REACTOR ANTINEUTRINOS IN THE SNO+ WATER PHASE AND DETECTOR R&D FOR LARGE-SCALE NEUTRINO DETECTORS

Tanner Kaptanoglu

A DISSERTATION
in
Physics and Astronomy
Presented to the Faculties of the University of Pennsylvania
in
Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

2020

Supervisor of Dissertation	**Graduate Group Chairperson**
Joshua R. Klein	Joshua R. Klein
Professor, Physics and Astronomy	Professor, Physics and Astronomy

Dissertation Comittee:

Christopher Mauger, Associate Professor of Physics and Astronomy
Joseph Kroll, Professor of Physics and Astronomy
Cullen Blake, Assistant Professor of Physics and Astronomy
Mark Trodden, Professor of Physics and Astronomy

To my family.

Acknowledgements

I first wish to thank my advisor Josh Klein, who has been a phenomenal advisor, mentor, collaborator, and friend. Josh's persistent optimism, refusal to give up, and insistence on quality have been a consistent source of inspiration during my time at Penn. There are few graduate students who can claim that they have had two wonderful advisers, but that has what my PhD felt like sitting next to Gene Beier, whose wisdom I am lucky to have received on many an occasion. I would also like to thank Christopher Mauger for many thoughtful discussions and Rick Van Berg for providing essential physics and electronics assistance and for patiently helping me debug an unknown number of issues in our lab. Lastly, I would never have made it to Penn for graduate school, if it was not for my incredible undergraduate advisor, Gabriel Orebi Gann.

At Penn I have been incredibly fortunate to have developed an amazing group of friends. I want to thank everyone in the neutrino group over my time at Penn, particularly the post-docs Nuno Barros, Ian Coulter, Logan Lebanowski, and Ben Land, who all taught me a great deal. A very special thanks to Andy Mastbaum, who provided daily guidance on all things SNO+, treated me as an essential colleague (even as a clueless first-year), and became a close friend. Lastly, this thesis would not have been possible without the continued support of Eric Marzec. We spent countless hours working together on electronics (both above and under the ground), playing Heroes III, drinking beer at Eddies, and discussing the Wheel of Time.

I would like to also thank the good friends I have made outside of the neutrino group, especially all of my fellow first-years. In particular, my friendship with Chris Lynn and Eric Horsley has made much of graduate school far my enjoyable. Thanks to Chris for always being up for playing soccer, board games, ping pong, and everything else we got up to, in spite of the fact that I am much better than you at all of those things. Thanks also to my furry-friend Ella, who only scratched and bit me a handful of times.

My family has always been an incredible system of support, kindness and love. Thanks to my mother for reading everything I have ever written, despite not understanding a word over the last six years. Thanks to my father for his relentless passion for physics, which inspired me to become a physics major many years ago. Thanks to my best-friends and brothers, Alan and Derek, who are incredibly supportive and have always been amazing about reaching-out to me. Thanks to my wonderful grandmother, aunt, and uncle for their support and for hosting me every Thanksgiving. Lastly, I will be eternally grateful to my girlfriend Lia, who has provided more love, kindness, and patience than I can express. During these often demanding and stressful times, I am so lucky to have her.

ABSTRACT

REACTOR ANTINEUTRINOS IN THE SNO+ WATER PHASE AND DETECTOR R&D FOR LARGE-SCALE NEUTRINO DETECTORS

Tanner Kaptanoglu
Joshua R. Klein

This dissertation presents two topics, both inspired by my work on the SNO+ experiment. First, I describe a search for reactor antineutrino interactions in the SNO+ water-filled detector using a 190.3 day dataset. Reactor neutrinos have never before been detected in a water Cherenkov detector, primarily due to the difficulty in detecting the low energy signal from neutron captures on free protons. However, enormous effort towards driving the trigger threshold to an unprecedentedly low level and reducing radioactive background levels have made this search possible in SNO+. The analysis described in this thesis uses a likelihood-ratio based approach to identify potential inverse beta decay interactions by selecting event pairs that are temporally and spatially coincident. This thesis presents the signal and background expectation and studies the background rates in carefully chosen sidebands. The signal window has has been partially unblinded, and a total of one event is observed, consistent with the expectation. Future efforts that build on the work presented in this thesis and include more SNO+ water-phase data are expected to result in the first ever significant detection of reactor antineutrinos in a water Cherenkov detector.

The second topic in this thesis focuses on R&D effort relevant for current and future large-scale neutrino detectors, such as SNO+, KamLAND-Zen, JUNO, and `THEIA`. Detectors hoping to observe low-energy solar neutrinos and/or neutrinoless double beta decay will need high efficiency PMTs, state-of-the-art liquid scintillator, and new and emerging technology. My work in bench-top characterization of several modern PMTs and liquid scintillator mixtures has been critical for understanding the expected detector response and sensitivty of SNO+, future upgrades to SNO+, and other liquid scintillator detectors. Additionally, a new instrument called the dichroicon was developed to provide spectral photon sorting for large-scale Cherenkov and liquid scintillator detectors. The dichroicon is characterized on the bench-top and in simulation and shows excellent promise for use in future detectors.

Table of Contents

Title **i**

Dedication **ii**

Acknowledgements **iii**

Abstract **v**

List of Tables **xiii**

List of Figures **xxii**

1 Introduction **1**

2 Physics of Massive Neutrinos **4**

2.1 Neutrino Oscillations in Vacuum 5

2.2 Neutrino Oscillations in Matter 10

2.3 Neutrino Mass 13

2.4 Neutrino Mass Hierarchy 15

2.5 Neutrino Interactions 17

3 Neutrino Sources and Experiments 23

3.1 Solar Neutrinos 23

3.2 Atmospheric Neutrinos 28

3.3 Reactor Neutrinos 29

3.3.1 KamLAND 30

3.3.2 Daya Bay, RENO, Double Chooz, and JUNO 33

3.3.3 Very Short Baseline Experiments 36

3.4 Geoneutrinos 37

4 The SNO+ Detector 39

4.1 Detector Design 39

4.2 Detector Principles 43

4.2.1 Water Target 43

4.2.2 Scintillator Target 45

4.3 Photomultiplier Tubes 48

4.4 Electronics and DAQ 50

4.5 Electronic Upgrades 57

4.5.1 XL3 57

4.5.2 MTC/A+ 58

4.5.3 TUBII 60

4.5.4 CAEN 61

4.5.5 Latch Board 62

4.6 Calibration Sources 63

4.6.1 ^{16}N 63

4.6.2 AmBe 64

4.7 Trigger Efficiency 64

5 Simulation and Reconstruction 72
- 5.1 Event Generation 73
- 5.2 Event Reconstruction 76

6 A Reactor Antineutrino Search with SNO+ 82
- 6.1 Sources of Antineutrinos 83
- 6.2 Event Rate 84
- 6.3 Detecting Reactor Antineutrinos in SNO+ 90
- 6.4 Data Selection 93
  - 6.4.1 Dataset and Livetime 93
  - 6.4.2 Run Selection 94
  - 6.4.3 Event Selection 94
    - 6.4.3.1 Low-level Cuts 95
    - 6.4.3.2 High-level Cuts 96
    - 6.4.3.3 Follower Cuts 98
  - 6.4.4 Summary of Cuts 98
- 6.5 Likelihood Method for Rejecting Backgrounds 98
  - 6.5.1 PDF Generation 99
  - 6.5.2 Likelihood Ratio 104
- 6.6 Calibration using the AmBe Source 104
  - 6.6.1 Calibrating the Neutron Detection Efficiency 107
    - 6.6.1.1 Change to the SNO+ Trigger System 112
    - 6.6.1.2 Relative Efficiency Correction 113
- 6.7 Signal Detection Efficiency 115
- 6.8 Backgrounds 118
  - 6.8.1 Accidentals 119
    - 6.8.1.1 Δt Sideband 121
  - 6.8.2 (α, n) 127

6.8.2.1 Measurement of the (α, n) Rate on the Acrylic 133
6.8.3 Atmospherics 140
6.8.3.1 Multiplicity Sideband 146
6.8.4 Geoneutrinos 149
6.8.5 Cosmogenics 150
6.9 Background Summary with $\Delta\mathcal{L}$ Cut 152
6.10 Partial Unblinding 154
6.11 Summary 155

7 R&D for SNO+ and Future Large-Scale Water Cherenkov and Scintillator Detectors 158
7.1 Large-area PMTs 160
7.1.1 Experimental Setup 160
7.1.2 Analysis 161
7.1.3 R5912-MOD PMT 165
7.1.4 11-inch ETEL PMT 173
7.1.5 R5912-200 PMT 174
7.2 Liquid Scintillator 177
7.2.1 Experimental Setup 179
7.2.2 `RAT` Model 181
7.2.3 Emission Timing 183
7.2.4 Light Yield 189
7.2.5 Scaling to Larger Volumes 191

8 The Dichroicon 195
8.1 Neutrinoless Double Beta Decay 195
8.2 Neutrinoless Double Beta Decay with SNO+ 200
8.3 The Dichroicon: Concept 205

8.4 The Dichroicon: Measurements 216
  8.4.1 Filter Characterization 217
  8.4.2 Measurements with Single Bandpass Filters 222
  8.4.3 Measurements with Single Dichroic Filters 230
    8.4.3.1 Results with a Large-area PMT 233
  8.4.4 Dichroicon Characterization 233
    8.4.4.1 PMT Calibration 233
    8.4.4.2 Experimental Setup 235
    8.4.4.3 DAQ and Data Analysis 238
    8.4.4.4 Simulation Models 240
    8.4.4.5 Cherenkov Source Results 243
    8.4.4.6 LAB+PPO Results 248
    8.4.4.7 Pulse Shape Discrimination 256
    8.4.4.8 Off-Axis Source 258
  8.4.5 Large-Scale Detector Simulation 264

9 Conclusion 267

A Data-Cleaning Cuts 270

B Electronics Calibration 273
  B.0.1 ECAL 273
  B.0.2 ECA 275
  B.0.3 PCA 275
  B.0.4 Dark Noise 276

C Hit Cleaning **277**
C.0.1 Cross-talk Cut 277
C.0.2 PMT Afterpulsing 279
C.0.3 Channel Flags 284

D Optical Calibration **287**
D.1 Optical Calibration Sources 287
D.1.1 Laserball 287
D.1.2 TELLIE 288
D.2 Results 288

E Reactor Complexes **291**

F The AmBe Source Run List **293**

G Atmospheric Neutrino Event Simulation with the `GENIE` Generator **296**
G.1 Verifying the `GENIE` MC 298

References **301**

List of Tables

2.1 The best fit mixing parameters from the NuFIT global fit (1) for the normal hierarchy (NH) and inverted hierarchy (IH), discussed in Section 2.4. Note that Δm^2_{3l}=$\Delta m^2_{31} > 0$ for the NH and $\Delta m^2_{3l} = \Delta m^2_{32} < 0$ for the IH. 10

5.1 The fitted values for the coefficients in Equation 5.1 used to model the antineutrino energy distribution from nuclear reactors. The values for ^{235}U, ^{239}Pu, and ^{241}Pu are from (2) and the values for ^{238}U are from (3). 73

6.1 Various cut variables and values for the prompt and delayed events. The efficiency of each cut on the reactor IBD events is detailed in Table 6.5. 99

6.2 The relative efficiency (data/MC) for two runs for different $\tilde{n}_{100}$ cut values. The AmBe calibration source in run 109133 is in the center and in 109153 it is located at (0, -4, 0) m. 111

6.3 The various corrections factors applied to the relative efficiency maps generated using the AmBe source. 114

6.4 The relative efficiency (data/MC) correction factors for four regions of the detector, after applying all cuts in Table 6.1, the additional AmBe criteria, and the likelihood ratio cut. Note these regions are mutual exclusive in descending order (i.e., z $\geq$ 4 includes all ρ, whereas $\rho \geq 3.5$ is for z positions between 4 and -4 m). The uncertainties are taken by using the standard deviation of the average as the statistical uncertainty, in quadrature with the systematic uncertainty from the trigger efficiency variation scaling (15%). 114

6.5 The breakdown of the signal efficiency of each cut on the prompt and delayed events. The events are simulated in a volume of 8500 mm. The details of the cuts are given in Table 6.1, some of which are reproduced in this table. An example likelihood ratio cut is included for completeness and does not indicate where the cut will be placed in the final analysis. No scale factors or sacrifices are applied to the efficiencies presented in this table, and additional scaling for differences in the neutron detection efficiency reduce the overall detection efficiency (see Table 6.6). 116

6.6 The various correction factors used to adjust the expected detection efficiency. The last two factors are specific to the reactor IBD events. 117

6.7 The breakdown of the effect of each cut on the prompt and delayed events in the data. The data-cleaning, pedestal cut, and fit valid cut are applied to all events in order to trim the data. The rest of the cuts are separated by prompt and delayed events. The details of the cuts are given in Table 6.1, some of which are reproduce in this table. Note that the Δt cut is shown twice; the $\Delta t < 5$ ms defines the coincidence and the second Δt cut rejects backgrounds, although in this case the cut is shown for the sideband analysis. An example likelihood ratio cut is included for completeness and does not indicate where the cut will be placed in the final analysis. 125

6.8 The various energies of the decay products and the associated branching ratios for the ^{13}C$(\alpha, n)^{16}$O interaction (with 5.3 MeV α particles). 130

6.9 The various energies of the decay products and the associated branching ratios for the ^{18}O$(\alpha, n)^{21}$Ne interaction (with 5.3 MeV α particles). 130

6.10 The neutron yields calculated using Equation 6.12 for the 5.3 MeV ^{210}Po α for acrylic (PMMA) and water. 131

6.11 The expected background from (α, n) interactions on the acrylic vessel and in the internal water in the 190.3 days of livetime. These values are provided prior to the cut on the likelihood ratio. The uncertainties on the interaction rates are detailed in (4) and the uncertainty on the counts passing the cuts include an additional 100% uncertainty on the excited-state cross-sections. The efficiency scale factors presented in Section 6.6 are not included in these numbers. 132

6.12 The number of (α, n) events that pass the cuts from Table 6.1 with the adjusted FV around the AV. The 100% uncertainty on the excited state cross-section is included in the expected counts/ROI/livetime. These numbers include the cut on $\Delta\mathcal{L}_{(\alpha,n)}$. Notably, the contribution from the contributions from the (α, n) interaction on oxygen are only about 5% of the total expected signal. In the sum the uncertainties on the inner and outer acrylic from each nuclide are taken as 100% correlated because they are dominated by the cross-section uncertainty. 134

6.13 The (α, n) signal and backgrounds for a Δt window of 0 to 500 μs, a prompt FV of 5700 to 6500 mm, and a delayed FV of 5500 to 6600 mm. The expected (α, n) events comes from Table 6.12. The accidental background is estimated using the mixed dataset in Figure 6.30. The fraction of nominal, F, is defined in equation 6.14. The factor F is used to scale the (α, n) MC as a background to the reactor antineutrino search. The uncertainty in F is dominated by the statistics in the observation. 138

6.14 The expected number of atmospheric interactions (in interactions/year) at solar minimum for each flavor and detector volume after applying oscillations. The atmospheric neutrinos are generated with an energy range of 0.1 to 10 GeV. Note the outer AV extends from the AV to a radius of 8500 mm. 142

6.15 About ten thousand of each neutrino flavor was simulated using the `GENIE` generator and propagated through `RAT`. The expected interactions per year are for a volume of 8500 mm, from Table 6.14. These numbers are given before the likelihood ratio cut is applied. The efficiency scale factors presented in Section 6.6 are not included in these numbers. The error bars are from cross-section and flux uncertainties, and are assumed to be 100% correlated between the various components. 143

6.16 Various cut variables and values for prompt and delayed event for the atmospheric sideband search. 148

6.17 The expected number of multiple coincidence events selected using the cuts outlined in Table 6.16. 149

6.18 The cosmogenically activated isotopes that undergo $\beta^- + n$ decays with endpoint kinetic energies above 1 MeV. Some of the decay energies are not precisely known. All isotopes other than ^{17}N are expected to contribute less than 10^{-10} events per year after the 20 second follower cut, and are included in this table as zero. 150

6.19 A breakdown of the backgrounds and signal expectation after an optimized cut of $\Delta\mathcal{L}$ < -12.35 and $\Delta\mathcal{L}_{atm}$ < -0.05. The second column gives the total number of counts per livetime for the signal and each background component, after applying the correction factors summarized in Table 6.6 (some of which are specific to the reactor neutrinos), propagating the uncertainties, and prior to applying the likelihood ratio cuts. All (α, n) backgrounds are scaled by 3.32 ± 1.12. The efficiency reduction from the likelihood cuts is given in the third column. The expected number of signal and background events after the likelihood ratio cuts is given in the fourth column. 153

6.20 Information about the three events in the 25% unblinded dataset (between runs 200004 and 202030) that pass a likelihood ratio cut of $\Delta\mathcal{L} < -10$. The first event is cut by both the optimized likelihood ratio cut, $\Delta\mathcal{L} < -12.35$, and the atmospheric likelihood ratio cut, $\Delta\mathcal{L}_{\mathrm{atm}} < -0.05$. The second event passes all cuts. The third event is cut by the optimized atmospheric likelihood cut. The values of the observables are given for the prompt event unless otherwise specified. 155

7.1 Voltage divider ratios between each stage. For example the ratio for the drop from the cathode to the first dynode is 11.5 times larger than the voltage drop between the first dynode and the grid. The total resistance across the entire base is 16.8 MΩ. 166

7.2 The SPE performance for the ZC2722 R5912-MOD PMT across gains spanning roughly 0.45×10^7 to 1.7×10^7. The coincidence rate changes by about 9% across the gain change and the TTS is constant within uncertainties. 170

7.3 The table of coincidence rates, TTS, late ratios, and peak to valleys of the three R5912-MODs tested as well as an R1408 and R5912 HQE for comparison. 172

7.4 Summary of the SPE results for the D784KFLB ETEL PMTs. The operating voltage for all PMTs was adjusted for a gain of 1×10^7, corresponding to a SPE peak of 1.6pC. This table includes the results for 19 PMTs. 174
7.5 The normalizations, N_i, from the fit to the LAB+PPO time profiles given in Equation 7.4. 186
7.6 The time constants, τ_i, from the fit to the LAB+PPO time profiles given in Equation 7.4. 186

8.1 List of some isotopes known to undergo double beta decay. Typically isotopes with higher Q-value and longer $2\nu\beta\beta$ half-life are desirable as they allow detectors to achieve lower backgrounds. Additionally, higher isotopic abundance means larger isotope mass can be easily be deployed without requiring expensive enrichment. Information for this table was partially taken from (5). 197
8.2 Backgrounds for 0.5% Te loaded SNO+ detector for 5 years of data taking with an ROI of 2.42 to 2.56 MeV and a FV of 3.3 m. 202
8.3 The details for the filters used for the dichroicon shown in Figure 8.11. The cut-on wavelength is given for an average incidence angle of 45°. 212
8.4 The details for the filters used for the dichroicon-2. The cut-on wavelength is given for an average incidence angle of 45°. 212
8.5 The central wavelength, FWHM, and transmission at the central wavelength for each of the bandpass filters. The tolerances on the central wavelengths of the filters are less than or equal to 3 nm. 224
8.6 The fit results for the data with each of the bandpass filters. The fit parameters are defined in Equation 8.8. P is the purity of the Cherenkov light selected in a prompt window and is defined in Equation 8.9. 229

8.7 The relative efficiencies of the R7600-U20 and R1408 PMTs. The $\mathrm{R}_{\mathbf{CEF}}$ measures the efficiency ratio between the R2257 PMT and the other two PMTs, after factoring out the expected difference in the quantum efficiencies. This effectively provides the value for the collection and front-end efficiencies of the other two PMTs. This factor is used when comparing expected numbers of detected photons for each PMT in the dichroicon measurements. 236

8.8 The Cherenkov source results for the R2257 and R7600-U20 aperture PMTs, with different configurations of the filters. C_{NORM} is defined in Equation 8.11. The errors are statistical only. The ✓indicates whether a given part of the setup was used. The first column corresponds to the central longpass dichroic filter, the second column to the absorbing longpass filter behind the dichroic filter, and the third column to the barrel of the dichroicon equipped with the shortpass dichroic filters. The results from data and the Chroma simulation are shown in the next two columns. To account for unmodeled inefficiencies in the simulation, the results for each PMT are scaled such that the case with no filters has the same C_{NORM} as data. 246

8.9 Results for the LAB+PPO source with the R2257 and R7600-U20 central PMTs. C^{*}_{NORM} is total Cherenkov light, normalized to the number of triggers and photocathode area, defined in Equation 8.13. P is the purity of the Cherenkov light selection in a prompt window, defined in Equation 8.12. R is $C^{*}_{\mathrm{NORM}} \times P$. The errors come from the uncertainties on the fit parameters. The ✓indicates whether a given part of the setup was used. The first column corresponds to the central longpass dichroic filter, the second column to the absorbing longpass filter behind the dichroic filter, and the third column to the barrel of the dichroicon equipped with the shortpass filters. 252

8.10 Results for the `Chroma` simulations of the LAB+PPO source with the R2257 and R7600-U20 central PMTs. C^*_{NORM} is total Cherenkov light, normalized to the number of triggers and photocathode area, defined in Equation 8.13. P is the purity of the Cherenkov light selection in a prompt window, defined in Equation 8.12. R is $C^*_{\text{NORM}} \times P$. The errors come from the uncertainties on the fit parameters. The ✓indicates whether a given part of the setup was simulated. The first column corresponds to the central longpass dichroic filter, the second column to the absorbing longpass filter behind the dichroic filter, and the third column to the barrel of the dichroicon equipped with the shortpass filters. To account for unmodeled inefficiencies in the simulation, the results for each PMT are scaled such that the case with only the longpass dichroic filter has the same C^*_{NORM} as data. 253

8.11 A summary of the results for the off-axis measurements with the R2257 PMT at the aperture of the dichroicon. The value of C^*_{NORM} for the central data is the same as given in Table 8.9. The C^*_{NORM} and Q_{TOT} values for the first off-axis measurement at y of 75 mm are consistent with a 50% reduction in the collection of the scintillation and Cherenkov light. At the more extreme incident angles the Cherenkov collection efficiency is reduced by a larger factor, as the source is outside the geometric field of view of the Winston cone. The Q_{TOT} values are given relative to the central data, which is normalized to one. 260

8.12 Comparison between LAB+PPO and LAB+PTP and the two different dichroicons for the R7600-U20 aperture PMT. For both scintillator cocktails, using the dichroicon-2 with shorter wavelength pass filter increased R. 262

E.1 All reactor complexes and the distance from the complex to the SNO+ detector, ordered by distance from the detector. 292

F.1 The full run list of all AmBe deployed source runs used. The source positions are determined using the rope tensions by the source manipulator system. 294

List of Figures

2.1 The normal and inverted neutrino mass hierarchies. The different colors indicated the probabilities of finding the α neutrino flavor in the ith neutrino mass eigenstate. This procedure is outlined in (6). The vertical scale in each band varies the value of δ_{CP} between 0 and 2π. This figure is from (7). 16

2.2 The tree-level Feynman diagrams for elastic scattering of a ν_e on an e^-. 18

2.3 The tree-level Feynman diagrams for elastic scattering of a $\bar{\nu}_e$ on an e^-. 18

2.4 The tree-level Feynman diagram for inverse beta decay interactions. 19

2.5 The upper panel shows the IBD cross-section for $\mathcal{O}(1/M)$ in the solid line and for the $\mathcal{O}(1)$ result in the short-dashed line. The long-dashed line shows results from (8). The bottom panel shows outgoing angle of the positron, $\langle\cos(\theta)\rangle$, for the same reaction. This Figure is from (9). 21

2.6 The total cross-section for neutrino energies above 100 MeV for the QE, RES, and DIS processes. This figure is from T. Leitner's thesis (10). 22

3.1 The *pp* chain of stellar thermonuclear reactions, reproduced from (11). 24

3.2 The *CNO* cycle of stellar thermonuclear reactions, reproduced from (11). 25

3.3 The neutrino energy spectra for the various components of the *pp* chain and CNO cycle along with the SSM uncertainties. This figure is from (12). 26

3.4 The flux of ^{8}B solar neutrinos detected as μ or τ flavor plotted the flux detected as e type. The diagonal dashed lines show the ^{8}B flux predicted by the SSM. The ES, CC and NC results are shown in the green, red, and blue bands respectively. The combined flux results are consistent with neutrino flavor transformation. This figure is reproduced from (13). 27

3.5 The survival probability for $\bar{\nu}_e$ as a function of distance from the reactor to the detector. The probability is averaged over the energy spectrum of the $\bar{\nu}_e$ emitted from the reactors, weighted by the cross-section of the IBD reaction. The values of θ_{13} and θ_{12} determine the amplitudes of the oscillations and the corresponding peaks in the survival probability are indicated. This figure is from (14). 31

3.6 The reactor antineutrino survival probability as a function of $\mathrm{L}_0/\mathrm{E}_{\bar{\nu}}$, from (15), where L_0 is the flux-average baseline of 180 km. The primary backgrounds have been removed from the data points. The blue curve shows the expectation based on the neutrino oscillation parameters determined by KamLAND. 32

3.7 (Left) Allowed parameter regions at 1, 2, and 3 σ for the GS98 (full regions and best fit marked by black star) and AGSS09 solar models (dashed, void contours with best fit marked by white dot) and for the KamLAND data (green contours and best fit indicated by green star). (Right) The $\Delta\chi^2$ dependence on Δm^2_{21} for the solar and KamAND analyses. These analyses are performed with a fixed $\sin^2(\theta_{13}) = 0.0224$. This plot and caption is reproduced from (1). 34

3.8 The L/E (km/MeV) spectrum for reactor neutrino oscillations assuming a distance of 60 km. The flux is integrated over the energy spectrum of the reactor antineutrinos and weighted by the IBD cross-section. This figure is motivated by the JUNO detector, presented in (16). 36

3.9 The energy spectrum of emitted antineutrinos in the ^{40}K, ^{238}U, ^{235}U, and ^{232}Th chains. This figure is from (17). 38

4.1 The SNO+ detector, showing the AV, PSUP and AV neck located in the cavity. Above the cavity is the deck, where the electronics and calibration systems are located. 41

4.2 The SNO+ rope-net system showing both the hold-up ropes from SNO and the hold-down ropes installed for SNO+. 42

4.3 Normalized differential cross section as a function of the angle of the outgoing electron, relative to the incoming neutrino direction. The ES interaction is strongly peaked in the direction pointing away from the incoming neutrino direction. This plot is from S. Seibert's thesis (18). 46

4.4 A schematic of the R1408 including the waterproof housing (left). The hexcell design to hold the PMTs and concentrators in the PSUP (right). 49

4.5 A SNO concentrators imaged in 2019 with considerable aging to the petals clearly visible. 50

4.6 Schematic of the SNO+ trigger system. Board upgrades or additions from SNO are highlighted in red. 56

4.7 The family of new electronics for the SNO+ detector. Not shown here is the Latch board developed primarily by T. LaTorre. 58

4.8 A comparison between the ^{16}N data and MC for a central ^{16}N deployment after calibrating the detection efficiency of the detector. This figure is from (19). 63

4.9 The CAEN baseline in ADC counts for a run with a lot of dropout. The Gaussian components of the fit correspond to one or more channels dropped out, shifting the CAEN baseline down. The fit shows the average amount of dropout at a given time is around two to three channels. Note that the CAEN baseline has been set to around 4050 counts with no dropout. This plot is from E. Marzec's thesis (20). 66

4.10 The analytic model in `RAT` compared to test-stand data for a 1 nhit N100 pulse. The noise on the N100 is largely caused by the environment on the teststand and is not representative of the noise in the detector. 68

4.11 The nhit monitor (data) and `RAT` (Monte Carlo) trigger efficiency curves for runs 114287 (left) and 200005 (right). The dashed line shows the underlying $\tilde{n}_{100}$ distribution for the simulated 2.2 MeV γs. The events are simulated out to the PSUP at 8.5 m. 69

4.12 The nhit monitor (data) and `RAT` (Monte Carlo) trigger efficiency curves for runs 200005. The left plot shows event simulated in the center and the right plot shows events simulated out to 5900 mm. The dashed line shows the underlying $\tilde{n}_{100}$ distribution for the simulated 2.2 MeV γs. 69

4.13 The trigger efficiency curve for three different TELLIE data sets. Run 207092 (black) used node 19 pulsed at 50 Hz. Run 207095 (red) used node 19 pulsed at 1 kHz. Run 207097 (green) used node 30 pulsed at 50 Hz. The different nodes are used to test the trigger efficiency for different parts of the detector. The different rates are used to test whether the rate impacts the determination of the trigger efficiency. The most recent nhit monitor, before these runs, is shown in blue. It is clear the nhit monitor over predicts the total trigger efficiency relative to the TELLIE results. 71

5.1 The predicted energy spectra for the antineutrinos emitted from the sum of all β-branches of the fission products ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu. The spectra are shown above the IBD threshold (1.8 MeV) shown in Figure 2.5, and are used in the `RAT` antineutrino generator. 74

5.2 The PMT hit time residuals for central ^{16}N calibration source data, calculated using the known calibration source position. 78

5.3 The PDF used in RAT when evaluating the likelihood of a hypothesized direction. The peak at 41° corresponds to the Cherenkov angle in water. 79

5.4 The total number of produced Cherenkov photons for electrons with energies from 0.7 to 9 MeV. 80

6.1 Map showing the SNO+ detector location and the location of six nearby reactor complexes in Canada and the US. The sizes of the nuclear stations are roughly scaled to the power of the complex. 86

6.2 The survival probability, $P_{ee}(E)$ given by Equation 3.8, for $\bar{\nu}_e$s with a distance of $L = 240$ km. 87

6.3 The energy spectrum of emitted antineutrinos, from Equation 5.1, is shown in black. The cross-section, given in Equation 2.51 and shown in Figure 2.5, is shown in red. The convolution of those curves produces the energy spectrum of detected antineutrinos shown in blue. The distributions are shown with arbitrary units and prior to oscillations being applied. This figure is reproduced from (3). 89

6.4 The total number of interactions per year in the AV as a function of neutrino energy. The black distributions are prior to (solid) and after (dashed) applying oscillations. The red curves show the same, but only for the CANDU reactor complexes. 90

6.5 Schematic representation of the IBD interaction in the SNO+ detector. The outgoing positron creates Cherenkov light before annihilating into two 511 keV γs. The neutron thermalizes and captures on hydrogen with a time-constant of about 200 μs. The neutron capture results in the emission of a 2.2 MeV de-excitation γ. 92

6.6 The prompt event (left) and delayed event (right) energy spectrum for data (black) and MC (red). 102

6.7 The prompt event (left) and delayed event (right) β_{14} spectrum for data (black) and MC (red). 102

6.8 The data (left) and MC (right) $(\mathrm{R}/\mathrm{R}_{AV})^3$ vs $u \cdot r$ for the prompt event. 103

6.9 The data (left) and MC (right) $(\mathrm{R}/\mathrm{R}_{AV})^3$ vs $u \cdot r$ for the delayed event. 103

6.10 The data (left) and MC (right) Δt vs Δr. 103

6.11 The $\Delta\mathcal{L}$ distribution for the simulated reactor IBD signal. 105

6.12 Δt distribution between the prompt and delayed AmBe events. The data is fit, excluding the first bin because of the retrigger data-cleaning cut, out to 1000 μs using Equation 6.6. 107

6.13 The left panels show the data (black) and MC (red) comparison for the prompt event in an AmBe central run. The right panels show the data (black) and MC (red) comparison for the delayed event in an AmBe central run. There are notable differences between the data and MC in the Δr and ITR distributions. 108

6.14 The left panels show the data (black) and MC (red) comparison for the prompt event in an AmBe central run. The right panels show the data (black) and MC (red) comparison for the delayed event in an AmBe central run. There are notable differences between the data and MC in the x, y, and z position distributions, which lead to the difference in the Δr distribution identified in Figure 6.13. 109

6.15 The $\Delta\mathcal{L}$ distribution for the AmBe calibration data and simulation. The right panel shows the distribution after arbitrarily broadening the reconstructed position distributions for the prompt and delayed events by 220 mm and 310 mm respectively. 110
6.16 The neutron detection efficiency for data (left) and MC (right) as a function of the prompt event position. The z-axis (color scale) shows the efficiency and the MC is generally predicts a higher efficiency than identified in the data. 112
6.17 The trigger efficiency comparison between the nhit monitor and the RAT trigger model for run 109133 (left) and run 200005 (right). The in-time hit ($\tilde{n}_{100}$) distribution is shown in the dashed line. The error bars are stat only. 113
6.18 The relative efficiency (data/MC) for detecting the delayed neutron capture event using a cut at $\Delta\mathcal{L} < -12$. The correction factors from Table 6.3 are applied. The lack of data at high z and large ρ is due to limited source deployements in that area, as can be seen in Appendix F. 115
6.19 The detection efficiency predicted by the MC after the cuts described in Table 6.1, including the $\Delta\mathcal{L}$ cut $<$ -12, which can be directly compared to Table 6.5. The runs are combined in three-day chunks to average over the statistical variations in the number of events simulated. The variations in the detection efficiency are well described by statistical fluctuations. 118
6.20 The Δt between all events that pass the cuts in 6.1. The event count is dominated by accidental backgrounds prior to the likelihood ratio cut, which results in the identified flat Δt distribution. 120
6.21 The prompt event rate by run-number. 121
6.22 The delayed event rates by run-number. 122

6.23 The observed number of counts in three-day chunks of livetime divided by the expected counts, based on the average prompt and delayed event rates in the three-day period. The fitted value of the flat constant is consistent with 1.0, indicating that the calculation of the expectation is predictive of the observation. 123

6.24 Likelihood ratio in the sideband region compared to the signal region after shifting the sideband distribution by $-S$. 124

6.25 The comparison between the signal MC and the data in the Δt sideband. Excellent separation between the signal and background is apparent. Neither of these distributions are shifted by S. There are 1991 events in the accidental (data) histogram, which is normalized to one to compare against the simulated signal events. 126

6.26 The data in the sideband compared directly to the 'mixed' dataset, generated by randomly choosing prompt and delayed candidates, pairing them, and assigning a Δt between 500 and 1000 μs. The mixed dataset matches the shape of the data nicely. 127

6.27 The reactor IBD signal MC compared directly to the mixed data. This plot shows qualitatively where the $\Delta\mathcal{L}$ cut will go in order to remove the accidental backgrounds. The sideband shift is not applied to either distribution. 128

6.28 The $\Delta\mathcal{L}$ distributions for reactor IBD events, ^{18}O$(\alpha, n)^{21}$Ne interactions in the internal water, and ^{13}C$(\alpha, n)^{16}$O interactions on the inner AV. $\Delta\mathcal{L}$ is calculated according to the procedure outlined in 6.5.1 using PDFs for the signal MC and for accidental backgrounds. 133

6.29 $\Delta\mathcal{L}_{(\alpha,n)}$ for the mixed dataset and for the $(\alpha, n)^{13}$C along the inner AV. 135

6.30 The $\Delta\mathcal{L}_{(\alpha,n)}$ for the data in the sideband (left) and signal (right) region compared to the mixed dataset. 136

6.31 The Δt distribution for the events that pass the (α, n) cuts. The fitted time constant is consistent with the neutron capture time on hydrogen. 137

6.32 The 15 observed events in the (α, n) fiducial volume with Δt between 0 and 500 μs. The $(\alpha, n)^{13}$C simulations along the inner and outer AV is shown in red, scaled arbitrarily so that the shape can be compared to the data. From top left to bottom right is the prompt energy, prompt $u \cdot r$, Δr, and the prompt $(R/R_{AV})^3$ position. 139

6.33 The Bartol-04 fluxes (21) averaged over zenith angle. This figure is reproduced from (22). 141

6.34 The $\Delta\mathcal{L}$ distribution for reactor IBD and atmospheric neutrino events. The atmospheric events are more 'signal-like' when using the accidental PDFs, so this parameter is not useful for rejecting the atmospheric neutrino background. 144

6.35 The signal (black) and atmospheric (red) Δt and Δr distributions. 145

6.36 The signal (black) and atmospheric (red) prompt energy and β_{14} distributions. 145

6.37 The atmospheric likelihood ratio, $\Delta\mathcal{L}_{\text{atm}}$, calculated using the Δr, prompt energy, and prompt β_{14} PDFs shown in Figures 6.35 and 6.36. 146

6.38 The neutron multiplicity for NC atmospheric interactions. The first bin corresponds to 0 generated neutrons. Note that this figure shows the final state particles in the neutrino interaction, as simulated by `GENIE`. More neutrons can be created if neutrons are knocked out of oxygen nuclei by the high energy neutrons and protons. 147

6.39 The cosmic muon flux as a function of depth, in units of meters water equivalent. SNOLAB is the second deepest underground laboratory; the Jinping underground laboratory, not shown on this plot, is the deepest. 151

6.40 The likelihood ratio distribution for events in the signal window for the 25% of the dataset that was unblinded, in linear (left) and log (right) scale. The three events with likelhood ratio values less than -10 are detailed in Table 6.20. 155

6.41 The likelihood ratio vs. the atmospheric likelihood ratio for events in the signal window for the 25% of the dataset that was unblinded. 156

7.1 Transit time spread of a SNO PMT (removed from the detector), measured relative to a fast trigger PMT using a source of Cherenkov light. The σ of the prompt peak shows the expected TTS of approximately 1.5 ns. The operating voltage of 2000V is typical for the R1408 PMTs and the dark rate of about 1 kHz is typical at room temperature. The coincidence rate is kept below 5% in these measurements to ensure the events are primarily SPE. The late ratio is the fraction of PMT hits outside of the prompt peak. 159

7.2 The R5912-MOD specifications provided by Hamamatsu photonics (23). 165

7.3 The charge distribution of the R5912-MOD PMT. Shown in the statistics box is some important characteristics of the SPE charge response. The Guassian fit to the SPE peak is shown in red. The χ^2/NDF is 70.68/59. 167

7.4 The transit time profile of the R5912-MOD. Shown in the statistics box is some important characteristic of the SPE time response. The Gaussian fit to the prompt light peak is shown in red. 168

7.5 The transit timing profile of the R5912-MOD broken down into the various components that make up the structure. 169

7.6 An example of an R5912-MOD PMT waveform with a triple lognormal fit shown in red. 171

7.7 A 2D histogram showing the time between the prompt light and the afterpulse plotted against the charge of the afterpulse. 173

7.8 The base used for the R5912-200 HQE PMTs (left). An R5912-200 HQE after custom potting, before going into the SNO+ detector (right). 175

7.9 The R5912-200 HQE detection efficiency curve compared to two SNO PMTs. Relative scaling between PMT efficiencies is from bench-top data at Penn. 176

7.10 The R5912-200 HQE SPE charge distribution compared to two SNO R1408 PMTs. 177
7.11 The R5912-200 HQE SPE transit time distribution. 178
7.12 The scintillator emission timing setup, which includes an R7600-U200 PMT optically coupled to the scintillation source, which is used as a fast trigger. In this picture, a red ^{210}Po disk source is deployed above a sample of liquid scintillator, held inside a UVT acrylic block. A second R7600-U200 PMT located about 30 cm away from the scintillation source. 180
7.13 The `RAT` model of the scintillation emission timing setup (left), corresponding to Figure 7.12, and the light yield setup (right). The blue lines show the paths of various optical photons. The 1-inch Hamamatsu trigger and measurement PMTs are shown for the timing setup and the R1408 PMT is shown in the light yield setup. 181
7.14 A R7600-U200 SPE waveform fit to a triple lognormal distribution (Equation 7.2). The corresponding best fit parameters are shown. 182
7.15 Simulated digitized waveforms for a R7600-U200 and R1408 PMTs. The R7600-U200 waveform is compared agianst a trigger threshold. For triggered events, the simulated waveforms for up to four channels are written to disk in .hdf5 format. 183
7.16 The β and α timing profiles compared for LAB+PPO (left) and Te-loaded LAB+PPO (right). For both scintillators the α particles causes a time profile with a longer tail, which can be used to discriminate between β and α particle excitation. The Te-loaded LAB+PPO has a quenched triplet state excitation relative to LAB+PPO, which reduces the number of photons emitted at larger times. 184
7.17 Example fits for LAB+PPO for β (left) and α (right) particles. The fit equation used is shown in Equation 7.4. 186

7.18 The f_{prompt} distribution for LAB+PPO and Te-loaded LAB+PPO, which show excellent discrimination between β and α particles. 187

7.19 A comparison between the data and `RAT` simulation for LAB+PPO with the ^{210}Po α source. 188

7.20 The LAB+PPO+bisMSB emission timing for data and simulation using a 1.4 ns re-emission constant in `RAT`. The fast-time component of the emission time is slightly slowed down due to additional absorption and re-emission, and this effect is well modeled in `RAT`. 189

