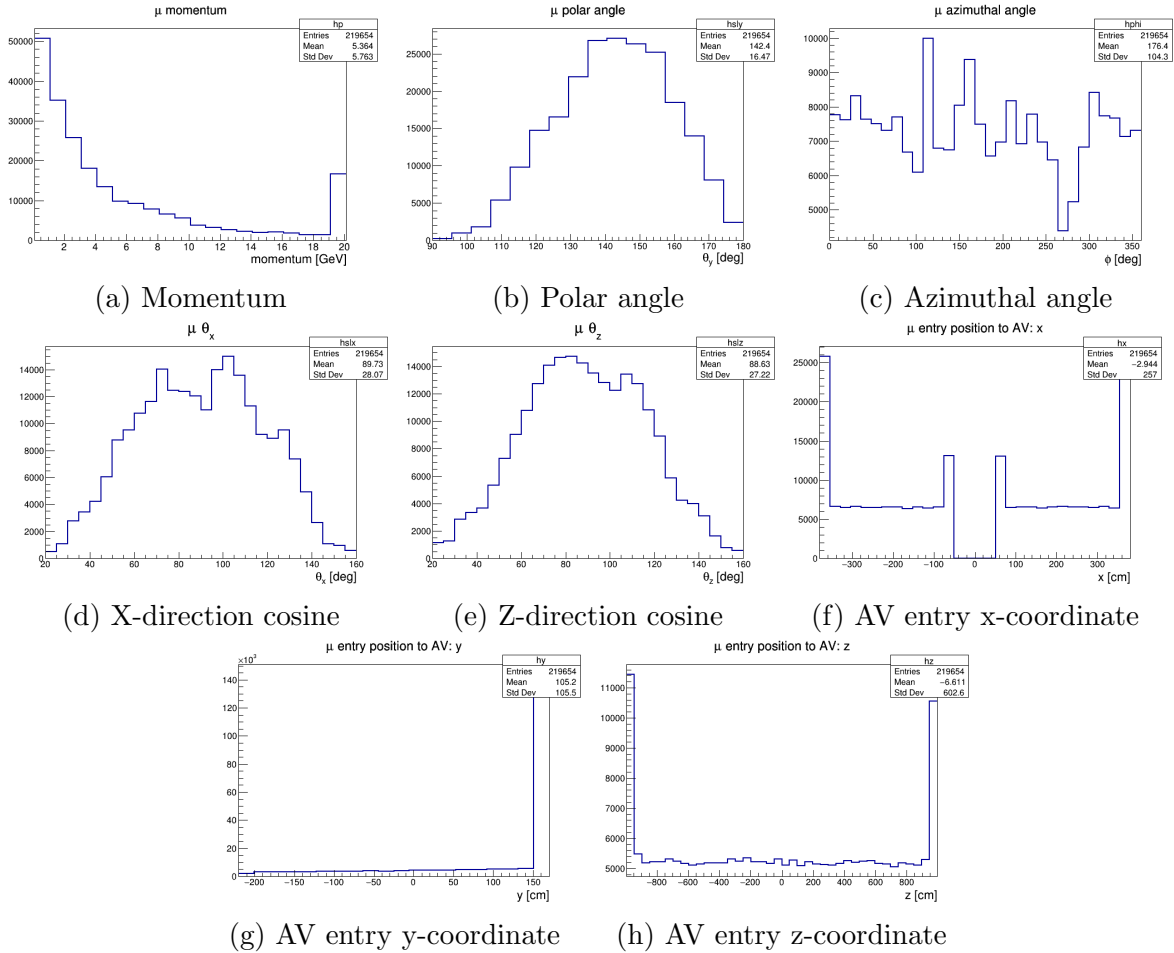


Figure 6.1: Truth-level kinematic information for cosmogenic muons, in the presence of 3-m concrete overburden, with the origin at the T600 center.

Recall from Section 4.3 that the position and time resolution for μ reconstructed CRT hits depends on the CRT subsystem and/or region. The fraction intercepted by each CRT subsystem and their respective rates are summarized in Tables 6.2 and 6.3. These tables are expanded into individual regions in Tables 6.4 and 6.5. This is important for

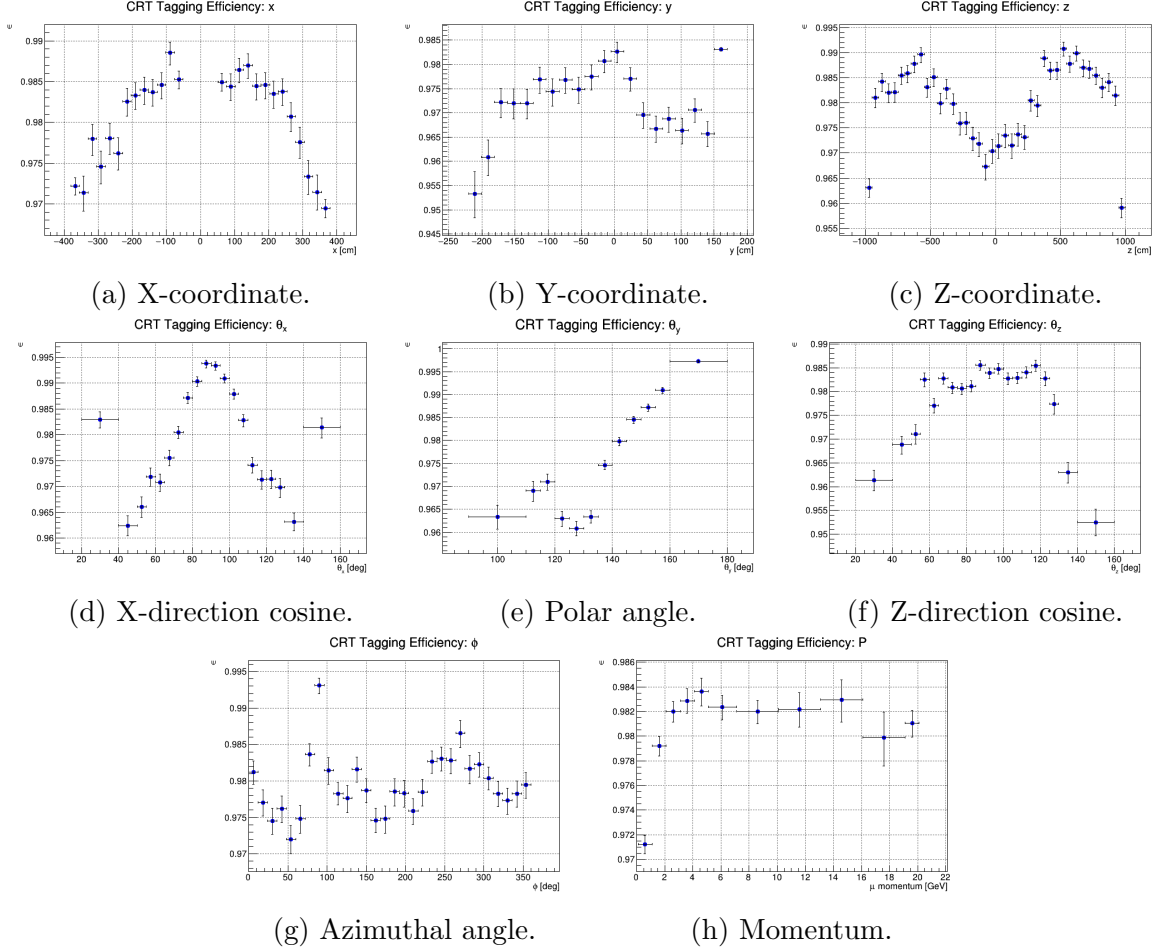


Figure 6.2: Truth-level CRT cosmogenic muon tagging efficiencies for different kinematic variables describing the muons at the point they enter the active volume. Here, 3-m concrete overburden is present and the origin is at the T600 center.

understanding how possible methods of cosmogenic background rejection facilitated by the CRT would be impacted by the worse spatial and time resolution for the Side CRT with respect to the Top CRT.