7.21 The Cherenkov source data for the R1408 PMT used for calibration prior to the light yield measurements. The overall efficiency of the PMT in `RAT` is scaled by a factor called the collection efficiency. The parameters are tuned in `RAT` until the Cherenkov source data and simulation agree, and is held fixed in the light yield simulations. For this measurement, a collection efficiency of 1.3 is selected based on the agreement with the data. 191

7.22 The R1408 PMT in data and simulation for the light yield setup. The scintillation efficiency, S, is tuned to 11900 photons/MeV. 192

7.23 The 30 L liquid scintillator setup with four PMTs. This setup is used as a large-scale test of the `RAT` scintillator model. In the `RAT` model the blue lines show optical photons from a ^{90}Sr β decay at the center of the 30 L volume. 193

7.24 The R1408 PMT data and MC for the 30L setup. The collection efficiency is set to 1.3 and the scintillation efficiency is set to 11900. 193

7.25 The number of expected PMTs that detect light (primarily SPE) per MeV, determined by simulating 1 MeV electrons at the center of the SNO+ detector using the optical model for Te-loaded LAB+PPO, without DDA. Adding the DDA improves the PE per MeV to about 450. 194

8.1 A schematic of the energies of the A = 76 isobars. The single-beta decay (shown in each case by the green arrow) is energetically forbidden between ^{76}Ge and ^{76}As. However, the double beta transition between ^{76}Ge and ^{76}Se (shown by the purple arrow) is allowed and proceeds with a very long half-life given in Table 8.1. This figure is from (5). 196

8.2 The Feynman diagrams for $0\nu\beta\beta$. 198

8.3 The energy spectra for $0\nu\beta\beta$ and $2\nu\beta\beta$ decay. The signature of $0\nu\beta\beta$ is a peak at the end-point of the $2\nu\beta\beta$ decay spectrum, which is shown spread by an arbitrary detector resolution. 199

8.4 The effective Majorana mass as a function of the lightest neutrino mass. The KamLAND-Zen limit is shown in the blue band. The width of the band correspond to different choices of nuclear matrix elements. The green and pink regions show the allow parameter space, which is obtained from scanning over the unknown Majorana phases for the inverted and normal hierarchies. The best limits as of 2016 for each isotope are shown in the inset. This figure is from (24). 201

8.5 The number of background counts per year expected for the SNO+ detector in the energy ROI from 2.42 to 2.56 MeV and a FV of 3.3 m. The backgrounds are dominated by ^{8}B solar neutrino interactions. There are also significant contributions from external γ-rays, internal Thorium chain (dominated by ^{212}BiPo decays), and $2\nu\beta\beta$ leakage into the ROI. 203

8.6 The expected reconstructed energy distribution for the backgrounds and a hypothetical $0\nu\beta\beta$ signal where $m_{\beta\beta} = 100$ meV. The FV used is 3.3 m. An asymmetric ROI of -0.5σ to 1.5σ around the mean of the signal is chosen to avoid the falling $2\nu\beta\beta$ background. 204

8.7 The radial distribution of the external backgrounds (which includes the ^{208}Tl from the AV, external water, ropes, and PMTs) compared to the $0\nu\beta\beta$ signal. The fiducial volume of 3.3 m corresponds to a value of $(R/R_{AV})^3 = 0.17$. The counts are normalized to 5-years of data-taking. 205

8.8 The time profile of the detected light for simulated 2.5 MeV electrons at the center of a SNO+-like detector, consisting of about 9000 PMTs with transit time spreads of 1.4 ns, a 6 meter radius acrylic vessel, and about 50% coverage. The Cherenkov light arrives promptly, but is difficult to identify due to the intrinsic resolution of the photodetectors and high light yield of the scintillator. 208

8.9 The quantum efficiency of the PMTs used in the various measurements compared to the Cherenkov emission spectrum and the emission spectra of the fluors PPO and PTP. The three emission spectra are arbitrarily scaled and show shape only. The quantum efficiency curves for the R7600-U20, R2257, and R1408 are taken from (25, 26, 27) respectively. The PPO and PTP emission spectra were taken from the PhotochemCAD database (28). 210

8.10 Simple schematics of two potential options for a system to detect light sorted by the dichroicon. The schematic on the left shows an option where a parabolic reflector is built around the dichroicon to detect the short-wavelength light. The schematic on the right an option using acrylic light guides to direct the short-wavelength light to one or more photodetectors. In this case, a pixelated light detector such as a large area picosecond photodetector (LAPPD) might be an ideal sensor for the dichroicon. In both designs the long-wavelength light is detected at the aperture of the dichroicon. Neither of these full designs are constructed, but are included to show potential detection schemes for the short- and long- wavelength light. The blue and red lines show possible photon tracks for short- and long-wavelength light respectively. 211

8.11 A head-on view of the dichroicon. The filters are held in a custom 3D printed plastic holder and can be easily swapped. The shortpass filters tile the barrel of the Winston cone and a central longpass filter is placed at the aperture. A small amount of black electrical tape is used to block a small gap between the filters and the holder at the top of the dichroicon. The outer diameter of the dichroicon is about 150 mm and the inner radius, where the long-pass filter is located, is about 50 mm. 213

8.12 The dichroic filter characterization setup. The normalization, transmission, and reflection PMTs are R7600-U200 PMTs. The dichroic filter is held on a rotating stage. 217

8.13 The transmission data for the 500 nm shortpass filter for incidence angles between 0 to 60°. The percent transmission was calculated using Equation 8.6. 220

8.14 The reflection data for the 500 nm shortpass filter for incidence angles between 15 to 60°. Note that the reflected data only extends to incident angles of 15° due to shadowing effects in the setup. The percent reflection was calculated using Equation 8.7. 221

8.15 The transmission through the 480 nm longpass filter, used at the aperture of the dichroicon, as a function of wavelength, for multiple incidence angles. 222

8.16 The transmission through the 480 nm longpass filter, used at the aperture of the dichroicon, as a function of wavelength, for 0° and 45° incidence angles, in air and in water. 223

8.17 A schematic of the experimental setup, showing the β source deployed above the LAB+PPO target and the locations of the mask, PMTs, and bandpass filter. The colored lines indicate example optical photon paths before and after the filter. The aperture is 1 cm in diameter. 224

8.18 The time emission profile, zoomed into the rise time of the LAB+PPO for different wavelength regimes. The light is selected using the bandpass filters listed in Table 8.5. The central wavelength and width of the bandpass filters are specified in the legends. The histograms are normalized to the peak of the scintillation light and shown in arbitrary units (A.U.). The Cherenkov light can be clearly identified at early times in the data for the filters longer than 450 nm. These histograms, as well as additional ones, are shown separately in Figure 8.19. 226

8.19 The time emission profile for each of the bandpass filters specified in Table 8.5. Each of the data sets are normalized to the peak of the scintillation light. The Cherenkov light becomes clearly separated above 450 nm. 227

8.20 The full fit for the 355 nm bandpass filter data (left) and 494 nm bandpass filter data (right) is shown in red. The Cherenkov and scintillation components of the fit are shown explicitly. Note the arbitrary Δt offset from 0 ns, mostly due to cable delays, was not removed for this plot, as it was for Figures 8.18 and 8.19. 228

8.21 The measured emission spectrum compared to the amount of scintillation light detected for each bandpass filter. The amount of scintillation light is scaled based on the total number of triggered events, the efficiency of the PMT, and the transmission of the filter. The shape extracted from the bandpass filter measurement agrees well with the emission spectrum. The noise above 550 nm and below 340 nm is from the dark noise of the spectrometer. 230

8.22 A schematic of the experimental setup with the dichroic filter. The setup is the similar to Figure 8.17, with the replacement of the bandpass filter with a dichroic filter and the addition of the reflection PMT. 231

8.23 The fit results for the longpass 506 nm dichroic filter using LAB+PPO. The transmitted light (left) shows clear Cherenkov and scintillation separation, while the reflected light (right) shows the LAB+PPO emission time-profile. 232

8.24 The fit results for the shortpass 500 nm dichroic filter using LAB+PPO. The transmitted light (left) shows the LAB+PPO emission time-profile, while the reflected light (right) shows modest Cherenkov and scintillation separation. 232

8.25 The data using the R5912-MOD PMT to detect the transmitted light through the longpass dichroic filter. The fit result shows a clear Cherenkov peak at early times. 234

8.26 A schematic showing the setup with the dichroicon and reflective cylinder. The R7600-U200 PMT is optically coupled to the acrylic or scintillator source and used as a fast trigger. The long-wavelength light is detected at the aperture of the dichroicon. The short-wavelength light is transmitted through the dichroicon, reflected off of the Mylar lining the cylinder, and detected by an R1408 PMT. The setup with the R2257 aperture PMT is identical to the one shown, except due to the length of the aperture PMT, the reflective cylinder is extended 150 mm. The back of the aperture PMT is covered in reflective foil. The area of the R1408 PMT outside of the reflective cylinder is masked off using felt. 238

8.27 A side view of the dark-box setup with the Cherenkov source. The dichroicon is shown with the R2257 PMT at the aperture. In front of the R2257 PMT is a 480 nm longpass dichroic filter. The barrel of the dichroicon consists of shortpass dichroic filters. The reflecting cylinder and R1408 PMT are not shown in this setup. The distances and size of the various important components is shown in Figure 8.26. 239

8.28 The full setup with the short-wavelength light detection system, which consists of a cylinder with a Mylar-lined reflecting interior that ends at an R1408 PMT. The part of the R1408 PMT outside the cylinder is masked off using black felt. 240

8.29 A direct view of the `Chroma` dichroicon model with the R2257 PMT at the center of the dichroicon. The outer diameter of the dichroicon is about 150 mm and the inner radius, where the long-pass filter is located, is about 50 mm. The two different colors in the barrel of the dichroicon indicate the two different types of short-pass filters, as detailed in Table 8.3. This figure is created by B. Land. 242

8.30 The full `Chroma` simulation setup for reproducing the Cherenkov source results. Geometry components with the same optical properties are colored similarly, however the colors are arbitrary. The sizes of the components and distances between the objects are identical to those used for the data, shown in Figure 8.26. This figure is created by B. Land. 243

8.31 The results for the R2257 central aperture PMT and the acrylic Cherenkov source. In black is data for the configuration with no filters or dichroicon. The blue shows the data with the longpass dichroic filter optically coupled to the R2257. The data with the full dichroicon added is shown in red. The corresponding `Chroma` results are shown in the dashed lines. These results are summarized in Table 8.8. The value of Δt is determined by cable delays and transit times through the PMTs and has no impact on the analysis. The results for the R7600-U20 look similar, but with a narrower TTS. 245

8.32 An example fit using Equation 8.8, to LAB-PPO data. The Cherenkov and scintillation components are shown separately, in addition to the total fit in red. The data were taken with the shortpass barrel of the dichroicon and the absorbing longpass filter behind the dichroic longpass filter at the aperture of the dichroicon. The data is normalized to 1.0. The Δt offset from zero is arbitrary and does not impact the fit. 250

8.33 The R2257 and R1408 waveforms for a single triggered event. This event has an early-time PMT pulse, corresponding to Cherenkov light, for the R2257 PMT and an approximately 5 PE pulse at the R1408 PMT. The waveforms are specifically selected by looking at the Δt histogram and selecting an event in the Cherenkov peak. The x-axis is labelled in 0.1 ns samples, and the offset between the two waveforms is showing the additional photon travel time and transit time through the R1408 PMT. 255

8.34 The dichroicon data for both the R7600-U20 PMT at the aperture and the R1408 PMT behind the dichroicon. The light detected by the R1408 is primarily scintillation light and the light detected by the R7600-U20 is primarily Cherenkov light. 256
8.35 The R1408 data for a ^{90}Sr β source and a ^{210}Po α source. The small bumps in the timing spectrum are due to the complicated PMT transit time distribution, which includes two different late-pulsing peaks. The difference between these distributions is used to discriminate between β and α particle excitation in liquid scintillator detectors. 258
8.36 The R7600-U20 data for a ^{90}Sr β source and a ^{210}Po α source. The lack of Cherenkov light for the below-threshold α particles can be clearly identified. 259
8.37 The data for the R7600-U20 at the aperture of the dichroicon using LAB+PPO and LAB+PTP targets, compared for the dichroicon and the dichroicon-2. 262
8.38 A visualization of a 1-kT right cylinder active volume of LAB+PPO instrumented with 13,350 dichroicons produced by `Chroma`. This figure was created by B. Land. 265
8.39 A `Chroma` event display showing a single 100 MeV electron event in LAB+PPO. Dichroicons are colored blue if a short-wavelength hit was detected, or red if a long-wavelength hit was detected. Both a long- and short-wavelength hit results in a magenta dichroicon. Despite all dichroicons detecting many short-wavelength scintillation photons, a clear long-wavelength Cherenkov ring can be seen in magenta. Every dichroicon detects at least one short-wavelength photon so there are no dichroicons colored red in this image. No selection criterion was used when choosing this event and most events appear similar during hand-scanning of the simulated data. This figure was created by B. Land. 266

A.1 (Left) The integral of the ESUM waveform for ^{16}N events as a function of nhit. The red lines show the range used by the data-cleaning cut. (Right) The integral of the ESUM waveform as a function of nhit for a physics run, which includes flashers and shark-fins that are removed by this cut. This figure is from Eric Marzec's thesis (20). 272

B.1 A schematic demonstrating the time-walk effect for two potential PMT pulses. Due to the constant discriminator threshold, the time measured depends on the size of the pulse. 276

C.1 The time-difference between a high charge hit (above a QHS of 50 ADC counts) and the hits in adjacent channels plotted against the QHS of the hit in the adjacent channel. There is clearly a large contribution from cross-talk at QHS around pedestal and Δt $>$ 5 ns. The left plot shows hits after ECA calibration but before PCA and the right plots shows after PCA calibration. The cross-talk cut is applied after ECA but before PCA. 279

C.2 (Left) The deposited charge at the R1408 PMT for an LED pointed directly at the PMT at various intensities. The LED intensity is adjusted in order to understand the afterpulsing rate as a function of number of electrons. The SPE peak for this PMT is around 1 pC, making it easy to roughly convert to number of PE. (Right) The time between the initial pulse on the PMT from the LED light and any afterpulses. There are clear afterpulsing peaks at around one and five μs. The number of afterpulses scales with number of prompt PEs. 281

C.3 The afterpulsing probability as a function of the average number of photoelectrons. The number of photoelectrons detected was tuned using a variable intensity LED, as shown in Figure C.2 (left). A probability over 100% indicates on average more than one afterpulse was detected for each initial pulse. The fit indicates that there is about a 0.9% chance of each individual PE creating an afterpulse. 282

C.4 The afterpulsing distribution for detector data showing the afterpulsing peaks at one and five μs, consistent with the benchtop data shown in C.2. 283

C.5 The PMT data bundle bitmap. 285

D.1 Internal and external water absorption coefficients/lengths (left/right vertical axis), in the black data points (29). 289

D.2 The acrylic vessel attenuation coefficiency/lengths (left/right vertical axis) (29). 289

D.3 The nhits ratio between data and MC for ^{16}N source positions (29). After optical calibrations, the number of PMTs hit in data and MC agree well regardless of event position in the detector. 290

F.1 The reconstructed ρ and z positions of the tagged prompt event, summing events over all AmBe runs shown in Table F.1. The z-axis (color scale) is showing the total number of events in each bin. 295

G.1 The $\nu/\bar{\nu}$ + ^{16}O cross sections used in the `GENIE` and `NEUT` generators, compared to data taken by T2K (30). 297

G.2 A visualization of the simplified SNO+ geometry used in the `GENIE` primary vertex generation. The detector model includes the acrylic vessesl with its neck (blue), and the light water (grey). The facets are an artifact of the visualization. This figure and caption are reproduced from A. Mastbaum's thesis (22). 298

G.3 The energy of the delayed events, compared for data and MC, shows the expected Michel electron energy spectrum. 299
G.4 The Δt between the prompt and delayed events shows the expected muon decay time. 300

Chapter 1

Introduction

Over the past several decades, neutrino detectors, utilizing neutrinos created in the sun, the atmosphere, nuclear reactors, and man-made beams, have discovered neutrino flavor oscillations, provided evidence of neutrino mass, and made precision measurements of many of the parameters that govern neutrino oscillations. The SNO detector cemented the discovery of neutrino flavor oscillation in 2001 (13), a measurement that implies that neutrinos have a non-zero mass. Subsequently, the mixing parameters have been measured by dozens of experiments, filling in a consistent picture of massive neutrino flavor oscillations. However, several critical questions remain unanswered. Do neutrino oscillations obey charge-parity symmetry? What is the ordering of the massive neutrino states? Is the neutrino a Majorana particle? To answer these questions, future experiments such as JUNO, Hyper-Kamiokande, and Theia are planning to build enormous, monolothic Cherenkov- or scintillation-based detectors. These detectors will succeed by building on the technology and experience developed from working on the current generation of large-scale scintillation and water detectors, such as SNO+. SNO+ will not only demonstrate a promising technology, but will also itself attempt to answer some of these important remaining questions.

This dissertation can largely be considered in three distinct parts. Chapters 2 and 3 describe the phenomonology of massive neutrino flavor mixing and the experiments that observe neutrino interactions. Chapters 4 through 6 detail the SNO+ detector, the associated simulation and event reconstruction software, and an analysis designed to search for reactor antineutrinos. Chapters 7 and 8 discuss significant research and development for SNO+ and future detectors, with a focus on improving technology for future large-scale scintillation-based detectors.

Antineutrinos from nuclear power reactors have been a critical tool for studying neutrino oscillations since KamLAND first measured $\bar{\nu}_e$ disappearance (31). Experiments have used this prolific source of antineutrinos to measure mixing parameters (32), (33), (34), (35) and to search for sterile neutrinos (36), (37), and future reactor-based neutrino experiments hope to uncover the mass heirarchy (16). The primary analysis presented in this dissertation is a search for antineutrinos from nuclear reactors using data taken with the SNO+ detector.

The SNO+ detector is impressively versatile and will run in three phases, with three different targets inside of the acrylic vessel: ultra-pure water, liquid scintillator, and tellurium-loaded liquid scintillator. The SNO+ detector began taking data with a water target in May of 2017. During two years of data-taking, finishing in July of 2019, the detector performed measurements of the ^{8}B solar neutrino flux (19), the neutron-proton capture cross-section (38), and put competitive limits on nucleon decay through invisible modes (39). This dissertation presents an analysis of the SNO+ water data developed to search for reactor antineutrinos.

Because of the imporance of neutron detection for reactor antineutrino and diffuse supernova background searches, an experiment utilizing a pure water target, capable of observing neutrons at a high efficiency, is of technological interest. Given the extremely low trigger thresholds and radioactive backgrounds in the SNO+ detector, it is the first pure water Cherenkov detector capable of detecting reactor antineutrinos. Notably, no special trigger is implemented for the data-taking; the standard trigger runs at a low enough threshold to maintain a high efficiency for detecting the 2.2 MeV γ-rays emitted when a neutron captures on hydrogen. This exceptionally low trigger threshold is critical for the analysis I present, and was made possible, in part, by my work on the trigger and front-end electronics. Chapter 6 will focus on the detection of reactor antineutrinos in the SNO+ water phase, which serves as a demonstration to the community that this measurement is possible without specialized technologies.

As of the writing of this thesis, the scintillator fill of the detector is onging. The physics program during scintillator phase includes measurements of low energy solar neutrinos, reactor antineutrinos, and geoneutrinos (40). Critical to the success of SNO+, and future experiments hoping to build on SNO+'s technology, is R&D effort on organic liquid scintillator and photomultiplier tubes (PMTs). In regard to this, Chapter 7 presents bench-top measurements of state-of-the art large-area PMTs and liquid scintillators. The bench-top liquid scintillator measurements, performed for both pure liquid scintillator and tellurium-loaded liquid scintillator, are compared to a detailed simulation to extract parameters for a scintillator model that is used to extrapolate to large-scale detectors. The performance and detailed understanding of the scintillator is even more important for SNO+'s final tellurium-loaded scintillator phase.

The primary goal of the SNO+ experiment is to perform a competitive search for neutrinoless double beta decay in ^{130}Te. The sensitivity is expected to be world-leading, but SNO+ will not cover a significant fraction of the parameter space. To achieve better sensitivity levels, the primary background from solar neutrino interactions must be reduced. Chapter 8 presents the developement of an instrument – the **dichroicon** – which provides this capability. Specifically, the dichroicon allows for the detection of Cherernkov light against the scintillation light background in a high light yield liquid scintillator detector by providing spectral sorting of photons. The dichroicon is an exciting and promising technology, that will hopefully be deployed in large-scale detectors, such as ANNIE or WATCHMAN, in the coming years.

Chapter 2

Physics of Massive Neutrinos

The existence of a neutral, weakly interacting particle called the 'neutron' was first proposed by W. Pauli in order to address the observed continuous energy spectrum in β decays. The particle was subsequently renamed the neutrino in 1932 by E. Fermi after the discovery of the neutron by J. Chadwick. E. Fermi also pointed out the neutrino could be massless in his formulation of a theory for β decays. The neutrino has since been the focus of enormous theoretical and experimental effort. Its elusiveness due to its weakly interacting nature makes its detection and complete understanding difficult. Indeed, when the particle was first proposed, W. Pauli is famously quoted as saying 'I have done something very bad today by proposing a particle that cannot be detected; it is something no theorist should ever do'.

Luckily, W. Pauli turned out to be incorrect – the neutrino was first detected in 1956 by F. Reines and C.L. Cowan using antineutrinos emitted from the Savannah River nuclear reactor, presented first in their paper titled simply 'The neutrino' (41). In 1957, motivated by discovery of oscillations in mesons ($K^0 \leftrightarrow \bar{K}^0$), B. Pontecorvo proposed the first theory of neutrino oscillations in which oscillation occurred between the neutrino and antineutrino (42). By 1969, Pontecorvo has developed a complete theory of two neutrino flavor mixing oscillations (43).

The first model independent observation of neutrino oscillation was performed by the Super-Kamiokande experiment in a measurement of atmospheric ν_μ interactions (44). Definitive evidence of neutrino oscillations was provided by the SNO experiment in 2002 (13). Since SNO, experiments measuring atmospheric, solar, and terrestrial neutrinos at a wide range of neutrino energies and distances from the source have confirmed the three-neutrino mixing model in which the three flavor neutrinos ν_e, ν_μ, and ν_τ are linear combinations of the massive neutrino states ν_1, ν_2, and ν_3. Neutrino mixing opens new questions into the nature of the neutrino, many that have yet to be addressed.

2.1 Neutrino Oscillations in Vacuum

Neutrinos (antineutrinos) in the standard model are included as massless, left-handed (right-handed), neutral, fermions that participate in interactions through the weak force. Neutrino oscillations, which directly imply that neutrinos have mass, are firmly beyond the scope of the current standard model. In the canonical theory of neutrino oscillations, a neutrino is created in a flavor state ϵ and propagates in the mass state j. The flavor state is described as a linear combination of mass states:

$$|\nu_\epsilon\rangle = \sum_j U^*_{\epsilon j} |\nu_j\rangle , \tag{2.1}$$

where U is the 3 x 3 unitary matrix. Using the unitarity of U, we can write the massive states in terms of the flavor states as:

$$|\nu_j\rangle = \sum_\epsilon U_{\epsilon j} |\nu_\epsilon\rangle . \tag{2.2}$$

The massive neutrino states are eigenstates of the free Hamiltonian, $\mathcal{H}$ and evolve according to the Schrödinger equation:

$$i\frac{d}{dt} |\nu_j(t)\rangle = \mathcal{H} |\nu_j(t)\rangle . \tag{2.3}$$

The solution to Equation 2.3 is plane waves, meaning the mass eigenstates evolve as:

$$|\nu_j(t)\rangle = e^{-iE_jt}\,|\nu_j\rangle\,, \tag{2.4}$$

with energy eigenvalues, E_j. Converting to the flavor basis using Equation 2.1 gives:

$$|\nu_\epsilon(t)\rangle = \sum_j U^*_{\epsilon j} e^{-iE_jt}\,|\nu_j\rangle\,, \tag{2.5}$$

where the neutrino is created with flavor ϵ, i.e. $|\nu_\epsilon(0)\rangle = |\nu_\epsilon\rangle$. The amplitude of a flavor transition, e.g. $\epsilon \rightarrow \gamma$ can then be written as:

$$\psi_{\epsilon\gamma}(t) = \langle\nu_\gamma|\nu_\epsilon(t)\rangle = \sum_j U^*_{\epsilon j} U_{\gamma j} e^{-iE_jt}, \tag{2.6}$$

which gives the transition probability:

$$P_{\nu_\epsilon\rightarrow\nu_\gamma}(t) = |\langle\nu_\gamma|\nu_\epsilon(t)\rangle|^2 = \sum_{j,k} U^*_{\epsilon j} U_{\gamma j} U_{\epsilon k} U^*_{\gamma k} e^{-i(E_j-E_k)t}. \tag{2.7}$$

Given that neutrinos are ultra relativistic with very small masses, we can expand the energy relation $E_j = \sqrt{\mathbf{p_j}^2 + m^2}$, giving $E_j \approx E + \frac{m_j^2}{2E}$, or $E_j - E_k = \frac{\Delta m_{jk}^2}{2E}$. Additionally, the neutrinos can be approximated as traveling at the speed of light, which allows a direct conversion from the time to distance, $t = L$. This approximation is referred to as the light-ray approximation, and is actually unjustified in a plane-wave treatment of oscillations. A full quantum treatment of oscillations using a wave packet desciption is given in (11). Rewriting the transition probability gives:

$$P_{\nu_\epsilon\rightarrow\nu_\gamma}(L,E) = |\langle\nu_\gamma|\nu_\epsilon(t)\rangle|^2 = \sum_{j,k} U^*_{\epsilon j} U_{\gamma j} U_{\epsilon k} U^*_{\gamma k} e^{-i\frac{\Delta m_{jk}^2 L}{2E}}. \tag{2.8}$$

It is clear from this equation that the phase of the oscillation is determined by the distance the neutrino has traveled, L, the mass-squared differences between the neutrino mass states, Δm_{jk}^2, and the energy of the neutrino, E. The amplitude of the oscillations is specified entirely by the constant elements in the mixing matrix U. An important consequence of Equation 2.8 is that oscillation experiments cannot probe the absolute masses of the neutrinos, only the differences between the mass states.

Equation 2.8 can be rewritten using the unitary relation and separating the real and imaginary parts of U:

$$\begin{aligned} P_{\nu_\epsilon \rightarrow \nu_\gamma}(L,E) = \delta_{\epsilon\gamma} - 4 \sum_{j>k} \mathrm{Re}\left[\mathrm{U}^*_{\epsilon \mathrm{j}} \mathrm{U}_{\gamma \mathrm{j}} \mathrm{U}_{\epsilon \mathrm{k}} \mathrm{U}^*_{\gamma \mathrm{k}}\right] \sin^2\left(\frac{\Delta \mathrm{m}^2_{\mathrm{jk}} \mathrm{L}}{4\mathrm{E}}\right) + \\ 2 \sum_{j>k} \mathrm{Im}\left[\mathrm{U}^*_{\epsilon \mathrm{j}} \mathrm{U}_{\gamma \mathrm{j}} \mathrm{U}_{\epsilon \mathrm{k}} \mathrm{U}^*_{\gamma \mathrm{k}}\right] \sin^2\left(\frac{\Delta \mathrm{m}^2_{\mathrm{jk}} \mathrm{L}}{2\mathrm{E}}\right). \end{aligned} \tag{2.9}$$

The oscillation probabilities where $\epsilon \rightarrow \gamma$ are typically referred to as *transition probabilities* and transitions within the same state, $\epsilon \rightarrow \epsilon$, are referred to as *survival probabilities.* For the survival probability, the imaginary term disappears and is given by:

$$P_{\nu_\epsilon \rightarrow \nu_\epsilon}(L,E) = 1 - 4 \sum_{j>k} |U_{\epsilon j}|^2 |U_{\epsilon k}|^2 \sin^2\left(\frac{\Delta m^2_{jk} L}{4E}\right). \tag{2.10}$$

For the case of antineutrinos the same derivation of the oscillation probabilities follows, with the exception that the mixing matrix is complex conjugated. In other words, the flavor antineutrinos are related to the mass states by:

$$|\bar{\nu}_\epsilon\rangle = \sum_j U_{\epsilon j} \, |\bar{\nu}_j\rangle \, . \tag{2.11}$$

The transition probability is then:

$$P_{\bar{\nu}_\epsilon \rightarrow \bar{\nu}_\gamma}(L,E) = \sum_{j,k} U_{\epsilon j} U^*_{\gamma j} U^*_{\epsilon k} U_{\gamma k} e^{-i \frac{\Delta m^2_{jk} L}{2E}}, \tag{2.12}$$

which can be rewritten as:

$$\begin{aligned}P_{\bar{\nu}_\epsilon\rightarrow\bar{\nu}_\gamma}(L,E) = \delta_{\epsilon\gamma} - 4\sum_{j>k}\mathrm{Re}\left[\mathrm{U}^*_{\epsilon\mathrm{j}}\mathrm{U}_{\gamma\mathrm{j}}\mathrm{U}_{\epsilon\mathrm{k}}\mathrm{U}^*_{\gamma\mathrm{k}}\right]\sin^2\left(\frac{\Delta\mathrm{m}^2_{\mathrm{jk}}\mathrm{L}}{4\mathrm{E}}\right) - \\ 2\sum_{j>k}\mathrm{Im}\left[\mathrm{U}^*_{\epsilon\mathrm{j}}\mathrm{U}_{\gamma\mathrm{j}}\mathrm{U}_{\epsilon\mathrm{k}}\mathrm{U}^*_{\gamma\mathrm{k}}\right]\sin^2\left(\frac{\Delta\mathrm{m}^2_{\mathrm{jk}}\mathrm{L}}{2\mathrm{E}}\right).\end{aligned} \tag{2.13}$$

Comparing to the transition probability for neutrinos in Equation 2.9, only the sign in front of the imaginary term changes. Because the imaginary term disappears in the survival probability, this implies that $P_{\nu_\epsilon\rightarrow\nu_\epsilon} = P_{\bar{\nu}_\epsilon\rightarrow\bar{\nu}_\epsilon}$. This is expected, as theories of neutrino oscillations are expected to obey CPT symmetry, which transforms the oscillation probability $\nu_\epsilon \rightarrow \nu_\gamma$ to $\bar{\nu}_\gamma \rightarrow \bar{\nu}_\epsilon$.

As a useful exercise, I consider the case of two neutrino mixing, which simplifies the formulas for the transition probabilities and is often a good approximation for many experiments. In general, a 2 x 2 unitary matrix U is described by four parameters: three phases and one mixing angle. However, the phases are unphysical because the neutrino states can be redefined with a phase transformation that cancels out the phase in U. Thus, the mixing matrix U can be written as:

$$U = \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix}. \tag{2.14}$$

In this case, it is straightforward to calculate the transition probability:

$$P_{\nu_\epsilon\rightarrow\nu_\gamma}(L,E) = \sin^2(2\theta)\sin^2\left(\frac{\Delta m^2 L}{2E}\right). \tag{2.15}$$

The oscillation length, L_{osc}, is defined by the distance the neutrino travels for the phase to become 2π. In the two neutrino mixing scenario we have $L_{osc} = \frac{4\pi E}{\Delta m^2}$. It is convenient for oscillation experiments to write Equation 2.15 with different units:

$$P_{\nu_\epsilon\rightarrow\nu_\gamma}(L,E) = \sin^2(2\theta)\sin^2\left(1.27\frac{\Delta m^2[\mathrm{eV}^2]\mathrm{L}[\mathrm{m}]}{2E[\mathrm{MeV}]}\right), \tag{2.16}$$

which gives an $L_{osc} = 2.47\frac{E[\mathrm{MeV}]}{\Delta m^2[\mathrm{eV}^2]}$ m. In this two neutrino framework it is straightforward to see that the amplitude of the oscillations is set by the value of the mixing angle θ, while the frequency of the oscillations is proportional to the mass difference between the neutrino states.

The 3 x 3 unitary mixing matrix U which connects the flavor to mass eigenstates is known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix and is given by:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{2.17}$$

The PMNS matrix is commonly parameterized in terms of three mixing angles θ_{12}, θ_{13}, and θ_{23} and a single phase δ_{CP} related to charge-parity (CP) violation. It can be written as:

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{23} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{12}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{pmatrix} \cdot P, \tag{2.18}$$

which can be split into the θ_{23}, θ_{13}, and θ_{12} components:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \cdot P, \tag{2.19}$$

where $s_{ij} \equiv \sin(\theta_{ij})$ and $c_{ij} \equiv \cos(\theta_{ij})$ and $P = \mathrm{diag}(1, \mathrm{e}^{\mathrm{i}\alpha}, \mathrm{e}^{\mathrm{i}\beta})$ is a matrix associated with the Majorana nature of neutrinos and contains additional Majorana phases α and β. If the neutrino are Dirac fermions then these phases are zero and $P = I$. The Majorana phases play no role in oscillations and cannot be measured in an oscillation experiment.

Dozens of oscillation experiments have contributed measurements of the various oscillation parameters. Global analyses of neutrino oscillation data (45), (46), such as that performed by the NuFIT framework (1), provides a method for combining the various experiments to determine the values of the six parameters in U. As of 2019, Table 2.1 shows the 3σ ranges for the best fit mixing parameters and Equation 2.20 shows the values in the

Parameter	Best fit (NH)	3σ range (NH)	Best fit (IH)	3σ range (IH)
θ_{12} [°]	33.82	31.61 – 36.27	33.82	31.61 – 36.27
θ_{13} [°]	8.61	8.22 – 8.99	8.65	8.27 – 9.03
θ_{23} [°]	49.6	40.3 – 52.4	49.8	40.6 – 52.5
δ_{CP} [°]	215	125 – 392	284	196 – 360
Δm^2_{21} [10^{-5} eV^2]	7.39	6.79 – 8.01	7.39	6.79 – 8.01
Δm^2_{3l} [10^{-3} eV^2]	2.53	2.43 – 2.63	-2.51	-2.61 – -2.41

Table 2.1: The best fit mixing parameters from the NuFIT global fit (1) for the normal hierarchy (NH) and inverted hierarchy (IH), discussed in Section 2.4. Note that Δm^2_{3l}=$\Delta m^2_{31} >$ 0 for the NH and $\Delta m^2_{3l} = \Delta m^2_{32} < 0$ for the IH.

PMNS matrix. The values reported are generated by global fits to oscillation data using NuFIT-4.0 and are presented in (1).

$$|U| = \begin{pmatrix} 0.797 - 0.842 & 0.518 - 0.585 & 0.143 - 0.156 \\ 0.233 - 0.495 & 0.448 - 0.679 & 0.639 - 0.783 \\ 0.287 - 0.532 & 0.486 - 0.706 & 0.604 - 0.754 \end{pmatrix} \tag{2.20}$$

2.2 Neutrino Oscillations in Matter

When propagating through matter, neutrinos undergo interactions with the constituent matter particles (p, n, and e). All three neutrino flavors participate in neutral current interactions, while only ν_e undergoes charged current interactions with electrons. Because the neutral current interactions affect all flavors equally, they contribute only a term that is proportional to the identity matrix to U, and are thus physically-irrelevant. However, as first pointed out by L. Wolfenstein (47) and later developed by Mikheyev and Smirnov (48), the asymmetry amongst the neutrino flavors in the charged current interactions causes a distortion to vacuum oscillations. Specifically, the effective potential from charged current scattering of ν_e is given by:

$$V_{CC} = \sqrt{2} G_F n_e, \tag{2.21}$$

where G_F is the Fermi constant and n_e is the number density of electrons in the medium. The evolution equation of the flavor transition amplitudes (defined in Equation 2.6) can be written in matrix form as:

$$i \frac{d}{dx} \Psi_\epsilon = \mathcal{H} \Psi_\epsilon, \tag{2.22}$$

where

$$\Psi_\epsilon = \begin{pmatrix} \psi_{\epsilon e} \\ \psi_{\epsilon \mu} \\ \psi_{\epsilon \tau} \end{pmatrix} \tag{2.23}$$

and $\mathcal{H}$ is the effective Hamiltonian matrix and x is the distance from the source. $\mathcal{H}$ is given by summing the vacuum Hamiltonian with an interaction term:

$$\mathcal{H} = \frac{1}{2E} (U M^2 U^\dagger + A), \tag{2.24}$$

where

$$M = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta m^2_{21} & 0 \\ 0 & 0 & \Delta m^2_{31} \end{pmatrix} \tag{2.25}$$

and

$$A = \begin{pmatrix} A_{CC} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \tag{2.26}$$

where $A_{CC} = 2EV_{CC}$. It is particularly interesting to consider the scenario where the initial state is an electron neutrino, in the framework of two neutrino mixing. Then, Equation 2.22 can be rewritten as:

$$i \frac{d}{dx} \begin{pmatrix} \psi_{ee} \\ \psi_{e\mu} \end{pmatrix} = \frac{1}{4E} \begin{pmatrix} -\Delta m^2 \cos(2\theta) + A_{CC} & \Delta m^2 \sin(2\theta) \\ \Delta m^2 \sin(2\theta) & \Delta m^2 \cos(2\theta) - A_{CC} \end{pmatrix} \begin{pmatrix} \psi_{ee} \\ \psi_{e\mu} \end{pmatrix}. \tag{2.27}$$

The Hamiltonian matrix can be diagonalized such that it takes the form:

$$\mathcal{H} = \frac{1}{4E} \begin{pmatrix} -\Delta m^2_{\text{eff}} & 0 \\ 0 & \Delta m^2_{\text{eff}} \end{pmatrix}, \tag{2.28}$$

where Δm_{eff} is the effective mass difference due to matter effects. With this transformation the unitary mixing matrix is:

$$U = \begin{pmatrix} \cos(\theta_{\text{eff}}) & \sin(\theta_{\text{eff}}) \\ -\sin(\theta_{\text{eff}}) & \cos(\theta_{\text{eff}}) \end{pmatrix} \tag{2.29}$$

with

$$\Delta m^2_{\text{eff}} = \sqrt{(\Delta m^2 \cos(2\theta) - A_{CC})^2 + (\Delta m^2 \sin(2\theta))^2} \tag{2.30}$$

and

$$\tan(2\theta_{\text{eff}}) = \frac{\tan(2\theta)}{1 - \frac{A_{CC}}{\Delta m^2 \cos(2\theta)}}. \tag{2.31}$$

Of particular interest is the value of A_{CC} that leads to maximal mixing:

$$A_{CC} = \Delta m^2 \cos(2\theta), \tag{2.32}$$

which corresponds to an electron number density of:

$$n_e = \frac{\Delta m^2 \cos(2\theta)}{2\sqrt{2} E G_F}. \tag{2.33}$$

This resonance would lead to large flavor transition amplitudes, and, because A_{CC} is positive, can only occur if $\theta > \pi/4$ and Δm^2 is positive. Additionally, the effective mixing parameters now depend on energy in the A_{CC} term. Because the potential flips signs for antineutrinos, the mixing in matter is different for neutrinos and antineutrinos. If the matter density is constant, the transition probability can be rewritten using the effective mixing parameters as:

$$P_{\nu_\epsilon \to \nu_\gamma}(x) = \sin^2(2\theta_{\text{eff}}) \sin^2\left(\frac{\Delta m^2_{\text{eff}} x}{2E}\right), \tag{2.34}$$

identical to the vacuum mixing in Equation 2.15, with the mixing angle and mass-squared difference replaced with the effective values in matter.

2.3 Neutrino Mass

With the discovery of neutrino oscillations and the revelation that neutrinos have mass, the origin of the neutrino mass has become a fundamental question to answer. Like the other fermions, the neutrinos could be Dirac particles. This would imply the neutrino mass is generated through the Higg's mechanism that gives masses to the quarks and leptons. The only extension necessary to the standard model is to add right-handed components ν_R of the neutrino fields. These right handed neutrinos would be *sterile* in that they would not interact through the weak force. The mass of Dirac neutrinos can be obtained by including ν_R in the standard model Higgs-lepton Yukawa Lagrangian and is found to be:

$$m_i = \frac{y_i^\nu V}{\sqrt{2}} \tag{2.35}$$

where V is the Higgs vacuum expectation value (246 GeV) and y_i^ν are the Higgs to neutrino Yukawa couplings with values necessary to generate the neutrino masses. This procedure is unsatisfying in that it does not address the smallness of the neutrino masses, relative to the mass of the other fermions. In particular, to account for the tiny neutrino masses $\mathcal{O}(0.1)$ eV, the Yukawa couplings need to be $\mathcal{O}(10^{-12})$.