When we add in detector effects, the tagging efficiency is reduced by about 2%, integrated over the full CRT system. The CRT trigger efficiency spans 95-97% for the different topologies discussed above. Thus, we have succeeded in reaching the tagging efficiency goal set during the SBN proposal of 95% or better.

With a high performance CRT system, how do we best apply the information provided by the CRT to reject cosmogenic backgrounds? Up to this point, I neglected complications

Table 6.2: The true fraction [%] of cosmogenic muons crossing the CRT including the first hit only, separated by CRT subsystem, in the presence of 3-m concrete overburden, is given for different topologies.

CRT region	All	CRT only	CRT + IV only	CRT + AV	CRT + FV
Top	74.5	74.0	69.8	76.4	78.7
Side	25.0	26.0	29.8	22.7	20.5
Bottom	0.4	0.04	0.4	0.9	0.8

Table 6.3: The true rates [kHz] of cosmogenic muons crossing the CRT including the first hit only, separated by CRT subsystem, in the presence of 3-m concrete overburden, is given for different topologies.

CRT region	All	CRT only	CRT + IV only	CRT + AV	CRT + FV
Top	32.7	16.4	3.0	13.3	9.4
Side	11.0	5.7	1.3	4.0	2.4
Bottom	0.2	0.01	0.02	0.2	0.05

that arise due to uncontained neutrino interactions which go on to produce a CRT hit.

This will be a central point of the discussion in the next section.

Table 6.4: The true rate [kHz] of cosmogenic muons crossing the CRT including the first hit only, separated by CRT subsystem and region, in the presence of 3-m concrete overburden, is given for different topologies.

CRT region	CRT only	CRT + IV only	CRT + AV	CRT + FV
Top, roof	14.6	2.7	11.5	7.8
Top, E/W rim	0.4	0.1	0.7	0.4
Top, N/S rim	0.5	0.05	0.1	0.1
Side, north	0.3	0.07	0.2	0.2
Side, south	0.9	0.2	0.6	0.2
Side, E/W north	0.7	0.2	0.5	0.4
Side, E/W center	0.7	0.2	0.5	0.4
Side, E/W south	0.9	0.2	0.5	0.4
Bottom	0.01	0.02	0.2	0.05

Table 6.5: The true fraction [%] of cosmogenic muons crossing the CRT including the first hit only, separated by CRT subsystem and region, in the presence of 3-m concrete overburden, is given for different topologies.

CRT region	CRT only	CRT + IV only	CRT + AV	CRT + FV
Top, roof	65.8	63.3	65.6	66.1
Top, E/W rim	2.0	2.0	4.2	5.0
Top, N/S rim	0.5	0.05	0.1	0.05
Side, north	0.3	0.07	0.2	0.1
Side, south	0.9	0.2	0.6	0.1
Side, E/W north	0.7	0.2	0.5	0.2
Side, E/W center	0.7	0.2	0.5	0.2
Side, E/W south	0.9	0.2	0.5	0.2
Bottom	0.01	0.02	0.2	0.05

6.2 Auto-Veto

The trigger system will trigger a readout of the TPC when there is a coincidence among a subset of PMTs above some photoelectron threshold in coincidence with the beam spill. Recent studies by the ICARUS trigger working group have shown that 1 in 50 BNB beam spills will have triggers due to cosmogenic activity coincident with the spill. In contrast, 1 in 180 triggers will be due to BNB neutrinos. From these rates, I calculate that 1 in 10,000 beam spills will have competing trigger sources. From the rates of cosmogenic muons presented in the previous section, 11 muons will cross the fiducial volume (FV),

on average, per 1-ms readout of the TPC. Well reconstructed muon tracks, especially tracks that enter and then exit from the TPC active volume (AV), can be unambiguously removed. However, as discussed in Section 2.5, this will not always be the case. Dedicated studies with beam-off data will be able to quantify this effect. Regardless of the outcome of such a study, we still must contend with muons that miss the AV entirely. The CRT provides a direct way to detect and remove these events.

The most straightforward approach is to use the CRT as a simple veto. We reject any neutrino event candidates with a CRT hit in time with the beam spill. This approach was first investigated by Umut Kose and Paula Sala from the CERN group in the SBN proposal. They performed a truth-level study with a simplified CRT geometry where each CRT region was approximated by two layers of scintillator slabs, each 1 cm thick and providing 4π coverage and about 99.9% tagging efficiency. For a simple veto, the cosmogenic muon rejection efficiency is the same as the CRT muon tagging efficiency. This study produced the results shown in Section 2.5. One potential problem with this approach is CRT induced neutrino signal loss, the so-called auto-veto fraction.

Table 6.6: For BNB neutrinos, auto-veto fractions [%] are summarized for different samples integrated over the full range in energy. The fractions are broken down by the volume containing the true neutrino vertex and whether the interaction is charged current or neutral current.

Sample	IV Only		AV Only		FV Only	
	CC	NC	CC	NC	CC	NC
ν_μ	20	3	11	1	9	1
Intrinsic ν_e	13	5	4	2	2	1
Oscillated ν_e	10	3	2	1	1	1

To estimate the auto-veto fraction, I generated one million BNB neutrino events each for ν_μ only, beam-intrinsic ν_e , and oscillated ν_e assuming the 3+1 oscillation phenomenology with $\sin^2 2\theta_{\mu e} = 0.013$ and $\Delta m_{41}^2 = 0.43 \text{ eV}^2$. If the CRT was incorporated into the TPC readout trigger at the hardware level, any CRT trigger would cause the neutrino event to be vetoed. Currently, it is not foreseen that the CRT will part of the TPC

hardware based trigger. Therefore, I will assume that CRT data can undergo a full reconstruction before a veto decision is made. If the CRT hit reconstruction efficiency is 100%, there will be no difference in the auto-veto fraction using CRT triggers or CRT hits. If the CRT hit reconstruction efficiency is less than 100%, then the auto-veto fraction would be higher than the results that follow if the CRT were to be integrated into the trigger system. The energy-integrated auto-veto fractions by neutrino interaction channel and neutrino location in the cryostat are summarized in Table 6.6.

I note that a simplifying assumption made for the SBN proposal was that there was no appreciable quenching of the scintillation light produced in the CRT scintillator. I showed that the auto-veto fraction is reduced when quenching is introduced. Using Birks' model, I showed that higher dE/dx particles, protons for example, became less likely to cause a CRT trigger. This led to a marginal reduction in the total auto-veto fraction, especially when CRT hits are restricted to occur within the beam spill. I introduced a similar quenching model into the CRT detector readout simulation as I used in my early studies, but I have disabled this model for the results presented here as it has not yet been thoroughly studied or validated with data. This renders a more conservative result.

The auto-veto fractions I have obtained are about half of those obtained in earlier studies by the CERN group. This is due to several factors. We have reduced CRT coverage lowering the geometric tagging efficiency from 99.9% to 97%. An important CRT region for auto-veto that is less important for tagging cosmogenic muons is the Side CRT North Wall. The Side CRT North Wall coverage was reduced by 44%. In addition, the bottom CRT has not been considered here as it is not useful for cosmogenic muon tagging so would not be included in a simple veto. Furthermore, realistic detector modeling further reduces the rate, due to attenuation, Poisson fluctuations and deadtime effects for example, that lower the integrated tagging efficiency to about 95%.

Aside from losses in signal efficiency, impacts on the reconstructed neutrino energy spectrum must be considered. In Figure 6.3, I show the dependence of the auto-veto

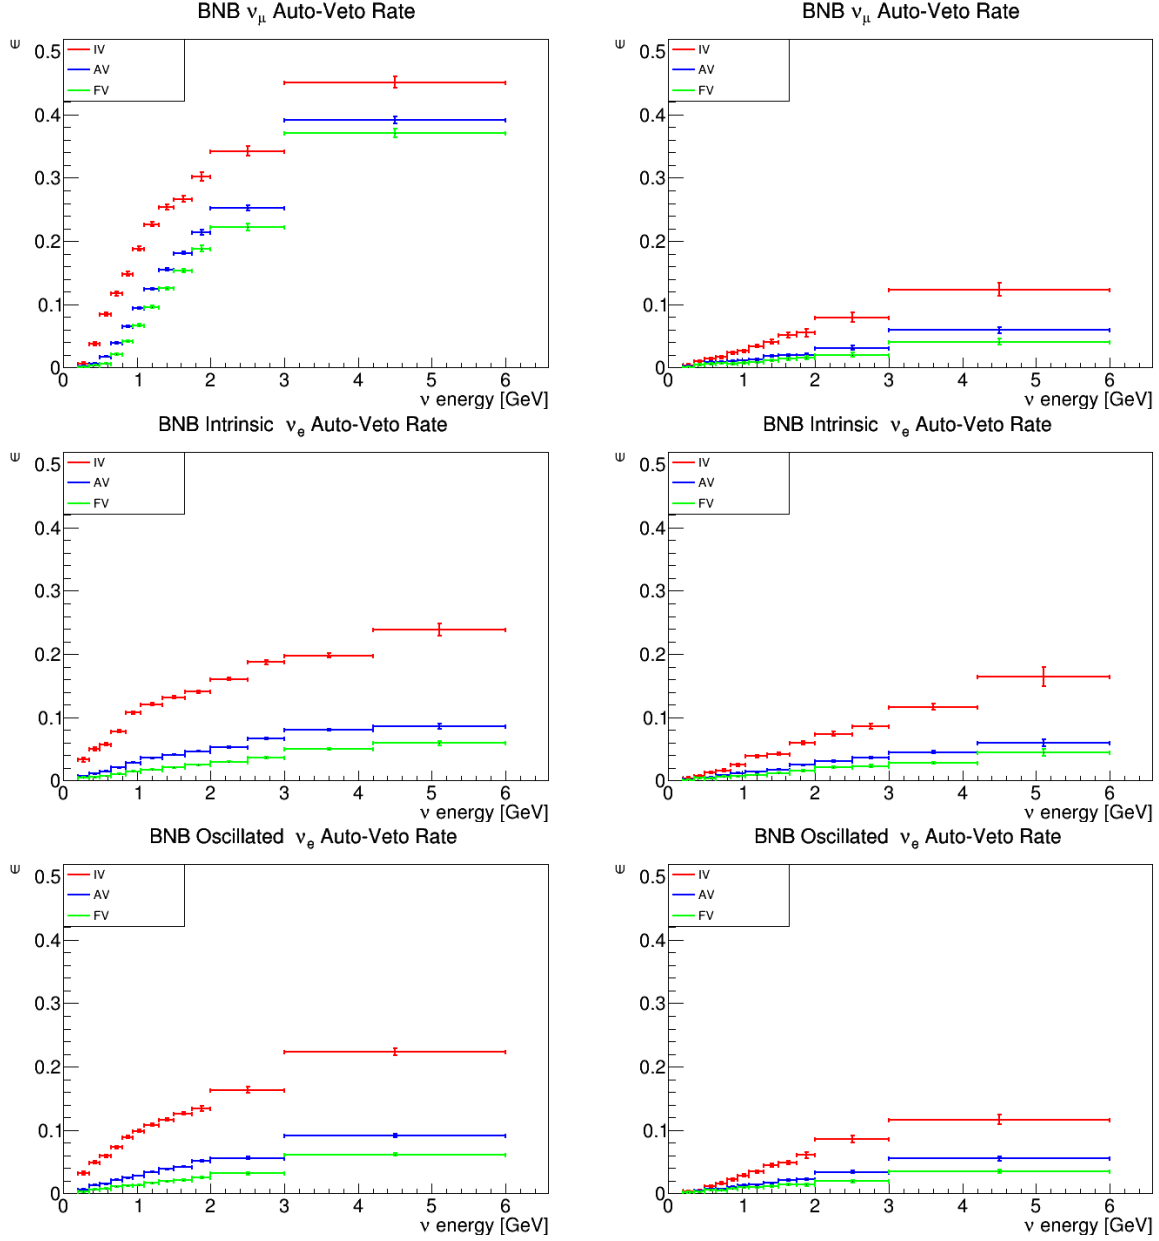


Figure 6.3: The auto-veto fraction is dependent on the location of the neutrino vertex, the neutrino energy and whether the interaction is charged-current or neutral-current as shown for the three different BNB neutrino samples: ν_μ (top row), intrinsic ν_e (middle row), and oscillated ν_e . Charged-current and neutral-current interactions are in the left and right columns respectively.

fraction on the true neutrino energy depending on whether the true vertex position was in the IV, AV, or FV and whether the interaction is CC or NC. As expected, as the neutrino energy increases, the likelihood that the neutrino secondaries escape the LAr

volume increases.