Because neutrinos are neutral, they do not have to be Dirac particles. To see this first consider the Dirac equation for chiral states ψ_L and ψ_R:

$$\psi = \psi_L + \psi_R \tag{2.36}$$

$$i\gamma^\mu \partial_\mu \psi_L = m\psi_R, \tag{2.37}$$

$$i\gamma^\mu \partial_\mu \psi_R = m\psi_L. \tag{2.38}$$

Setting $m = 0$ yields the Weyl equations:

$$i\gamma^\mu \partial_\mu \psi_L = 0, \tag{2.39}$$

$$i\gamma^\mu \partial_\mu \psi_R = 0. \tag{2.40}$$

The Weyl equations are special in that the solutions are two-component spinors, rather than four-components spinors (the more general solution to the Dirac equation). In the SM the neutrino, as a massless particle, is included as a left-handed Weyl spinor. It is reasonable to ask whether there are other conditions (besides $m = 0$) in which the solution to the Dirac equation is a two-component spinor. As discovered by Majorana, the answer is yes, under the condition that:

$$\psi_R = C\overline{\psi_L}^T \tag{2.41}$$

where C is the charge-conjugation operator, which converts particles to antiparticles, and $\overline{\psi} \equiv \psi^\dagger\gamma^0$. Additionally, the charge-conjugate field is defined such that $\psi^C = C\overline{\psi}^T$. Substituting for ψ_R in Equation 2.36 give:

$$\psi = \psi_L + C\overline{\psi_L}^T = \psi_L + \psi_L^C \tag{2.42}$$

Taking the charge conjugate of this field gives:

$$\psi^C = (\psi_L + \psi_L^C)^C = \psi_L^C + \psi_L = \psi \tag{2.43}$$

If the Majorana condition in Equation 2.43, in which the charge conjugation of the field is equal to itself, is satified, then a massive fermion can be described by a two-component spinor, analogous to the Weyl spinors. Of course, this condition can only hold for neutral particles, meaning that neutrinos are the only fermions that could be Majorana particles. Due to the Majorana relation, a Majorana particle and antiparticle must obey the same equations of state and must then be the same particle. In other words, a Majorana particle is its own antiparticle. Because a lepton number of +1 is assigned to neutrinos and a lepton number of -1 is assigned to antineutrinos, it is clear that a Majorana neutrino would not obey lepton number conservation.

It is interesting to consider the possible allowed states for Dirac and Majorana neutrinos with a given momentum $\mathbf{p}$ and helicity h. Dirac neutrinos have four possible states: $\nu(\mathbf{p}, h)$, $\bar{\nu}(\mathbf{p}, h)$, $\nu(\mathbf{p}, -h)$, and $\bar{\nu}(\mathbf{p}, -h)$. Note that the direction of the momentum can be flipped with a 180° rotation, which leaves h unchanged, so the states with $-\mathbf{p}$ do not represent unique states. In contrast, for the Majorana case, there are two possible states: $\nu(\mathbf{p}, h)$ and $\nu(\mathbf{p}, -h)$, as charge conjugation does not change the state. This means that Majorana neutrinos have half of the possible number of degrees of freedom as Dirac particles. Most extended standard model theories include neutrinos as Majorana particles, in part because there are only two independent components. Theories such as the type I see-saw mechanism more naturally explain the smallness of the neutrino mass, for example by including a heavy sterile right-handed neutrino partner, as described in (49).

In the case the neutrino is Majorana, the discussed differences in the oscillation patterns for neutrinos and antineutrinos is actually differences due to the chirality of the states. Because of the unknown nature of the neutrino, throughout this dissertation the term antineutrino is used to refer to either Dirac antineutrinos or Majorana neutrinos with left-handed chirality. The most promising method for discovering Majorana neutrinos is through neutrinoless double beta decay, which is discussed in Section 8.1 and reviewed in (50).

2.4 Neutrino Mass Hierarchy

As discussed in Section 2.1, oscillation experiments are sensitive only to the differences in mass squared, $\Delta m^2_{ij} = m^2_i - m^2_j$. From Table 2.1 the best fit to the 2018 oscillation data gives $\Delta m^2_{21} = 7.39 \times 10^{-5}$ eV^2 and $\Delta m^2_{31} \approx \Delta m^2_{32} = 2.53 \times 10^{-3}$ eV^2. Due to matter effects from within the sun, the sign of $\Delta m^2_{21} > 0$ is known from solar neutrino experiments. However, the sign of Δm_{31} is still unknown. The resulting picture of possible neutrino mass hierarchies (or orderings) is shown schematically in Figure 2.1. The normal ordering (NO) or normal hierarchy (NH) is defined by $m_3 > m_2 > m_1$ whereas the inverted ordering (IO) or inverted hierarchy (IH) is defined by $m_2 > m_1 > m_3$.

Despite not knowing the ordering, there is additional information that can be extracted from the current picture. At least two of the neutrino masses are larger than $\sqrt{\Delta m^2_{21}} \sim$

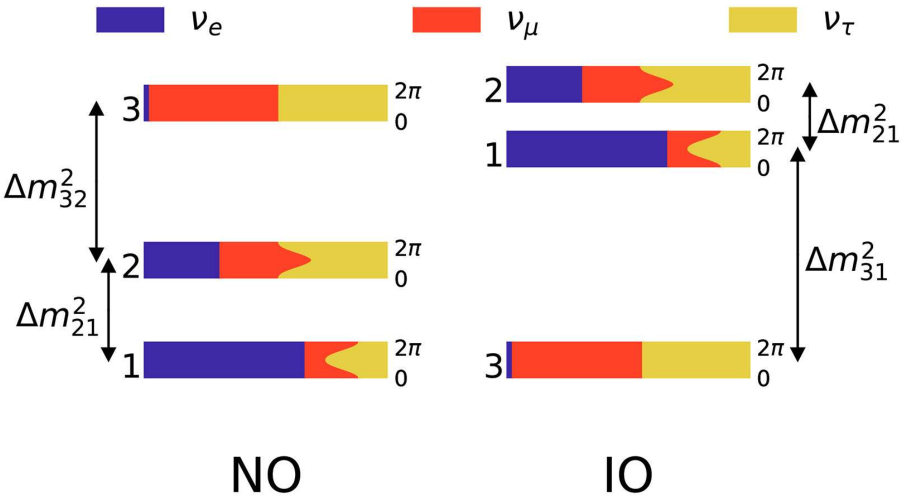

Figure 2.1: The normal and inverted neutrino mass hierarchies. The different colors indicated the probabilities of finding the α neutrino flavor in the ith neutrino mass eigenstate. This procedure is outlined in (6). The vertical scale in each band varies the value of δ_{CP} between 0 and 2π. This figure is from (7).

8.6 meV, and at least one mass is larger than $\sqrt{\Delta m^2_{31}} \sim 50$ meV. Additionally, constraints on the sum of the three neutrino mass can be set at $\sum m_\nu > 60$ meV (NO) and $\sum m_\nu > 100$ meV (IO). This is important as cosmological measurements provide strong limits on $\sum m_\nu$ (51) and could eventually indicate the neutrino mass ordering.

The most promising method for determining the mass ordering is through long-baseline oscillation experiments, using either accelerator or atmospheric neutrinos. In these experiments, as the neutrinos travel through matter, the transition probabilities are enhanced in the case of normal ordering and suppressed in the case of inverted ordering. This effect is opposite for antineutrinos. In a simple two-flavor mixing framework (e.g., $\nu_e \leftrightarrow \nu_\mu, \nu_\tau$), the effective θ_{13} angle is given by:

$$\sin^2(2\theta_{13}^{\text{eff}}) = \frac{\sin^2(2\theta_{13})}{\sin^2(2\theta_{13}) + \left(\cos(2\theta_{13}) \pm \frac{2EV_{CC}}{\Delta m^2_{31}}\right)^2}, \tag{2.44}$$

where V_{CC} is given in Equation 2.21. The plus sign corresponds to antineutrinos and minus sign corresponds to neutrinos and the sign of Δ_{31} is fixed by the ordering (positive for normal and negative for inverted). The matter effects will be particularly important when the resonance condition:

$$\Delta m^2_{31} \cos(2\theta_{13}) = 2\sqrt{2} G_F n_e E \tag{2.45}$$

is satisfied. For $\Delta m^2_{31} \sim 2.5 \times 10^{-3}$ eV^2 and distances of several hundreds of kilometers through the Earth, the resonance occurs for neutrino energies on the order of 1 GeV. This effect can be measured in the T2K and NOνA experiments, which generate high energy neutrino beams pointed at a detector, placed several hundred kilometers away. Recent global fits using new T2K and NOνA data tend to prefer the normal ordering (1), (45), (52).

2.5 Neutrino Interactions

Neutrino interactions are described impressively well by the standard model, and thus far no deviations have been observed. The simplest way in which neutrinos interact with matter is through scattering off of electrons. Low energy neutrinos of all three flavors can interact through elastic scattering (ES) on an electron. For example, for a ν_e:

$$\nu_e + e^- \to \nu_e + e^-, \tag{2.46}$$

which has no energy threshold because the initial and final state are identical. This process occurs similarly for a $\bar{\nu}_e$:

$$\bar{\nu}_e + e^- \to \bar{\nu}_e + e^-. \tag{2.47}$$

The charged current (CC) and neutral current (NC) tree-level Feynman diagrams are shown in Figures 2.2 and 2.3. Note that for low energy ν_μ and ν_τ interactions, only the NC diagram contributes as there is not enough energy to create a μ^- or τ^-.

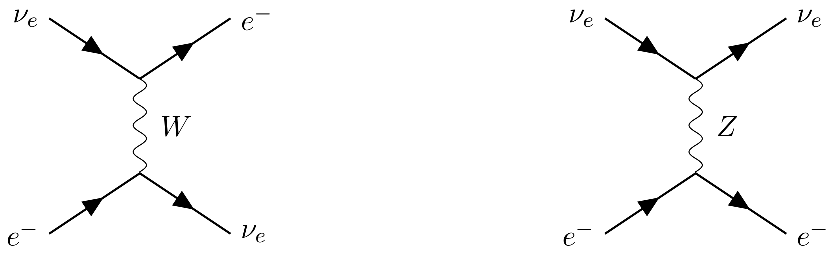

(a) The ν_e - e^- CC ES interaction.

(b) The ν_e - e^- NC ES interaction.

Figure 2.2: The tree-level Feynman diagrams for elastic scattering of a ν_e on an e^-.

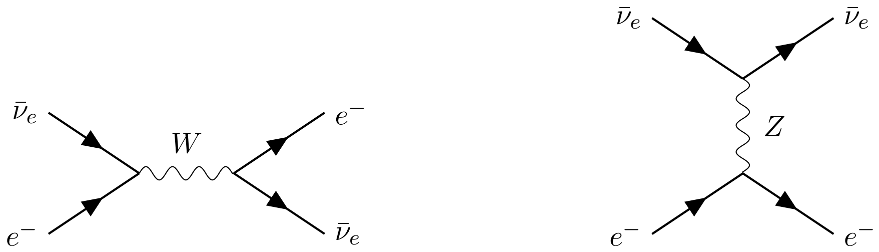

(a) The $\bar{\nu}_e$ - e^- CC ES interaction.

(b) The $\bar{\nu}_e$ - e^- NC ES interaction.

Figure 2.3: The tree-level Feynman diagrams for elastic scattering of a $\bar{\nu}_e$ on an e^-.

Neutrinos can also interact via quasi-elastic CC interactions with components of the nucleus via:

$$\nu_l + n \rightarrow p + l^-, \tag{2.48}$$

and

$$\bar{\nu}_l + p \rightarrow n + l^+. \tag{2.49}$$

In particular, at low energies below around 100 MeV, Equation 2.49 is referred to as inverse beta decay (IBD), and is shown in Figure 2.4. This is the process through which low energy (several MeV) antineutrinos are most commonly detected, primarily because of it has a more clear detection signature and a significantly larger cross-section than $\bar{\nu}_e$ ES. Chapter 6 discusses a search for reactor antineutrinos using a water target in the SNO+ detector, which utilizes this IBD interaction. At low energies these CC interactions only occur for $\bar{\nu}_e$.

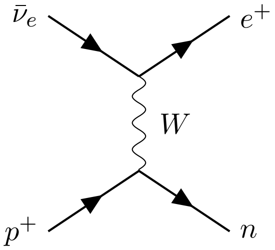

Figure 2.4: The tree-level Feynman diagram for inverse beta decay interactions.

The differential cross-section for IBD reactions is given at zeroth order in $1/M$, the nucleon mass, in (9) as:

$$\left(\frac{d\sigma}{d\cos(\theta)}\right)^0 = \frac{\sigma_0}{2}\left[(f^2+3g^2)+(f^2-g^2)v_e^0\cos(\theta)\right]E_e^0 p_e^0, \tag{2.50}$$

where E_e is the positron energy, f is the vector coupling constant and g is the axial-vector coupling constant, set equal to 1 and 1.26 respectively at zeroth order, v_e is the velocity of the positron, θ is the outgoing angle of the positron, and p_e is the momentum of the positron, all given in the lab frame. σ_0 is an energy-independent normalizing constant. Integrating over $\cos(\theta)$ gives the total cross-section:

$$\sigma_{\text{tot}}^0 = \sigma_0(f^2+3g^2)E_e^0 p_e^0 = 0.0952\left(\frac{E_e^0 p_e^0}{1\text{MeV}^2}\right)\times 10^{-42}\text{cm}^2 \tag{2.51}$$

This equation can be rewritten in terms of the neutron lifetime, τ_n as:

$$\sigma_{\text{tot}}^0 = \frac{2\pi^2}{f_{\text{p.s}}\tau_n m_e^5}E_e^0 p_e^0, \tag{2.52}$$

where $f_{\mathrm{p.s}} = 1.7152$ is a phase space factor and m_e is the positron mass. The proportionality of the IBD cross-section to the neutron beta decay rate is interesting, but not unexpected, as the IBD diagram in Figure 2.4 is related to neutron decay via crossing symmetry.

At zeroth order the energy of the positron is given by summing the energies of the outgoing and incoming particles:

$$E_e^0 = E_\nu + m_p - m_n. \tag{2.53}$$

The first order correction to the positron energy is given by:

$$E_e^1 = E^0 \left[1 - \frac{E_\nu}{M}(1 - v_e^0 \cos(\theta))\right] - \frac{y^2}{M} \tag{2.54}$$

where $y^2 = [(m_n - m_p)^2 - m_e^2]/2$. In addition to the energy of the positron, the outgoing angle of the positron is an important observable for detectors. Again from (9) the angular distribution is given by:

$$\langle \cos(\theta) \rangle = \frac{1}{3} v_e a(E_\nu). \tag{2.55}$$

Expanding in zeroth order in $1/M$ gives:

$$\langle \cos(\theta) \rangle^0 = \frac{1}{3} v_e a^0 \approx -0.034 v_e^0, \tag{2.56}$$

where $a^0 \approx -0.10$, is given completely in terms of f and g. At first order it can be shown this is extended to:

$$\langle \cos(\theta) \rangle^1 \approx -0.034 v_e^0 + 2.4 \frac{E_\nu}{M}. \tag{2.57}$$

The angle is slightly backwards relative to the incoming neutrino direction, but is fairly close to flat across the neutrino energies between 2 to 10 MeV that are of interest in the reactor neutrino search. Figure 2.5 shows the results for σ_{tot} and $\langle \cos(\theta) \rangle$ as a function of neutrino energy from 0 to 10 MeV.

At higher energies, neutrinos and antineutrinos can interact with the entire nucleus through three primary processes: quasi-elastic scattering (QE), resonant production (RES)

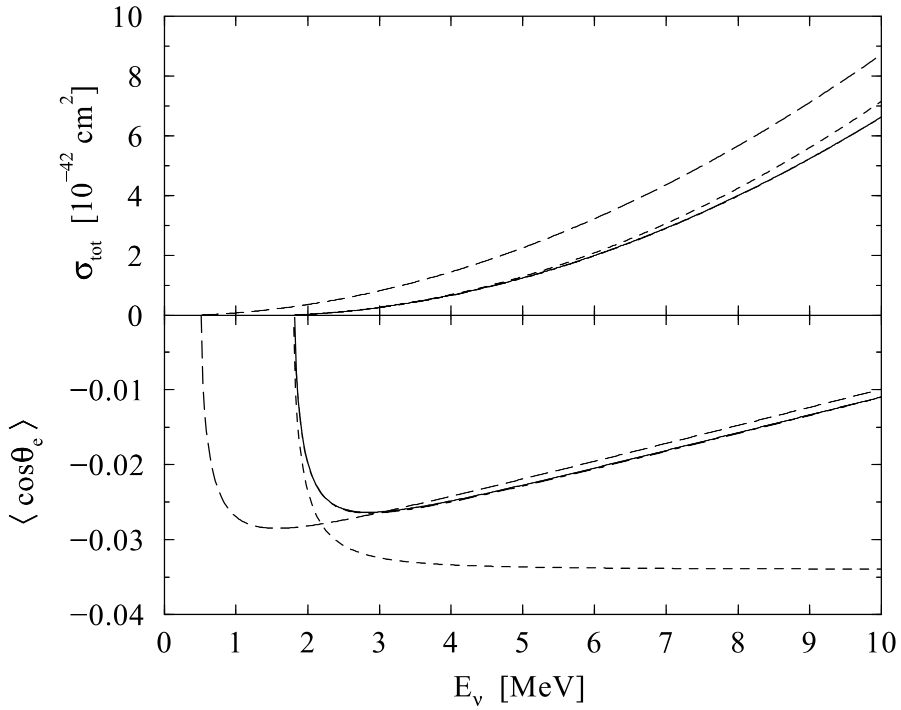

Figure 2.5: The upper panel shows the IBD cross-section for $\mathcal{O}(1/M)$ in the solid line and for the $\mathcal{O}(1)$ result in the short-dashed line. The long-dashed line shows results from (8). The bottom panel shows outgoing angle of the positron, $\langle\cos(\theta)\rangle$, for the same reaction. This Figure is from (9).

and deep inelastic scattering (DIS). In resonant production, the neutrino interacts with the proton or neutron and produces a Δ or π in the final state. In deep-inelastic scattering the neutrino interaction can produce many final state hadrons. The cross-sections of each of these processes as a function of neutrino energy is shown in Figure 2.6 and a detailed treatment of these complicated processes is given in (10).

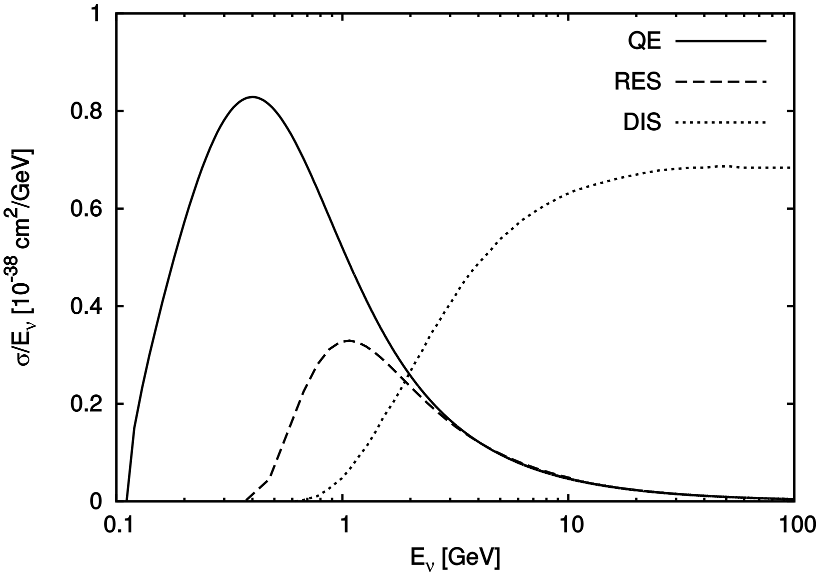

Figure 2.6: The total cross-section for neutrino energies above 100 MeV for the QE, RES, and DIS processes. This figure is from T. Leitner's thesis (10).

Chapter 3

Neutrino Sources and Experiments

Historically, neutrino experiments have detected neutrinos emitted from six different sources: the sun, the atmosphere, man-made beams, terrestrial fission reactors, radioactivity in the Earth's crust and mantle, and supernova. Several of these sources and associated experiments are discussed in this section. Most relevant for the analysis presented in Chapter 6 is Section 3.3 describing reactor antineutrinos. While important for many experiments, beam neutrinos and supernova neutrinos are not discussed in this chapter.

3.1 Solar Neutrinos

Standard solar models (SSMs) are constructed to predict the Sun's observed luminosity, radius, and surface heavy-element-to-hydrogen ration (53). At current, the solar models are crucial tools used for predictions of the pressure oscillations of the Sun (helioseismology) and in the prediction of the solar neutrino fluxes. Precision comparisons between measurements of solar neutrinos and the predictions from the solar models allows for the determination of oscillation parameters and to test different models of neutrino propagation (54). Historically, solar models were critical for comparing observations against expected fluxes, and it was these comparisons that gave the first indications of neutrino oscillations.

The sun produces enormous numbers of neutrinos in the thermonuclear fusion reactions within the solar core. The nuclear reactions that occur in the sun are the *pp* chain and CNO cycle, shown in Figures 3.1 and 3.2. The reactions in the *pp* and CNO chains produce neutrinos with energies primarily below 10 MeV, as shown in Figure 3.3. Notably, the reaction that produce neutrinos at the highest energies are the ^{8}B and hep reactions; however, the rate of hep neutrino emission is approximately 0.15% of the ^{8}B flux and has never been measured (22).

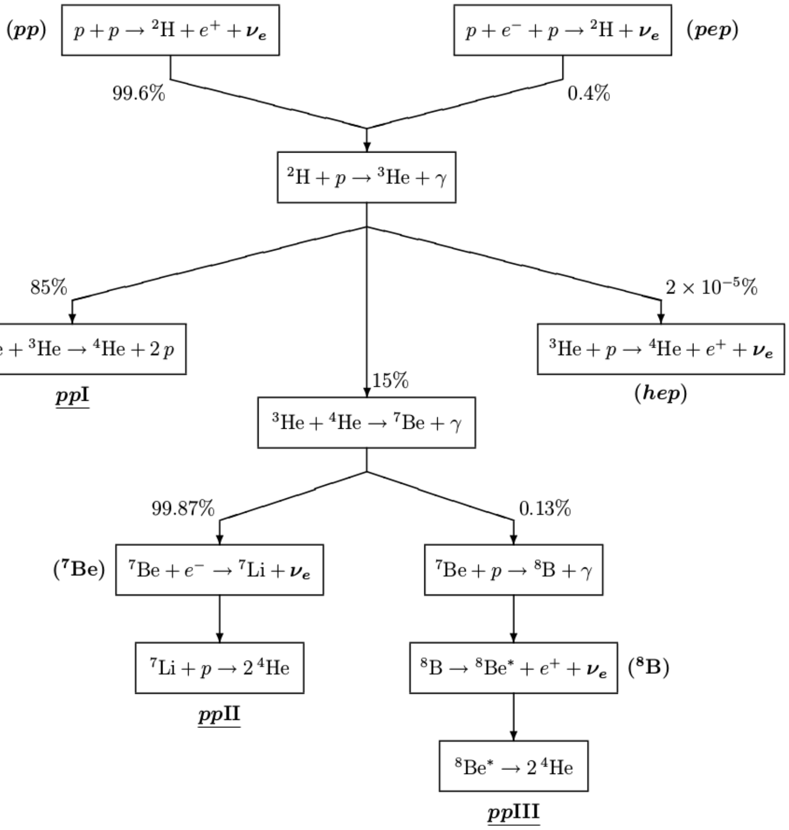

(pp)
$p + p \to {}^2\mathrm{H} + e^+ + \nu_e$
$p + e^- + p \to {}^2\mathrm{H} + \nu_e$
(pep)
99.6%
0.4%
${}^2\mathrm{H} + p \to {}^3\mathrm{He} + \gamma$
85%
$2 \times 10^{-5}\%$
${}^3\mathrm{He} + {}^3\mathrm{He} \to {}^4\mathrm{He} + 2\,p$
ppI
15%
${}^3\mathrm{He} + p \to {}^4\mathrm{He} + e^+ + \nu_e$
(hep)
${}^3\mathrm{He} + {}^4\mathrm{He} \to {}^7\mathrm{Be} + \gamma$
99.87%
0.13%
(⁷Be)
${}^7\mathrm{Be} + e^- \to {}^7\mathrm{Li} + \nu_e$
${}^7\mathrm{Be} + p \to {}^8\mathrm{B} + \gamma$
${}^7\mathrm{Li} + p \to 2\,{}^4\mathrm{He}$
ppII
${}^8\mathrm{B} \to {}^8\mathrm{Be}^* + e^+ + \nu_e$
(⁸B)
${}^8\mathrm{Be}^* \to 2\,{}^4\mathrm{He}$
ppIII

Figure 3.1: The *pp* chain of stellar thermonuclear reactions, reproduced from (11).

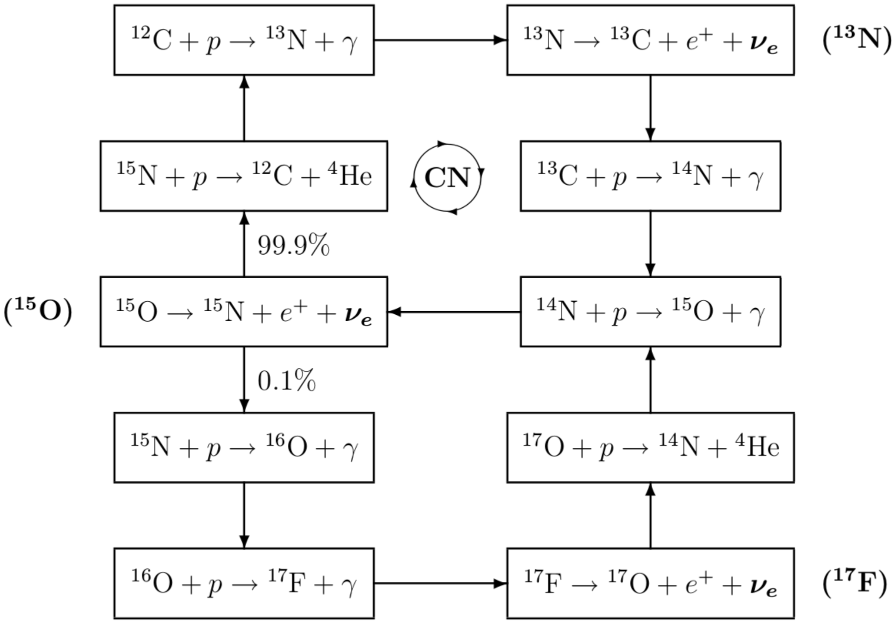
$^{12}C + p \rightarrow {}^{13}N + \gamma$
$^{13}N \rightarrow {}^{13}C + e^+ + \nu_e$
(^{13}N)
$^{15}N + p \rightarrow {}^{12}C + {}^{4}He$
CN
$^{13}C + p \rightarrow {}^{14}N + \gamma$
99.9%
(^{15}O)
$^{15}O \rightarrow {}^{15}N + e^+ + \nu_e$
$^{14}N + p \rightarrow {}^{15}O + \gamma$
0.1%
$^{15}N + p \rightarrow {}^{16}O + \gamma$
$^{17}O + p \rightarrow {}^{14}N + {}^{4}He$
$^{16}O + p \rightarrow {}^{17}F + \gamma$
$^{17}F \rightarrow {}^{17}O + e^+ + \nu_e$
(^{17}F)

Figure 3.2: The *CNO* cycle of stellar thermonuclear reactions, reproduced from (11).

Solar neutrinos were first detected in 1970 by the Homestake experiment (55), which continued to take data for the next 25 years (56). The Homestake experiment measured a deficit of about one-third of the SSM prediction in the number of detected solar neutrinos, which was dubbed the ‘solar neutrino problem’. The solar neutrino problem was confirmed by the SAGE (57), GALLEX (58), and GNO (59) experiments. The most definitive evidence of neutrino flavor transformation as the solution to the solar neutrino problem was provided by the SNO experiment, which utilized heavy water to measure all three neutrino flavors (13).

The SNO experiment utilized a heavy water target, which provided a unique method for detection of all three neutrino flavors through the following reactions:

$$\mathrm{CC}: \nu_{\mathrm{e}} + \mathrm{d} \rightarrow \mathrm{e}^- + \mathrm{p} + \mathrm{p} \tag{3.1}$$

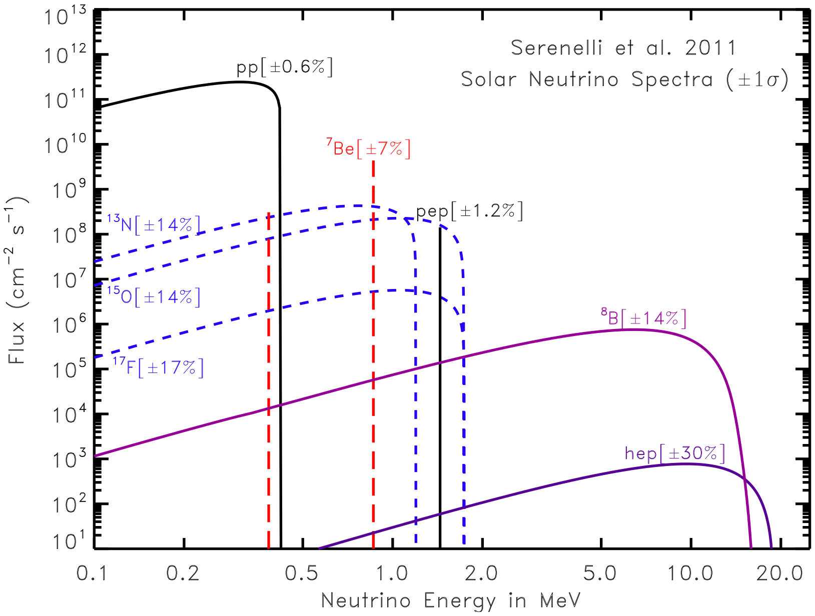

Figure 3.3: The neutrino energy spectra for the various components of the *pp* chain and CNO cycle along with the SSM uncertainties. This figure is from (12).

$$\mathrm{NC}: \nu_{\mathrm{x}} + \mathrm{d} \rightarrow \nu_{\mathrm{x}} + \mathrm{n} + \mathrm{p} \tag{3.2}$$

$$\mathrm{ES}: \nu_{\mathrm{x}} + \mathrm{e}^{-} \rightarrow \nu_{\mathrm{x}} + \mathrm{e}^{-}, \tag{3.3}$$

for $x \in (e, \mu, \tau)$. Typically in a water detector only the elastic scattering channel is measured, as there are no deuterons for the neutrino to interact with. All of the neutrino created in the sun are electron type, so the ν_μ and ν_τ are created through oscillations. The SNO flux results were presented in (13) separately for the CC, NC, and ES channels. The NC reaction on deuterium was particularly important, as it is equally sensitive to all three neutrino flavors and thus provides the most robust check against the SSM. The elastic scattering process is sensitive to the different flavors, but the total cross-section for ν_e is a factor of about six larger than for ν_μ and ν_τ. The CC reaction on deuterium is only sensitive to ν_e, but is important because it provides the most direct measure of the energy of the electron neutrinos. The ratio $\frac{\phi_{CC}}{\phi_{NC}} = 0.34$ clearly shows the deficit in the CC channel. Additionally, the NC current channel is consistent with the expected flux from the SSM. The full results are shown in Figure 3.4, where the fluxes have been converted to electron ϕ_e and non-electron $\phi_{\mu,\tau}$ types with a change of variables. The results are compared directly to the SSM and are shown to be consistent in the framework of neutrino oscillations.

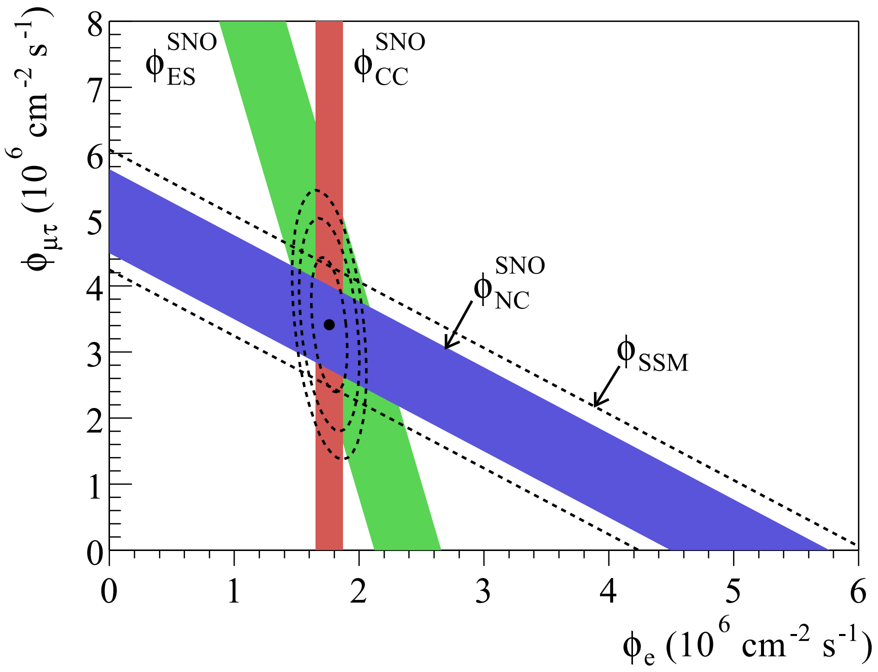

Figure 3.4: The flux of ^{8}B solar neutrinos detected as μ or τ flavor plotted the flux detected as e type. The diagonal dashed lines show the ^{8}B flux predicted by the SSM. The ES, CC and NC results are shown in the green, red, and blue bands respectively. The combined flux results are consistent with neutrino flavor transformation. This figure is reproduced from (13).

In the context of this thesis, solar neutrinos are expected to be the dominant background for SNO+'s $0\nu\beta\beta$ search. This realization led to significant R&D efforts to help reduce this background in large-scale scintillation-based $0\nu\beta\beta$ detectors, discussed in Chapter 8.

3.2 Atmospheric Neutrinos

Neutrinos are created when cosmic rays, primarily protons, interact with nuclei in the atmosphere. In these interactions, mesons, such as pions and kaons, are produced, which decay mainly into muons. Neutrinos are produced both in the pion and kaon decays and in the subsequent muon decay:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \tag{3.4}$$

$$K^+ \rightarrow \mu^+ + \nu_\mu \tag{3.5}$$

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu \tag{3.6}$$

which occurs similarly for π^-, K^- and μ^-. The neutrinos generated in these interactions are primarily in the range of 100 MeV to 10 GeV.

The Kamiokande experiment, originally built to study proton decay, made some of the first measurements of atmospheric neutrinos. The Kamiokande experiment measured the CC ν_μ and ν_e interactions in their water Cherenkov detector, and had good separation between μ^- and e^- events. Kamiokande found a deficit in the ratio of muon to electron neutrino interactions (60):

$$R_{\mu/e} = 0.60^{+0.07}_{-0.06} \tag{3.7}$$

which is now explained by ν_μ disappearance ($\nu_\mu \rightarrow \nu_\tau$). The Super Kamiokande experiment confirmed these results (44) and cast the results in the framework of neutrino oscillations.

3.3 Reactor Neutrinos

Electron-flavor antineutrinos ($\bar{\nu}_e$) are produced in nuclear reactors around the Earth at energies up to about 10 MeV. These reactors provide the most intense source of man-made antineutrinos, which are particular useful for studying oscillations. The first detection of antineutrinos was performed in the Cowan-Reines neutrino experiment, which observed $\bar{\nu}_e$s from the Savannah River nuclear reactor (41). Since then, reactor antineutrino experiments have been used to measure four of the neutrino mixing parameters given in Table 2.1. In general, reactor antineutrino oscillation experiments look for the disappearance of $\bar{\nu}_e$s. That is, the $\bar{\nu}_e$s oscillate to $\bar{\nu}_\mu$s and $\bar{\nu}_\tau$ which do not undergo charged current interactions at these low energies.

The $\bar{\nu}_e$s primarily interact in the detectors via inverse beta decay (IBD) reactions, shown in Figure 2.4. They also elastically scatter off of electrons, shown in Figure 2.3, but the cross-section is about 150 times smaller than the IBD reaction and the process is more difficult to detect. The reactor antineutrino detectors typically measure oscillation parameters using both the flux and energy spectrum of the detected antineutrinos. The IBD cross-section and energy threshold important for these searches is given in Section 2.5. The interaction rate and detection mechanism for these antineutrinos is discussed in Section 6.2. This section will focus on the history of experiments that use antineutrinos from nuclear reactors and the model parameters that were measured.

In the three flavor neutrino model the survival probability of the $\bar{\nu}_e$ is given by expanding the sum in Equation 2.12:

$$\begin{aligned} P_{\bar{\nu}_e\to\bar{\nu}_e} =& 1 - \cos^4(\theta_{13})\sin^2(2\theta_{12})\sin^2(\Delta_{21}) - \\ & \sin^2(2\theta_{13})\left(\cos^2(\theta_{12})\sin^2(\Delta_{31}) + \sin^2(\theta_{12})\sin^2(\Delta_{32})\right), \end{aligned} \tag{3.8}$$

where $\Delta_{ij} = \frac{1.27L[\mathrm{m}]\Delta \mathrm{m_{ij}}[\mathrm{eV}^2]}{E[\mathrm{MeV}]}$. Note that the angle θ_{23} and δ_{CP} are the only mixing parameters that do not enter this transition probability. Thus, reactor oscillation experiments are sensitive to the four other mixing parameters: θ_{12}, θ_{13}, Δm^2_{21} and Δm^2_{31}. Given the unknown mass hierarchy, described in Section 2.4, reactor experiments typically report $\Delta m^2_{ee} \approx \Delta m^2_{31} \pm 0.023 \times 10^{-3}\text{eV}^2$ where the plus sign corresponds to the normal hierarchy and the minus sign corresponds to the inverted hierarchy. The oscillation are not sensitive to matter effects because of the relatively small amount of matter the antineutrinos travel through between the source and the detector.

Figure 3.5 shows the survival probability as a function of distance from the reactor to the detector, integrated over the energy spectrum of the $\bar{\nu}_e$ energy spectrum, weighted by the IBD cross-section. This figure indicates that around 1.5 km a dip in the survival probability occurs, which is driven by the θ_{13} component. After tens of kilometers, a large dip in the survival probability, driven by the θ_{12} oscillation, dominates.

The Kamioka Liquid Scintillator Antineutrino Detector (KamLAND) experiment, described in Section 3.3.1, is a large scintillator detector located an average flux-weighted distance of 180 kilometers from the nuclear reactors, making it sensitive primarily to the θ_{12} oscillations. The Daya Bay, RENO, and Double Chooz experiments, described in Section 3.3.2, are located around 1 km from nuclear reactor facilities and are sensitive to the θ_{13} oscillations.

3.3.1 KamLAND

The KamLAND experiment measures the oscillation of $\bar{\nu}_e$ emitted from distant nuclear reactors. In particular, the Kashiawazaki station was the highest power-station in the world, at about 8 GW, while KamLAND was taking data. The flux-averaged distance of the nuclear reactors is 180 km from the KamLAND detector. The first measurement of reactor $\bar{\nu}_e$ disappearance was reported in 2004 (32).

The KamLAND detector, located in the Kamioka mine, consists of 1 kton of liquid scintillator inside of an 18 m containment vessel. The vessel is surround by a buffer oil that provides shielding from external backgrounds. A total of 1879 photomultiplier tubes view

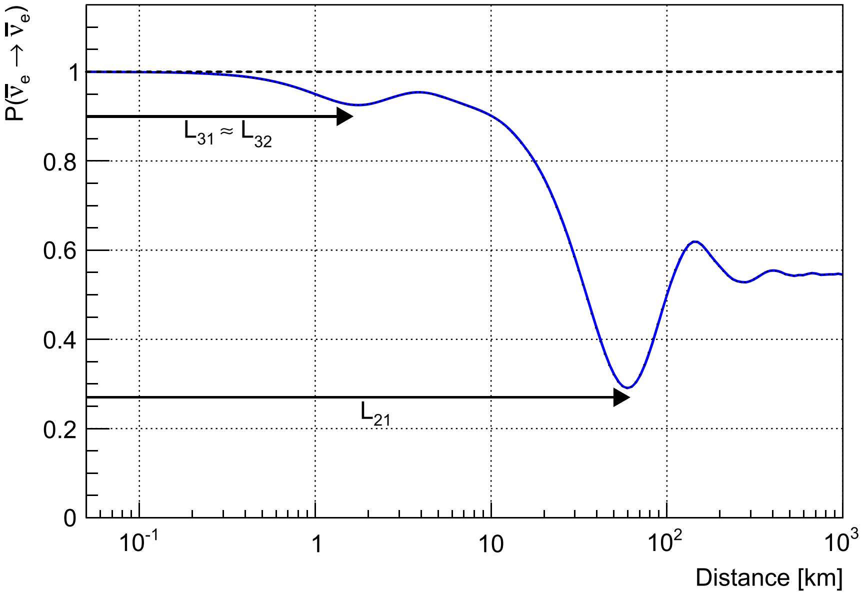

Figure 3.5: The survival probability for $\bar{\nu}_e$ as a function of distance from the reactor to the detector. The probability is averaged over the energy spectrum of the $\bar{\nu}_e$ emitted from the reactors, weighted by the cross-section of the IBD reaction. The values of θ_{13} and θ_{12} determine the amplitudes of the oscillations and the corresponding peaks in the survival probability are indicated. This figure is from (14).

the target with a coverage of approximately 34%. The excellent energy resolution provided by the liquid scintillator provides an accurate determination of the positron energy, which relates directly to the energy of the antineutrino (given in Equations 2.53 and 2.54). It is through a spectral analysis that KamLAND is able to perform the most precise measurement of Δm^2_{21}, which can be compared to solar neutrino experiments. The mixing parameters measured by KamLAND are

$$\Delta m^2_{21} = 7.58 \pm^{0.14}_{0.13} \text{ (stat)} \pm 0.15\text{(syst)} \times 10^{-5}\text{eV}^2 \tag{3.9}$$

and

$$\theta_{12} = 0.56 \pm^{0.10}_{0.07} \text{(stat)} \pm^{0.10}_{0.06} \text{(syst)} \tag{3.10}$$

from (15). The KamLAND survival probability as a function of $L_0/E_{\bar{\nu}_e}$ with the corresponding best fit is shown in Figure 3.6.