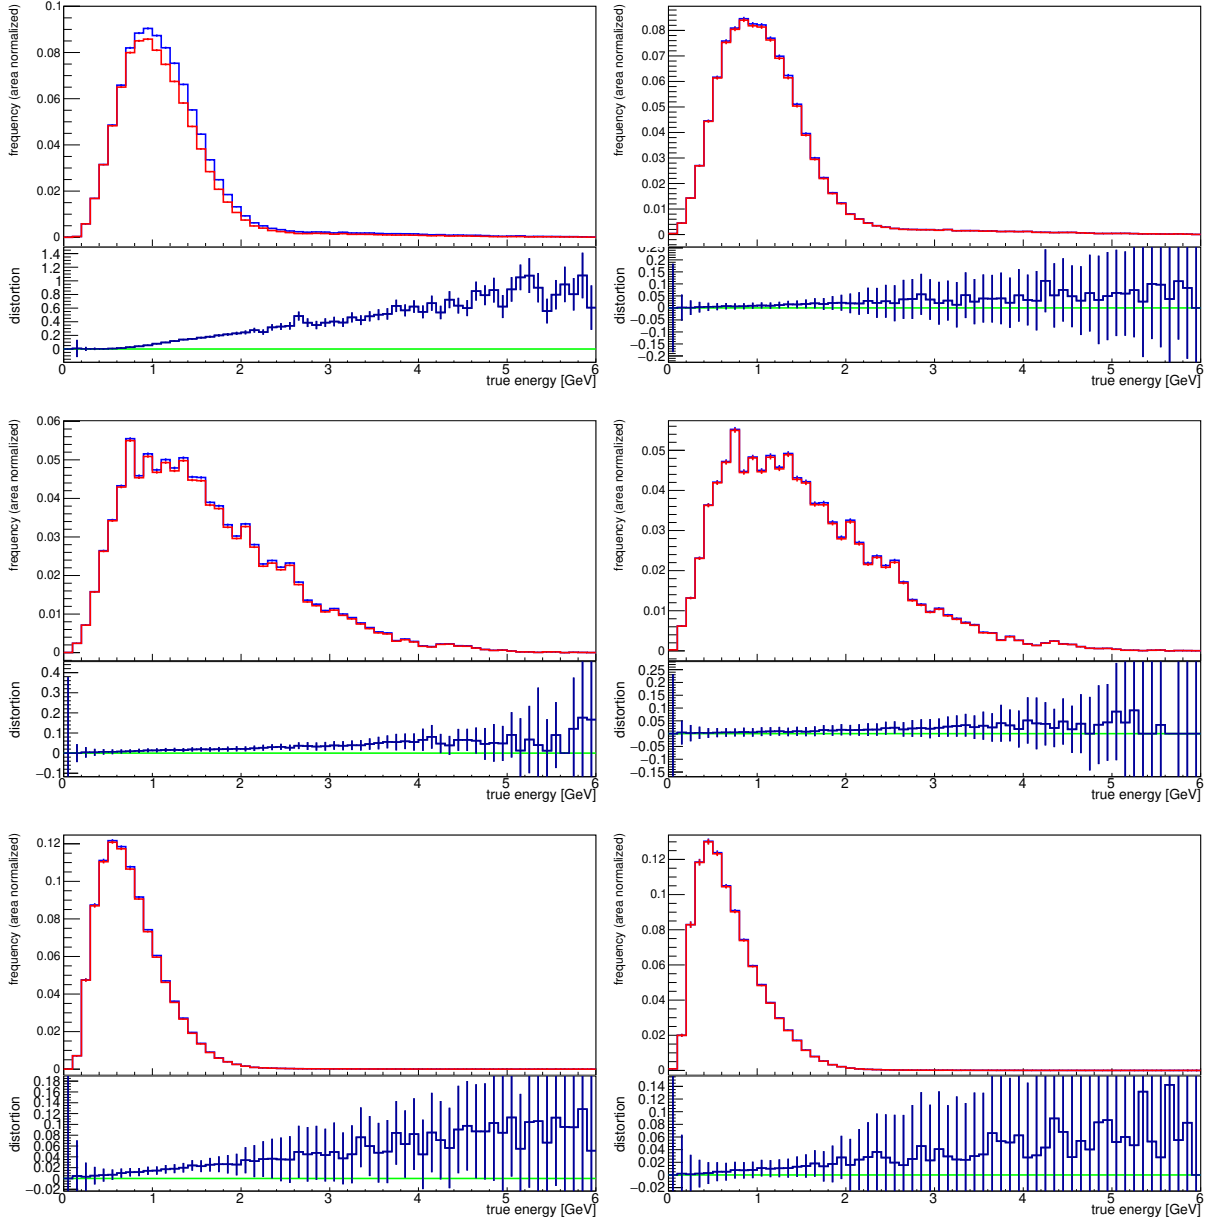


Figure 6.4: The neutrino energy dependence of the auto-veto fraction leads to distortion of the neutrino energy spectrum as shown by comparing the true spectrum before (blue) and after auto-veto (red) for the three different BNB neutrino samples: ν_μ (top row), intrinsic ν_e (middle row), and oscillated ν_e . Charged-current and neutral-current interactions are in the left and right columns respectively. The spectra are shown in the upper subpanels while the bin-by-bin fractional differences are shown in the lower subpanels.

To understand the impact on the reconstructed neutrino energy spectrum, I compare the true neutrino energy spectrum before and after auto-veto. The results for events

with the true neutrino vertex position located in the FV separated by CC or NC for the three different neutrino samples are shown in Figure 6.4. The full and auto-vetoed spectra are normalized to the integral of the full spectrum. The bin-by-bin fractional differences are shown in the histograms below the spectra. For the ν_μ CC, intrinsic ν_e CC, and oscillated ν_e CC samples, the bin distortion at the peak is 20%, 2%, and 1% respectively. The distortion increases at higher energies in each sample, up to 80%, 10% and 8% respectively. Depending on the interplay between the auto-veto induced distortion and systematics introduced by other event selection methods, this level of distortion may be insignificant.

Taking signal efficiency losses and neutrino energy spectral distortion into account, we might decide that these affects are acceptable. Cosmogenic backgrounds are not considered to be an important background for the ν_μ analysis. If this is verified with beam-off data, we could decline to veto any ν_μ candidates.

Though the auto-veto fraction for ν_e CC events is relatively low (about 2%), we should try to reduce it. To avoid the auto-veto effect incurred by the simple veto approach, since the CRT is not part of the hardware trigger, I also consider using the CRT as part of higher level reconstruction. Rather than veto every event with a CRT hit in time with the beam spill, we could use the CRT hits to reduce the fiducial volume under certain conditions. For example, we saw in Section 2.5 how TPC methods alone could reject 99.9% of cosmogenic photons in the ν_e analysis, though this was not enough. The remaining cosmogenic contributions to the photon background are primarily due to secondary photons produced by muons passing in close proximity to, but not entering, the active volume. CRT hits provide additional information that can be used to reject these backgrounds. Depending on the location of the CRT hit, we can define a cone centered on the CRT hit that extends into the fiducial volume. If we remove the intersected volume from the analysis, we can reject backgrounds with a significant reduction in auto-veto. This requires some combination with TPC reconstruction, which is beyond the scope of this work.

In this section, I showed that the current simulation predicts the auto-veto fraction to be low and may be negligible. The CRT tagging efficiency for cosmogenic muons remains at or above the target set in the SBN proposal. In the next section, I continue the discussion of using the CRT as part of higher level reconstruction.

6.3 Time-of-Flight Veto

Both the CRT and the PDS are expected to achieve nanosecond-level time resolution. We can combine the Photon Detection System (PDS) with the CRT to possibly improve upon the CRT only method in terms of reducing losses in neutrino efficiency and avoiding possible distortion in the reconstructed neutrino energy spectrum due to auto-veto. By




Figure 6.5: An illustration of the time-of-flight veto concept shows a muon entering from the top CRT before entering into the active liquid argon volume (AV). The time-of-flight is given by the delay between the CRT hit and the first PMT signal, approximated here by the shortest path to the PMT closest to the AV entry point .

matching reconstructed objects between the two systems, we can distinguish between incoming and outgoing tracks from the argon volume by calculating a time-of-flight (TOF), which is why I call this method the TOF veto.

I start with a truth-level study, calculating the earliest possible time of arrival as illustrated in Figure 6.5. The true TOF for cosmogenic muons is shown in Figure 6.6a.

There are two modes in the distribution: the dominant mode is due to muons that enter from the top CRT where we have a few meters of space between the CRT and the argon volumes; the subdominant mode is due to muons that enter from the side CRT, which is located just outside of the cryostat. There is a clear low edge at about 6 ns.




Figure 6.6: The TOF distribution for cosmogenic muons from the CRT to the earliest possible light to arrive at a PMT in truth (a) or using the full simulation and reconstruction chain for the CRT and applying a 2-ns smearing to the PMT arrival time (b) demonstrate the feasibility of the TOF veto method.

To understand the impact of moving to the full CRT simulation and reconstruction chain, I replace the true CRT hit time with the reconstructed time. Then, to understand the possible impact from the expected PMT time resolution, I apply a 2-ns Gaussian smearing to the true photon arrival time at the PMT. The resulting TOF distribution is shown in Figure 6.6b. The modal structure of the distribution is no longer resolvable. In this case, 98% of cosmogenic muons present a TOF larger than 0 ns.

The last step before considering the realistic scenario with a full combined CRT-PDS simulation and reconstruction with cosmogenic backgrounds is to analyze single particle samples with the full simulation and reconstruction chain for both the CRT and PDS. This serves as a sanity check. In Figure 6.7, you can see the TOF distributions for single, vertical, 4-GeV muons beginning either above the top CRT and propagating downward into the TPC or in the TPC center and propagating upward through the top CRT. We

see a clear separation between incoming and outgoing tracks with an absolute value of the most probable TOF being of 19 ns with an 3 ns RMS.




Figure 6.7: A simple test of the combined CRT-PDS reconstruction chain using vertical muons starting either downgoing from above CRT or upgoing from the TPC center show the absolute TOF provides a clean separation with a 19 ns mean and a 3 ns RMS.

In order to match CRT hits to optical flashes, I calculate the time difference between CRT hits and optical flashes and choose the minimum difference as the TOF. With exclusive neutrino samples, there is no problem with ambiguity. However, cosmogenic samples are complicated with several interactions uniformly distributed over the TPC readout window. As I showed in Section 6.1, about 50% of cosmogenic muons that generate a CRT hit do not enter the active volume. Other cosmogenic interactions will generate optical flashes, and this introduces ambiguity in matching without any TPC reconstruction.