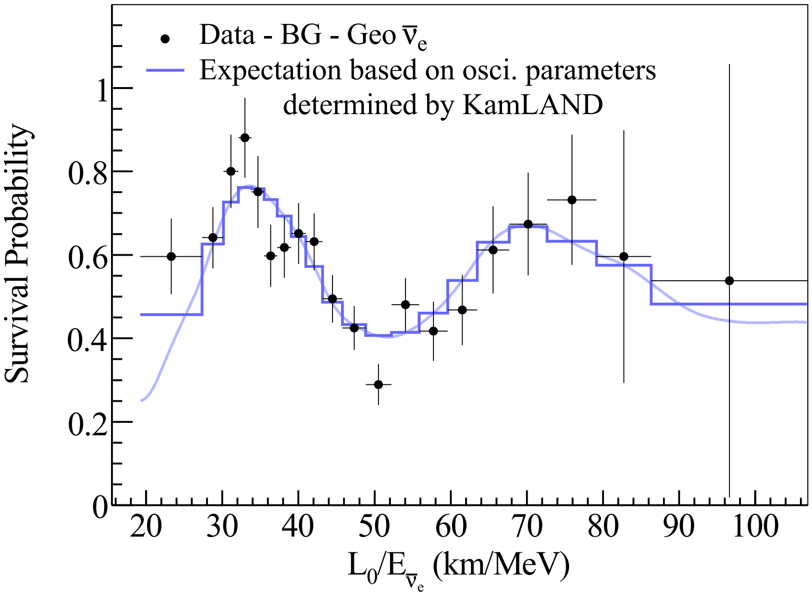

Figure 3.6: The reactor antineutrino survival probability as a function of $L_0/E_{\bar{\nu}}$, from (15), where L_0 is the flux-average baseline of 180 km. The primary backgrounds have been removed from the data points. The blue curve shows the expectation based on the neutrino oscillation parameters determined by KamLAND.

The solar neutrino experiments are also sensitive to Δm^2_{21} through its effect on the MSW transition region. The ^{8}B solar neutrino survival probability has been measured down to 3.5 MeV by the Super-Kamiokande-IV (61) and SNO (62), (63) collaborations. Combining the results, the best fit solar Δm^2_{21} is:

$$\Delta m^2_{21} = 4.8 \pm^{1.5}_{0.8} \times 10^{-5}\text{eV}^2, \tag{3.11}$$

which is in clear tension with the KamLAND result in Equation 3.9. Figure 3.7 shows a global solar analysis in comparison with the KamLAND data, performed by the NuFit collaboration (1). In this framework, the best fit Δm^2_{21} of KamLAND is at $\Delta\chi^2_{\text{solar}} = 4.7$. This tension is of particular interest in the community, and there have been various attempts to explain the discrepancy as non-standard neutrino interaction and propagation (64), (65), (66) or as vacuum-enhanced neutrino mixing (20). During the scintillator phase, SNO+ will be able to detect, at considerable rates, both low-energy solar neutrinos and reactor antineutrinos. These measurements will potentially shed light on this interesting tension.

3.3.2 Daya Bay, RENO, Double Chooz, and JUNO

As illustrated in Figure 3.5, antineutrinos with the energy spectrum produced in nuclear reactors have a dip in survival probability at 1.1 km. The size of the dip can be used to measure the value of θ_{13}. The Palo Verde and Chooz reactor antineutrino experiments took data from about 1998 to 2000 with detectors placed about 1 km from the source. Neither experiment found evidence of neutrino oscillations (67), (68). This led to predictions that θ_{13} could be very small or equal to zero (69) and to a new generation of experiments searching for oscillations driven by θ_{13}. The Double Chooz experiment in France, the Daya Bay experiment in China, and the RENO experiment in South Korea became the set of experiments searching for this oscillation. Importantly, these experiments chose to build near detectors (several hundred meters from the reactor) to measure the antineutrino flux prior to oscillation.

The detector technology of all three experiments was quite similar. All three opted to use liquid scintillator targets inside of acrylic vessels, surrounded by PMTs. The scintillator is loaded with gadolinium to improve the neutron detection efficiency. The primary difference between the experiments is in the expected rate of interactions due to the distance from and power of the nearby nuclear reactors. Daya Bay built eight detectors in total, two at each near site and four at the far site, in order to increase the signal statistics and to compare

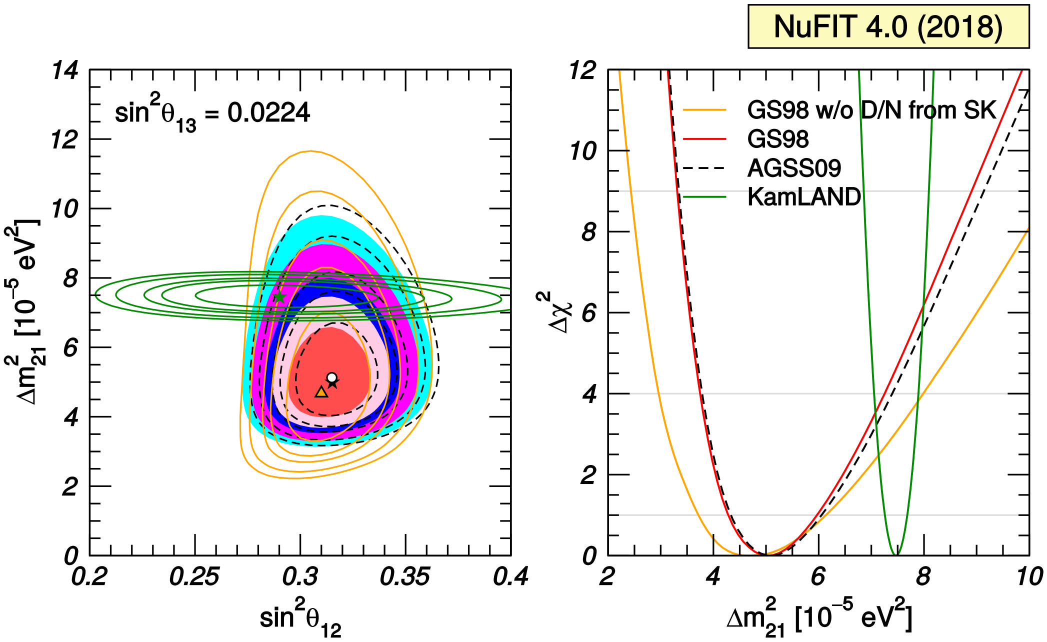

Figure 3.7: (Left) Allowed parameter regions at 1, 2, and 3 σ for the GS98 (full regions and best fit marked by black star) and AGSS09 solar models (dashed, void contours with best fit marked by white dot) and for the KamLAND data (green contours and best fit indicated by green star). (Right) The $\Delta\chi^2$ dependence on Δm^2_{21} for the solar and KamAND analyses. These analyses are performed with a fixed $\sin^2(\theta_{13}) = 0.0224$. This plot and caption is reproduced from (1).

efficiencies and systematics. Overall, this meant that Daya Bay had the highest signal rate of the three detectors and was able to perform the most sensitive measurement of θ_{13}.

All three experiments began collecting data in 2011. The Double Chooz detector published the first results, using the rate and energy spectrum of the detected events to measure $\sin^2(2\theta_{13}) = 0.087 \pm 0.041(\text{stat}) \pm 0.030(\text{syst.})$, which excluded $\theta_{13} = 0$ at almost 2σ (33). In 2012 Daya Bay published a $>5\sigma$ discover of $\bar{\nu}_e$ disappearance (34), measuring $\sin^2(2\theta_{13}) = 0.092 \pm 0.016(\text{stat}) \pm 0.005(\text{syst.})$. Quickly afterwards, RENO published first results consistent with Daya Bay results and at a similar significance (35). Notably, all three experiments later published results using the neutron capture signal on hydrogen (70), (71), (72) (as opposed to the primary signal from neutron capture on gadolinium), which is an important signal for the analysis discussed in Chapter 6.

The measurement of θ_{13}, particularly given its relatively large value around 8.6°, has consequences for future measurements of the neutrino mass hierarchy and CP violation in the neutrino sector. Several experiments, including Hyper Kamiokande and DUNE, hope to persue these measurements using neutrinos created as a beam using accelerators. Relevant for this section is the JUNO experiment, which plans to measure the mass hierarchy using a 20 kton (fiducial) mass liquid scintillator detector (16) to observe reactor antineutrino interactions. The sensitivity to the mass hierarchy for a reactor neutrino experiment is most clearly understood by rewriting Equation 3.8 replacing Δ_{32} with Δ_{ee}:

$$\begin{aligned} P_{\bar{\nu}_e \to \bar{\nu}_e} =& 1 - \frac{1}{2}\sin^2(2\theta_{13})\left[1 - \sqrt{1 - \sin^2(\theta_{12})\sin^2(\Delta_{21})}\cos(2|\Delta_{ee}| \pm \phi)\right] \\ & - \cos^4(\theta_{13})\sin^2(2\theta_{12})\sin^2(\Delta_{21}), \end{aligned} \tag{3.12}$$

where

$$\sin(\phi) = \frac{\cos^2(\theta_{12})\sin(2\sin^2(\theta_{12})\Delta_{21}) - \sin^2(\theta_{12})\sin(2\cos^2(\theta_{12})\Delta_{21})}{\sqrt{1 - \sin^2(2\theta_{12})\sin^2(\Delta_{21})}}, \tag{3.13}$$

and

$$\Delta m^2_{ee} = \cos^2(\theta_{12})\Delta m^2_{31} + \sin^2(\theta_{12})\Delta m^2_{32}. \tag{3.14}$$

The ± term in Equation 3.12 is a plus sign for the normal hierarchy and a minus sign for the inverted hierarchy. Note that the ϕ term only enters as the argument of a cos term, such that the hierarchy only affects the frequency of the oscillation and not the amplitude. The manifestation of this term in the survival probability is visualized in Figure 3.8, which shows L/E for a 60 km baseline and the energy spectrum of the reactor antineutrinos, weighted by the IBD cross-section. This figure illustrates that the mass hierarchy can be measured using reactor antineutrinos, but only if the detector expects to have a high event rate and excellent energy resolution.

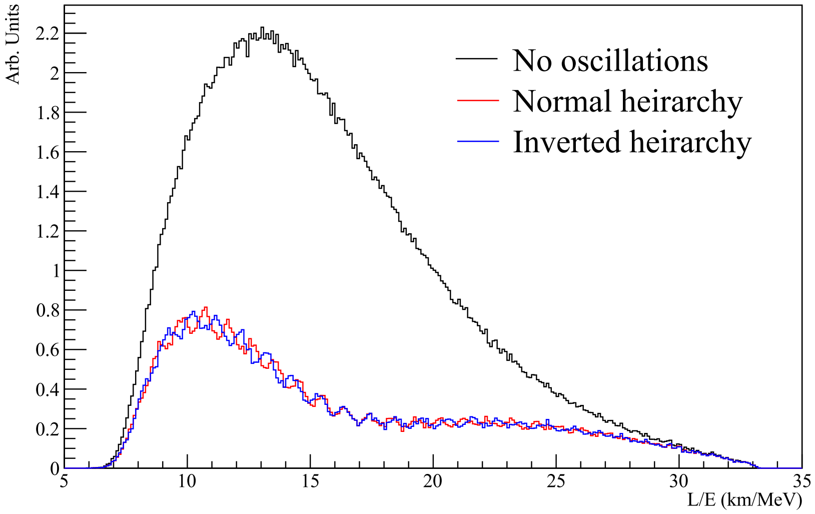

Figure 3.8: The L/E (km/MeV) spectrum for reactor neutrino oscillations assuming a distance of 60 km. The flux is integrated over the energy spectrum of the reactor antineutrinos and weighted by the IBD cross-section. This figure is motivated by the JUNO detector, presented in (16).

3.3.3 Very Short Baseline Experiments

Very short baseline reactor experiments are designed primarily to test oscillations into a fourth sterile neutrino flavor. These searches are motivated by results from the LSND and MiniBooNE experiments, which both looked for $\bar{\nu}_\mu \to \bar{\nu}_e$ oscillations and MiniBooNE additionally observed $\nu_\mu \to \nu_e$ oscillations. Both experiments observed excesses relative to expectation that have been interpreted in a model with one additional sterile neutrino with a $\Delta m \sim 1$ eV^2 (73).

In this 3+1 model, the survival probability at very short distances is given by:

$$P_{\bar{\nu}_e \to \bar{\nu}_x} = 1 - \sin^2(2\theta_{14}) \sin^2 \left(\frac{1.27\Delta m^2_{14}[\text{eV}^2]\text{L}[\text{m}]}{E[\text{MeV}]} \right). \tag{3.15}$$

At a very short baseline, around 10 m, L/E $\sim$ 3 m/MeV, which gives sensitivity to $\Delta m^2_{14} \sim$ 1 eV^2.

A host of these experiments are being built to search for sterile neutrinos by placing a detector very close to a nuclear reactor. The DANSS experiment uses segmented plastic scintillator placed underneath the core of the Kalinin Nuclear Power Plant (37). The NEOS experiment has 108 L of 0.5% gadolinium loaded liquid scintillator, installed 23 m from the center of the Hanbit Nuclear Power Reactor core in Korea (74). The NEUTRINO-4 experiment uses liquid scintillator as a target near the SM-3 research reactor in Russia (75). The SoLid experiment consists of cubes of plastic scintillator installed 6 m from the SCK CEN BR2 research reactor in Belgium (76). The STEREO detector uses gadolinium loaded liquid scintillator placed 9.4 to 11 meters from the core of the High Flux Reactor in France (77). The PROSPECT experiment uses rectangular volumes consisting of ^{6}Li-doped liquid scintillator installed near the High Flux Isotope Reactor at Oak Ridge National Laboratory (36). Thus far, no convincing detection of sterile neutrino oscillations has been observed in these detectors.

3.4 Geoneutrinos

The Earth produces radiogenic heat through the decay of long-lived isotopes, with half-lives comparable or longer than the Earth's age. In the decay of these radioactive elements, antineutrinos are emitted. The primary isotopes are ^{238}U, ^{235}U, ^{232}U, and ^{40}K. The decay chains of each of these isotopes, proceeding to a stable nucleus, produces several antineutrinos and dozens of MeV of energy. The study of geoneutrinos would provides a useful measurement of the composition of Earth's interior and insight into the radiogenic heat contribution to the Earth's total surface heat flux (17).

Thus far, geoneutrinos have only been measured by the KamLAND (78) and Borexino detectors (17). As illustrated in Figure 3.9, the geoneutrinos are low in energy and relatively few of them are above the IBD threshold, making them difficult to detect. Additionally, reactor antineutrinos provide a background for a search for geoneutrinos, and visa versa (32). SNO+ hopes to be the third detector to detect antineutrinos from the Earth during the scintillator phase; however, the flux and detection efficiency is far to small to hope to detect geoneutrinos in the water phase. Geoneutrinos provide a negligible source of backgrounds to the reactor antineutrino search presented in Section 6.

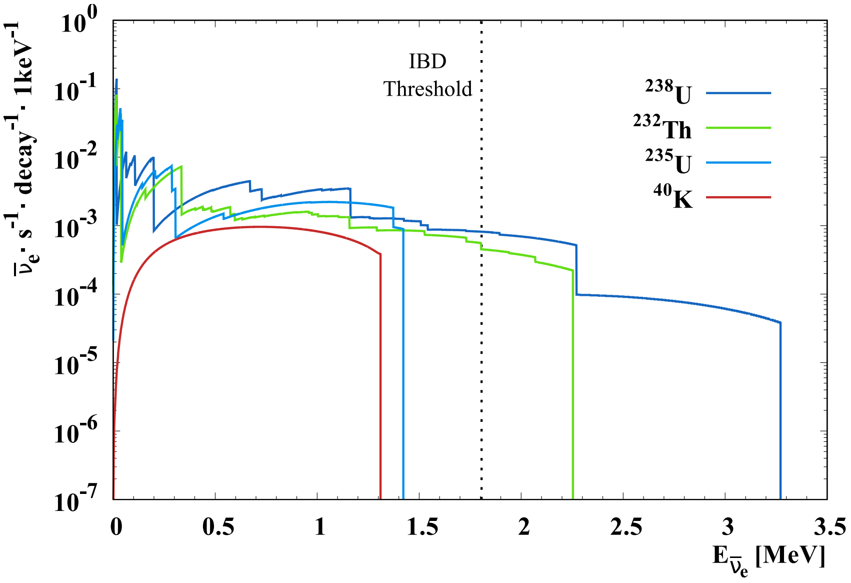

Figure 3.9: The energy spectrum of emitted antineutrinos in the ^{40}K, ^{238}U, ^{235}U, and ^{232}Th chains. This figure is from (17).

Chapter 4

The SNO+ Detector

The Sudbury Neutrino Observatory (SNO) is a solar neutrino experiment that took data from 1999 to 2006. Using heavy water as a target, the ^{8}B solar neutrino measurement performed by SNO culminated in the resolution of the solar neutrino problem by providing definitive evidence of matter-enhanced flavor oscillation. This effort ultimately led to the 2015 Nobel Prize in physics for Arthur McDonald along with Takaaki Kajita for Super-Kamiokande. The SNO+ experiment aims to build upon the infrastructure and success of SNO with the primary goal of measuring neutrinoless double beta decay. A full description of the SNO detector can be found in (79).

4.1 Detector Design

The SNO+ experiment re-uses the majority of the SNO infrastructure. The detector is shown schematically in Figure 4.1 and consists of an acrylic vessel (AV) 6 meters in radius surrounded by photomultiplier tubes (PMTs) located on a support structure (PSUP) that is 8.5 meters in radius. About 9400 inward facing PMTs view the target volume inside of the AV and about 100 outward-looking (OWL) PMTs are used as a veto to reject events that originate in the outer volume. The AV has a 7 m tall cylindrical neck that connects to the deck and allows for the deployment of calibration sources. Both the AV and PSUP are located in a large cavity that is filled with ultra-pure water (UPW). The entire detector is located at the 6800-foot level of the Creighton nickel mine in Sudbury, Canada, which provides a 6020 m (water equivalent) rock overburden protection from cosmic-rays. SNO+ has taken a phased approach to deploying the target medium inside of the AV.

First, the detector was filled with ultra-pure water and took data in this configuration from May, 2017 to July, 2019. The primary goal of the water phase was to calibrate the detector response, ensure stable running, and test the response of the new electronics and calibration systems. Using the water phase data, competitive limits on invisible nucleon decay were set (39) and a ^{8}B solar neutrinos measurement with very low backgrounds was performed (19). Additionally, using the AmBe calibration source, discussed in Section 4.6.2, a measurement of the neutron-proton capture cross-section was achieved (38). A search for reactor antineutrinos using the water phase data is presented in Chapter 6.

The second phase of SNO+ is a liquid scintillator phase, which replaces the water in the AV with a liquid scintillator called linear alkyl benzene (LAB). LAB was chosen due to its compatibility with acrylic, ease of handling, and excellent optical properties. A fluor called 2,5-Diphenyloxazole (PPO) is added to the LAB to increase the light output and to wavelength shift the emitted light into a region visible to the PMTs. There are several key advantages of liquid scintillator over water. First, the LAB+PPO has a high light yield relative to water (about two orders of magnitude more light in the PMT sensitivity region), drastically improving the energy and position resolution. Second, the intrinsic radioactive backgrounds of scintillator are lower than water and can be purified and cleaned using a variety of methods. This was demonstrated by the Borexino collaboration which achieved very low levels of internal radioactivity, in a large part due to their purification system (80).

Several upgrades were made necessary by the addition of liquid scintillator into the AV. The scintillator is considerably less dense than water and the entire acrylic vessel would be pushed upward in the cavity. A hold-down rope-net system was installed to counteract

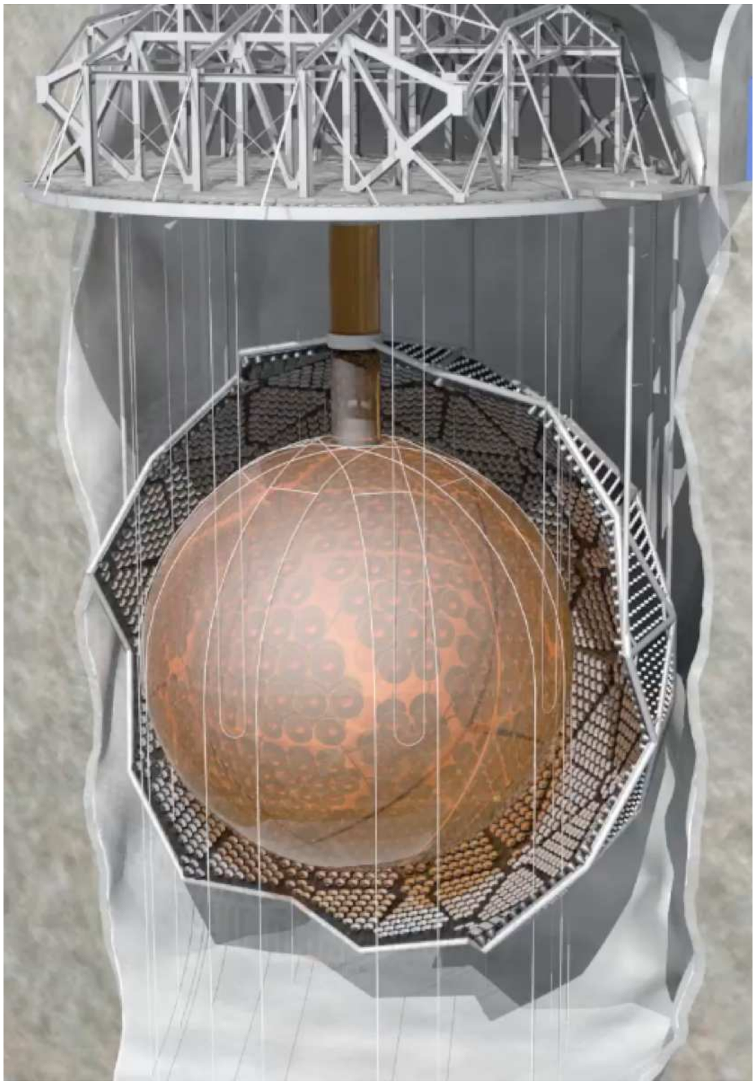

Figure 4.1: The SNO+ detector, showing the AV, PSUP and AV neck located in the cavity. Above the cavity is the deck, where the electronics and calibration systems are located.

the buoyant force in order to ensure that the AV does not float up into the deck. The installation of this hold-down ropes system required the removal of about 60 PMTs from the PSUP.

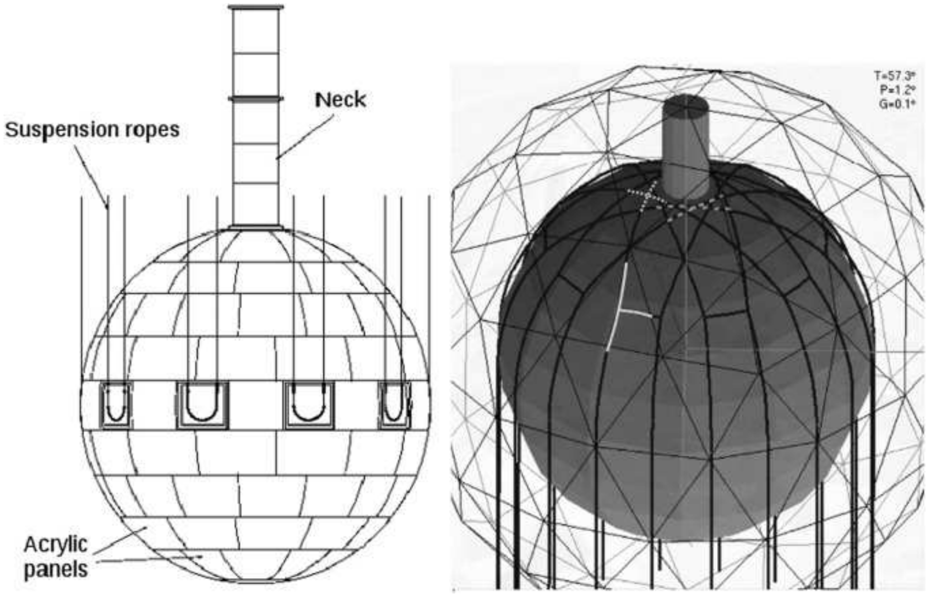

Figure 4.2: The SNO+ rope-net system showing both the hold-up ropes from SNO and the hold-down ropes installed for SNO+.

In order to purify the liquid scintillator, an underground scintillator purification plant was built. The plant is a specialized implementation of standard petrochemical purification processes and was built with strict cleanliness and vacuum-tight requirements. Additionally, a new sealed cover gas system, designed to reduce radon ingress into the detector through the interface at the neck, was installed.

The final phase of SNO+ requires the loading of tellurium into the liquid scintillator. Intense R&D has been conducted by the SNO+ collaboration in order to understand metal-loading of liquid scintillator, and two additional underground plants have been designed and built for the cleaning and loading of the Te. The final goal will be a 0.5% Te loading (by mass) into the liquid scintillator. The isotope ^{130}Te is known to undergo two-neutrino double beta decay and exists at about 34% natural abundance in Te. The search for neutrinoless double beta decay with the SNO+ detector is discussed in much greater detail in Section 8.1.

As of the writing of this thesis, the detector is currently filling with liquid scintillator and is about half full. The liquid scintillator is added directly on top of the water, which is drained out from underneath. The analysis presented in this thesis will use data taken in the water-filled detector; however, a large fraction of the livetime considered does contain data with a small amount of LAB in the neck of the detector.

4.2 Detector Principles

As discussed, the SNO+ detector will change the target volume inside the AV in phases. The first phase is a water phase, the second phase is a scintillator phase, and the final phase is the Te-loaded scintillator phase. The physics principles behind the water and scintillator targets are discussed in this section.

4.2.1 Water Target

While the primary intention behind filling the SNO+ detector with water was for calibration and testing the detector response, there is interesting physics that can be done in the water-filled detector. Indeed, there is a rich history of using enormous, deep underground detectors filled with water in neutrino physics. Super-Kamiokande is filled with about 50,000 tons of ultra pure water which is used as a target to detect solar and atmospheric neutrino interactions. This led to evidence of oscillation of atmospheric neutrinos as well as precision measurements of the ^{8}B solar neutrino flux (44, 81, 82). The SNO detector deployed a heavy water target in order to open additional interaction channels, first proposed by Herb Chen (83), that ultimately led to the most direct evidence of neutrino flavor transformation (13).

The detection mechanism is fairly simple in water. An incoming neutrino interacts with the target in the detector, producing charged particles in the final state. Those charged particles create Cherenkov radiation which is viewed by an array of photodetectors, most commonly PMTs. Cherenkov light is produced when a relativistic charged particle has speed greater than the speed of light in the medium, $\frac{c}{n}$. For electrons this corresponds to about 265 keV of kinetic energy. The Cherenkov light moves out spherically at the phase velocity of the medium, leading to a cone of emission in the direction of the particle motion. The half-angle of this cone is given by

$$\cos(\theta_c) = \frac{1}{n\beta} \tag{4.1}$$

where n is the index of refraction of the medium. For these relativistic particles in water, where $\beta \sim 1$ and $n = 1.33$, the Cherenkov angle is about 41°.

In a pure water detector, like the first phase of SNO+, the primary interaction mechanism for neutrinos with energy below 10 MeV is through neutral current (NC) and charged current (CC) elastic scattering (ES). The Feynman diagrams for these processes are shown in Figure 2.2. For antineutrinos ($\bar{\nu}$) from nuclear reactors, the Earth, or supernova, the primary interaction mechanism is through inverse beta decay (IBD), where the $\bar{\nu}_e$ scatters off of a proton, creating a positron and neutron. The Feynman diagram for this pure CC processes is shown in Figure 2.4. Relativistic electrons are also created by the decay of radioactive elements in and around the detector, either by direct β decay or by the creation of γ-rays, which can Compton scatter or pair produce. Neutrons are detected when they capture on hydrogen, produce a 2.2MeV de-excitation γ. Particles such as alphas from radioactive decays, protons below 1.4 GeV, and muons below 160 MeV are invisible in a water Cherenkov detector.

The number of Cherenkov photons produced by these charged particles is given by:

$$\frac{dN}{d\lambda dx} = \frac{2\pi\alpha z^2}{\lambda^2}\left(1 - \frac{1}{\beta^2 n(\lambda)^2}\right), \tag{4.2}$$

where N is the number of produced photons, λ is the wavelength of the emitted photon, α is the fine-structure constant, and z is the charge of the particle in units of the electron charge. The Cherenkov spectrum is proportional to $1/\lambda^2$ meaning that the majority of the light is emitted in the UV but with a significant fraction emitted in the visible between 350 to 600 nm. Because the UV light if often absorbed by the detector medium, the visible light is critical to detect with a high efficiency.

An important aspect of Cherenkov light is the directionality of the photons. This is contrasted in the next section in the discussion of the isotropic scintillation light. For the ES events, the differential cross-section is highly directional – the direction of the scattered electron is strongly correlated with the direction of the incoming neutrino. Figure 4.3 illustrates the directionality of the ES interaction. This is critical for the discussion of rejecting ^{8}B ES scatter events in the SNO+ double beta decay search in Chapter 8.

4.2.2 Scintillator Target

Liquid scintillator targets have additionally been an important tools in neutrino particle physics and have been used by numerous experiments. Interestingly, Cowan and Reines used a detector that includes tanks of both water and liquid scintillator in their neutrino experiment that led to the first detection of the neutrino (41). The LSND experiment detected neutrions from a beam at Los Alamos using a liquid scintillator, and reported results consistent with a hypothesized sterile neutrino (84). The KamLAND (31), Daya Bay (34), and RENO (35) experiments utilized liquid scintillator detectors to measure oscillation parameters by detecting antineutrino from nearby nuclear reactors. Borexino (85) deployed a liquid scintillator target to detect the ^{7}Be, ^{8}B, and pep solar neutrinos. Future detectors such as `THEIA` (86), Jinping (87), and JUNO (16) hope to deploy enormous volumes of liquid scintillators or water-based liquid scintillators (on the scale of 20 to 50 ktonnes) for future measurements of solar, supernova, and reactor neutrinos and potentially for measurements of $0\nu\beta\beta$ decay. The largest currently running liquid scintillator detector is NOvA, which

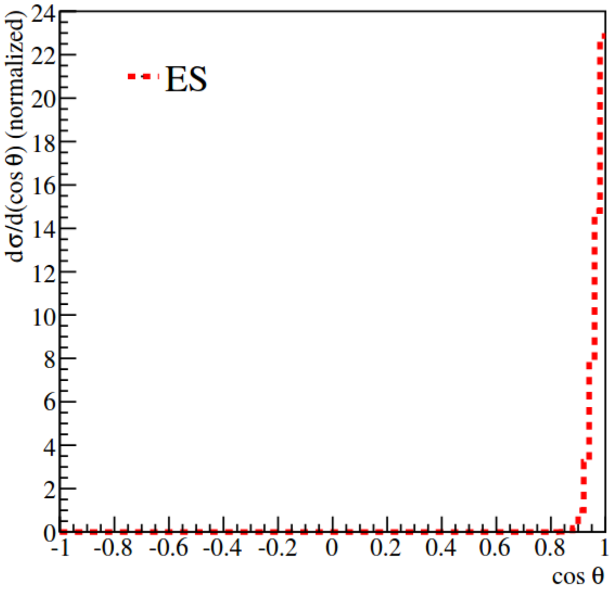

Figure 4.3: Normalized differential cross section as a function of the angle of the outgoing electron, relative to the incoming neutrino direction. The ES interaction is strongly peaked in the direction pointing away from the incoming neutrino direction. This plot is from S. Seibert's thesis (18).

uses 13 kt of liquid scintillator contained in plastic extrusions connected to wavelength shifter fibers (88). Recently, the PROSPECT experiment placed a detector with ^{6}Li-doped liquid scintillator within 10 m of a reactor core, primarily as a test of the sterile neutrino hypothesis as a solution to the LSND anomaly (36).

For SNO+, the primary advantages in moving from a water to an organic liquid scintillator detector is the high light yield and radiopurity of liquid scintillator. The high light yield is critical in order to achieve the required energy resolution necessary for performing a competitive measurement of $0\nu\beta\beta$ decay.

The scintillation processes itself is a sequence in which the ionizing radiation excites the electrons of the solvent into various excited states. As the molecule de-excites it fluoresces, producing relatively large amounts of UV or optical photons. Often a fluor or wavelength shifter is dissolved into the scintillator, which can accept energy from the solvent both radiatively and non-radiatively and increase the total light output.

An empirical formula called Birk's law gives the light yield per path length as a function of the energy loss per path length:

$$\frac{dL}{dx} = S\frac{\frac{dE}{dx}}{1 + k_B\frac{dE}{dx}}, \tag{4.3}$$

where L is the total number of produced photons, S is the scintillation efficiency (units of photons per energy), and k_B is Birk's constant, which depends on the material. The deviation from linearity is due to quenching effects, captured in the non-zero value of k_B (89). The scintillation efficiency and Birk's constant has been measured for LAB+PPO, and the total light output is expected to be more than 10000 photons/MeV.

In addition to the high light yield, the intrinsic levels of radioactivity in liquid scintillator can be reduced to extremely low levels. Borexino has demonstrated chemical techniques for distilling, filtering, and applying nitrogen stripping to liquid scintillator, which reduces the uranium and thorium chain contents to about two orders of magnitude lower than SNO levels (80). This excellent background reduction enabled Borexino's low energy solar neutrino program. The SNO+ scintillator plant was designed and constructed with many of the same goals and techniques as Borexino.

Charged particles such as αs and protons are visible in a scintillation detector, which opens new detection channels and adds possible backgrounds. For example α-decays from ^{210}Po, ^{212}Po, and ^{214}Po provide a source of backgrounds for the low energy solar neutrino and $0\nu\beta\beta$ analyses. Additionally, particles such as αs and βs are quenched differently, meaning both the emission-time distribution and the expected number of produced photons per MeV is different between the two particles. For example, in LAB+PPO a 10 MeV α interaction looks (very roughly) like a 1 MeV β in terms of total light output.

The scintillation light emission occurs over a much longer time period than Cherenkov light. The liquid scintillator can be excited to both singlet and triplet states and the de-excitation time depends on the state. The singlet state consists of fast de-excitation on the order of several nanoseconds, while the triple state can take hundreds of nanoseconds to de-excite. Notably, different particles excite the liquid scintillator molecules into different fractions of singlet and triplet states. For example, α particle interactions excite the scintillator into relatively more triplet states than β interactions, causing the emission timing to have a longer tail for αs than βs. This difference can be used to reject α backgrounds. A measurement of the scintillation decay time profiles for the SNO+ liquid scintillator is given in Section 7.2.3.

Similarly to water, the mechanism for detecting solar neutrinos in a liquid scintillator detector is through elastic scattering. However, direction reconstruction in a liquid scintillator detector is difficult, and solar neutrino interactions are distinguished using their energy spectrum. For the SNO+ $0\nu\beta\beta$ search, elastic scattering of electrons from solar neutrinos is a dominant background, discussed in Section 8.2.

4.3 Photomultiplier Tubes

The SNO+ detector deploys approximately 9400 R1408 Hamamatsu PMTs, the same as those used in the SNO detector. These PMTs were originally selected for their relatively low dark rate (about 500 Hz at 16° C), reasonable charge resolution, and relatively good timing. Over the last two decades, new and improved large-area PMTs have been designed. Bench-top measurements of several of these PMTs is discussed in Section 7.1. In the interim, between the end of SNO and the beginning of data-taking with SNO+, PMTs that were no longer operating at the end of SNO were removed, had their bases replaced or fixed, and then were introduced back into the detector. This effort regained several hundred working PMTs for SNO+.

The PMT base, which distributes the high voltage to the dynode stack, must be protected from the environment in order to operate. Waterproof housing was designed using soft-silicon gel and the R1408 PMT with this housing is shown in Figure 4.4 (left). The PMTs are placed into 27 cm diameter light concentrators, imaged in Figure 4.5. The light concentrators consist of reflective petals arranged in the shape of a Winston-cone, an ideal non-imaging light collector (90). Over the lifetime of SNO and SNO+ the concentrator petals have aged and dulled, reducing the overall reflectivity. During the SNO data-taking the photocathode coverage, including the reflectivity of the concentrators was 54%. The SNO+ coverage is slightly worse due to decreases in the reflectivity of the concentrator petals.

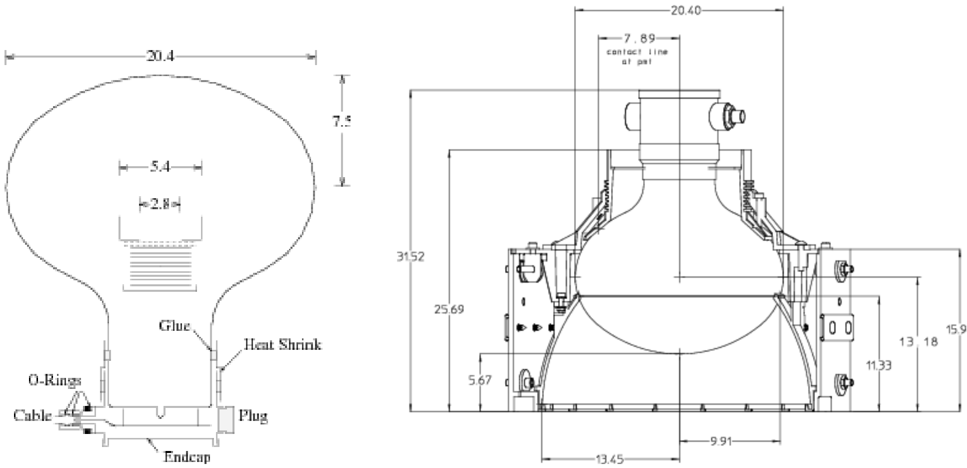

Figure 4.4: A schematic of the R1408 including the waterproof housing (left). The hexcell design to hold the PMTs and concentrators in the PSUP (right).

In addition to the standard inward looking PMT channels, there are about 50 PMTs with the last dynode tapped to read out a second, low-gain signal. There are also 5 PMTs without concentrator petals, intentionally removed to directly test the effectiveness off the reflective petals. There are four non-R1408, high quantum efficiency PMTs installed in

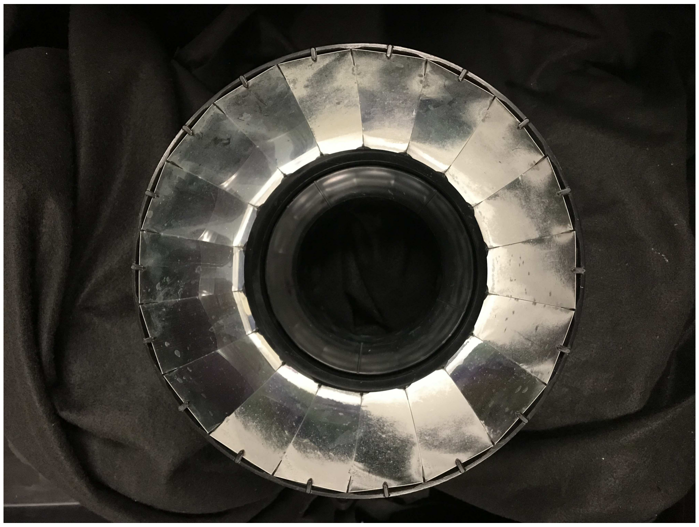

Figure 4.5: A SNO concentrators imaged in 2019 with considerable aging to the petals clearly visible.

the detector, discussed in Section 7.1.5. The 100 OWL PMTs are located along the PSUP pointing outward and are not contained within concentrators. There are four additional neck PMTs located at the top of the neck, which are used to reject events that occur in the neck of the detector.

4.4 Electronics and DAQ

The electronics chain designed to read out and trigger on signals from the PMT array make up the data acquisition (DAQ) system. The SNO+ DAQ inherits much of its infrastructure from SNO, with several notable changes. The electronics for SNO were designed for sub-nanosecond time-precision with a large dynamic range, but only for relatively few expected PMT signals per 400 ns event, at a reasonably low rate of tens of Hz. In order to deal with the higher occupancy and trigger rates expected with a liquid scintillator target, several of the electronics board required upgrading. These upgrades are described in Section 4.5.

The SNO+ electronics system consists of 19 crates each containing 16 front-end cards (FEC) and 16 PMT interface cards (PMTICS). On each FEC, four daughterboards (DBs) are attached. The DBs house custom integrated circuits for eight PMT channels that perform the important signal processing. Thus, each crate is responsible for the analog processing and trigger generation for 512 PMT signals.

The PMTs in each crate are powered by a single high voltage box which supplies about 60 mA of current across the 512 channels. The high voltage supply connects directly into the SNO+ backplane, which distributes the high voltage to the PMTICs. 32 75Ω waterproof RG51-like coaxial cables connect the PMTIC to the PMTs. These cables travel 32 meters down to the PMT carrying the high voltage to the PMT base, where it is distributed to the dynode stack. The typical high voltage supplied to the PMTs is 2000 V, but it varies from 1800 V to 2350 V across the crates.

The analog signal from the PMT travels up the same cables back to the PMTICs where a HV-blocking capacitor picks off the DC component. The PMTIC also supplies other important features, such as disconnect relays that allow one to disconnect eight sets of channels at a time, and overvoltage and breakdown protection diodes. Additionally, the PMTIC has flexible feed-through resistors that allow tunable high voltage adjustments for each PMT. The high voltage set at the supply and the value of these feed resistors is tuned such that the gain of the PMTs on the crate is close to 1×10^7. The mapping from high voltage to gain for each PMT was measured in dedicated setup at Queen's University which characterized more than 9,800 PMTs for SNO (91).

The signal from the PMT is passed directly from the PMTIC to the FECs on a connector that bypasses the backplane. On the FECs, the signals are distributed to one of the four DBs, depending on the location of the PMT connection. On the DBs the signals are passed into a four-channel custom discriminator (SNOD), where they are each compared against programmable thresholds. Typically these ‘channel thresholds’ are set to correspond to one-quarter of the peak voltage of an average SPE signal. For channels that cross threshold a PMT ‘hit’ is created. The total number of PMT hits in a triggered event is often referred to as the ‘nhit’ of the event.

Each PMT signal crossing threshold is passed to a custom integrator (SNOI) where the signal is integrated across three paths to produce three charge values. These charge values are high gain, long integration (QHL), high gain short integration (QHS), and low gain short integration (QLX). The SNOI also shapes the input PMT signal and outputs two versions with different gains to be used in triggering: ESUMH and ESUML. Simultaneously, on a custom CMOS chip a TAC is started on the leading edge of the signal from the SNOD. The TAC is stopped by a centrally created global trigger, discussed later. The CMOS chip also creates two square fixed-current pulses used for triggering, one of length about 100 ns (N100) and the other with length about 20 ns (N20). The N100 and N20 can be disabled at the channel level for broken channels. The CMOS chip also provides sixteen analog memory cells per measurement for storage.

The time from the TAC, the three charge values, and the crate, FEC number (card number), channel number, and cell number (CCCC), and some error flags (discussed in Appendix C.0.3) are the entirety of the analog-processing information stored by the SNO electronics. A direct map of the CCCC to PMT location is stored in a database and used to physically locate the origin of the signal.

The N100, N20, and ESUM signals are the primary trigger signals for the detector. There exist separate trigger signals for the OWLs that are separate but handled equivalently. Each trigger signal generated on the DBs is passed onto the FECs, where the bottom and top 16 channels are summed and passed onto the backplane. The backplane provides a path to the crate trigger cards (CTCs) where each trigger signal is summed separately for the entire crate. That is, in each crate there is a crate-level sum of N100, N20, ESUMH, and ESUML.

The signal for each trigger for each of the 19 CTCs is passed to the master trigger cards – analog (MTC/A+s), located in a central timing rack. Each MTC/A+ is primarily responsible for a detector-wide analog sum of the trigger signal with three different gains, each of which is compares against a programmable threshold. For trigger sums that cross threshold, a 20-ns digital signal called the raw trigger is created. Details about the MTC/A+ are described in Section 4.5.2. There are a total of seven MTC/A+s in the timing rack, one for each trigger signal. In principle all of the MTC/A+s are identical.

Only a small fraction of the 21 total trigger signals (seven MTC/A+s with three gain paths) are used. The most important trigger signal for triggering on physics events in the detector during the SNO+ water phase is the N100. This is largely due to the size of the detector – it takes about 60 ns for light to transit the entire detector, so the N100 signal allows all unreflected Cherenkov photons from an energy deposition to be counted toward a possible trigger. Of the three N100 signals, the high gain path is referred to as N100L ('L' is for low-threshold) is the most important and runs with the lowest threshold due to the better signal-to-noise on this signal compared to the lower gain paths. The other gain paths allow for a large dynamic range, but are not as useful in Cherenkov detector where the total number of hits in an event is fairly small. The medium gain path N100 is referred to as N100M and the low gain path is referred to as N100H. The N20 has a similar naming scheme, and we only use the high gain ESUMH signal during the water phase.

The raw trigger out from the MTC/A+ is created for each individual gain path that crosses threshold. These raw trigger signals are sent to the master trigger card – digital (MTC/D). The MTC/D itself has a trigger mask that decided which raw trigger signals to accept as a valid trigger signal. For triggers that are masked in, the MTC/D creates a global trigger (GT), which effectively ends the trigger cycle and indicates that the detector should read out its data. The GT is fanned out to all crates in the detector, which in turn pass the GT down the backplane to every individual channel. Simultaneous with the GT, the MTC/D also creates a signal called lockout (LO), a 420 ns signal during which the MTC/D will not emit another global trigger. This effectively creates a 420 ns trigger window in which to readout the PMT hits. Additionally, the MTC/D keeps track of the trigger type that causes the GT, as well as any other triggers that arrives within roughly a 20 ns window around the GT, in a trigger word. A counter called the global trigger ID (GTID) is incremented at the MTC/D for each generated GT. The GTID is used to make sure the entire system remains in sync. At the crate side, the GT is distributed to each of the channels, which stops the TAC ramps and indicates to each channel that it should read out its data. The channels hold their data waiting for a GT for a length of time known as GTValid. If a GT does not arrive before GTValid expires, the channels charge and TAC values are cleared. The channels independently keep track of the GTID, which is checked against the master trigger system every 2^{16} GTIDs to ensure the channels are in-sync.

The channels that have valid hits are buffered into the FEC memory. The data is then read out from each FEC by the XL3s. Each individual crate has its own XL3, which is an upgraded board for SNO+ (described in Section 4.5.1). The XL3 then ships the data for the crate to a central DAQ computer over Ethernet. The central DAQ computer runs several servers that handle the data-readout and talk to the hardware. The `xl3-server` and `mtc-server` handle sending commands to and receiving replies from the XL3 and MTC respectively. The data-server receives PMT data from the XL3s and trigger data from the MTC/D and relays that data to any subscribed clients. The most important client is the SNO+ event builder, often referred to simply as ‘the builder’. The builder uses the GTID stored in the PMT data and the trigger data to combine the PMT and trigger data into an ‘event’. Once the builder is satisfied that all PMT hits have been saved to the correct event, the builder writes the data to disk in a custom ‘zdab’ format. Data is written to the same file until a maximum file size is reached, upon which the builder will create a new file. The files are indexed by runs, which are typically set to one hour for physics running, and are used to conveniently chunk the data together. The runs themselves have important information associated with them, such as the run-number and the run-type (marking the type of running such as physics, calibration, etc.) and any common activity occurring during the run (such as recirculation, people on deck, etc.).

The absolute time of the event and relative time between events are recorded by two separate clocks. A 10 MHz clock is distributed by a commercial GPS system located on surface. This 10 MHz clock is connected to the detector via a 4 km long fiber optics cable, and an underground board called the UG Board provides the signal relay to the MTC/D. The 10 MHz clock provides a 'universal' time and is synced to the GPS once per hour. A 100 MHz clock runs on the trigger utility board mk.II (TUBii, described in Section 4.5.3) and is plugged directly into the MTC/D, where it is converted to 50 MHz. This clock is used for inter-event timing and provides 20 ns resolution.

An auxiliary commercial CAEN v1720 waveform digitizer runs in the same crate as the MTC/D and is used to digitize the summed trigger signals, discussed in 4.5.4. The raw triggers from the MTC/A+ are plugged in to the MTC/D through a board, called the Latch board, described in Section 4.5.5. Figure 4.6 shows a detailed schematic of the SNO+ trigger system. The entire path from PMT to the MTC/D is shown with the new or upgraded boards from SNO to SNO+ highlighted in red.

The MTC/D provides a pulser than can produce forced global triggers at a fixed rate. This pulsed global trigger (PGT) provides a zero bias trigger that is used for a measurements of the dark noise rate (described in Section B.0.4) and CMOS dropout rate (described in Section 4.7. In combination with the PGT, the MTC/D can emit a pedestal signal to all crates in the crate pedestal mask (set on the MTC/D) that tells the crates to read out the data for any channels in the channel pedestal mask (set by the XL3). This bypasses the discriminator and forces channels to read out their data. The pedestal signals are used for eectronics calibrations (such as the ECAL and ECA, described in Appendix B) and by the nhit monitor for determination of the trigger efficiency.

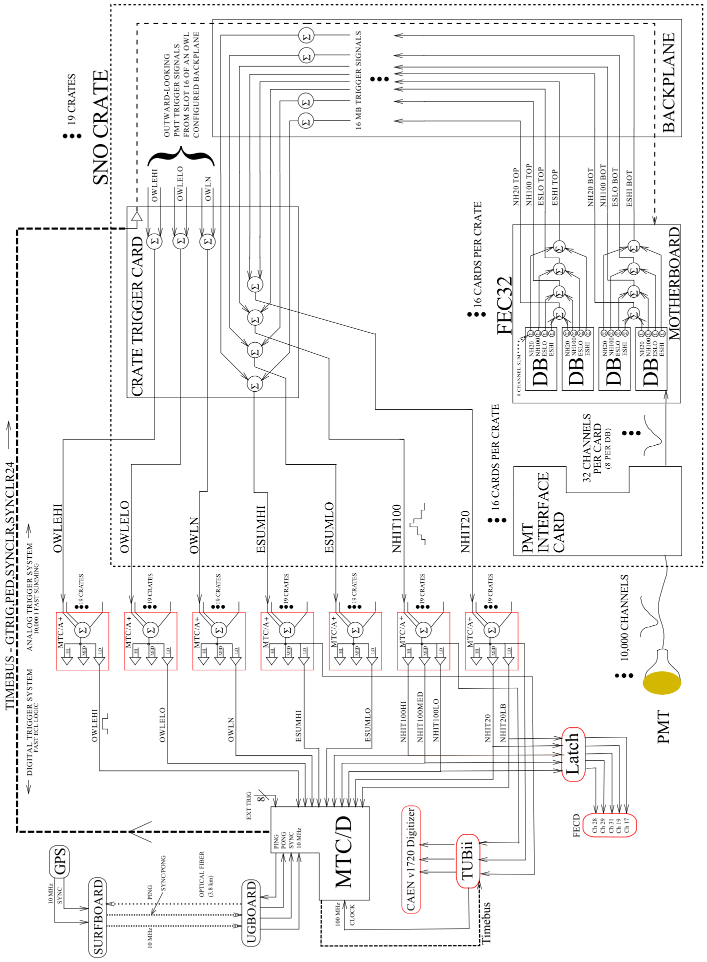
SNO CRATE
19 CRATES
CRATE TRIGGER CARD
OUTWARD-LOOKING PMT TRIGGER SIGNALS FROM SLOT 16 OF AN OWL CONFIGURED BACKPLANE
OWLEHI
OWLELO
OWLN
16 MB TRIGGER SIGNALS
BACKPLANE
NH20 TOP
NH100 TOP
ESLO TOP
ESHI TOP
NH20 BOT
NH100 BOT
ESLO BOT
ESHI BOT
16 CARDS PER CRATE
FEC32
DB
MOTHERBOARD
32 CHANNELS PER CARD (8 PER DB)
16 CARDS PER CRATE
PMT INTERFACE CARD
10,000 CHANNELS
PMT
OWLEHI
OWLELO
OWLN
ESUMHI
ESUMLO
NHIT100
NHIT20
MTC/A+
TIMEBUS - GTRIG,PED,SYNCLR,SYNCLR24
ANALOG TRIGGER SYSTEM
DIGITAL TRIGGER SYSTEM
NHIT100HI
NHIT100MED
NHIT100LO
NHIT20
NHIT20LB
Latch
FECD
Ch 28
Ch 29
Ch 31
Ch 19
Ch 17
EXT TRIG
8
MTC/D
PING
PONG
SYNC
10 MHz
CAEN v1720 Digitizer
TUBii
100 MHz CLOCK
Timebus
GPS
10 MHz
SYNC
SURFBOARD
PING
SYNC/PONG
OPTICAL FIBER (3.8 km)
10 MHz
UGBOARD

Figure 4.6: Schematic of the SNO+ trigger system. Board upgrades or additions from SNO are highlighted in red.

4.5 Electronic Upgrades

Several electronics boards were added or upgraded to the SNO electronics chain. As mentioned, these electronics are required due to the higher trigger rates and average PMT occupancy expected by using a liquid scintillator target.

4.5.1 XL3

The data-readout system for SNO consisted of the XL1s and XL2s, which provided centralized readout of the crates, but was limited to communicate and receive data from one crate at a time. These boards were replaced with the XL3s, which allowed for asynchronous, parallel read out of the crates, drastically increasing the possible read out speed. The bandwidth limit of the SNO electronics was very roughly 250 kB/s, which is easily achievable with the XL3s. An XL3 is located in each crate to provide the read out of the front-end electronics, independent of any central computer. The XL3 hosts a Xilinx ML403 which acts as a carrier board for a Virtex-4 FPGA. The ML403 also contains a PowerPC processor, providing optimal flexibility. The FEC readout is run from the firmware on the Virtex-4 FPGA, which is configured to pull data across the backplane. C code running on the PowerPC processor requests the data from the FPGA and buffers it in local memory. The data is then pushed (as opposed to requested) by the XL3s to a central DAQ computer over ethernet using standard TCP/IP protocols. This upgrade from the VME-like protocol used by SNO also largely increased the bandwidth of the system. The lightweight IP libararies run on the PowerPC and allow the XL3 to push data over Ethernet at a maximum rate of around 14 MB/s, corresponding to a detector wide rate of 266 MB/s. During the water phase, the higher bandwidth has meant SNO+ has been able to trigger at significantly higher trigger rates with thresholds much lower than SNO ever achieved.

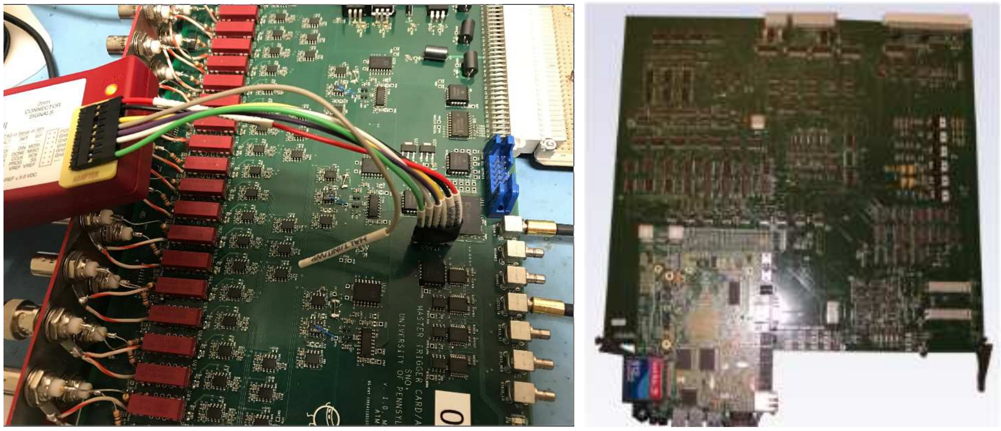

(a) The MTC/A+ trigger board, developed primarily by A. Mastbaum.

(b) The XL3 read out board, developed primarily by R. Bonventre.

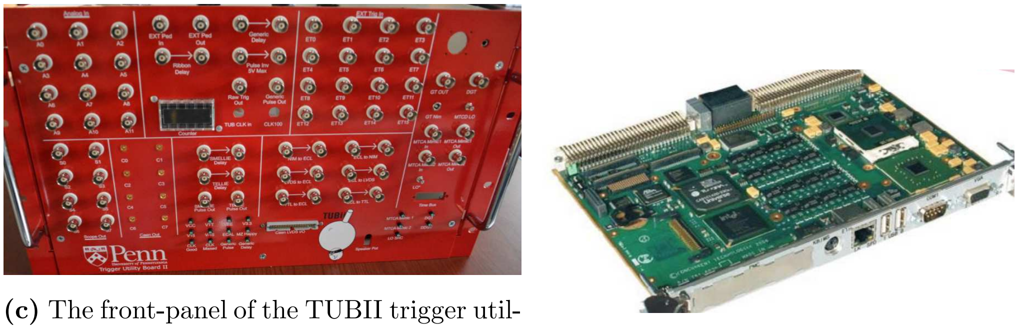

(c) The front-panel of the TUBII trigger utility board, developed primarily by E. Marzec.

(d) The CAEN v1720 waveform digitizer.

Figure 4.7: The family of new electronics for the SNO+ detector. Not shown here is the Latch board developed primarily by T. LaTorre.

4.5.2 MTC/A+

For SNO+ the MTC/A, which performed the detector wide trigger sum for SNO, is replaced with the MTC/A+. Similarly to the other upgrades, the motivation for the development for the MTC/A+ is the large number of expected PMT hits. The MTC/A has a dynamic range of about 1000 PMT hits, which might not be sufficient for the expect light yield. The MTC/A+ provides different gains such that SNO+ can use the full 10000 hit range for triggering. Additionally, the MTC/A+ provides better than 1 nhit resolution for the low-threshold water data-taking, which has been critical for achieving our very low trigger thresholds discussed in Section 4.7.

Three primary additional upgrades are the baseline restoration, crate disconnects, and programmable trigger logic. Baseline restoration on the MTC/A+ is an active LRC feedback loop that restores the baseline to the nominal value with a time constant on the order of tens of seconds. This alleviate issues associated with trigger baseline shifts that can be caused by variations in the otherwise DC-coupled analog sum baseline. The slow time constant does not integrate the fast signals from the PMTs.

The crate disconnects are a set of controllable relays that allows one or more crates to be disconnected from the trigger sum. This allows fast debugging of issues with the trigger signal by identifying the problem crate. The control is critical for running the detector during periods of time when the trigger signal from one or more crates is not working.

The programmable trigger logic on the MTC/A+ is implemented in firmware running on a Xilinx XC2C512 CPLD. This allows for the trigger logic to be reprogrammed as desired. The input into the CPLD is the raw triggers for each of the three gains paths, and the CPLD can perform logical operations with those signals. The form of logic currently implemented on the CPLD has two different types of triggering. The standard trigger logic is a raw trigger is emitted for the corresponding gain path if the analog sum crosses threshold. Additionally, a raw trigger is emitted for a gain path if at the end of the LO window the analog sum is still above threshold. This is often referred to as retriggering, as a global trigger is created (if the trigger is masked in at the MTC/D) immediately after the previously triggered event.

There are several other options for triggering with the MTC/A+. A forced retriggering scheme would fire raw triggers for N sequential window after a trigger threshold crossing.

This can be gain path specific and the number of raw triggers emitted can be programmed to any value. Forced retriggers have not been used during normal data-taking, but might be useful for background rejection in the scintillator phase or for performing specific measurements of the detector response, such as PMT afterpulsing.

A delayed copy of the global trigger (DGT) is created on TUBII and cabled back to the CPLD, where it can used in the trigger logic. Potential logic using this signal includes providing a retrigger if the signal is still high at DGT (or crosses threshold again at DGT), which might be half the length of LO. This could be useful for triggering specifically on β-α coincidence events during the scintillator phase. Finally, a gated low energy trigger could be implemented on the CPLD, where a path with a high threshold enables a path with a low threshold for some amount of time if it crosses threshold.

4.5.3 TUBII

The Trigger Utility Board Mk.II (TUBii) replaces the SNO trigger utility board (TUB) and adds many additional features. Besides providing for additional triggering logic, TUBii performs the following roles:

1. Analog shaping of the trigger signals sent to the CAEN
2. Removal of the -5V baseline from the MTC/A+ for the trigger signals sent to the CAEN and control-room oscilloscope (used for monitoring purposes)
3. Production of the DGT signal discussed in Section 4.5.2
4. Interface for the calibration systems (including the laserball, TELLIE, and SMELLIE) in which it provides synchronization with the trigger system and tunable pulsers and delays
5. A variety of different logic converters (eg, ECL to TTL)

6. A measurement of the MTC/A+ baseline values used for monitoring
7. 'MTC/A+ Mimic' circuity, providing functionality very similar to the trigger implemented in the CPLD on the MTC/A+s
8. Audio output of the global trigger rate, used for monitoring
9. Flexible trigger logic using a Xilinx MicroZed, which carries and FPGA

The triggering logic on TUBII is designed to use as input any of the raw trigger signals from the MTC/A+. Nearly any logical combination of the trigger signals can be used to create a global trigger for the detector. For example, TUBII can take as input an N20 raw trigger, where the threshold has been set relatively low, and an N100 raw trigger, where the threshold has been set relatively high. This trigger was implemented as a method for detecting γ-rays in the external volume during scintillator filling, where the low energy decays of ^{14}C in the scintillator make it difficult to measure the events in the external water volume. Due to the extended scintillation time-profile, the ^{14}C events will trigger both the N20 and N100 thresholds, while the external γ-rays often only cause the N20 to trigger in this scheme. Logic on TUBII is implemented to provide a global trigger in situations when the N20 raw trigger is received and the N100 raw trigger is not received, within some time window around the N20. This provides a sample of triggered events that should have a high purity of external events for measurement of the external backgrounds. There are many other possibilities for triggering on TUBII, which provides more flexibility than the firmware on the CPLD on the MTC/A+. Supernova burst triggers and BiPo coincidence triggers are among those that are often discussed.

4.5.4 CAEN

The CAEN v1720 is an 8 channel, 12-bit, 250 MS/sec waveform digitizer with a dynamic range of 2 V. The CAEN digitizer replaces the SNO analog measurement board (AMB), which provided analog signal processing of the ESUM trigger signal. The output for the high gain path of the N100, N20, and ESUMH are plugged into a CAEN (the path goes through TUBII). The CAEN digitization window is set to 420 ns, to be consistent with the width of LO, with a sample width of 4 ns. The digitized waveforms are read out as part of the trigger event and the CAEN keeps track of its own GTID to ensure that it stays in sync. These digitized waveforms are primarily used for data-cleaning and for monitoring the health of the trigger signals.

4.5.5 Latch Board

The latch board provides a trigger fanout for the MTC/A+ which is critical functionality in the determination of the trigger efficiency, described in Section 4.7. The motivation for building the latch board is the limitation in the information provided in the trigger word of the MTC/D. The trigger word on the MTC/D is generated such that only raw triggers that arrive roughly 10 ns around the raw trigger that caused the global trigger are latched into the word.

The latch board takes as input the raw triggers from the N100L, N100M, N100H, N20L, and N20M. It outputs a copy of the raw trigger to be used by the MTC/D for detector triggering. Additionally, the latch board outputs a negative-gong 5V pulse to a calibration FEC called the FECD that is located in crate 17, slot 15. The FECD is a spare FEC with no PMTs attached that provides additional calibration channels. The FECD is read out by the XL3 like a normal FEC for calibration sources, such as the ^{16}N trigger PMT. In this case it is used for the latch board to mark whether each of the raw triggers from the MTC/A+ fired.

4.6 Calibration Sources

4.6.1 ^{16}N

The ^{16}N calibration source is the primary high energy calibration source for the SNO+ water phase. The ^{16}N is produced via (n, p) reaction on ^{16}O using 14 MeV neutrons from a Deuterium-Tritium generator on deck. The ^{16}N is transported from the deck into the source deployed in the detector using a high pressured gas-stream down the umbilical. In the source chamber the ^{16}N beta-decays with a half-life eof 7.13 seconds to ^{16}O, which de-excites and produces a primarily 6.13 MeV γ. The β interacts with a small piece of plastic scintillator inside the source, which produces scintillation light that is detected by a PMT located in the neck of the source. The signal from that neck PMT is used to tag the event. The γ leaves the source and interacts in the detector. This interaction is used as to calibrate the total detection efficiency and the energy response of the detector. Full details of the ^{16}N source can be found in (92) and (93).

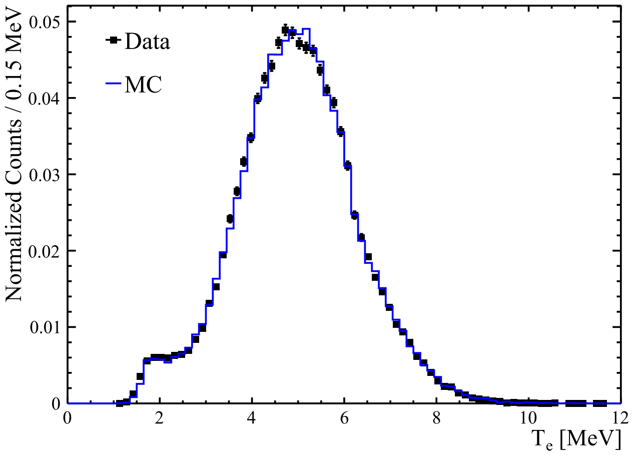

Figure 4.8: A comparison between the ^{16}N data and MC for a central ^{16}N deployment after calibrating the detection efficiency of the detector. This figure is from (19).

4.6.2 AmBe

The Americium-Beryllium (AmBe) calibration source is a power source produced in 2005 and stored underground. The ^{241}Am nuclei undergo α-decay with a half-life of 432 years and a ^{9}Be target absorbs the emitted αs, producing ^{12}C and a neutron. The ^{12}C nucleus is produced in an excited state about 60% of the time, which de-excites via a 4.4 MeV γ. The majority of the neutrons leave the source and thermalize and capture on hydrogen in the detector. The capture produces a 2.2 MeV γ with the production of ^{2}H. The AmBe calibration source provides a low energy calibration of the reconstruction and is an ideal source for the antineutrino analysis described in Chapter 6. Additionally, the AmBe source allows for a measurement of the neutron detection efficiency, which is found to be around 50% with a small variation across the detector (38). An analysis of the AmBe source data is described in Section 6.6 and more details about the source can be found in (93).

4.7 Trigger Efficiency

The trigger efficiency in the SNO+ detector is the probability that the detector triggers given a number of 'in-time' PMT hits. The trigger signals have been described in Section 4.4. The primary trigger for the SNO+ detector for the majority of the data-taking during the water phase is the N100L. This is only not true for runs before run 104613. These runs will not be used in the analysis presented in this thesis due to the complicated behavior of the trigger that was not well-understood. A detailed look at the trigger efficiency for these runs is given in (94) and is not included here.

As already discussed, the SNO+ trigger system discriminates on a detector-wide analog sum of the trigger signal. For the water phase, the N100 – specifically the high-gain path on the N100 MTC/A+ called the N100L – has the lowest trigger threshold and is considered the 'primary' trigger for the detector. The trigger threshold is set to correspond to about 7 PMT hits. However, given the various inefficiencies, noise in the system, and several other trigger features that will be discussed, the efficiency for triggering at 7 hits is not 100%. Indeed, there is a non-zero efficiency for triggering the detector down to about 5 hits and does not reach 100% until about 9 hits.

Whether or not the detector triggers is effectively a function of the number of in-time hits, denoted $\tilde{n}_{100}$. Because the trigger signals are 100 ns in length, the definition of in-time hits is the number of PMT pulses that cross their respective discriminator threshold within 100 ns of one another and have their N100 trigger signals enabled. Because of the rise-time of the square pulses this is not precisely true because each individual pulse does not reach its maximum voltage until several ns after it is created. The rise-time comes both from the creation of the pulse at the CMOS chip, but also from the bandwidth of the MTC/A+.

Perhaps most importantly, the trigger response is affected by an error in the design of the CMOS chip, called ‘dropout’. In some situations, the N100 and N20 for a channel latches up, without going back to baseline, effectively lowering the trigger threshold by one hit. This lasts until that channel gets a real PMT hit crossing the discriminator threshold, so is a more dangerous issue for channels that are running at low rates. The total amount of dropout can be measured by looking at the trigger sum on the CAEN. In situations where one or more channels are dropped out the CAEN baseline should shift by the corresponding amount. A plot showing the CAEN baseline for the N100 trace (in ADC counts) from E. Marzec’s thesis (20) is shown in Figure 4.9. In that plot, a fit composed of the sum of several Gaussian distributions indicates that the CAEN traces are composed of several different baselines, corresponding to different amounts of dropout. This run had a particularly large amount of dropout and there are components that show seven channels dropped out in the fit to the baseline.

The trigger efficiency is monitored on an approximately hourly basis using a method called the nhit monitor. The nhit monitor sends pedestals to a single crate at a rate of 1 kHz, and then, one at a time, enables channels in the pedestal mask. The nhit monitor starts with a single channel on the crate receiving pedestals and increments to 50 channels.

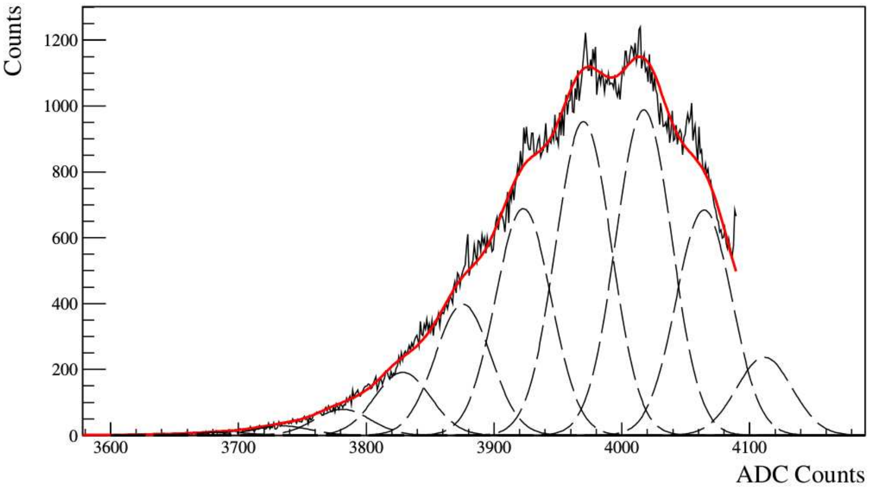

Figure 4.9: The CAEN baseline in ADC counts for a run with a lot of dropout. The Gaussian components of the fit correspond to one or more channels dropped out, shifting the CAEN baseline down. The fit shows the average amount of dropout at a given time is around two to three channels. Note that the CAEN baseline has been set to around 4050 counts with no dropout. This plot is from E. Marzec's thesis (20).

At each step 1000 pedestals are sent and, for each pedestal event, it is checked whether the N100L triggered the detector. The pedestal hits are always exactly in-time and the number of channels firing are known exactly. This allows one to extract the trigger efficiency curve by using the total number of triggered events divided by the total number of pedestal events for a giving number of in-time hits. The trigger efficiency can also be measured using laserball or TELLIE data, as will be discussed later. Note that given limitations in the latching of the trigger word of the MTC/D, the N100L, even if the threshold is crossed on the MTC/A+, might not be included in the trigger word. This is mitigated by the latch board, discussed in Section 4.5.5, which splits the raw trigger out from the MTC/A+ to the FECD, where a hit on a specific channel corresponds to the N100L firing in that event. This method is robust to the relative timing of the triggers and reliably marks whether each the various trigger threshold were crossed.

Understanding the trigger efficiency curve is critically important if events are being analyzed that have nhits in the range where the efficiency is not 0 or 100%. For the solar neutrino and nucleon decay analyses (19, 39) the energy cut was such that the events were 100% certain to trigger the detector and thus we did not need to worry about the trigger efficiency. However, for the reactor antineutrino analysis, the low energy neutron capture events will be affected by inefficiencies in the trigger. Thus, understanding our modeling of the trigger around threshold is critically important for these analyses. In (38) the AmBe source was used to understand the trigger efficiency for 2.2 MeV γs produced from neutron captures on hydrogen.

The trigger model used in `RAT` uses a detailed model of the entire trigger system and includes in-situ measurements of the rise-time of the trigger signals, noise in the trigger system, variation in the size of the nhit pulses, and the width of the nhit pulses. The model of the nhit pulse in `RAT` is compared to data taken on the teststand in Figure 4.10. Most importantly, the model of the dropout discussed is used as direct input into `RAT` and the baseline of the trigger. For each run the average amount of dropout, R, is measured and `RAT` draws from a Poisson distribution with that average to determine the dropout for each event. Additionally, `RAT` includes a model of the baseline restoration by subtracting the average dropout, such that the trigger baseline (B) for any event is given by:

$$B = \text{Poisson}(R) - R. \tag{4.4}$$

In order to test the `RAT` model, simulated 2.2 MeV γs are used to probe the relevant nhit range. Because the amount of dropout is measured run-by-run, the simulation is done using a specific run-number. Many runs are tested with various amount of dropout but only two runs are presented here. In Figure 4.11 the trigger efficiency distribution is compared between the nhit monitor (labeled as data in the plots) and the output of the `RAT` simulation. The calculation of $\tilde{n}_{100}$ is done using the trigger efficiency processor in the

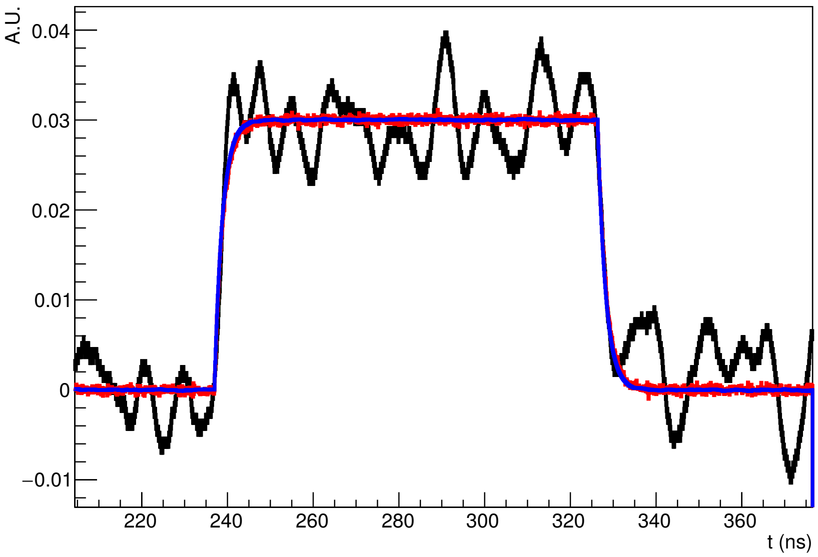

Figure 4.10: The analytic model in `RAT` compared to test-stand data for a 1 nhit N100 pulse. The noise on the N100 is largely caused by the environment on the teststand and is not representative of the noise in the detector.

same way for both the nhit monitor and for `RAT`. This calculation includes the rise-time and realistic trigger widths in its calculation of $\tilde{n}_{100}$. Figure 4.11 shows events simulated out to the PSUP for both runs, while Figure 4.12 shows events simulated only in the center (left) and out to a fiducial volume of 5900 mm (right).

It is clear from these plots that the model in `RAT` does a reasonable job of reproducing the trigger efficiency curve extracted from the nhit monitor. This can be quantified by convolving the underlying $\tilde{n}_{100}$ distributions (dashed line in Figures 4.11, 4.12) with the trigger efficiency curves. Doing this for run 200005, the total efficiency for detecting the 2.2 MeV γs using the `RAT` trigger efficiency curve is 51.5% and using the nhit monitor trigger efficiency curve is 52.5%. This is about a 2% systematic uncertainty, under the assumption

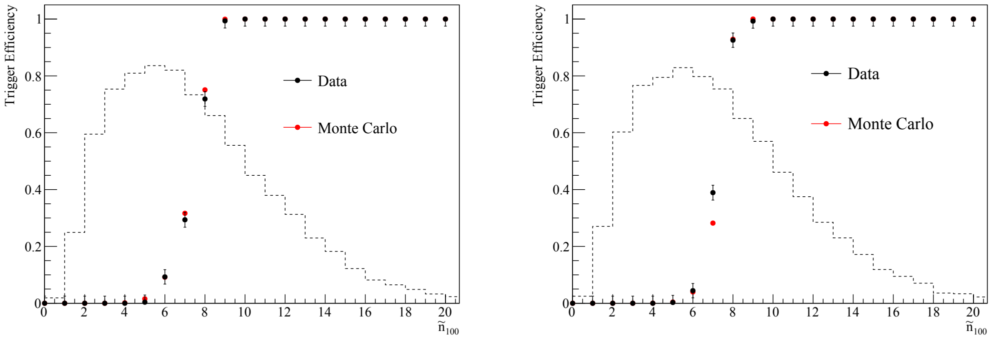

Figure 4.11: The nhit monitor (data) and `RAT` (Monte Carlo) trigger efficiency curves for runs 114287 (left) and 200005 (right). The dashed line shows the underlying $\tilde{n}_{100}$ distribution for the simulated 2.2 MeV γs. The events are simulated out to the PSUP at 8.5 m.

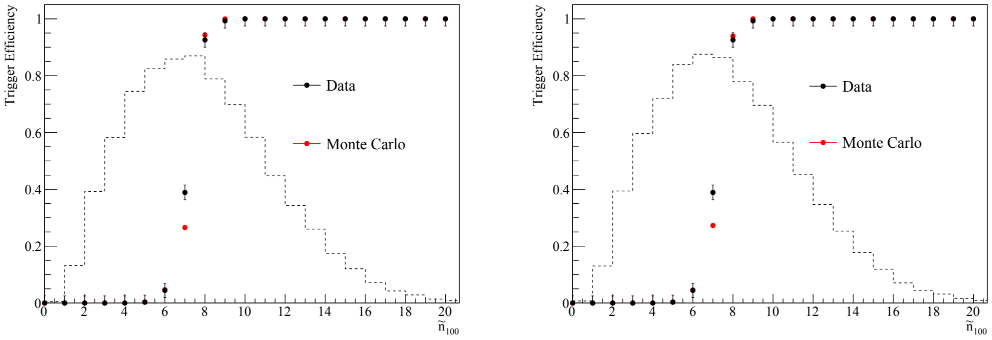

Figure 4.12: The nhit monitor (data) and `RAT` (Monte Carlo) trigger efficiency curves for runs 200005. The left plot shows event simulated in the center and the right plot shows events simulated out to 5900 mm. The dashed line shows the underlying $\tilde{n}_{100}$ distribution for the simulated 2.2 MeV γs.

that the nhit monitor is reflective of the ‘true’ trigger efficiency. This analysis was run over a large set of physics runs and the maximum systematic difference between the nhit monitor and `RAT` was 4%.

There are several issues with the assumption that the nhit monitor data is representative of the true trigger efficiency. The nhit monitor runs on the same channels in crate 4 every time and does not probe the response of the other crates or channels. The nhit monitor does run during physics runs, allowing it to snap shot the state off the detector at a time in which we are taking physics data, but it increases the total rate by 1 kHz (about doubling the average trigger rate) while doing so. This might affect the baseline restoration on the MTC/A+s effectively correcting any changing baseline, which might not be representative of the detector response during physics running.

In order to check the nhit monitor results, dedicated runs using the TELLIE system were taken to measure the trigger efficiency. These runs were taken with a mean occupancy of around 10 nhit in order to probe the relevant turn-on of the trigger efficiency. The results for three of these runs are shown in Figure 4.13 and are compared to the nhit monitor results. The TELLIE runs are performed on two fibers at two different rates. The runs on the same fiber at 50 Hz and 1 kHz show consistent results, suggesting there is probably little rate dependence in the trigger efficiency curve. However, the runs on the different fibers do show systematic differences in the trigger efficiency curve, suggesting there is a significant effect from differences in the electronics behavior across the detector.