Further complicating matters, about 30% of cosmogenic muons will generate multiple CRT hits. A track that crosses the Top CRT, passes through the argon volume, and exits through the Bottom CRT will generate two CRT hits. The TOF reconstruction would result in one incoming and one outgoing match. The outgoing match would be identified as most likely originating from a neutrino interaction and would therefore not be rejected. The solution to this problem is to incorporate CRT tracks into the TOF matching algorithm. Only the CRT hit with the earliest time stamp is used for the TOF match while the remaining CRT hits associated with the track are ignored.

In examining the candidate TOF distribution for cosmogenic muons, I observe a

uniform distribution for cases with $|\text{TOF}| > 70$ ns. This is the expectation for random coincidences and these matches can be rejected. For this analysis, the TOF limit set to ± 40 ns.

The other factor that can limit the efficacy of the TOF veto is the optical flash reconstruction efficiency. Optimization studies of the trigger system are currently ongoing; however, preliminary results show trigger efficiencies for charged-current neutrino events to be better than 99%. Trigger efficiencies for cosmogenic interactions are similar to those for ν_μ CC interactions since, with 3-m concrete overburden, most cosmogenic interactions are due to muons.

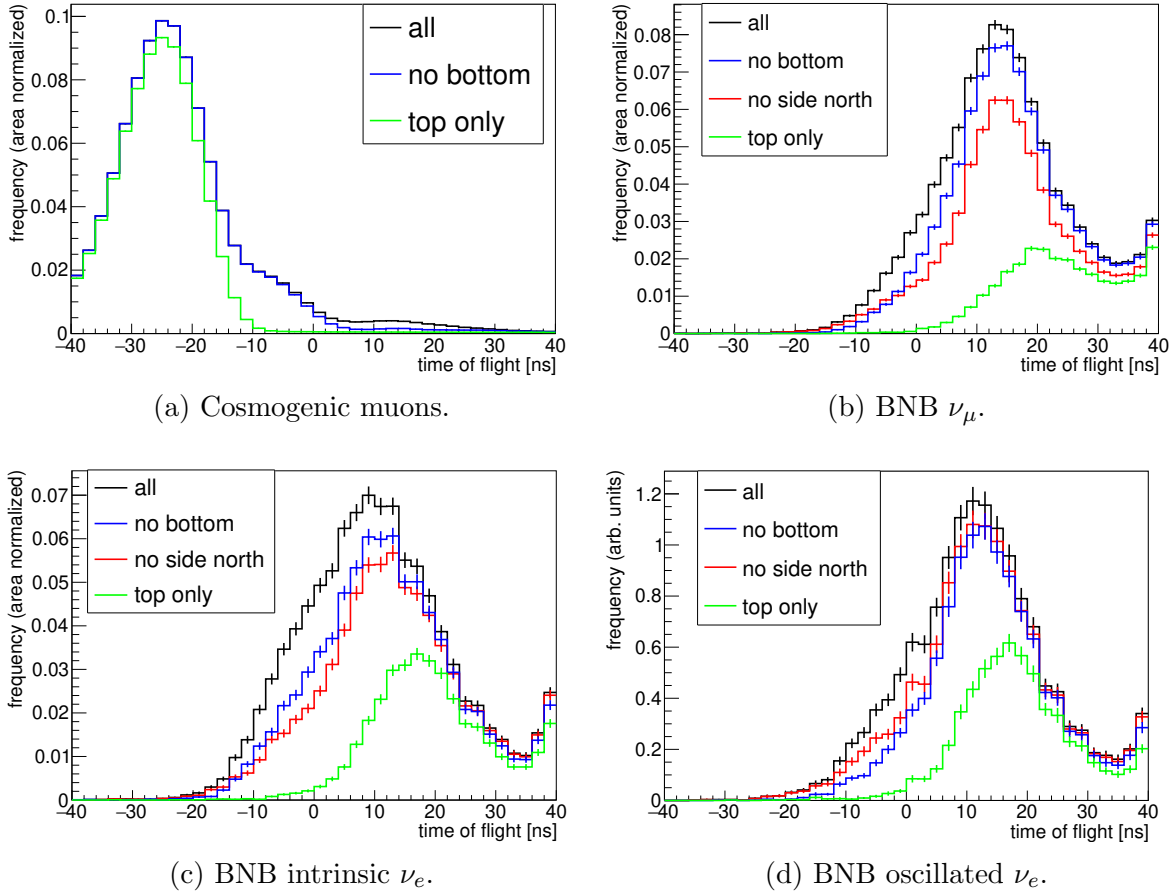


Figure 6.8: Shown here are the TOF distributions for cosmogenic particles (a) and BNB neutrino samples (b,c,d) with $\sin^2 2\theta=0.013$, $\Delta m^2=0.43$ eV².

TOF distributions obtained with the full combined CRT-PDS reconstruction chain

and the matching procedure described above are shown in Figure 6.8. The vast majority of cosmogenic matches have TOF values below 0 ns, while the opposite is true for the neutrino samples. I have included distributions using the full CRT system, ignoring hits from the bottom CRT, neglecting the Side CRT North Wall, and using the Top CRT only. Not surprisingly, with the few-meter gap and 3-ns time resolution of the Top CRT, the distributions using the Top CRT only outperform the rest. For this analysis, I will use the distributions with Bottom CRT hits excluded. Some optimization must be done in setting the TOF cut to reject cosmogenic backgrounds and avoid auto-veto. To be worthwhile, the TOF veto method should achieve a similar level of rejection power as for the simple veto method but with a lower auto-veto rate.

To compute the efficiency and purity, I will adopt some numbers from the SBN proposal. See Section 2.5 for more details. These values will be exclusively for the ν_e CC appearance analysis using the global best fit point to the LSND allowed region, $\sin^2 2\theta=0.013$, $\Delta m^2=0.43 \text{ eV}^2$. Recall that the CRT is most crucial for values of Δm^2 below about 1 eV^2 . After dE/dx and distance-to-the-muon cuts, 204 cosmogenic background events remain, including a $15\%/\sqrt{E}$ energy resolution. After fiducial volume cuts and an applied threshold of 200 MeV, we have 498 ν_e CC candidates.

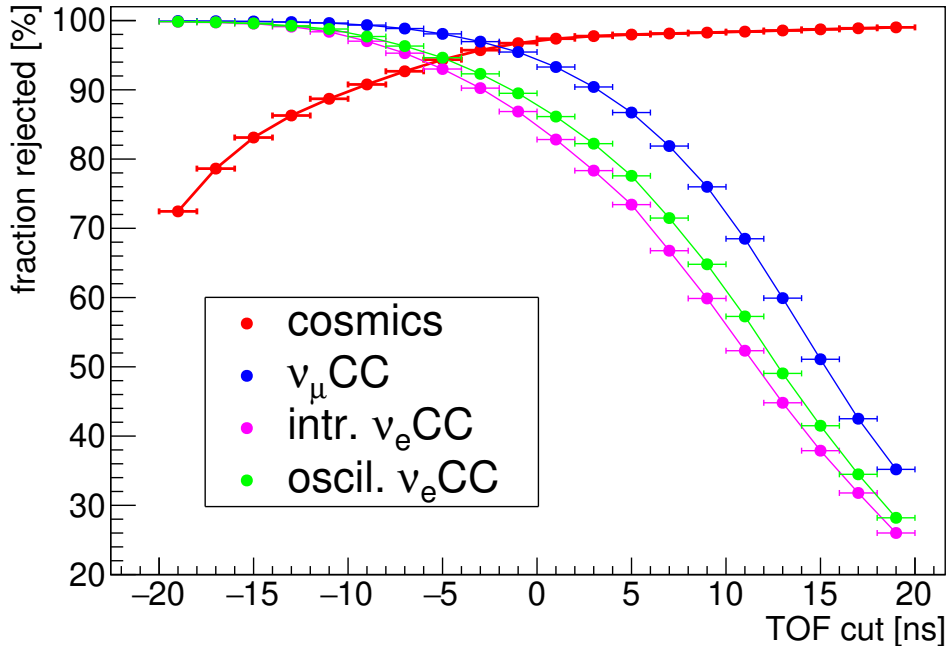
To understand the performance of the TOF veto method, I will define efficiency and purity in terms of cosmogenic background events and auto-veto losses. Contributions from other backgrounds, the beam intrinsic ν_e for example, will not be included. The efficiency is defined as the $1 - \epsilon_{\text{auto-veto}}$, and the purity is defined as $N_{\text{signal}}/(N_{\text{signal}} + N_{\text{background}})$.