The difference between the nhit monitor results and the TELLIE results are not fully understood, but are consistent with the results found in (94). The nhit monitor only probes a single crate, and its possible that crate 4's response is particularly good. Additionally, the pedestal hits all come exactly in-time, whereas the TELLIE hits can be spread in-time due to optical reflections and the PMT transit time. Regardless of the reason, this difference can be taken as a systematic uncertainty. For the TELLIE run that is maximally different from the nhit monitor in Figure 4.13 the total predicted trigger efficiency differs by 15%. This dominates over the run-by-run uncertainty. This systematic is important for the analysis described in Chapter 6.

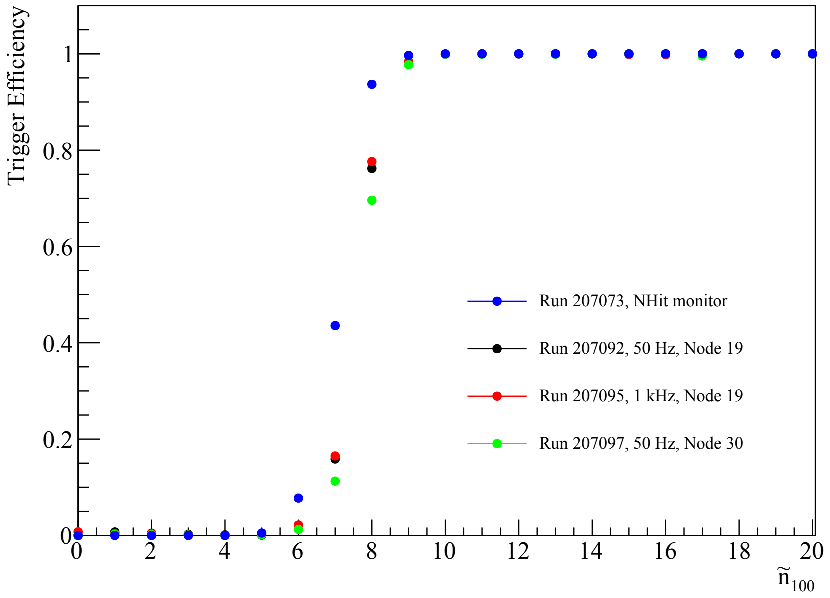

Figure 4.13: The trigger efficiency curve for three different TELLIE data sets. Run 207092 (black) used node 19 pulsed at 50 Hz. Run 207095 (red) used node 19 pulsed at 1 kHz. Run 207097 (green) used node 30 pulsed at 50 Hz. The different nodes are used to test the trigger efficiency for different parts of the detector. The different rates are used to test whether the rate impacts the determination of the trigger efficiency. The most recent nhit monitor, before these runs, is shown in blue. It is clear the nhit monitor over predicts the total trigger efficiency relative to the TELLIE results.

Chapter 5

Simulation and Reconstruction

SNO+ relies on a detailed Monte Carlo simulation for predictions of efficiencies and signal expectations, and for constructing systematic uncertainties. The simulation software used by the SNO+ experiment is called `RAT`, a `GEANT4`-based (95) simulation and analysis package originally developed by Stan Seibert (96). `RAT` uses custom photon tracking and a detailed detector model, including full simulation of the PMTs, electronics, and DAQ response. `RAT` relies on a dynamic database, `RATDB`, containing information about the detector with adjustable parameters that can be easily tuned by the user. `RATDB` contains both static tables, not expected to change as the detector runs, such as the PMT locations, as well as run-by-run tables that are uploaded every run with the detector-state information for that run. In addition to event simulation, `RAT` provides a flexible framework for reading in data, such that the exact same software is used for the simulated and real data. The output of `RAT` is a `ROOT` (97) file. For the most part, the files used in the analysis presented in this thesis are flat `ROOT` trees called ntuples.

5.1 Event Generation

The first step in the simulation of a Monte Carlo (MC) event is the generation of an event vertex containing the particle type(s), time(s), energy(ies), position(s), and direction(s). For this discussion, the generation of the vertex for reactor antineutrino interactions will be used as an example due to its relevance for the analysis presented in Chapter 6. The reactor antineutrinos interact via IBD, producing a positron and neutron.

The reactor antineutrino energy distribution is generated using a useful phenomenological description, fit to the shape of the spectrum that is predicted by modeling the β-branches of the four fission products. The spectrum from each of the four primary fission isotopes – ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu – is modeled as an exponential to a 5th order polynomial of the form:

$$S_{fit} = \exp\left(\sum_{\mathrm{p}=1}^{6} \alpha_{\mathrm{p}} \mathrm{E}_{\nu}^{\mathrm{p}-1}\right) \tag{5.1}$$

given in (2) and (3). The fitted parameters used are shown in Table 5.1 and the associated spectra are shown in Figure 5.1.

Isotope	α_1	α_2	α_3	α_4	α_5	α_6
^{235}U	4.367	-4.577	2.100	-0.5294	0.06186	-0.002777
^{238}U	0.4833	0.1927	-0.1283	-0.006762	0.002233	-0.0001536
^{239}Pu	4.757	-5.392	2.563	-0.6596	0.07820	-0.003536
^{241}Pu	2.990	-2.882	1.278	-0.3343	0.03905	-0.001754

Table 5.1: The fitted values for the coefficients in Equation 5.1 used to model the antineutrino energy distribution from nuclear reactors. The values for ^{235}U, ^{239}Pu, and ^{241}Pu are from (2) and the values for ^{238}U are from (3).

In addition to generating the energy distribution, the other non-trivial vertex information to generate is the direction of the incoming antineutrino. There are dozens of possible nuclear reactors running at various powers that can contribute to the incoming flux, each

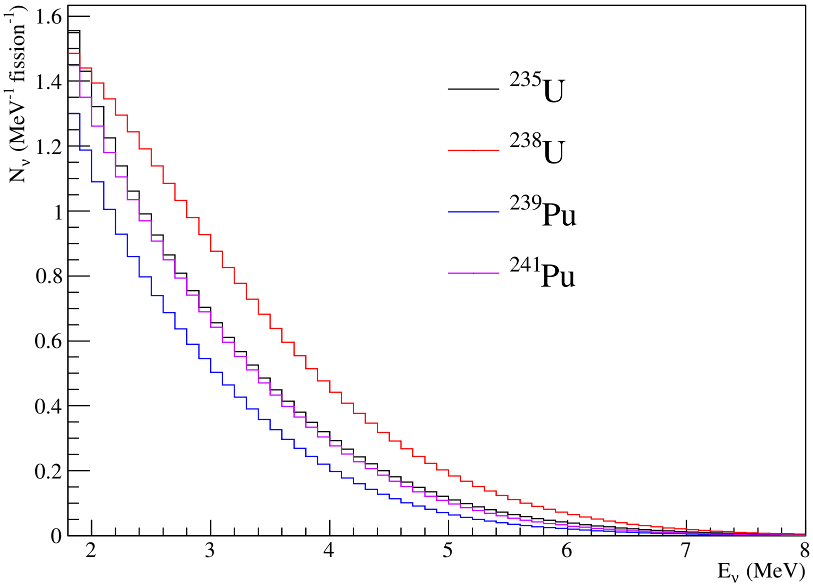

Figure 5.1: The predicted energy spectra for the antineutrinos emitted from the sum of all β-branches of the fission products ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu. The spectra are shown above the IBD threshold (1.8 MeV) shown in Figure 2.5, and are used in the `RAT` antineutrino generator.

with different distances to SNO+ (shown in Table E.1). The source reactor, for each generated event, is selected based on its fractional contribution to the total expected flux for SNO+ (the details for the rate of emission are given in Section 6.2). Then, the neutrino's direction is set such that it is incoming from the selected reactor complex.

The vertex of the interaction is distributed uniformly in the selected detector volume, where the primary outgoing particles (the neutron and the positron) are generated. The energy and direction of the positron are related directly to the energy of the neutrino according to Equations 2.54 and 2.57 respectively. After generating the positron's vertex information, the magnitude and direction of the neutron momentum is set by conservation of momentum.

Once the final state particles' vertex information are finalized, the propagation and production of photons and secondary particles are simulated by `GEANT4`. `GEANT4` provides detailed and verified methods for the simulation of physics processes as the particles are transported through the detector. Each particle is moved through the detector in steps, the length of which are randomly sampled using the mean free path for the next physics interaction. `GEANT4` provides base classes for both electromagnetic physics and hadronic physics interactions and also handles the decay of unstable particles.

The Cherenkov production of light is handled by generating a Poisson distributed number of photons according to the Cherenkov formula given in 4.2. The photons are distributed evenly along the track segment and uniformly in azimuth with respect to the charged particle's direction. The wavelengths of the photons are generated between 200 to 800 nm.

The photons are propagated through the detector, where they can be absorbed or scattered in each of the detector components. The absorption and scattering length of the internal water, acrylic vessel, and external vessel are calibrated using in-situ calibration sources, described in Appendix D. Once the photons reach the PMTs, a 'grey-disk' model is used to simulate the reflectivity of the concentrators and PMT glass. The grey disk model is a substitute for the full 3D PMT simulation, which speeds up the time of the simulation due to the complicated geometry of the full 3D PMTs. The PMT and concentrator are replaced by a disc at the top of the concentrator position, and the disc surface properties are calibrated as a function of wavelength according to optical calibration data. For photons that are absorbed at the grey disk (rather than reflected) an effective efficiency for creating a photoelectron (PE) is simulated. This efficiency comes from ex-situ measurements of the PMT properties, but is calibrated using in-situ ^{16}N source deployment.

A generated PE enters the trigger and DAQ simulation. The first step is the front-end simulation, where the PE is converted to a PMT pulse. The pulses for each channel are compared against a discriminator threshold, which is drawn from the channel hardware status (CHS) `RATDB` table, containing information about the state of the hardware for every run. About 25% of the pulses fall below the discriminator threshold, which adds an additional efficiency loss that is separate from the efficiency included at the grey disk. PMT pulses that pass the discriminator thresholds enter the trigger simulation, where corresponding N20, N100, and ESUM trigger signals are generated. The trigger simulation includes many realistic components, such as electronics noise and finite rise-times. The trigger signals are summed detector wide, and each trigger sum is compared against the corresponding threshold set at the MTC/A+. The trigger simulation includes important features such as the length of the trigger window, set by lockout, and keeps track of the global trigger ID. For events that cause a trigger signal that crosses threshold, a global trigger is generated and the event is built into the simulated dataset.

For each triggered event, the simulated PMT information contains hit-times and charges, analogous to the values output by the ECA and PCA processes that are applied to the data. PMT dark current is simulated based on a run-by-run measurement of the hit-counts on each individual PMT. The calibrated PMT information is passed into the reconstruction algorithms, described in Section 5.2. The final output of the simulation is a `ROOT` file containing the event, PMT, and reconstruction information, that is identical in most ways to the data. The simulated files additionally contain the ‘truth’ information about the event, which is often useful for validating aspects of the reconstruction. The tracks of the particles and photons through the detector can be stored, but are often discarded to save space. A second output, called the ‘ntuple’ file, is a flat `ROOT` tree, that strips away the PMT hit information in order to save space and speed up the analysis time.

5.2 Event Reconstruction

Event reconstruction is applied to data and simulated events identically. The reconstruction uses the pattern and time of the hit PMTs to generate the most probable time, position, direction and energy of the event. This reconstruction process assumes that the event is from Cherenkov light caused by an electron. Channels that are identified as offline or poorly calibrated are not used in the reconstruction algorithm for either data or simulation.

The time and position of the event are determined by minimizing the time residuals against a PDF. The residuals are generated by varying the vertex position, and, for each PMT hit, calculating:

$$t_{res} = t_{PMT} - t_{transit} - t_0 \tag{5.2}$$

where t_{PMT} is the calibrated hit-time of the PMT, $t_{transit}$ is the photon travel time given a hypothetical event position, and t_0 is the time of the event. The position and time of the vertex are reconstructed by minimizing the total t_{res} distribution for all PMT hits against a PDF that is generated offline using simulated events. An example time residuals distribution is shown for the ^{16}N calibration data deployed at the center of the detector in Figure 5.2.

With the reconstructed event position $\mathbf{r}$, the event direction, $\mathbf{d}$, is reconstructed by calculating, for each hit PMT:

$$\cos(\theta_i) = \mathbf{d} \cdot (\mathbf{r}_i - \mathbf{r}), \tag{5.3}$$

where $\mathbf{r_i}$ is the location of the ith hit PMT. Summing over the hits, the θ distribution is minimized against a PDF, shown in Figure 5.3, to generate a reconstructed direction. An effect of these methods is a bias in the fitted position along the direction the particle was traveling, because the vertex can be moved along this axis without degrading the fit. This ‘drive’ is corrected for, based on Monte Carlo estimates of the effect, as a function of position and direction.

The energy of the event is determined after the time, position and direction have been reconstructed, using a method called energy response (energyRSP). The time residuals are generated using the reconstructed position and time, and only PMT hits occurring within a prompt 18 ns window are selected to be used in the energy estimation. The

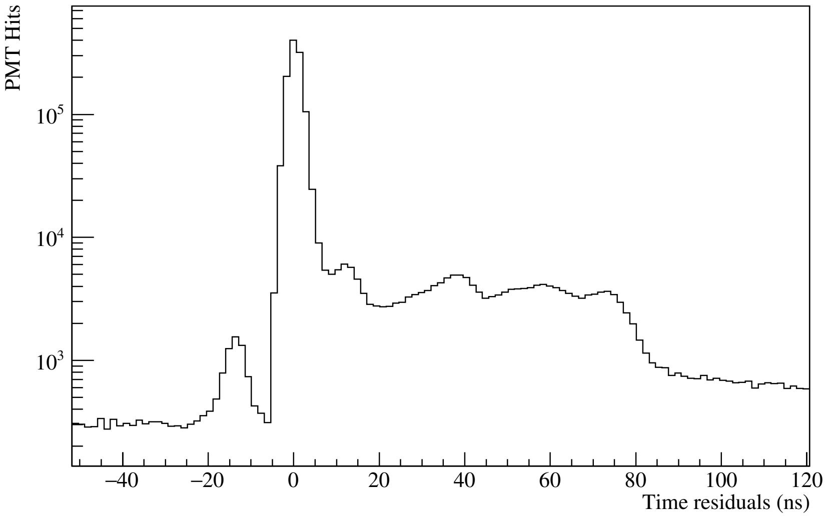

Figure 5.2: The PMT hit time residuals for central ^{16}N calibration source data, calculated using the known calibration source position.

energyRSP method works by accounting for a total detection efficiency which includes the PMT quantum and collection efficiencies, the angular response of the grey disk model, the solid angle acceptance, the PMT dark noise rate, the Cherenkov angular dependence, and the transmission, attenuation, and Raleigh scattering probabilities. This total efficiency is determined primarily through the use of look-up tables generated by offline simulations. The number of Cherenkov photons as a function of electron energy, as shown in Figure 5.4, is scaled by the total detection efficiency to be converted to a total number of prompt PMT hits. The energy is adjusted until the scaled number of Cherenkov hits matches the detected number of prompt hits.

The reconstruction time, position, direction, and energy of the event are used to calculate several quantities. These are listed below and are often used to classify the goodness-of-fit, to reject poor reconstructions, and to remove background events.

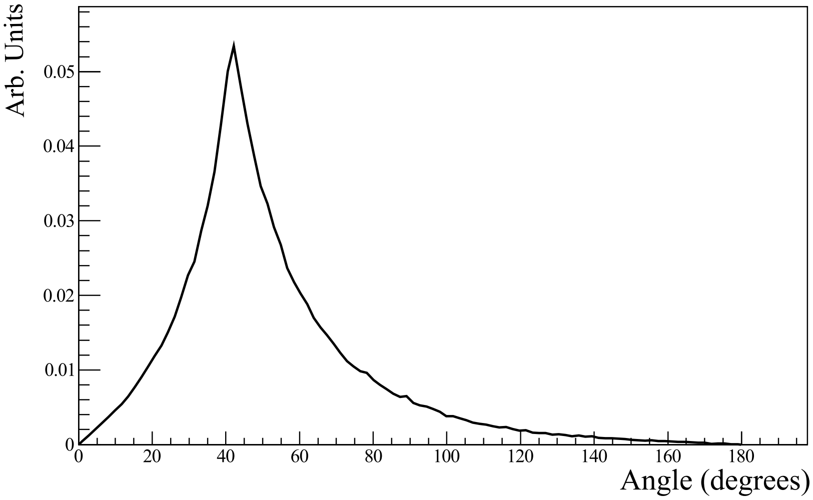

Figure 5.3: The PDF used in RAT when evaluating the likelihood of a hypothesized direction. The peak at 41° corresponds to the Cherenkov angle in water.

In-time ratio (ITR): The total number of PMT hits with time residuals between -2.5 and 5.0 ns divided by the total number of PMTs hits. ITR provides discrimination between events from Cherenkov light, which produces hits in a very prompt window, and other types of events (such as those from instrumental noise).

$\boldsymbol{\beta_{14}}$: An isotropy parameter defined as:

$$\beta_{14} = \beta_1 + 4\beta_4 \tag{5.4}$$

where

$$\beta_k = \sum_i \sum_{j \neq i} P_k(\cos(\theta_{ij})) \tag{5.5}$$

is the kth Legendre polynomial evaluated as the cosines of the angles between all pairs of hit PMTs i and j. The combination selected for β_{14} was chosen during SNO for its discrimination power and Gaussian-like distribution (98). β_{14} is used to separate more

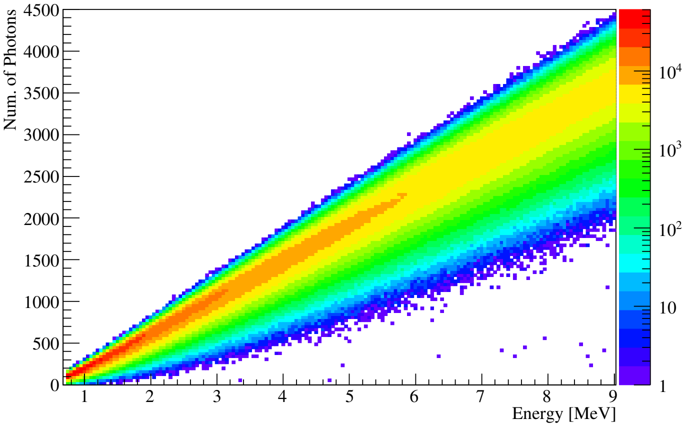

Figure 5.4: The total number of produced Cherenkov photons for electrons with energies from 0.7 to 9 MeV.

isotropic events, such as those originating from γ-rays, from less isotropic events, such as those from electrons.

$\mathbf{u} \cdot \mathbf{r}$: The dot product of the reconstructed direction and position vectors. $\mathbf{u} \cdot \mathbf{r}$ provides discrimination against external events, which are typically at high radius and inward pointing.

Position Error (σ_R): The estimated error on the fitted position, as returned by the fitting method. Larger values indicate more poorly estimated positions.

Position figure of merit ($\mathbf{R}_{FOM}$): The value of the maximum log likelihood of the position fit, divided by the total number of hits. Smaller values indicate that the maximum log likelihood was relatively small and that the fit might be poorly estimated.

Energy U-test: The energy U-test calculates a test statistic by ordering the hit probabilities returned by energyRSP for each PMT in descending order. Then each PMT is ranked according to the index in this list. For N PMTs hit in the event, the U-test is calculated as:

$$U = \frac{\sum_i^N \mathrm{R_i} - \frac{\mathrm{N(N+1)}}{2}}{N(N_{\mathrm{act}} - N)}, \tag{5.6}$$

where R_i is the rank of the PMT and N_{act} is the total number of inward-looking online channels. Small values of U indicate a possible poor energy estimation.

Energy G-test: The energy G-test sums over the hit PMTs to calculate:

$$G = \frac{1}{N} \sum_i \log\left(\frac{1}{E_i}\right), \tag{5.7}$$

where E_i is the expected number of hits on channel i, as calculated by the energyRSP method. Large values of G indicate a possible poor energy estimation.

Energy Z-factor: The energy Z-factor uses the medians and the median absolute deviations of the hit-probabilities returned by the energy RSP method:

$$Z = 1 - \frac{3(\sigma_{\mathrm{med_p}} + \sigma_{\mathrm{med_n}})}{\mathrm{med_p} - \mathrm{med_n}}, \tag{5.8}$$

where med_p is the median of the probabilities for the hit channels, med_n is the median of the probabilities for unhit, active channels and $\sigma_{\mathrm{med_p}}$ and $\sigma_{\mathrm{med_n}}$ are the associated median absolute deviations.

Chapter 6

A Reactor Antineutrino Search with SNO+

The detection of neutrons is notoriously difficult in a pure water detector. Indeed, in his technical report on loading water with gadolinium, M. Vagins writes: “Water Cherenkov detectors have been used for many years as inexpensive, effective detectors for neutrino interactions and nucleon decay searches. While many important measurements have been made with these detectors a major drawback has been their inability to detect the absorption of thermal neutrons” (99). Using the SNO+ detector, I present a search for reactor antineutrinos, which have yet to be observed in a pure water Cherenkov detector, largely due to the difficulty of neutron detection. The neutrons liberated by antineutrino interactions thermalize in the detector and capture on a free proton, producing a 2.2 MeV de-excitation γ-ray. The low energy nature of the de-excitation γ makes it difficult to detect in water Cherenkov detectors, where the trigger thresholds have traditionally been set around 5 MeV. However, due to advanced, purposefully designed triggering and front-end electronics, as well as excellent detector cleanliness, SNO+ hopes to be the first water Cherenkov detector to observe reactor antineutrino interactions.

This chapter begins in Section 6.1 with a brief discussion of antineutrino sources other than nuclear reactors. Section 6.2 lays out a calculation of the expected reactor antineutrino interaction rate in the SNO+ detector. Section 6.3 describes the detection process for reactor antineutrinos in water Cherenkov detectors generally, and SNO+ specifically. Then, Section 6.4 details the dataset used, the corresponding livetime, and the event selection process. A likelihood-based analysis method for rejecting backgrounds is detailed in Section 6.5. The neutron detection efficiency from the AmBe calibration source is discussed in Section 6.6. The calculated event rate after cuts, including the analysis efficiency and the calibration corrections, is given in Section 6.7. The various backgrounds are discussed in Section 6.8 and Section 6.9. Lastly, Section 6.11 summarizes the results.

6.1 Sources of Antineutrinos

In addition to nuclear reactors, which will be discussed in much more detail in this chapter, there are several other sources of $\bar{\nu}_e$s. Antineutrinos can be created in the atmosphere at high energies (primarily above 100 MeV), as discussed in Section 3.2. When these neutrinos interact with the nuclei in the detector, they can produce secondary particles, some of which turn out to be backgrounds for the reactor antineutrino search (see Section 6.8.3). The diffuse supernova neutrino background (DSNB) is a weak glow of MeV scale neutrinos that were created in distant core-collapse supernova (100). These neutrinos and antineutrinos will appear isotropic and constant in time, but have never been detected due to their extremely low density. The Super-Kamiokande detector is currently filling with gadolinium loaded water (called SuperK-Gd) to perform the first ever detection of the DSNB (101). Antineutrinos from the uranium and thorium chain reactions within the Earth are a background for the KamLAND reactor antineutrino measurement, as discussed in Section 3.4, but these geoneutrinos are too low in energy and expected flux to be detected in the SNO+ water dataset. By far the largest source of MeV scale antineutrinos originate in nearby nuclear reactors.

6.2 Event Rate

The rate of reactor antineutrino ($\bar{\nu}_e$) inverse beta decay (IBD) interactions in the SNO+ detector is determined by three things: the rate of antineutrino production in various nuclear reactors, the oscillation of the antineutrinos on the path to the detector, and the cross-section of the reaction.

Nuclear reactors produce energy through the nuclear fission of four primary isotopes: ^{235}U, ^{239}Pu, ^{238}U, and ^{241}Pu. The neutron-rich daughters β decay, producing $\bar{\nu}_e$s isotropically. An average of ∼6 $\bar{\nu}_e$s are produced per fission, leading to about $\sim 2 \times 10^{20}$ $\bar{\nu}_e$s being emitted per second (102) for a 1 GW reactor (thermal power). The antineutrinos are all created in a pure flavor state, but oscillate on their path to the detector. The IBD reaction, shown in Figure 2.4, is sensitive only to the electron flavor, so oscillations between the reactor core and the detector decrease the flux.

The rate of emission of the $\bar{\nu}_e$s can be calculated using:

$$\frac{d^2N(E,t)}{dEdt} = \frac{W_{th}(t)}{\sum_i f_i(t)e_i} \sum_i f_i(t)S_i(E)c_i^{ne}(E,t) + S_{snf}(E,t). \tag{6.1}$$

where the sums are over the four primary fissile isotopes. $W_{th}(t)$ is the thermal power of the reactor, which can change throughout the year. f_i gives the fraction of fissions due to the ith isotope. The average thermal energy released per fission is given by e_i and is approximately 200 MeV on average for all four isotopes (103). The $\bar{\nu}_e$ yields per fission are given by $S_i(E)$ and are taken from (104) for ^{235}U, ^{241}Pu, ^{239}Pu and from (105) for ^{238}U. $S_i(E)$ is shown in Figure 5.1 for all four isotopes. The factor $c_i^{ne}(E,t)$ accounts for sub-dominant corrections to the energy spectrum due to non-equilibrium effects of long-lived fission fragments. This factor is estimated by Daya Bay as impacting the total flux by 0.6 ± 0.2% (106), and is neglected in the event rate calculation. $S_{snf}(E,t)$ accounts for contributions from spent nuclear fuel rods that were removed and stored nearby, which is estimated by Daya Bay as 0.3 ± 0.3% of the flux (106), and is neglected in the event rate calculation.

The most imporantant reactors for the study are the reactors closest to SNO+, which are all located in Ontario. The Bruce, Darlington, and Pickering Nuclear Generating Stations are all Canada Deuterium Uranium (CANDU) reactors, a type of pressurized heavy water nuclear reactor. The total power of the reactor complexes is 6.4 MW, 3.5 MW, and 3.1 MW respectively. The Bruce Nuclear Generation Station is the largest fully operational nuclear generating station by total reactor count, the number of currently operating reactors, and the nominal capacity power (note the Kashiwazaki-Kariwa power station in Japan has suspended operation since 2011 and the Kori power station in South Korea is not fully operational as of the writing of this thesis). These three complexes contribute approximately 60% of the total flux. The remainder of the flux comes from various reactors across the globe, the largest contribution of which is from three US reactors: R.E. Ginna Nuclear Power Plant, Nine Mile Point Nuclear Generating Station, and Davis-Besse Nuclear Power Station. A map showing the distances between SNO+ and the CANDU and US reactors is shown in Figure 6.1. A list of all reactor complexes is given in Appendix E in Table E.1, all of which are included in the flux estimation.

The operating power of the reactors directly impacts the total $\bar{\nu}_e$ flux. The initial flux calculations are performed using the quoted maximum operating power from each reactor complex and assuming that all reactors are operating at full-capacity year round. To account for offline periods and reactors operating below full capacity, a correction is calculated using the thermal power data published yearly by the International Atomic Energy Agency (107), and made easily available by researchers at the Laboratory for Nuclear Technologies Applied to the Environment at INFN (108), (109). For the Canadian reactors, the hourly electrical data is taken from (110), which provides more accurate and frequent information on the reactor operating conditions. Overall, a factor of 0.844 is identified as an appropriate scaling using this information.

In addition to the changing thermal powers, the relative abundance of the fissile isotopes fluctuates with time, changing the number and energy of the flux of $\bar{\nu}_e$s. This is a relatively

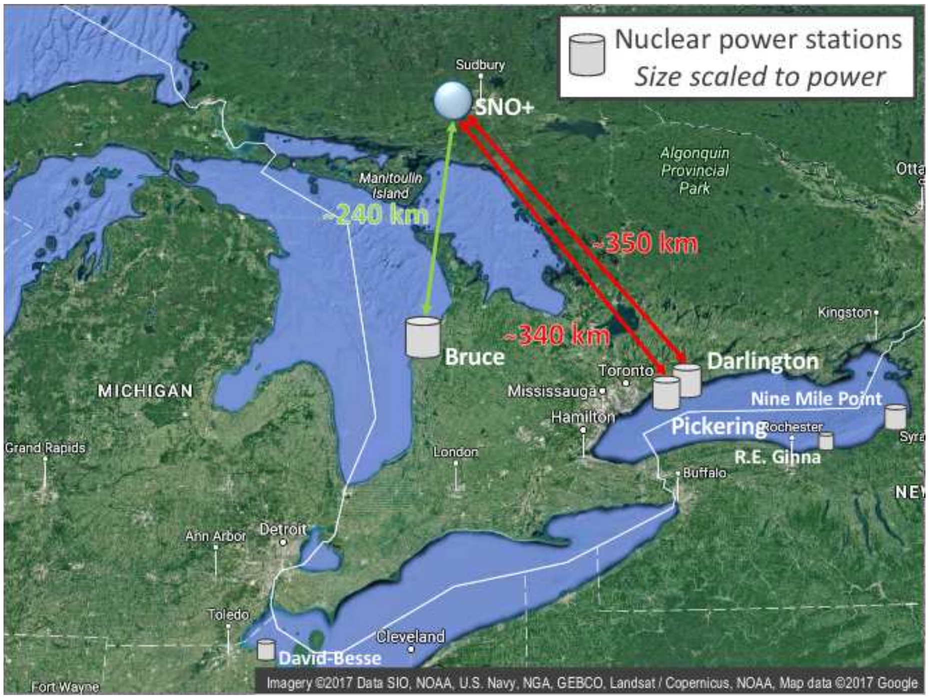
Nuclear power stations
Size scaled to power
Sudbury
SNO+
Algonquin Provincial Park
Otta
Manitoulin Island
~240 km
~350 km
~340 km
Kingston
Bruce
Darlington
Toronto
Mississauga
Hamilton
Nine Mile Point
Pickering
Rochester
R.E. Ginna
Syra
MICHIGAN
Grand Rapids
London
Buffalo
NE
Ann Arbor
Detroit
Toledo
Cleveland
David-Besse
Fort Wayne
Imagery ©2017 Data SIO, NOAA, U.S. Navy, NGA, GEBCO, Landsat / Copernicus, NOAA, Map data ©2017 Google

Figure 6.1: Map showing the SNO+ detector location and the location of six nearby reactor complexes in Canada and the US. The sizes of the nuclear stations are roughly scaled to the power of the complex.

small effect, and due to the constant online-refueling of the CANDU reactors, the flux and spectrum from those complexes should change very little with time. The relative fission fractions are taken as constant values: 56.8%, 7.8%, 29.7% and 5.7% for ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu respectively (2) for the pressurized/boiling water reactors. For the pressurized heavy water reactors (Canadian) the constant fraction values are: 52%, 5%, 42%, 1% for ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu respectively (111).

Incorporating these factors, and using a specific run-time and energy spectrum (Figure 5.1), Equation 6.1 can be used to generate the number of emitted antineutrinos per reactor complex. A point-like approximation is made for the distance, L, from the source to the detector, which effectively scales the flux reaching the detector by $1/L^2$. Following this, neutrino oscillations are applied using the full three-flavor survival probability given in Equation 3.8. Oscillation parameters from the 2018 PDG global fit are used (112). Uncertainties in oscillation parameters propogate about a ±4% unceraинty in the expected IBD rate. Figure 6.2 shows this survival probability as a function of neutrino energy for a distance of $L = 240$ km, which is approximately the distance of the Bruce Nuclear Generating Station to the SNO+ detector.

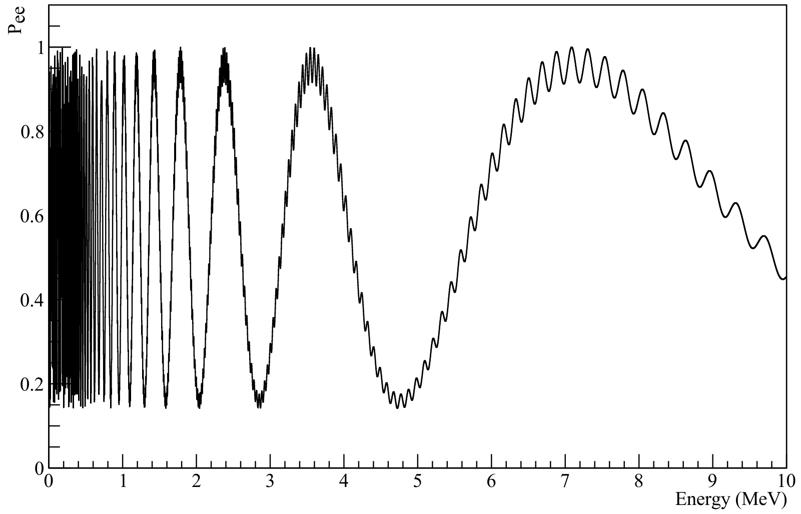

Figure 6.2: The survival probability, $P_{ee}(E)$ given by Equation 3.8, for $\bar{\nu}_e$s with a distance of $L = 240$ km.

The convolution of the neutrino energy spectrum of arriving antineutrinos and the energy spectrum of the cross-section, given in Equation 2.51, yields the energy spectrum of detected antineutrinos. This process is shown schematically in Figure 6.3. Of note is the IBD energy threshold, given in the laboratory frame (proton rest frame) as:

$$E_{\nu}^{thr.} = \frac{(M_n + m_e)^2 - M_p^2}{2(M_n + m_e)} = 1.806\,\mathrm{MeV}, \tag{6.2}$$

where M_n is the neutron mass, m_e is the positron mass, and M_p is the proton mass. Integrating the spectrum of detected antineutrinos over time and energy yields the total number of interactions per target.

To calculate the number of free protons, the number density of hydrogen is calculated using:

$$n_H = \frac{\rho w_H N_A}{m_H}, \tag{6.3}$$

where ρ is the density of water, w_H is the mass fraction of hydrogen in water, N_A is Avogadro's number, and m_H is the molar mass of hydrogen. Using the average cavity temperature of 15°C, Equation 6.3 yields a number density of 0.668×10^{29} m^{-3}. Taking the difference in temperature between the top and the bottom of the cavity, which is about 3°C, as a systematic on the total number of hydrogen atoms yields an expected ±0.045% uncertainty (113). Multiplying by the volume of the acrylic vessel gives 5.04×10^{33} total free protons. Accounting for this, the total number of interactions per year in the AV as a function of neutrino energy is shown in Figure 6.4. This figure also illustrates that the effect of oscillations is about a 55% reduction in the $\bar{\nu}_e$ flux and that about 60% of all detected antineutrinos originate in the CANDU reactors.

Integrating the over entire energy spectrum yields an expected 137.7 (391.5) events/year interacting within the AV (PSUP), before correcting for the realistic reactor powers. In this calculation, a correction is applied for the arbitrary flux scaling, 27610, used in the generation of the events. Lastly, it is necessary to account for the fact that absolute measurements of $\bar{\nu}_e$ fluxes from reactors, performed by Daya Bay and several others, indicates a deficit in the number of detected antineutrinos relative to the model prediction. A scaling of 95.2% is used to account for this, which is taken from (114). Applying this factor and the reactor

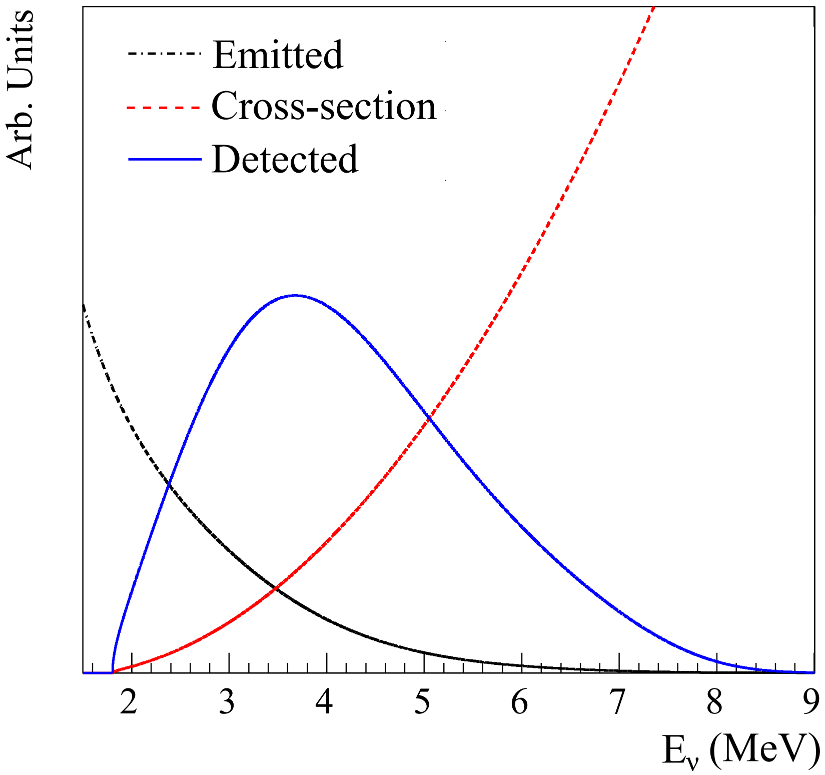

Figure 6.3: The energy spectrum of emitted antineutrinos, from Equation 5.1, is shown in black. The cross-section, given in Equation 2.51 and shown in Figure 2.5, is shown in red. The convolution of those curves produces the energy spectrum of detected antineutrinos shown in blue. The distributions are shown with arbitrary units and prior to oscillations being applied. This figure is reproduced from (3).

power scaling, the total number of expected interactions per year is 110.4 (313.9) within the AV (PSUP).

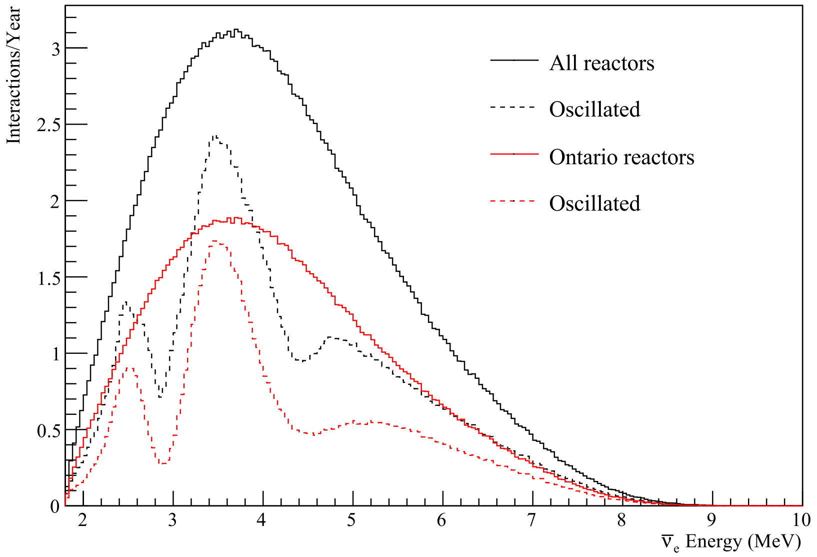

Figure 6.4: The total number of interactions per year in the AV as a function of neutrino energy. The black distributions are prior to (solid) and after (dashed) applying oscillations. The red curves show the same, but only for the CANDU reactor complexes.

6.3 Detecting Reactor Antineutrinos in SNO+

The SNO+ search for reactor antineutrinos presented in this thesis will look for IBD reactions on the water target; however, SNO+ will peform a more sensitive search once filled with liquid scintillator, which will drastically improve the neutron detection efficiency and energy resolution. This search using water is intended to demonstrate the feasibility for detecting reactor antineutrinos using a pure water detector. Importantly, SNO+ has collected data with the lowest trigger threshold ever achieved in a large-scale water Cherenkov detector. This low threshold implies that the work presented here is the lowest energy analysis ever performed for a water Cherenkov detector. Because of the importance of neutron detection for reactor antineutrino and DSNB searches, an experiment capable of performing these measurements without any doping of the water is of technological interest. Notably, no special trigger is implemented for the data-taking; the standard trigger runs at a low enough threshold to maintain good trigger efficiency for the 2.2 MeV γs.

Figure 6.5 shows the IBD interaction schematically. The outgoing positron creates prompt Cherenkov light before annihilating and producing two 511 keV γ-rays, which can Compton scatter and produce more Cherenkov light. The liberated neutron will thermalize in the detector through collisions with free protons. On average it takes about 10 μs, approximately 20 collisions, to lower the energy of the neutron from several MeV to room temperature, around 0.025 eV. Once thermalized, the neutron will continue to scatter until it is eventually captured on hydrogen. The neutron capture time-constant is around 200 μs in water (115), (38). After capturing, a 2.2 MeV γ is emitted, which Compton scatters and produces Cherenkov light. The signal from the positron is referred to as the 'prompt' event and the signal from the neutron capture is referred to as the 'delayed' event. Importantly, it is through enforcing the detection of this coincidence pair that drives down backgrounds and enables this search.

As previously discussed, the difficulty in detecting IBD interactions in a pure water Cherenkov detector is in the low energy nature of the signal. Water Cherenkov detectors, such as SNO and Super-Kamiokande, often run with trigger thresholds well above 3 to 4 MeV, and low energy radioactive backgrounds in the water make it difficult to reduce the trigger threshold. The Super-Kamiokande detector demonstrated a technique for detecting neutrons with an efficiency of less than 20% using a dedicated trigger (116). Due to the various electronics upgrades described in Section 4.5 and the cleanliness of the detector, SNO+ has been running with a trigger threshold of around 7 PMT hits (see Section 4.7). With this threshold the detector triggers on approximately 50% of the 2.2 MeV γs that interact in the detector, with a small dependence on the position of interaction.

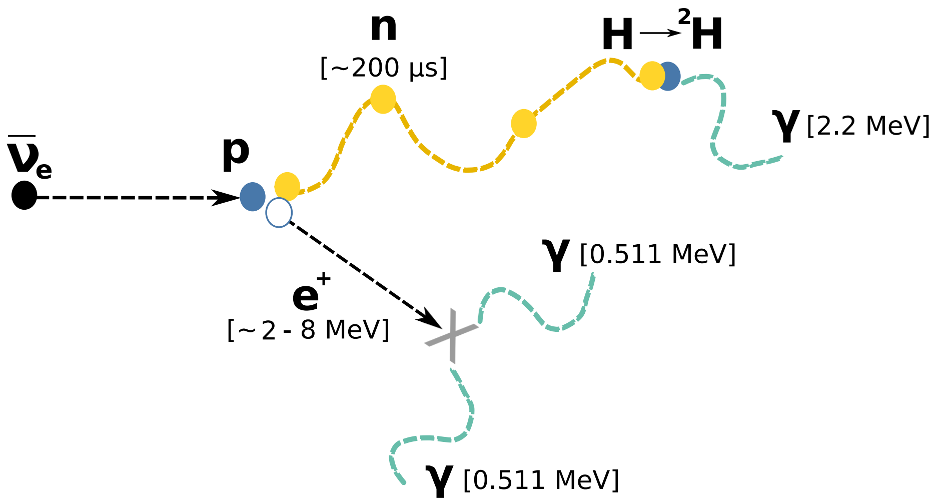

Figure 6.5: Schematic representation of the IBD interaction in the SNO+ detector. The outgoing positron creates Cherenkov light before annihilating into two 511 keV γs. The neutron thermalizes and captures on hydrogen with a time-constant of about 200 μs. The neutron capture results in the emission of a 2.2 MeV de-excitation γ.

Because neutron detection is particularly difficult in water Cherenkov detectors, typically a small amounts of material with both a large neutron capture cross-section and high energy de-excitation γ-rays is loaded into the water. For example, the SNO detector mixed NaCl into the heavy water for its 'salt' phase. ^{35}Cl has a large neutron cross-section (100 times that of hydrogen) and releases a high energy (8.6 MeV) γ cascade after neutron captures (117). The Super-Kamiokande detector is currently loading with Gd(SO_4)$_3$, which provides a large capture cross-section and high energy γ cascades from Gd (118, 119). The EGADS prototype detector in Kamioka mine demonstrated that the Gd-loading technique performs well at large scales (120). With significantly improved neutron detection, Super-Kamiokande with Gd hopes to perform the first detection of the DSNB. Due to the success of this program, future detectors such as WATCHMAN and ANNIE also hope to use Gd-loaded water to boost their neutron detection efficiency (121), (122).

6.4 Data Selection

The selection of quality data proceeds across several stages. Chunks of data are split into runs (most commonly one hour long), which are rejected or approved according to a set of data-quality criterion, called ‘run selection’, described in Section 6.4.2. In the runs, individual events are rejected or approved by examining the low-level PMT information as well as the reconstructed quantities. The PMT hits are selected for use in the reconstruction according to the their calibration status and whether they pass the hit-cleaning process described in Appendix C.

6.4.1 Dataset and Livetime

The dataset used for this analysis was taken between October 24th, 2018 and July 10th, 2019. The runs used are numbered between 200000 and 207718. Runs are selected to be used in this analysis based on a set of criterion described in Section 6.4.2. The total length of the dataset selected is 198 days. Data taking is performed in ‘physics’ mode, with detector settings appropriate for performing physics analyses; however, significant amounts of time are also spent calibrating the detector and performing maintenance of the various detector components and electronics.

The dataset selected is particularly low background because of the successful installation of the sealed nitrogen cover-gas system in September of 2018. This led to a reduction of radioactivity levels in the detector and makes the selected runs a particularly important fraction of the full SNO+ water dataset. Data taken earlier, such as that used for the nucleon decay (39) and solar neutrino (19) papers, is not considered in this thesis, but will be eventually incorporated into the analysis. Notably, in the dataset used for this thesis, there is a very small amount of the liquid scintillator linear alkyl benzene (LAB) in the neck of the detector. Various checks were performed to ensure this does not affect the data.

Note that the LAB did not contain any fluors, and the emission spectrum of the LAB is not expected to transmit through the acrylic vessel.

The livetime of the dataset includes corrections to the total data taking time, which remove dead-time introduced through various data-cleaning cuts, such as the muon follower cut. The dead-time is added up, accounting for overlaps between the various cuts, and yields a total livetime of 190.3 days.

6.4.2 Run Selection

Run selection is a set of checks performed to ensure the quality of the data taken during the run. Often short bursts of triggers due to trigger baseline shifts or PMT breakdowns will cause data loss during a run. Rather than attempt to maintain small fractions of runs in which an issue occurred, the entire run is removed from the dataset to simplify the data-selection processes. Other common reasons a run might be rejected is due to activity in the deck clean room, the run length is less than 30 minutes, one of the crate's high voltage is turned off, or the trigger signals are disabled for an entire crate. A full list of all of the run selection criterion is given in (123). For the data taken between October 24th, 2018 and July 10th, 2019, with run numbers between 200000 and 207718, approximately 85% of the data taken in physics running mode is selected to be used in this analysis.

6.4.3 Event Selection

The event selection process identifies triggered events that pass a set of of 'low-level' and 'high-level' criterion. The low-level cuts remove events based on the PMT information, such as the time, charge, and location of the PMTs that fired. These cuts are primarily designed to remove events caused by instrumental noise. The high-level cuts use the reconstructed quantities to determine if the events are signal-like, based on distributions extracted from the simulation.

6.4.3.1 Low-level Cuts

The low-level cuts consist primarily of a set of algorithms called data-cleaning, which identify events caused by instrumental effects. Common instrumentals that must be removed are called 'flashers', which are caused by small sparks in the base of the PMT, emitting light in the detector. Similarly, 'shark-fins', named for their shape on the ESUMH trigger signal, are caused by sparks at the PMT but do not lead to light emission in the detector. Flashers and shark-fins can occur at significant rates and must be removed from the dataset. These events are primarily tagged using the high charge PMT hit, corresponding to the location of the spark, and the surrounding cross-talk hits. However, several orthogonal methods are used to ensure 100% rejection of these events. One such method uses the CAEN waveforms to identify events with a particularly large ESUM trace (high charge) and small nhit trace (few PMT hits), which is typical for a flasher. Additionally, other types of electronics effects, such as dry-end breakdowns, are removed using the location of the PMT hits in crate-space. The data-cleaning bit mask used for this analysis is 0xFB0000017FFE, and each of the cuts in the mask are described in Appendix A.

Importantly, several of the data-cleaning cuts are considered 'livetime' cuts in the sense that they cut all events in a window of time, based on a tagged initial event or event pair. For example, notable livetime cuts include the muon follower cut, the atmospheric cut, and the retrigger cut. The muon follower cut removes all events 20 seconds after a tagged muon, which ensures that any decays of relatively short-lived cosmogenically activated isotopes are removed from the dataset. Similarly, the atmospheric cut tags charged-current atmospheric neutrino interactions, using the Michel electron follower to tag the event. A 20 second window is removed after these events. The retrigger cut is particularly important for this analysis, because it removes any event within 3 μs of another event. Often flashers and other instrumentals occur in short bursts and can be removed using the retrigger cut. This cut has a clear impact on the analysis, as the Δt between the prompt positron and delayed neutron can fall within this 3 μs window. Because data-cleaning is only applied to data and not simulation, a Δt > 3 μs must be included in the analysis.

The other low-level cut is applied during the data processing. Because of the high rates of low nhit events and the significant processing time required to perform reconstruction, it was not feasible to run the reconstruction algorithms on every event. Based on initial benchmarking, a threshold of 15 nhits was selected. Additionally, all events within a one ms coincidence window after an event with nhit 15 or higher are reconstructed. This ensures that the neutron capture events, which often fall below the 15 nhit threshold, are reconstructed. This process effectively enforces a 15 nhit threshold on the prompt positron event. This coincidence processor is not used on the simulation events, so a 15 nhit threshold on the prompt event must be applied in the analysis. In all cases, the definition of nhits includes only channels that are calibrated and with hit-cleaning applied. Lastly, the events are required to have greater than or equal to 6 hits in order to reconstruct, which only affects the delayed event.

6.4.3.2 High-level Cuts

The high-level cuts consist of selections made on reconstructed quantities, all of which are described in Section 5.2. For both the prompt and delayed event, a valid fit is required, which indicates that the position, direction, and energy of the event have been successfully identified by the reconstruction algorithms. Then, cuts on these reconstructed values are chosen to have small signal sacrifice. The background rejection is discussed in more detail in Section 6.8, and this section will present the cut values along with their primary justification.

An energy region of interest (ROI) for the prompt event is selected between 2.5 MeV and 9 MeV. From Figure 6.4 it is clear this will have a small signal sacrifice. The lower energy cut is largely degenerate with the 15 nhit cut enforced in the processing, and provides a rejection of low energy radioactive backgrounds with minimal signal sacrifice. The high energy cut is selected to reduce possible high energy backgrounds, primarily coming from atmospheric neutrino interactions that can produce de-excitation γs or Michel electrons. A high energy cut at 4 MeV is applied to the delayed event. The delayed event has no lower energy cut in order to maintain a high signal efficiency. Note that a bug in the energyRSP algorithm, discussed in Section 5.2, was identified post data-processing, and a correction factor is applied in the analysis using a method in `RAT-6.18.1` called ReconCorrector.

The position of the event, after correcting for the AV offset of 108 mm (converting to a coordinate system where (0,0,0) is the center of the AV), is restricted to a fiducial volume (FV) of 5700 mm. The FV reduces the background from external γ-rays and the (α, n) background along the AV. The β_{14} variable, a measure of the isotropy of the PMT hits, is used to discriminate between the positron signal and γ-ray backgrounds, as the latter will often Compton scatter several times and produce PMT hit distributions that are fairly isotropic. β_{14} provides a strong handle on rejecting γ backgrounds, particularly for the prompt event. The cut values are chosen to maintain a high signal efficiency, and the shape of the β_{14} distribution is used to further reject backgrounds, as described in Section 6.5. The ITR, σ_R, R_{FOM}, energy G-test, energy U-test, and the energy Z-factor are used to reject events with bad fits. These figures of merit are somewhat correlated but each provides additional background rejection power (124). The strongest handle on identifying the reactor IBD signal against the random, accidental backgrounds, discussed more in Section 6.8.1, is the temporal and spatial coincidence between the prompt and delayed events. The Δt calculation uses only the trigger time, from the 50 MHz clock, because any effect from using the fitted times is negligibly small. The Δr distribution is calculated using the prompt and delayed reconstructed positions:

$$\Delta r = \sqrt{(x_p - x_d)^2 + (y_p - y_d)^2 + (z_p - z_d)^2}. \tag{6.4}$$

Cuts on both the Δt and Δr are selected to maintain signal efficiency while removing accidental backgrounds. Similarly to β_{14}, the shape of the Δt and Δr distributions are used to further reject backgrounds.

6.4.3.3 Follower Cuts

As discussed, several of the data-cleaning cuts, described in Appendix A, remove all events in chunks of time after an important event. However, none of these livetime cuts are general enough to tag various types of possible charge-current atmospheric neutrino events, which often produce multiple hadrons in the final state. In situations where there are multiple neutrons in the final state, a neutron-neutron coincidence can produce a distinct background. To reduce this background, a cut that removes 0.1 seconds after any event with nhit > 300 is introduced. Additionally, for neutral current events, with no prompt lepton to act as a tag, the total number of neutrons in a window around the prompt event can be used as a handle. To reject these events, a neutron multiplicity count after each prompt candidate is constructed by counting the number of events that pass the delayed event cuts in a one ms window. Any event with a multiplicity of more than one is rejected. Both this cut and the high-nhit follower cut are described in more detail in the context of the atmospheric neutrino background in Section 6.8.3.

6.4.4 Summary of Cuts

A summary of the cuts described in Section 6.4.3 is given in Table 6.1. These cuts are performed prior to a likelihood ratio cut described in Section 6.5.

6.5 Likelihood Method for Rejecting Backgrounds

The set of analysis cuts, given in Table 6.1, does well to trim the dataset, removing the majority of the random or 'accidental' coincidences. However, even after these cuts, the accidental coincidences dominate the remaining dataset. A method that generates a likelihood ratio used to reject these accidental coincidences is presented in this section. The random, accidental coincidences are described further in Section 6.8.1.

Table 6.1: Various cut variables and values for the prompt and delayed events. The efficiency of each cut on the reactor IBD events is detailed in Table 6.5.

Variable	Cut Value Prompt	Cut Value Delayed
Data-cleaning	0xFB0000017FFE	0xFB0000017FFE
NHits Clean	> 15	
Fit Valid	True	True
Energy	$\geq$ 2.5 MeV	$\leq$ 4.0 MeV
	$\leq$ 9.0 MeV	
Position	$\leq$ 5700 mm	$\leq$ 5700 mm
ITR	> 0.5	> 0.4
β_{14}	$\geq$ -0.6	$\geq$ -0.6
	≤ 1.6	≤ 1.6
σ_R	$\leq$ 600 mm	$\leq$ 900 mm
R_{FOM}	≥ 9.9	≥ 9.9
Energy G-Test	> 0.0	> 0.0
	< 1.45	< 1.45
Energy U-Test	> 0.5	> 0.5
Energy Z-Factor	< 1.0	< 1.0
Δt		≥ 3 μs
		≤ 500 μs
Δr		$\leq$ 3000 mm
Multiplicity		$= 1$

6.5.1 PDF Generation

As discussed in Section 6.4.3.2, there are several important reconstructed quantities that can be used to distinguish the reactor IBD signal from the accidental backgrounds. Many of these variables are already used in cuts, detailed in Table 6.1, but the shape of the distribution is also an important handle for background rejection. Various probability distribution functions (PDFs) are given in Figures 6.6 to 6.10. The PDFs provide, for a given bin, the likelihood that the variable would equal that sample. The PDFs have been selected to span the ranges inside of the cut values from Table 6.1, for example the prompt energy PDF spans 2.5 to 9 MeV. Both the signal and background the PDFs consist of 100000 events.

The signal PDFs are generated using reactor IBD Monte Carlo generated inside of the PSUP volume (sphere with radius of 8500 mm). All events selected for the PDFs pass the quality cuts in Table 6.1. The prompt events correspond to the signal from the positron and the delayed events correspond to the signal from the neutron capture.

The accidental background PDFs are generated using data according to the following process:

1. Prompt candidates are selected by identifying events that pass only the prompt cuts listed in Table 6.1.
2. Delayed candidates are selected by identifying events that pass only the delayed cuts listed in Table 6.1, except the Δt cut. The Δr cut is applied by choosing a position for the prompt event based on a distribution of prompt event positions (also taken from data).
3. Prompt and delayed event candidates are paired at random and a Δt is assigned by drawing from a flat distribution with values between 0 and 1000 μs. This ensures that random prompt and delayed candidates are paired together, and eliminates possible correlations between true event pairs in the data.
4. Delayed events above and below 15 nhit are treated distinctly to ensure that the delayed event PDFs sample the appropriate distributions, which are biased towards higher energy events by the method in which the data is processed.

This method for generating PDFs for the random, accidental backgrounds ensures that true coincidences in the data are not included in PDFs. This random re-ording of the events also allows the creation higher statistics PDFs, because no Δt condition is necessary.

The prompt energy PDFs, shown on the left in Figure 6.6, show that the accidental backgrounds are primarily low energy and are almost completely gone by 4 MeV. The signal PDF displays the expected positron energy distribution. The delayed energy distributions, the right plot in Figure 6.6, shows relatively less separation between the low energy signal and the background. The prompt and delayed β_{14} PDFs are shown in Figure 6.7, which provides particularly good discrimination for the prompt event. This is expected as β_{14} is used to characterize the isotropty of the event, which distinguishes the positron signal from the γ backgrounds.

Two dimensional PDFs that show the position, cubed and normalized by the AV radius as $(R/R_{AV})^3$, plotted against $u \cdot r$ are given for data and simulation for the prompt event in 6.8 and for the delayed event in 6.9. It is clear that the accidental PDFs are dominated by external events at high radius and that are inward pointing ($u \cdot r \approx -1$). These events are primarily γ-rays from the AV, ropes, external water, and the PMTs. The signal is flat in $u \cdot r$ and $(R/R_{AV})^3$, up to effects from the Δr requirement (i.e., prompt events at the edge of the fiducial volume are less likely to pass the Δr cut because there is less fiducial volume available for the delayed event). In the background PDFs, the externals dominate more strongly as backgrounds to the prompt event, whereas the background PDF for the delayed event is slightly flatter.

Lastly, the two dimensional Δt and Δr PDFs for data and MC are shown in Figure 6.10. Using the shape of these distributions is the strongest handle at rejecting accidental background. The two dimensional PDF utilizes the small correlation between the time difference and spatial separation of the positron and neutron.

The PDFs are generated using the first several hundred runs, corresponding to the first couple days of livetime. These runs will still be used in the final analysis. Any time variation

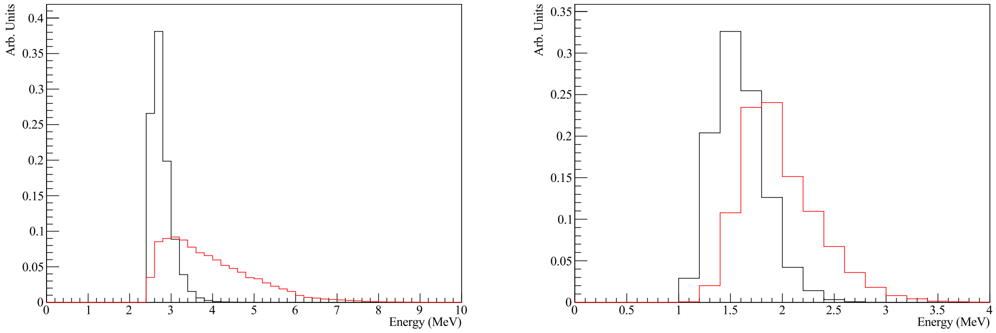

Figure 6.6: The prompt event (left) and delayed event (right) energy spectrum for data (black) and MC (red).

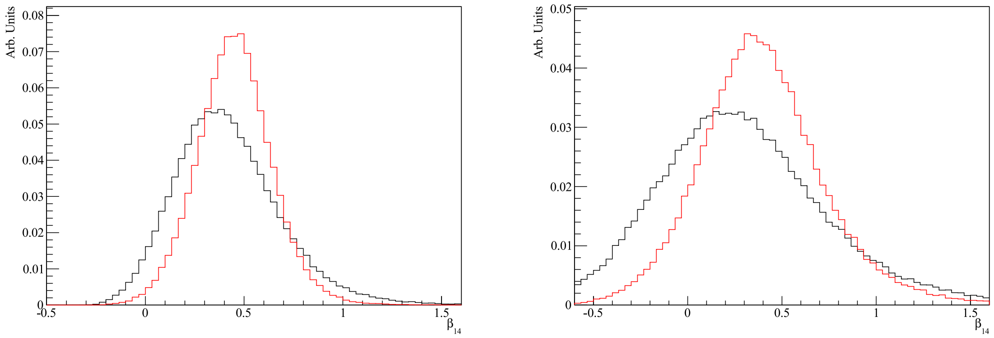

Figure 6.7: The prompt event (left) and delayed event (right) β_{14} spectrum for data (black) and MC (red).

to the shape of the PDFs is not included, but no significant time-variation is identified in any of the observables.

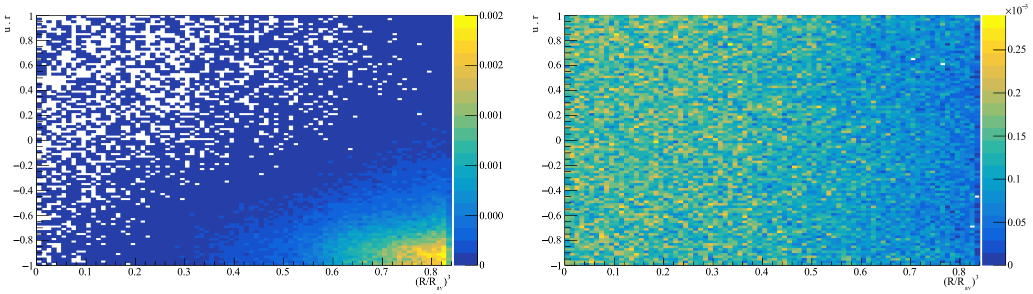

Figure 6.8: The data (left) and MC (right) (R/R$_{AV}$)^3 vs $u \cdot r$ for the prompt event.

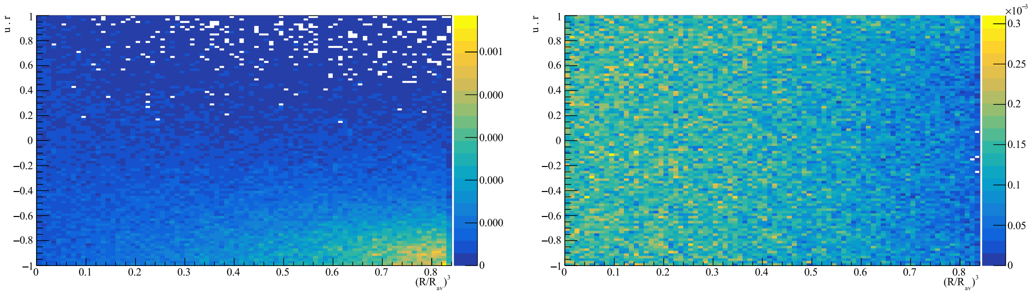

Figure 6.9: The data (left) and MC (right) (R/R$_{AV}$)^3 vs $u \cdot r$ for the delayed event.

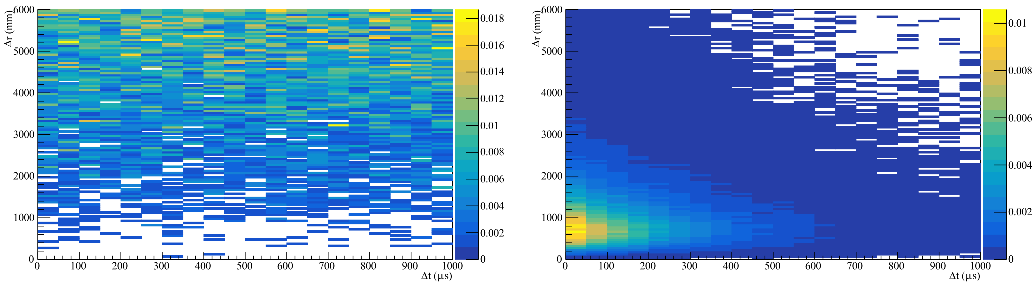

Figure 6.10: The data (left) and MC (right) Δt vs Δr.

6.5.2 Likelihood Ratio

Two likelihood are calculated using the PDFs for the prompt and delayed events. For coincidences of events that pass the cuts in Table 6.1, the likelihood for each measured observable is extracted by reading the probability from the correspoding bin in the PDF. The likelihood that the event is a reactor IBD event, $\mathrm{L}_{\bar{\nu}_e}$, is determined using the signal PDFs from simulation. The likelihood the event is an accidental background, $\mathrm{L}_{\mathrm{acc.}}$, is determined similarly using the PDFs generated from the data. The sum of the log of the likelihoods is calculated, and the ratio of the summed log-likelihoods:

$$\Delta\mathcal{L} = \log\left(\frac{\mathrm{L}_{\bar{\nu}_\mathrm{e}}}{\mathrm{L}_{\mathrm{acc.}}}\right) = \log(\mathrm{L}_{\bar{\nu}_\mathrm{e}}) - \log(\mathrm{L}_{\mathrm{acc.}}) \tag{6.5}$$

is used as the ultimate discriminant between signal and accidental backgrounds.

The $\Delta\mathcal{L}$ distribution for the simulated reactor IBD signal is shown in Figure 6.11. Given the definition of $\Delta\mathcal{L}$, signal-like events have values less than 0 and accidental-like events have values greater than 0. The small fraction of events in Figure 6.11 with values of $\Delta\mathcal{L}$ greater than 0 typically have large values of Δt and Δr, making them look more accidental-like.

6.6 Calibration using the AmBe Source

Prior to estimating the expected signal efficiency, the simulation must first be calibrated. The AmBe source provides an ideal dataset to use as a calibration for this analysis. The calibration run list, given in Appendix F, includes a (y, z) scan of the detector at various source positions out to 5.5 m from the center of the AV. The x position is kept fixed at zero in the scan. Note that these runs are taken in January of 2018, prior to the runs used in the data-set.

The AmBe source, introduced in Section 4.6.2, produces a 4.4 MeV γ, from the de-excitation of ^{12}C, and a neutron. The neutrons primarily thermalize in the detector, capture on hydrogen, and produce a 2.2 MeV γ. These signals span the relevant energy ranges

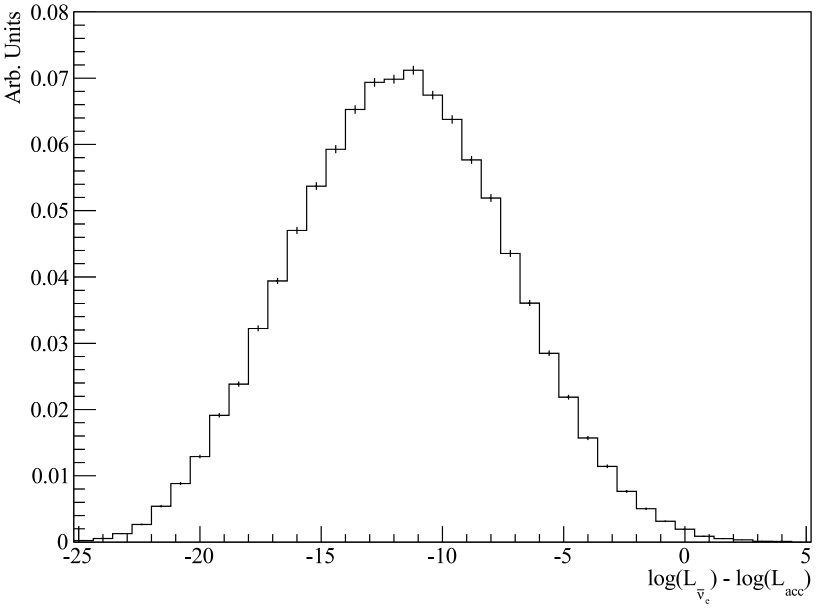

Figure 6.11: The $\Delta\mathcal{L}$ distribution for the simulated reactor IBD signal.

for this analysis and, importantly, consist of a coincidence of a prompt event and a delayed neutron capture. Unlike the ^{16}N source, the AmBe source events are not tagged, so particular care must be taken in understanding the purity of the event selection. The coincidence requirement is the most important factor in driving background leakage down. However, in addition to the cuts applied in Table 6.1, events in both the AmBe data and simulation are selected using:

1. The prompt event position must be within 1.8 m of the source
2. Δt $< 100\ \mu$s
3. Upper energy cut reduced to 6.0 MeV for the prompt event

The first two selections are designed to reduce accidental backgrounds and should not be strongly correlated with the other observables (i.e., selecting events at small values of Δt does not bias the selection). The upper energy cut is reduced to 6.0 MeV for the AmBe analysis because of a distinct high energy tail present in the MC but not the data (likely from poorly modeled ^{16}N production). To characterize the purity of the selection, the Δt distribution for the coincidences is fit out to 1 ms using:

$$f = N \times e^{-t/\tau} + B, \tag{6.6}$$

where N is normalization, τ is the neutron capture time constant, and B is the flat background component. Figure 6.12 shows the data with the corresponding fit for events that pass all analysis cuts but with the Δt window extended to 1 ms. From the fit values shown in Figure 6.12, the fraction of accidental events falling within the 3 to 100 μs window is 2.3%. Additionally, this same method is applied to the data after removing the data-cleaning cuts for the delayed event. The fit after removing data-cleaning indicates an increase in the total number of tagged coincidences by 1.7%, accounting for the increase in the accidental background rate.

Figure 6.13 shows important distributions compared between data and MC for the central AmBe run 109133. Notably the Δr distribution used in the likelihood calculation, shows significant disagreement between data and MC. This is caused by the broader reconstructed positions for both the prompt and delayed event in the data, as seen in Figure 6.14. Additionally, there is particularly poor agreement between data and MC in the ITR and position FOM distributions for both the prompt and delayed events. These general features are consistent across different source positions. The likelihood ratio, $\Delta\mathcal{L}$, is generated for the AmBe data and MC using these distributions and shown in Figure 6.15. The right panel shows the comparison after arbitrarily broadening the reconstructed position distributions in the MC for the prompt and delayed events by 220 mm and 310 mm respectively. This yields better agreement for the reconstructed x, y, and z distributions for the prompt and

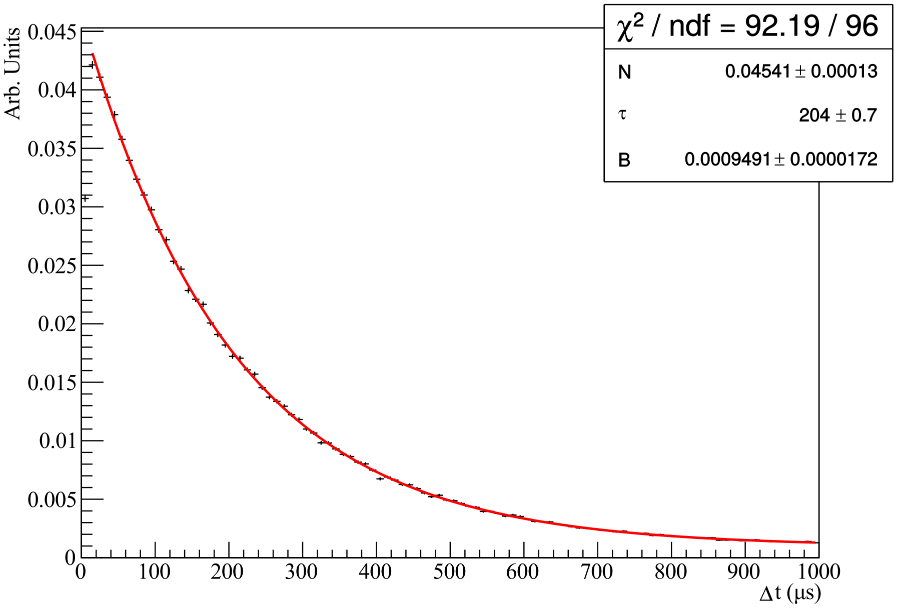

Figure 6.12: Δt distribution between the prompt and delayed AmBe events. The data is fit, excluding the first bin because of the retrigger data-cleaning cut, out to 1000 μs using Equation 6.6.

delayed events (top three panels in Figure 6.14) between data and MC, and thus better agreement in the Δr distribution. The $\Delta\mathcal{L}$ calculated after this arbitrary position broadening agrees better with the data, indicating the primary driver of the difference between data and MC in the $\Delta\mathcal{L}$ distribution is the shape difference of the Δr distribution.

6.6.1 Calibrating the Neutron Detection Efficiency

The most critical difference between the data and MC is in the detection efficiency of the 2.2 MeV γ-ray from the neutron capture on hydrogen. The detection efficiency of this low energy interaction is largely affected by the trigger efficiency, which is not

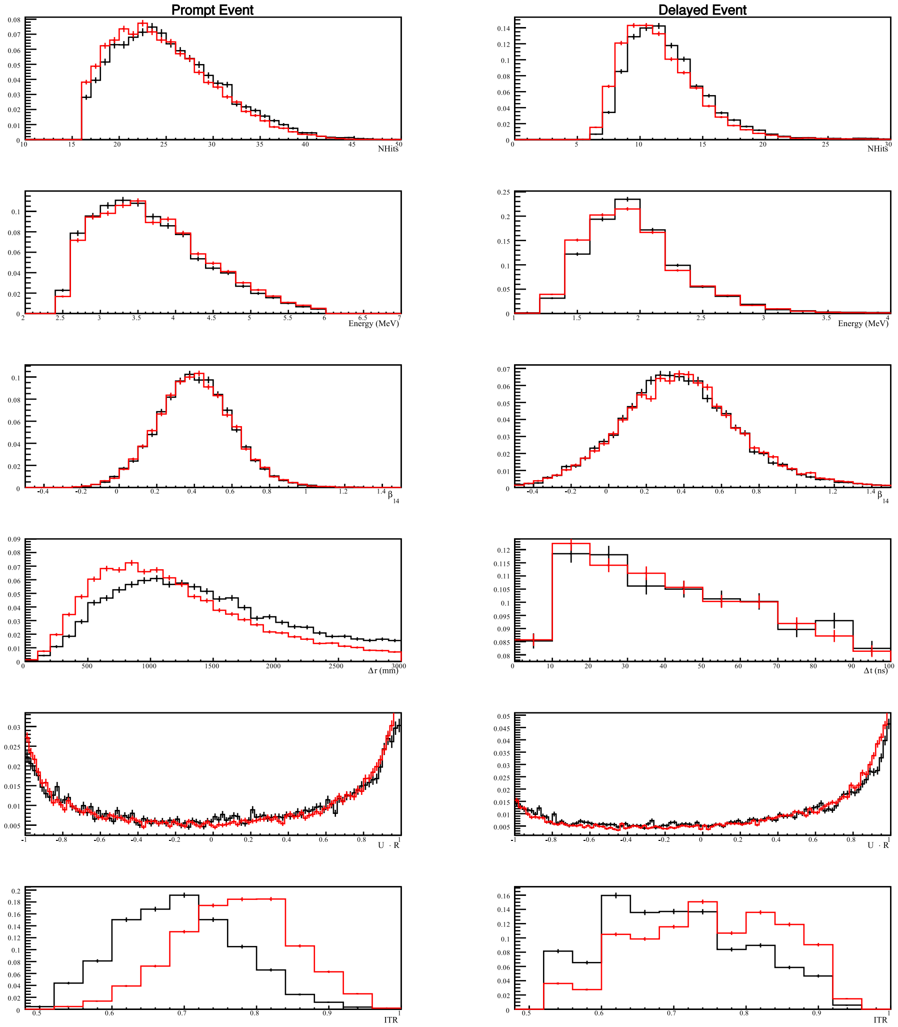

Figure 6.13: The left panels show the data (black) and MC (red) comparison for the prompt event in an AmBe central run. The right panels show the data (black) and MC (red) comparison for the delayed event in an AmBe central run. There are notable differences between the data and MC in the Δr and ITR distributions.

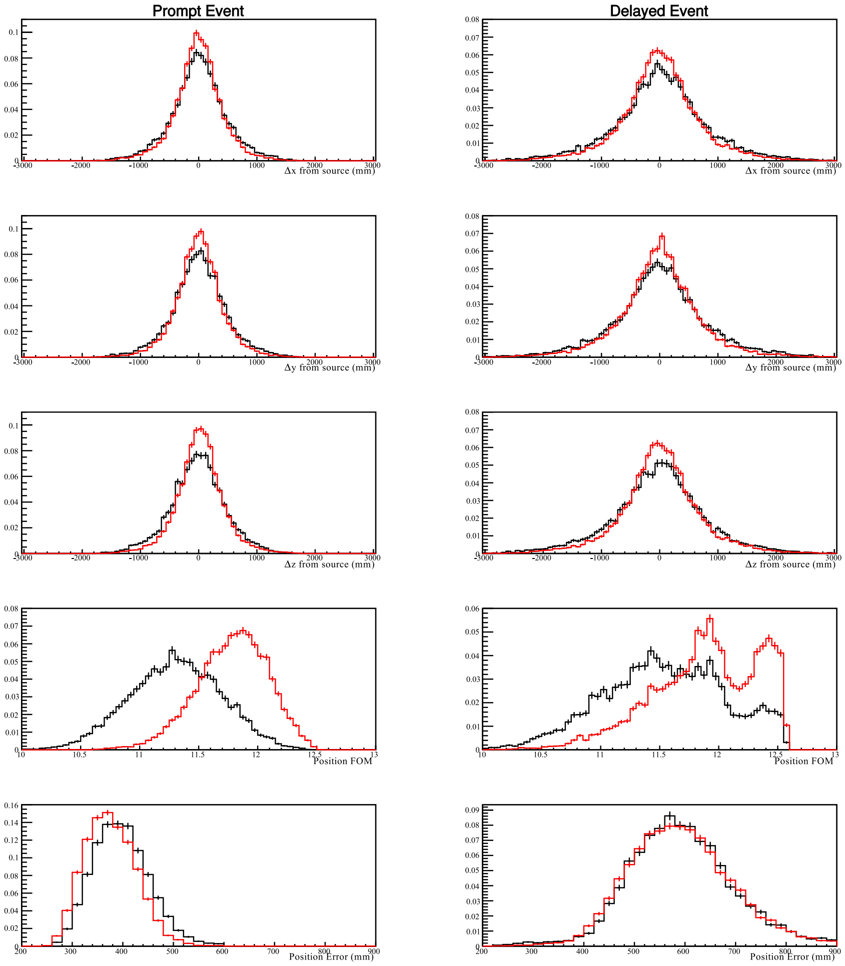

Figure 6.14: The left panels show the data (black) and MC (red) comparison for the prompt event in an AmBe central run. The right panels show the data (black) and MC (red) comparison for the delayed event in an AmBe central run. There are notable differences between the data and MC in the x, y, and z position distributions, which lead to the difference in the Δr distribution identified in Figure 6.13.

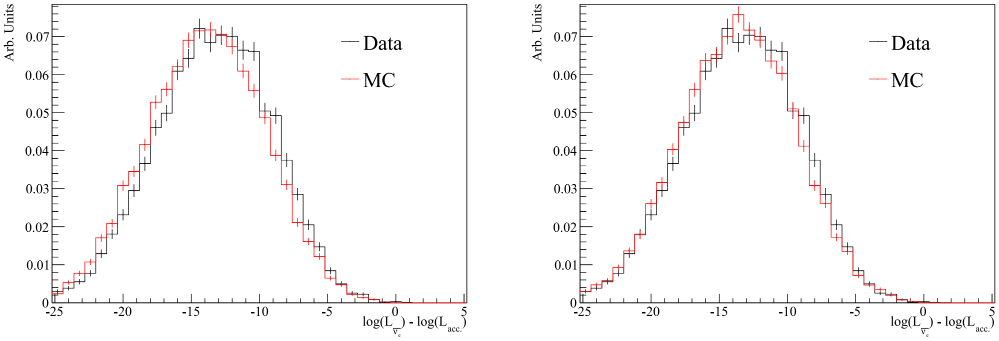

Figure 6.15: The $\Delta\mathcal{L}$ distribution for the AmBe calibration data and simulation. The right panel shows the distribution after arbitrarily broadening the reconstructed position distributions for the prompt and delayed events by 220 mm and 310 mm respectively.

modeled perfectly in the simulation. Section 4.7 described the measurement of the trigger efficiency, often comparing the nhit monitor to the model of the trigger in `RAT`. Notably, a measurement of the trigger efficiency using TELLIE showed about a 15% lower efficiency than that predicted by `RAT`. This is qualitatively confirmed by investigating a central AmBe run, where the detection efficiency in MC is 15.2% and the detection efficiency in MC is 12.2%. Note that these overall efficiencies should not be mistaken for the efficiency of the reactor antineutrino selection as there are additional criterion for the AmBe processing used to purify the sample (e.g., Δt $< 100\ \mu$s). The important quantity is the relative predicted efficiency, which in this case is $\frac{12.2}{15.2} = 80.2\%$. In other words, the MC is over-predicting the measured neutron detection efficiency by about 20% for events near the center of the detector. Table 6.2 investigates the effect of the trigger efficiency on the relative efficiency directly by scanning over a cut on the in-time hits, $\tilde{n}_{100}$. As the cut value on $\tilde{n}_{100}$ increases, differences between the data and MC are reduced, until after about $\tilde{n}_{100} > 10$, after which the trigger in both the data and MC should be 100% efficient. This scan makes clear that the trigger efficiency accounts for the vast majority of the difference between data and MC.

Table 6.2: The relative efficiency (data/MC) for two runs for different $\tilde{n}_{100}$ cut values. The AmBe calibration source in run 109133 is in the center and in 109153 it is located at (0, -4, 0) m.