In Figure 6.9a, the TOF rejection efficiency is shown for cosmogenic muons as well as for each neutrino sample as a function of the TOF cut value. Note that this does not include cosmogenic muons rejected with no reconstructed TOF, those outside of the 40-ns gate. The way the TOF is used is to identify an auto-veto event and retain the corresponding candidate neutrino event in the analysis. If the TOF value falls below the TOF cut or no TOF is reconstructed, the candidate neutrino event is rejected. The


efficiencies shown for the neutrino samples are in terms of identifying auto-veto events.

Now I turn to the purity and efficiency for the ν_e CC appearance sample. The result is shown in Figure 6.9b. For comparison, I show the efficiency, in terms of loss from auto-veto, and purity from cosmogenic backgrounds by using the simple veto method. For all values of the TOF cut between -20 ns and 20 ns, the efficiency with the TOF veto exceeds that for the simple veto by up to 1%, about 5 events. At a TOF cut of 3 ns, the cosmogenic rejection efficiency becomes flat while the purity for the ν_e sample also becomes flat. The purity for the TOF veto is about 1% lower than for the simple veto. For this same choice of the TOF cut, the efficiency for the TOF veto is about 1% better than the simple veto.

I have demonstrated the feasibility of the TOF veto method. I have shown that we can achieve sufficient time resolution to distinguish incoming from outgoing tracks from the TPC with the level of ambiguity at the level of a few percent. This method achieves similar performance as with the simple veto with the benefit of a significant reduction in the auto-veto fraction for the ν_μ CC sample.



(a) Background and auto-veto rejection.



(b) Efficiency and purity for oscillated ν_e CC.

Figure 6.9: The performance of the TOF veto method is shown in terms of separating cosmogenic muons from neutrinos (a) and in terms of the purity and efficiency of the ν_e CC appearance sample in terms of cosmogenic backgrounds and auto-veto respectively. Oscillation parameters used here are $\sin^2 2\theta=0.013$, $\Delta m^2=0.43$ eV².

6.4 Impact on Sterile Neutrino Search

I have shown that the CRT is a powerful tool for cosmogenic background rejection, but there are some caveats. Depending on the method of applying the CRT, we could introduce inefficiencies or systematic errors into the sterile neutrino analysis. In this section, I use the baseline sensitivities presented in Chapter 2 as a benchmark. The Oscillation Sensitivities Working Group is still in the beginning stages of incorporating the full reconstruction chain for each SBN experiment into the sensitivities and systematics analysis. In addition, the techniques I have developed for ICARUS need to be implemented in a similar way for SBND. Once the reconstruction code has been tested for each experiment in a common, shared analysis framework, the full impact on the oscillation analysis will be determined.

For the sensitivities presented in the SBN proposal, a flat rejection factor of 95% was applied to all cosmogenic backgrounds, assumed to have been provided by some combination of the CRT with the PDS. In Section 2.5, I showed that this level of rejection was required to maintain 5σ significance below Δm_{41}^2 values of about 1 eV^2 in testing the sterile neutrino hypothesis along the LSND 99% C.L. contour. This result took into account an approximate 3% loss in the ν_e CC signal sample due to auto-veto. However, the impact on systematics due to distortion of the reconstructed neutrino energy spectrum was not taken into account.

In this chapter, I have shown that the simple and TOF veto approaches should be able to maintain maintain 95% rejection efficiency. As an improvement over the simple veto, the TOF veto brings the auto-veto fraction to negligible levels for both ν_μ CC and ν_e CC. While a relative low-mass squared difference (0.43 eV^2) was used in my analysis, the BNB intrinsic ν_e sample has a larger mean energy. The effect of increasing the mass-squared difference is to re-weight the parent energy distribution inherited from the ν_μ sample, biasing it toward higher energies. I showed that the auto-veto rate for both the intrinsic

and oscillated ν_e samples is low. Thus, my result should be robust across the region of interest. With the assumptions of the SBN proposal supported by this fully simulated and reconstructed result, it appears that SBN is still well positioned to achieve 5σ significance in the low Δm^2 region.

Here, I have presented two different methods for cosmogenic background rejection. Cosmogenic backgrounds are complicated and difficult to model, although we expect the dominant effects to be well modeled. The results from my studies are encouraging, but they need to be validated with data. Depending on the outcome, one method may be better suited to the oscillation analysis than the other; multiple options provide some insurance. Since it is not anticipated that the CRT will be incorporated into the TPC trigger system, either option can be considered in the oscillation analysis. There is reason to be cautiously optimistic about the results I have presented here given the outcome of the data-Monte Carlo comparison study presented in Chapter 5.

Chapter 7

Conclusion and Outlook

7.1 Summary

It is an exciting time for the Short-Baseline Neutrino Program. MicroBooNE has completed its data and is completing its flagship low-energy excess analysis; ICARUS is about to begin physics running; SBND is under construction and is on track to start commissioning in about one year. As a collaboration, we are poised to make several world-leading cross-section measurements of neutrinos on argon. We are in a position to make a 5σ discovery or exclusion of the sterile neutrino interpretation of the short-baseline neutrino anomalies in both ν_μ -disappearance and ν_e -appearance channels. Of course, this will depend on exactly what nature has in store for us. If sterile neutrinos do exist, with a mass-squared difference of the order of 1 eV^2 , we must contend with several sources of background, especially from cosmogenic muons.

For lighter sterile neutrinos, the Cosmic Ray Tagging System is essential for maintaining 5σ sensitivity below values of Δm_{41}^2 of about 1 eV^2 . I played a crucial role in developing the Side CRT subsystem based on repurposed MINOS scintillator modules. This work included testing 173 modules, characterizing their response to cosmogenic muons in terms of light output and attenuation, developing a new optical readout, and installing and commissioning the first 25% of the system.

Using the data acquired at the Side CRT test stand, I developed a data-driven model for the CRT detector response including light output, attenuation, and electronics effects, with the goal of making the CRT simulation as realistic as possible with the available data. I tested the simulation for consistency with the input model. I analyzed the first

data from the production Side CRT system. Using this data, I led a study with my CSU colleagues to perform a first data-Monte Carlo comparison that tested the accuracy of the simulation without tuning. These results show that the CRT detector simulation, while not perfect, produces results that are qualitatively similar to the data. The most probable light yield predicted by the simulation agrees with the data to within 25%. Furthermore, the CRT trigger rates agree within 5-10% between the simulation and the data. This directly impacts any conclusions drawing regarding the tagging efficiency of the CRT. From this, I concluded that the current simulation is suitable for developing tools for cosmogenic background rejection.

In addition to the detector simulation, I developed a suite of CRT reconstruction tools both for low-level, general-purpose reconstruction as well as high-level tools incorporating the Photon Detection System. I have shown that we can reject 95% of cosmogenic muons with the simplest application of the CRT system, as a simple veto, whereby any neutrino event candidate with a CRT trigger or reconstructed hit in time with the beam spill is rejected. I also showed that this method of cosmogenic background rejection induces neutrino sample loss at the level of a few percent but leads to a small distortion of the reconstructed neutrino energy spectrum. This effect is known as auto-veto.

To mitigate the auto-veto problem, I incorporated additional information from the Photon Detection System (PDS). Using the nanosecond-level timing of both the CRT and the PDS, I matched CRT hits to reconstructed scintillation flashes in the TPC. With this extra information, we distinguish incoming from outgoing tracks in the TPC and reject backgrounds with only minimal adverse impact on the neutrino samples. In the particular algorithm configuration I used to generate my results, I was able to obtain 97.5 % purity for the ν_e CC appearance sample ($\sin^2 2\theta=0.013$, $\Delta m^2=0.43 \text{ eV}^2$) with an efficiency of 99.5% while reducing the auto-veto fraction for the ν_μ CC sample by 85%.

7.2 Next Steps

The tools I developed for CRT simulation, reconstruction and analysis are primed and ready for real data with some exceptions. The detector simulation for the CRT was only tested with side CRT data. There is no reason to expect the models to be completely accurate in describing the other CRT subsystems; however, the current model should produce reasonable results. As we saw with the data-Monte Carlo comparison, the detector simulation for the Side CRT could use some tuning. As data for the other subsystems become available, separate light yield and attenuation models should be implemented and tested. The reconstruction algorithms should remain valid however.

The remainder of the CRT system will be installed and commissioned this year. There is still some tuning of the current system configuration to be done including equalizing SiPM gains and discriminator thresholds. The tools developed for this task will be extended to the full system.

Additional validation should be done with the full CRT system. A study should be performed similar to the data-Monte Carlo study presented in Chapter 5. With the full system, we should quickly amass a sample of CRT-reconstructed muon tracks. We should be able to make a more firm comparison with the simulation prediction for the cosmogenic muon angular distribution. In addition, we can project these CRT tracks into the TPC active volume. This will be a useful tool for further studies as I explain below.

By the time the full system comes online, the production version of the CRT DAQ should be available to facilitate merging the CRT data streams with the other subsystems' data streams via a central event builder server. To understand the synchronization between the data streams several cross-checks will need to be made. Cosmogenic muons provide a convenient source for cross-system calibration. By matching CRT tracks to TPC tracks and matched scintillation flashes, we can ensure that synchronization has been achieved. This is an important prerequisite to applying the CRT to cosmogenic

background rejection.

If the collaboration opts to use the simple veto approach, this method could begin use at this stage. If instead the collaboration opts to use the time-of-flight veto method, some other studies should first be performed. Using the flash matching only, the photon library should be validated within the active volume. It is difficult to validate the photon visibility due to ionization in the inactive argon volume using the TPC alone however. Using the CRT, reconstructed CRT tracks in particular, we can select precisely these events to test the light yield and photon visibility in the inactive volume where the light yield is expected to be larger than inside the drift field. Despite the results of my simulation based study showing that flashes due to activity in the inactive volume are negligible, we need to ensure that we observe non-negligible impact on CRT-to-flash matching with data. If these types of flashes do occur at a significant rate, my analysis should be updated to take this case into account. In fact, it might even be beneficial as muons missing the active volume that produce photons that go on to convert in the active volume are a dangerous background indeed.

While there is much work yet to be done, we have many effective tools in place. With our current team, the tasks I have outlined above could be reasonably completed by the start of return of the BNB after the summer shutdown at the end of 2021. By this time, we should be well positioned to understand the cosmogenic background and the best way to mitigate it. Then, the hunt for new physics can begin.

References

- [1] C. Patrignani *et al.* (Particle Data Group). *Neutrino Cross Section Measurements*. Chin. Phys. C,40, 100001 (2016) and 2017 update.
- [2] I. Brivio and M. Trott. *The Standard Model as an Effective Field Theory*. Physics Reports 793, 1-98 (2019).
- [3] André de Gouvêa. *Neutrinos Have Mass - So What?*. Modern Phys. Letters A 2004 19:38, 2799-2813.
- [4] J.N. Bahcall. *Solar Neutrino Cross Sections and Nuclear Beta Decay*. Phys. Rev. 135, B137 (1964).
- [5] R. Davis, Jr., D.S. Harmer, and K.C. Hoffman. *Search for Neutrinos from the Sun*. Phys. Rev. Lett. 20, 1205 (1968).
- [6] C. Giganti, S. Lavignac, M. Zito. *Neutrino Oscillations: the Rise of the PMNS Paradigm* arXiv:1710.00715 [hep-ex] (2017).
- [7] M. Tanabashi *et al.*(Particle Data Group). *Neutrino Masses, Mixing, and Oscillations*. Phys. Rev. D98, 030001 (2018) and 2019 update.
- [8] A.B. Balantekin. *Nuclear Physics and Astrophysics of Neutrino Oscillations*. JPS Conf. Proc. 14, 010701 (2017).
- [9] IceCube Collaboration. *The IceCube Neutrino Observatory: Instrumentation and Online Systems*. arXiv:1612.05093 [astro-ph.IM] (2016).
- [10] Daya Bay Collaboration. *A Precision Measurement of the Neutrino Mixing Angle θ_{13} Using Reactor Antineutrinos at Daya Bay*. arXiv:hep-ex/0701029 (2007).

- [11] C. Palomares (for Double-Chooz Collaboration). *Double-Chooz Neutrino Experiment*. arXiv:0911.3227 [hep-ex] (2009).
- [12] S-B Kim and the RENO Collaboration. *RENO: reactor experiment for neutrino oscillation at Yonggwang*. J. Phys.: Conf. Ser. 120 052025 (2008).
- [13] M. Abbes *et al.* *The Bugey 3 Neutrino Detector*. Nucl. Inst. and Meth. A 374 P164-187 (1996).
- [14] Mikhail Danilov (On behalf of the DANSS Collaboration). *Recent results of the DANSS experiment*. arXiv:1911.10140 [hep-ex] (2019).
- [15] J.Ashenfelter *et al.* (PROSPECT Collaboration). *The PROSPECT reactor antineutrino experiment*. Nucl. Inst. and Meth. A 922, P287-309 (2019).
- [16] N. Allemandou *et al.* (STEREO Collaboration). *The STEREO Experiment*. JINST 13 (2018) P07009.
- [17] Youngju Ko *et al.* *NEOS Experiment*. 2019 J. Phys.: Conf. Ser.1216 012004
- [18] V.N.Gavrin *et al.* *Sage: The Soviet-American Gallium Solar Neutrino Experiment*. Nuclear Physics B - Proceedings Supplements 28(1) P75-81 (1992).
- [19] GALLEX Collaboration. Phys. Lett. B 284(4) P376-389 (1992).
- [20] M. Aker *et al.* (KATRIN Collaboration). *First Operation of the KATRIN Experiment with Tritium*. Eur. Phys. J. C 80, 3, 264 (2020).
- [21] M. Auger *et al.* *The EXO-200 Detector, Part I: Detector Design and Construction*. JINST 7 (2012) P05010.
- [22] I. Nutini *et al.* (CUORE Collaboration). *The CUORE Detector and Results*. J. Low Temp. Phys. 199, 519–528 (2020).

- [23] M. Agostini *et al.* *Search of Neutrinoless Double Beta Decay with the GERDA Experiment*. Nuclear and Particle Physics Proceedings 273-275, P1876-1882 (2016).
- [24] M.A. Acero *et al.* (NOvA Collaboration) *First Measurement of Neutrino Oscillation Parameters using Neutrinos and Antineutrinos by NOvA*. Phys.Rev.Lett. 123 (2019) 15, 151803.
- [25] K. Abe, R. Akutsu, A. Ali, *et al.* (T2K Collaboration). *Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations*. Nature 580, 339–344 (2020).
- [26] J.N. Abdurashitov *et al.* *Measurement of the Response of a Ga Solar Neutrino Experiment to Neutrinos from a ^{37}Ar Source*. Phys. Rev. C 73, 045805 (2006).
- [27] P. Anselmann *et al.* *First Results from the ^{51}Cr Neutrino Source Experiment with the GALLEX Detector*. Phys. Lett. B 342, 440 (1995).
- [28] W. Hampel *et al.* *Final Results of the ^{51}Cr Neutrino Source Experiments in GALLEX*. Phys. Lett. B 420, 114 (1996).
- [29] C. Giunti and M. Laveder. *Statistical Significance of the Gallium Anomaly*. Phys. Rev. C 83, 065504 (2011).
- [30] S. Gariazzo, C. Giunti, M. Laveder, and Y.F. Li. *Updated Global 3+1 Analysis of Short-Baseline Neutrino Oscillations*. JHEP 06 (2017) 135.
- [31] A. Aguilar *et al.* (LSND Collaboration). *Evidence for Neutrino Oscillations from the Observation of $\bar{\nu}_e$ Appearance in a $\bar{\nu}_\mu$ Beam*. Phys.Rev. D64, 112007 (2001).
- [32] A.A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration). *Significant Excess of Electronlike Events in the MiniBooNE Short-Baseline Neutrino Experiment*. Phys. Rev. Lett. 121, 221801 (2018).

- [33] M. Dentler *et al.*. *Updated Global Analysis of Neutrino Oscillations in the Presence of eV-Scale Sterile Neutrinos*. Journal of High Energy Physics (JHEP) 08 (2018) 010.
- [34] R. Acciarri *et al.*. *A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam*. arXiv:1503.01520 [physics.ins-det].
- [35] C. Rubbia. *The Liquid Argon Time Projection Chamber: a New Concept for Neutrino Detectors*. CERN-EP-INT-77-08, CERN-EP-77-08.
- [36] J.N. Marx and D.R. Nygren. *The Time Projection Chamber*. Physics Today 31, 10, 46 (1978).
- [37] Brookhaven National Laboratory. *Liquid Argon Properties (Tables and Calculators)*. <https://lar.bnl.gov/properties/>.
- [38] M. Babicz *et al.*. *Propagation of scintillation light in Liquid Argon*. JINST 15 (2020) 03, C03035.
- [39] C. Anderson *et al.* (ArgoNeut Collaboration). *The ArgoNeuT Detector in the NuMI Low-Energy Beam Line at Fermilab*. Journal of Instrumentation (JINST), 2012 JINST Vol. 7 P10019.
- [40] MicroBooNE Project. *The MicroBooNE Technical Design Report*. CD3b Review (February 2012).
- [41] B. Abi *et al.* (DUNE Collaboration). *Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume I Introduction to DUNE*. arXiv 2002.02967 [physics.ins-det].
- [42] S. Amerio *et al.* (ICARUS Collaboration). *Design, Construction and Tests of the ICARUS T600 Detector*. Nucl. Inst. and Meth. A 527, P329-410 (2004).
- [43] M. Antonello *et al.*. *Precise 3D Track Reconstruction Algorithm for the ICARUS T600 Liquid Argon Time Projection Chamber Detector*. AHEP (2013) 260820.

- [44] A. Ankowski *et al.* (ICARUS Collaboration). *Energy Reconstruction of Electromagnetic Showers from π^0 Decays with the ICARUS T600 Liquid Argon TPC*. arXiv 0812.2373 [hep-ex].
- [45] C. Rubbia *et al.* *Underground Operation of the ICARUS T600 LAr-TPC: First Results*. JINST 6 P07011 (2011).
- [46] M. Antonello *et al.* *Muon Momentum Measurement in ICARUS-T600 LAr-TPC via Multiple Scattering in Few-GeV Range*. JINST 12 P04010 (2017).
- [47] A.A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration). *The Neutrino Flux prediction at Mini-BooNE*. Phys.Rev. D79, 072002 (2009).
- [48] M.G. Catanesi *et al.* (HARP Collaboration). *Measurement of the Production Cross-Section of Positive Pions in the Collision of 8.9-GeV/c Protons on Beryllium*. Eur.Phys.J. C52, 29–53 (2007).
- [49] M. Babicz *et al.* *Test and characterization of 400 Hamamatsu R5912-MOD photomultiplier tubes for the ICARUS T600 detector*. JINST 13, P10030 (2018).
- [50] L. Bagby *et al.*. *New Read-Out Electronics for ICARUS-T600 Liquid Argon TPC. Description, Simulation and Tests of the New Front-End and ADC System*. JINST 13, P12007 (2018).
- [51] MicroBooNE Collaboration. *First Neutrino Interactions Observed with the MicroBooNE Liquid-Argon TPC Detector*. MICROBOONE-NOTE-1002-PUB (2015).
- [52] R. Acciarri *et al.* (ArgoNeuT Collaboration). *First Observation of Low Energy Electron Neutrinos in a Liquid Argon Time Projection Chamber* Phys. Rev. D 95, 072005 (2017).
- [53] P.A.N. Machado, O. Palamara, D.W. Schmitz. *The Short-Baseline Neutrino Program at Fermilab*. Annu. Rev. Nucl. Part. Sci. 2019. 69:363–87.

- [54] Pattison *et al.* *High resolution NaI(Tl) counters.*
<http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2842777>
- [55] A. Pla-Dalmau, A.D. Bross, V.V. Rykalin. *Extruding Plastic Scintillator at Fermilab.*
<https://lss.fnal.gov/archive/2003/conf/fermilab-conf-03-318-e.pdf>
- [56] M. Auger *et al.* *Multi-Channel Front-End Board for SiPM Readout.* JINST, 11, P10005 (2016).
- [57] Kuraray Wavelength Shifting Fibers. <http://kuraraypsf.jp/psf/ws.html>
- [58] SensL C-Series SiPM data sheet. <http://sensl.com/downloads/ds/DS-MicroCseries.pdf>
- [59] Hamamatsu SiPMs.
<https://www.hamamatsu.com/us/en/product/optical-sensors/mppc/index.html>
- [60] Particle Data Group. *Cosmic Rays.* <http://pdg.lbl.gov/2019/reviews/rpp2019-rev-cosmic-rays.pdf>
- [61] M. Kostin *et al.* (MINOS Collaboration). *The Magnetized Steel and Scintillator Calorimeters of the MINOS Experiment.* Nucl.Instrum.Meth.A 596 (2008) 190-228.
- [62] Simon Nieswand, Masterarbeit in Physik. *Measurement of the Exit Characteristics of Light from Optical Multimode Plastic Fibres.* Rheinisch-Westfälischen Technischen Hochschule – Aachen (2014).
- [63] R. Pordes, E. Snider (LArSoft Collaboration). *The Liquid Argon Software Toolkit (LArSoft): Goals, Status and Plan.* PoS(ICHEP2016)182 2017.
- [64] C. Andreopoulos *et al.* *The GENIE Neutrino Monte Carlo Generator.* Nucl. Instrum. Meth. A 614 (2010) 87-104.

- [65] C. Haggmann, D. Lange, D. Wright. *Cosmic-Ray Shower Generator (CRY) for Monte Carlo Transport Codes*. Nuclear Science Symposium Conference Record, 2007. NSS '07. IEEE Volume: 2.
- [66] D. Heck *et al.* *CORSIKA: a Monte Carlo Code to Simulate Extensive Air Showers*. Forschungszentrum Karlsruhe Report FZKA 6019 (1998).
- [67] S. Agostinelli *et al.* (Geant4 Collaboration). *Geant4—a Simulation Toolkit*. Nucl. Instrum. Meth. A 506 (2003) 250-303.
- [68] B. Baller. *Liquid Argon TPC Signal Formation, Signal Processing and Reconstruction Techniques*. JINST 12 P07010 (2017).
- [69] J. S. Marshall, M. A. Thomson. *The Pandora Software Development Kit for Pattern Recognition*. arXiv:1506.05348 [physics.data-an] 2015.
- [70] B. Jones. *Optical Simulations in LArSoft - Technical Manual*. <https://microboone-docdb.fnal.gov/cgi-bin/RetrieveFile?docid=2313> (2012).
- [71] K. Biery, E. Flumerfelt, J. Freeman, W. Ketchum, G. Lukhanin, R. Rechenmacher. *artdaq: DAQ Software Development Made Simple*. Journal of Physics: Conference Series (2017).