$\tilde{n}_{100}$ Cut Value	Rel. Eff (%), 109133	Rel. Eff (%), 109153
None	80.3	79.0
6	82.0	80.6
7	89.1	87.4
8	94.0	91.1
9	97.1	93.4
10	100	96.4
11	103	102

To account for the poor modeling of the trigger efficiency, a correction must be applied to the simulated events when estimating the signal (and background) efficiencies. This correction is investigated as a function of the reconstructed position of the prompt event. Figure 6.16 shows the neutron detection efficiencies for data (left) and MC (right) as a function of the prompt event $\rho = \sqrt{x^2 + y^2}$ and z positions. Both the data and MC predict qualitatively similar effects, where the efficiency decreases at large and small values of z, but increases with r towards the central-edge of the fiducial volume. As mentioned, the data also shows consistently smaller efficiencies than predicted by the MC. However, there are several corrections that must be applied prior to calculating the final relative efficiency (data/MC) map. The most important correction, which handles a change to the trigger system made between the AmBe runs and the analysis runs, is discussed in Section 6.6.1.1. A summary of the corrections and the final relative efficiency map are given in Section 6.6.1.2

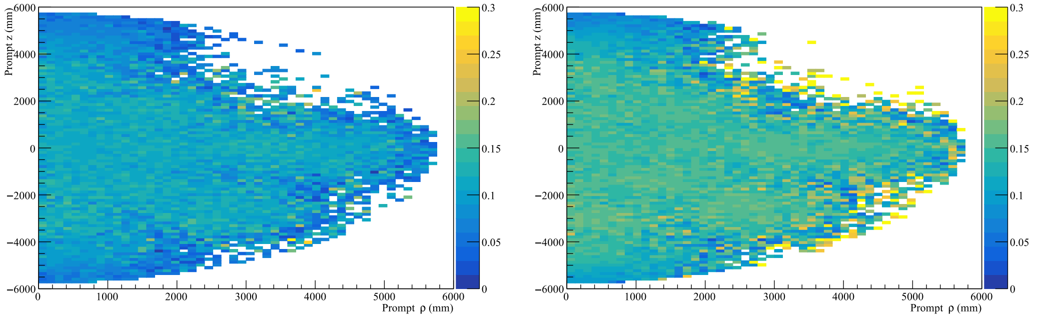

Figure 6.16: The neutron detection efficiency for data (left) and MC (right) as a function of the prompt event position. The z-axis (color scale) shows the efficiency and the MC is generally predicts a higher efficiency than identified in the data.

6.6.1.1 Change to the SNO+ Trigger System

As noted before, the AmBe dataset (January 2018) was taken several months prior to the beginning of the analysis dataset (October 2018). In June of 2018, between the AmBe and analysis datasets, an upgrade to the N100 MTC/A+ was installed, which was designed to immprove the rise-time of the trigger signal. This section focuses on the fact that the trigger efficiency changed with the new MTC/A+. A dynamic model in RAT accounts for this change; in a simulation of 2.2 MeV γs distributed throughout the AV the trigger efficiency predicted for run 109133 is 55.4% and the trigger efficiency predicted for run 200005 is 50.0%. This decreased efficiency is largely accounted for by the improved rise-time, yielding fewer detected events at low nhit. The trigger efficiency curves from the `RAT` model and the nhit monitor, with the corresponding in-time hits distributions, is shown in Figure 6.17.

Importantly, from Figure 6.17 the relative difference in predicted efficiency between the `RAT` model and the nhit monitor flips between run 109133 and 200005. That is, in run 109133 the trigger model predicts a higher efficiency than the nhit monitor, while in run 200005

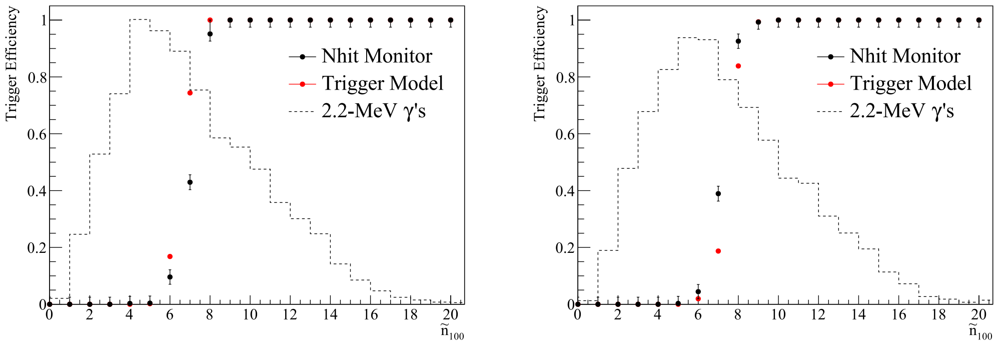

Figure 6.17: The trigger efficiency comparison between the nhit monitor and the RAT trigger model for run 109133 (left) and run 200005 (right). The in-time hit ($\tilde{n}_{100}$) distribution is shown in the dashed line. The error bars are stat only.

the trigger model predicts a lower efficiency than the nhit monitor. This trend, identified for these two runs, is consistent across a wide range of runs. Generally, `RAT` predicts a total efficiency ∼10% higher than the nhit monitor predicts during the AmBe runs and ∼5% lower than the nhit monitor predicts during the analysis runs. Thus, during the AmBe runs it is expected that `RAT` over predicts the total efficiency. This over prediction has already been observed in the relative efficiency calculated for the central AmBe run. Any correction applied to the analysis runs must account for the fact that the relative trigger behavior has changed. To do this, the relative efficiency in the AmBe data is scaled up by 15%. Because the nhit monitor is not a particularly reliable method for extracting the true trigger efficiency, a 100% uncertainty is taken on this scaling.

6.6.1.2 Relative Efficiency Correction

The factors that are included as scaling to the relative efficiency are shown in Table 6.3. The trigger efficiency correction and uncertainty, detailed in Section 6.6.1.1, dominates the correction factor. After applying these scale factors and after all cuts, including the likelihood ratio cut at < -12, the relative efficiency (data/MC) is shown in Figure 6.18. Because the AmBe source does not extend to high radii, there are empty bins near the edge of the fiducial volume. To simplify the application of the correction factors, the relative efficiency is averaged in four regions of the detectors, shown in Table 6.4. Applying this region-by-region scaling to the simulated reactor IBD events leads to a decrease in the total number of expected events of 92.1 ± 21.3%. Note this correction factor is used for the signal and all backgrounds except the (α, n) events along the AV, which have a different radial distribution. For those events, a scaling of 87.4 ± 24.0% is calculated.

Table 6.3: The various corrections factors applied to the relative efficiency maps generated using the AmBe source.

Scaling	Value
Accidentals	-2.32 ± 0.04%
Data-cleaning sacrifice	+1.73 ± 0.03%
Trigger efficiency changes	+15 ± 15%

Table 6.4: The relative efficiency (data/MC) correction factors for four regions of the detector, after applying all cuts in Table 6.1, the additional AmBe criteria, and the likelihood ratio cut. Note these regions are mutual exclusive in descending order (i.e., z $\geq$ 4 includes all ρ, whereas $\rho \geq 3.5$ is for z positions between 4 and -4 m). The uncertainties are taken by using the standard deviation of the average as the statistical uncertainty, in quadrature with the systematic uncertainty from the trigger efficiency variation scaling (15%).

Position (m)	Correction
z $\geq$ 4	97.8 ± 18.4
z $\leq$ -4	87.6 ± 18.1
$\rho \geq 3.5$	85.4 ± 26.2
$\rho < 3.5$	98.1 ± 17.5

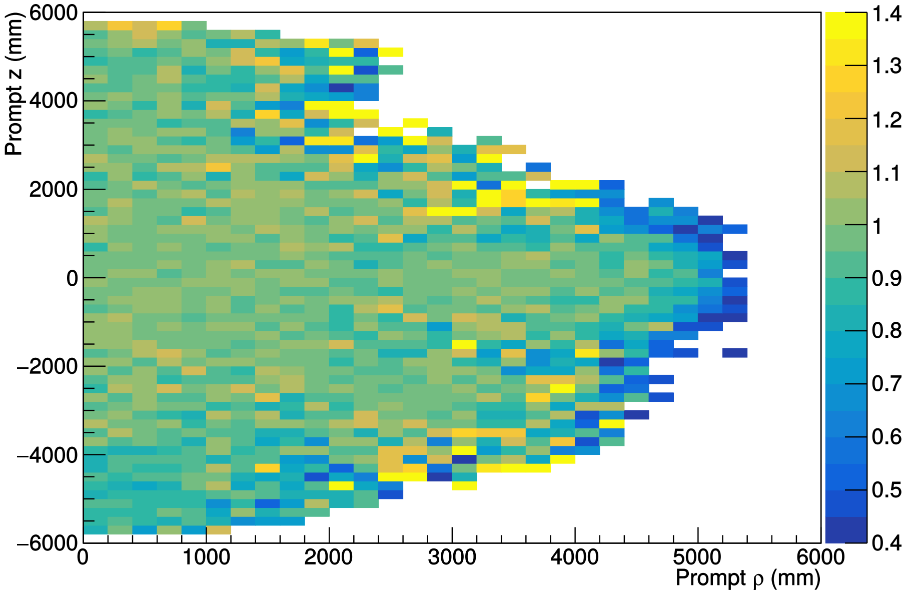

Figure 6.18: The relative efficiency (data/MC) for detecting the delayed neutron capture event using a cut at $\Delta\mathcal{L} < -12$. The correction factors from Table 6.3 are applied. The lack of data at high z and large ρ is due to limited source deployements in that area, as can be seen in Appendix F.

6.7 Signal Detection Efficiency

The interaction rate of reactor $\bar{\nu}_e$ interactions per year is given in Section 6.2 as 110.4 (313.9) in the AV (PSUP) regions. The simulation of reactor $\bar{\nu}_e$ occurs within the PSUP region, in a spherical volume with radius 8.5 m. The cuts from Table 6.1 are applied directly to the simulation, which yields an expected efficiency of 3.905% for events simulated within the PSUP region, as shown in Table 6.5 (note that this efficiency includes the fiducial volume of 5.7 m). This yields an expected 6.41 events in 190.3 days of livetime, prior to the cut on the likelihood ratio.

Table 6.5: The breakdown of the signal efficiency of each cut on the prompt and delayed events. The events are simulated in a volume of 8500 mm. The details of the cuts are given in Table 6.1, some of which are reproduced in this table. An example likelihood ratio cut is included for completeness and does not indicate where the cut will be placed in the final analysis. No scale factors or sacrifices are applied to the efficiencies presented in this table, and additional scaling for differences in the neutron detection efficiency reduce the overall detection efficiency (see Table 6.6).

Cut	Value	Efficiency (%)	Cumulative Efficiency (%)
Prompt			
Trigger		76.99	76.99
NHits	> 15	52.08	40.01
Fit valid	True	94.05	37.71
Energy	≥ 2.5 MeV ≤ 9.0 MeV	81.34	30.67
Fiducial Volume	≤ 5700 mm	42.12	12.92
Classifiers	see Table 6.1	98.67	12.75
FOMs	see Table 6.1	99.59	12.70
Multiplicity	$= 1$	100.00	12.70
Delayed			
Trigger		47.570	6.057
Fit valid	True	95.96	5.812
Energy	< 9 MeV	99.88	5.801
Fiducial Volume	≤ 5700 mm	81.86	4.752
Classifiers	see Table 6.1	98.00	4.657
FOMs	see Table 6.1	97.45	4.539
Δt	$\geq 3\ \mu$s $\leq 500\ \mu$s	90.76	4.119
Δr	≤ 3000 mm	94.79	3.905
$\Delta\mathcal{L}$	$<$ -12	50.91	1.988

The sacrifice of several of the cuts should be included in the efficiency estimate. First, the data-cleaning sacrifice is estimated using the AmBe data by comparing the total number of coincidence events (accounting for an increase in the accidental rate) before and after applying data-cleaning. The sacrifices are 1.73 ±0.03% and 2.35±0.04% for the delayed and prompt events respectively. The sacrifice of the multiplicity cut is estimated by applying it to randomly selected prompt candidate events and is 0.2%. The relative neutron detection efficiency scaling of 92.1 ± 21.3%, calculated in Section 6.6.1.2, is also applied. All of the various scale factors are summarized in Table 6.6. This reduces the total expected efficiency to 3.44 ± 0.73%. The large uncertainty is dominated by the 100% uncertainty taken on the trigger efficiency scaling. Multiplying this efficiency by 313.9 expected interactions in the PSUP volume yields 10.82 ± 2.30 expected coincidence events per year or 5.65 ± 1.20 events in 190.3 days of livetime. Note that this will be reduced when the likelihood ratio cut is applied. Figure 6.19, which breaks the predicted efficiency down in 3-day periods, shows that the time variations in the efficiency are well-described by the MC stats. The fitted efficiency from Figure 6.19 is consistent with the value presented in Table 6.5 after the likelihood cut.

Table 6.6: The various correction factors used to adjust the expected detection efficiency. The last two factors are specific to the reactor IBD events.

Correction	Scale Factor
Data-cleaning, prompt	0.976 ± 0.0004
Data-cleaning, delayed	0.983 ± 0.0003
Multiplicity	0.998 ± 0.0002
AmBe relative scaling	0.921 ± 0.213
Model scaling (114)	0.952 ± 0.027
Realistic reactor powers	0.844 ± 0.084

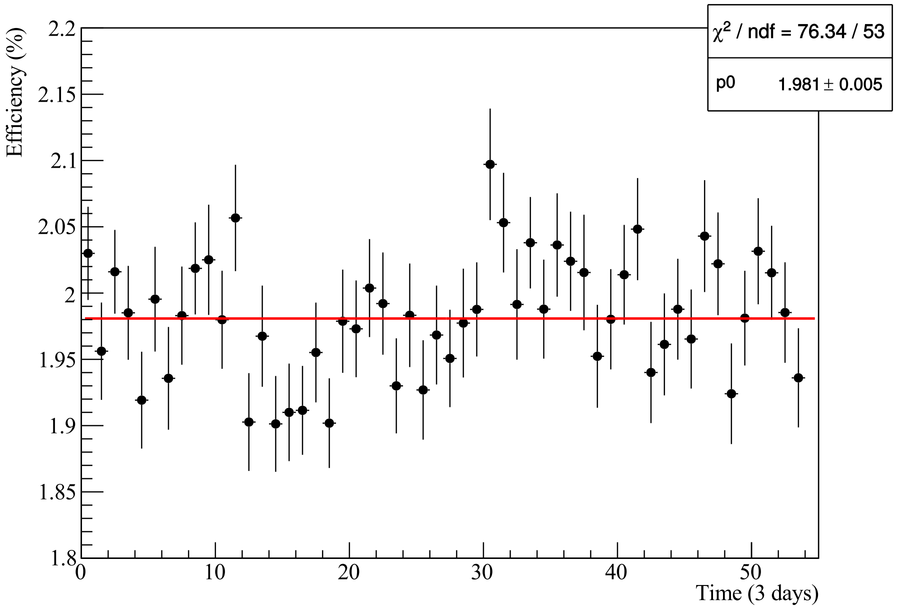

Figure 6.19: The detection efficiency predicted by the MC after the cuts described in Table 6.1, including the $\Delta\mathcal{L}$ cut $<$ -12, which can be directly compared to Table 6.5. The runs are combined in three-day chunks to average over the statistical variations in the number of events simulated. The variations in the detection efficiency are well described by statistical fluctuations.

6.8 Backgrounds

The primary backgrounds to the antineutrino search are accidental coincidences, (α, n) events, and atmospheric neutrino interactions. Geoneutrinos and cosmogenically activated isotopes from muons are negligible contributions to the total background.

6.8.1 Accidentals

The term accidental is used to described the coincidence of two uncorrelated events that, by chance, happen to both occur in the selected time-window. Given the rate of radioactive decays in the external and internal volumes, particularly at low energy, there is a significant number of events that can pile-up and mimic the coincidence signal. Specifically, ^{214}Bi and ^{208}Tl decays in the internal and external water, and along the acrylic vessel, ropes, and PMTs yield relatively high rates of events below 4 MeV. The external volumes have higher level of radioactivity, resulting in a larger number of random backgrounds along the outer edge of the fiducial volume. This background is unique in that it is the only background that does not consist of truly correlated coincidences (with a delayed neutron capture). Figure 6.20 shows the Δt distribution for the entire dataset, for events that pass the cuts detailed in Table 6.1, prior to the likelihood ratio cut. A flat distribution is expected for accidentals in this Δt window, based on the rates of the prompt-like and delayed-like events. This figure indicates that the event count is dominated by accidentals, which will need to be drastically reduced by the likelihood ratio cut. A total of 1991 events are identified between 3 and 500 μs and a total of 1981 events are identified between 500 and 1000 μs.

To understand the accidental coincidences in the data, the rates of prompt and delayed events are investigated separately. This process uses the cuts in Table 6.1, but applied to the prompt and delayed candidates separately (removing the Δt condition). The Δr cut for the delayed event is applied by selecting at random a position for the prompt event (from the prompt event position PDF). The prompt and delayed event rates for the entire run range are shown in Figures 6.21 and 6.22 respectively. The event rate jump for the prompt event around run 203000 is caused by a period of detector maintenance that led to more online channels. This raised the average nhit distribution, leading to more events passing the 15 nhit cut. While an energy cut would calibrate for the larger number of online channels, the nhits cut on the prompt event causes the event rate to depend on the total

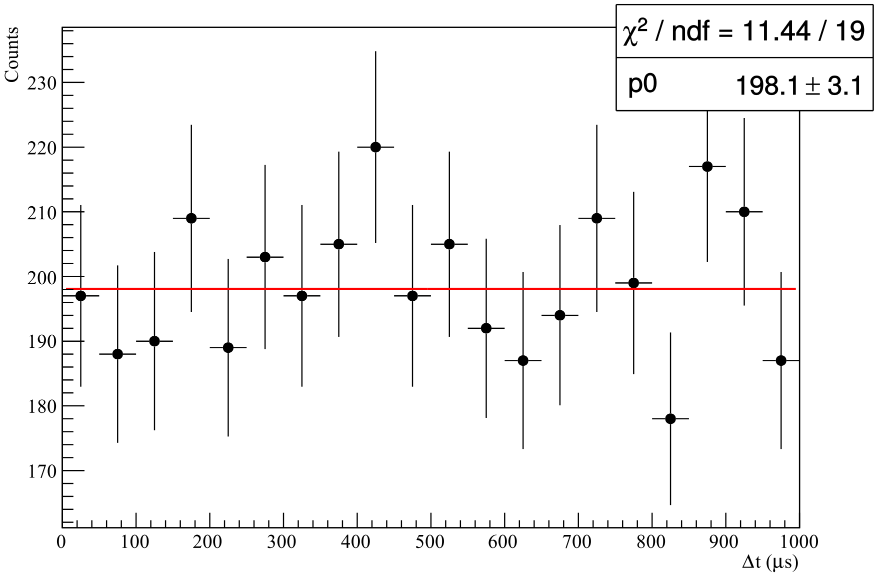

Figure 6.20: The Δt between all events that pass the cuts in 6.1. The event count is dominated by accidental backgrounds prior to the likelihood ratio cut, which results in the identified flat Δt distribution.

number of online channels. Because there is no nhit cut on the delayed event, there is no corresponding rate change.

Using these rates, the expected number of accidental coincidence events, $N_{acc.}$, can be calculated using:

$$N_{acc.} = r_p \times r_d \times \Delta t \times T, \tag{6.7}$$

where r_p is the rate of the prompt events, r_d is the rate of the delayed events, Δt = 497 μs is the coincidence window, and T is the livetime. Using Equation 6.7, the total number of expected accidental backgrounds, prior to the application of likelihood ratio cut, is calculated. This calculation can be directly compared to an observation of $N_{acc.}$, using a sideband window in Δt between 500 and 1000 μs (a sideband is chosen to ensure blindness).

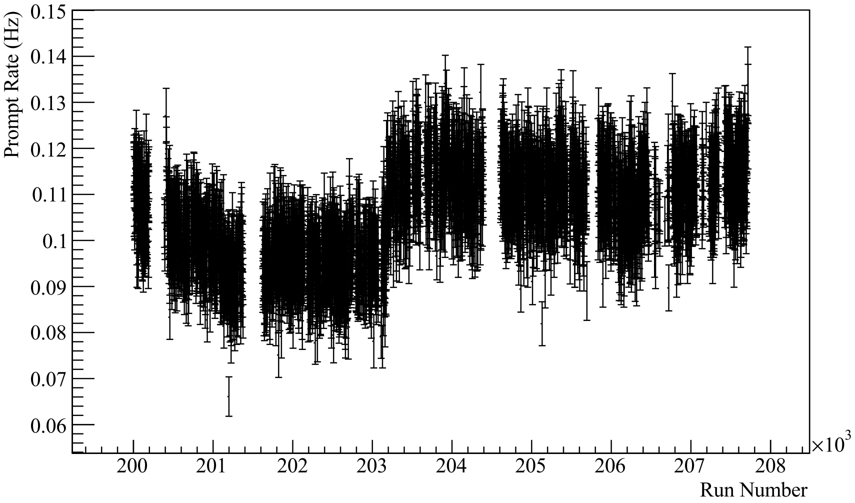

Figure 6.21: The prompt event rate by run-number.

Figure 6.23 shows this comparison, broken up in three-day periods, and indicates that the simple calculation is predictive of the observation. Note, the calculation of the expected rate is not used in the prediction for the accidental background, and is shown here to demonstrate the accidentals are well-understood.

6.8.1.1 Δt Sideband

To estimate the number of accidental events in the signal ROI, a sideband in Δt is selected. The signal region extends from 3 to 500 μs, so a window between 500 and 1000 μs is selected to look for accidental coincidences. The reactor IBD signal efficiency in this window (for events simulated in the PSUP) is about 0.34%, yielding an expected 0.49 events in 190.3 days. Thus, there is very little expected signal contamination in this window designed to measure the accidentals backgrounds.

This sideband region is used tune the cut on the likelihood ratio distribution without

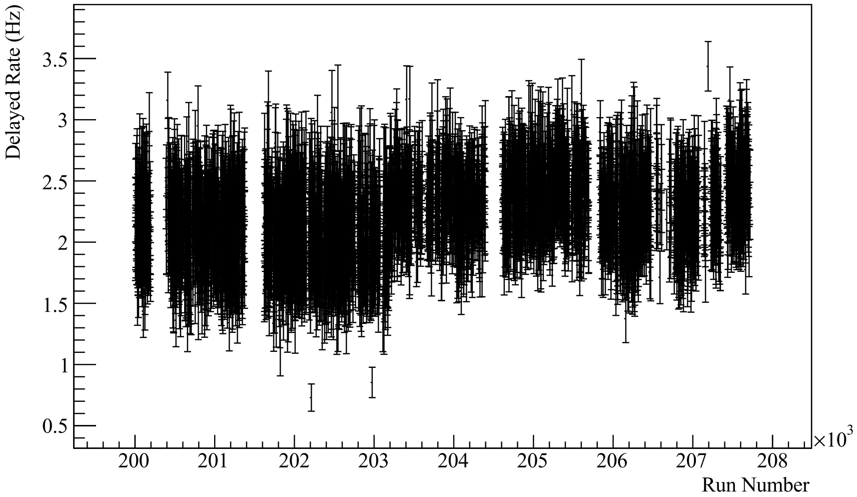

Figure 6.22: The delayed event rates by run-number.

unblinding the signal region. Because the Δt between the events is used in the calculation of the likelihood ratio, $\Delta\mathcal{L}$, the value of $\Delta\mathcal{L}$ shifts more 'accidental-like' in the sideband (events with larger Δt). This is seen in Figure 6.24, which shows $\Delta\mathcal{L}$ compared for the signal and sideband regions. Good agreement between the distributions is identified after shifting the entire $\Delta\mathcal{L}$ distribution for the events in the sideband. Because the accidental PDF is flat in Δt, the shift value, S, is given as:

$$S = \log\left(\int_{3}^{500} \exp^{-t/204.0}\right) - \log\left(\int_{500}^{1000} \exp^{-t/204.0}\right) = 2.43. \tag{6.8}$$

This indicates the sideband can be used to tune the cut on $\Delta\mathcal{L}$. The cut value is then shifted by S when looking at events in the signal region.

Running the analysis over the data in the Δt sideband, using the cuts described in Table 6.1, yields 1991 observed events or 10.5 events/day. The breakdown of the number of events passing each set of cuts is shown in Table 6.7. The resulting likelihood distribution for the

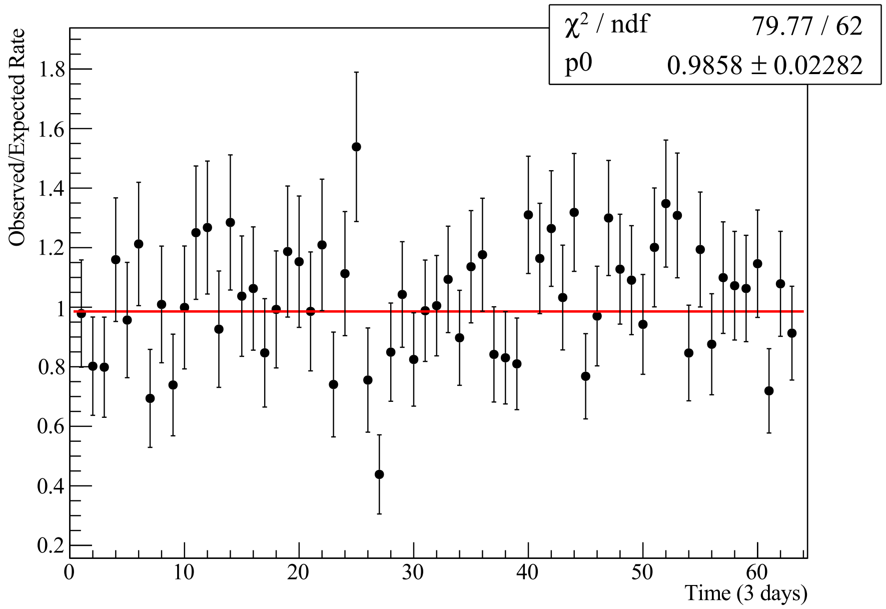

Figure 6.23: The observed number of counts in three-day chunks of livetime divided by the expected counts, based on the average prompt and delayed event rates in the three-day period. The fitted value of the flat constant is consistent with 1.0, indicating that the calculation of the expectation is predictive of the observation.

signal MC and the data is shown in Figure 6.25. It is clear that the likelihood separates the signal and background effectively. Indeed, there are no events in the sideband below a value of $\Delta\mathcal{L} < -9$, corresponding to $\Delta\mathcal{L} < -11.43$ in the signal window.

In order to make a robust prediction for the accidental background rate, given a cut on $\Delta\mathcal{L}$, a dataset of prompt candidate events (passing only the prompt cuts) and delayed candidate events (passing only the delayed cuts) is created. These candidate events pass all of the various cuts except no Δt or Δr cut is applied. Then, for each prompt candidate, a corresponding delayed candidate is selected at random, and the positions of the events are

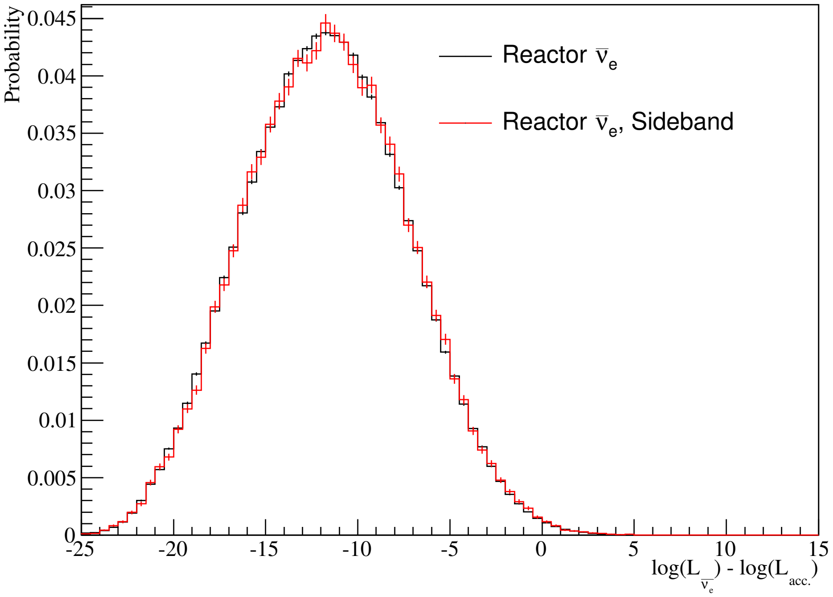

Figure 6.24: Likelihood ratio in the sideband region compared to the signal region after shifting the sideband distribution by $-S$.

used to calculate Δr. For events that pass the Δr cut, a uniform Δt is assumed (between 0 and 500 μs for the signal region or 500 and 1000 μs for the sideband region), and $\Delta\mathcal{L}$ is calculated. This calculated likelihood ratio for the 'mixed' dataset is shown compared to the actual coincidences in the sideband in Figure 6.26. As expected, the $\Delta\mathcal{L}$ distribution in the sideband is well described by the mixed dataset. The $\Delta\mathcal{L}$ distribution for the mixed data is compared to the distribution for the IBD signal in Figure 6.27. It is clear a cut value around -10 will remove the vast majority of the accidental backgrounds, while removing about 50% of additional signal efficiency.

Ultimately, the mixed dataset is used to generate a prediction for the accidental background rate after the likelihood ratio cut. The mixed dataset provides a high statistics

Table 6.7: The breakdown of the effect of each cut on the prompt and delayed events in the data. The data-cleaning, pedestal cut, and fit valid cut are applied to all events in order to trim the data. The rest of the cuts are separated by prompt and delayed events. The details of the cuts are given in Table 6.1, some of which are reproduce in this table. Note that the Δt cut is shown twice; the $\Delta t < 5$ ms defines the coincidence and the second Δt cut rejects backgrounds, although in this case the cut is shown for the sideband analysis. An example likelihood ratio cut is included for completeness and does not indicate where the cut will be placed in the final analysis.

Cut	Value	Event Count	Efficiency (%)	Cumulative Efficiency (%)
Total		1.248e10		
Data-cleaning	0xFB0000017FFE	7.304e9	58.53	58.53
Fit valid	True	1.331e8	1.823	1.067
Prompt				
NHits	> 15	6.561e7	49.28	0.5258
Fit valid	True			
Energy	≥ 2.5 MeV	3.343e7	50.95	0.2679
	≤ 9.0 MeV			
Fiducial Volume	≤ 5700 mm	2.121e6	6.344	0.01700
Classifiers	see Table 6.1	1.926e6	90.83	0.01524
FOMs	see Table 6.1	1.804e6	93.64	0.01446
Delayed				
Δt	< 5 ms	235537	13.01	1.887e-5
Energy	< 9 MeV	235537	100	1.887e-5
Fiducial Volume	≤ 5700 mm	80910	34.39	6.483e-6
Classifiers	see Table 6.1	67011	82.82	5.370e-6
FOMs	see Table 6.1	52816	78.82	4.232e-6
Δt	≥ 500 μs	26410	50.00	2.116e-6
	≤ 1000 μs			
Δr	≤ 3000 mm	1983	7.509	1.589e-7
Multiplicity	$= 1$	1981	99.90	1.587e-7
$\Delta\mathcal{L}$	$<$ -12	0	0	0

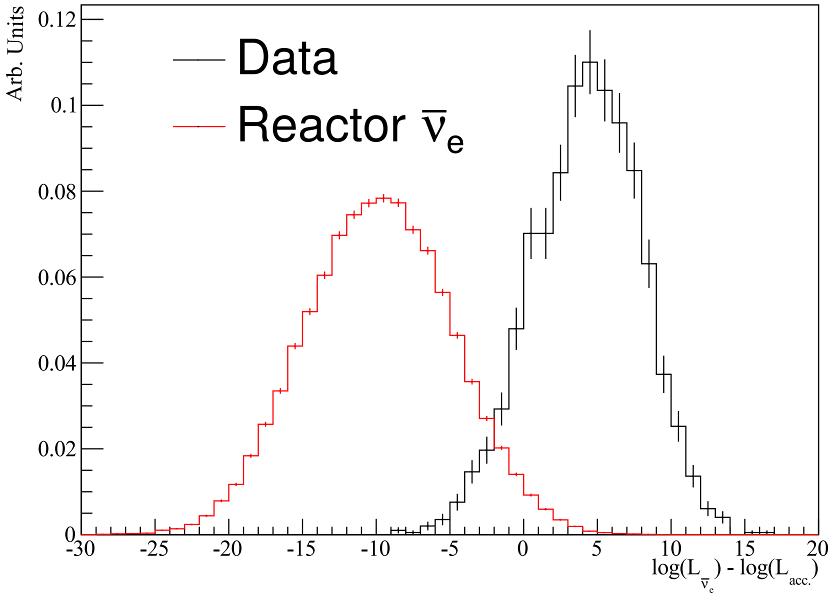

Figure 6.25: The comparison between the signal MC and the data in the Δt sideband. Excellent separation between the signal and background is apparent. Neither of these distributions are shifted by S. There are 1991 events in the accidental (data) histogram, which is normalized to one to compare against the simulated signal events.

sample of accidental coincidences and has been shown to recover the shape of the observed $\Delta\mathcal{L}$ distribution in the sideband. For example, based on Figure 6.27, the optimized $\Delta\mathcal{L}$ cut might be located around -12. In the mixed dataset, a total of 46 out of 136515 events pass this cut, yielding an efficiency of 0.031%. Note that even with this high statistics mixed dataset, the uncertainty on the efficiency is dominated by statistical limitations. In other words, the fractional uncertainty in the efficiency in this example is $\sqrt{46}/46 = 14.7\%$. In the sideband window 1991 events are observed prior to the likelihood cut, which we multiply by the efficiency to calculated the expected background. Thus, a total of 0.62 ± 0.09

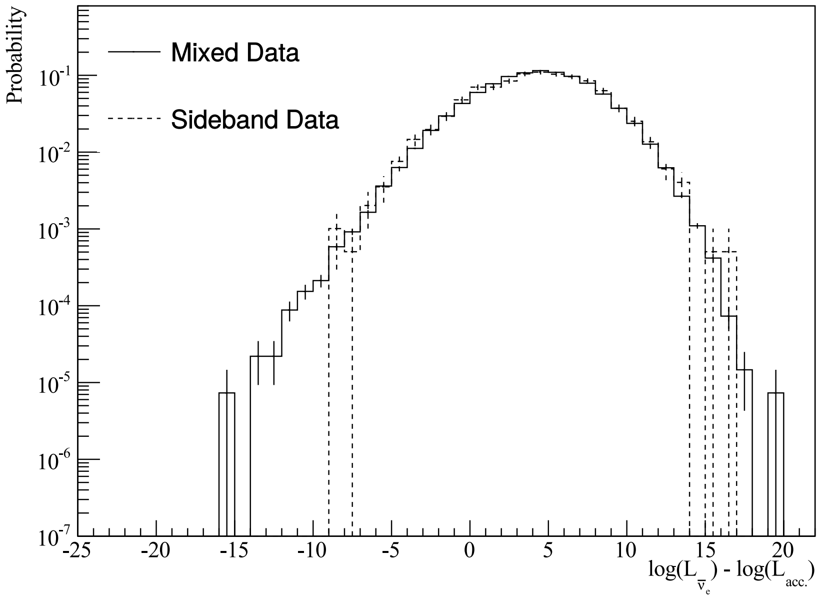

Figure 6.26: The data in the sideband compared directly to the 'mixed' dataset, generated by randomly choosing prompt and delayed candidates, pairing them, and assigning a Δt between 500 and 1000 μs. The mixed dataset matches the shape of the data nicely.

accidental events are expected to pass this example likelihood cut. After shifting the $\Delta\mathcal{L}$ values of the events in the sideband by S there are 0 events that pass this cut.

6.8.2 (α, n)

Background coincidence signals can be produced via the interaction of energetic α particles with the ^{13}C and ^{18}O atoms in the water and AV. In the uranium and thorium decay chains there are about a dozen different isotopes that α-decay, producing α particles with energies in the range of 5 to 10 MeV. By far the most abundant isotope in the detector that produces α particles is ^{210}Po. The primary source of the ^{210}Po is the radon deposition on

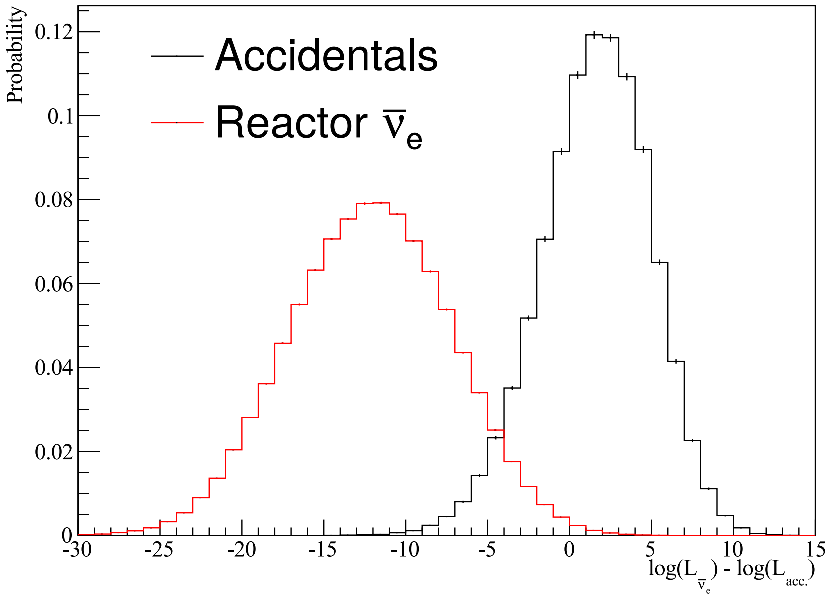

Figure 6.27: The reactor IBD signal MC compared directly to the mixed data. This plot shows qualitatively where the $\Delta\mathcal{L}$ cut will go in order to remove the accidental backgrounds. The sideband shift is not applied to either distribution.

the AV surface during the construction of SNO and in the period while the SNO detector was empty (prior to filling for SNO+). The radon daughters along the AV surface can leach off of the vessel into the water during data-taking. Because the α particles are invisible in a water Cherenkov detector, an ex-situ measurement of the expected rate of α-decays is required.

The initial activity of ^{210}Po on the inner AV surface was measured in 2013, prior to filling SNO+ with water, and found to be 1.9 ± 0.4 Bq/m^2 at the AV bottom and 3.2 ± 0.4 Bq/m^2 at the AV pipes. The average activity across the vessel is about 2.4 ± 0.8 Bq/m^2, or about 1150 Bq on the inner AV surface (125). The total rate leached into the internal water at a given time, t, is given by:

$$R_L = R_S(1 - e^{Lt}), \tag{6.9}$$

where R_L is the rate of events leached into the water, R_S is the initial surface activity rate, and L is the leaching rate of ^{210}Po into the water. The rate of leaching of the ^{210}Po into the water depends on the temperature of the water, and for water at 18.4°C the leaching rate is $1.05^{+0.18}_{-0.19} \times 10^{-3}$ day^{-1} (126). Using this model, half of the total surface activity leaches into the water after two years. Additionally, the ^{210}Po concentration in the internal water is affected by the water recirculation though the underground water purification plant, which reduces the concentration of ^{210}Po by a factor of 0.41 per pass (4).

The (α, n) reactions important for the SNO+ water phase are:

$$^{13}\mathrm{C} + \alpha \to {}^{16}\mathrm{O} + \mathrm{n} \tag{6.10}$$

and

$$^{18}\mathrm{O} + \alpha \to {}^{21}\mathrm{Ne} + \mathrm{n}. \tag{6.11}$$

The former reaction is an important background for KamLAND, and is discussed in (127), (128), (129), (130). The ^{16}O and ^{21}Ne are often produced in excited states, which lead to de-excitation products, most commonly γ-rays. The de-excitation γ-rays produce a prompt signal and the neutron will capture on hydrogen, producing a delayed signal. Note that (α, n) reactions on the small amount of natural deuterium in the detector do not contribute to the overall background. This is largely because excited states of deuterium are not produced in interactions with α particles with energies less than 6.6 MeV. For the ^{210}Po decay, the α particle is produced at 5.3 MeV. For this energy, the $(\alpha, n)^{13}$C interactions can only produce ^{16}O in the ground, first, and second excited states. In the $(\alpha, n)^{18}$O interactions, the ^{21}Ne can reach up to the fifth excited state, but the de-excitation γs are at lower energy. The different excited states are summarized in Tables 6.8 and 6.9.

Table 6.8: The various energies of the decay products and the associated branching ratios for the ^{13}C$(\alpha, n)^{16}$O interaction (with 5.3 MeV α particles).

Excited State	Decay Product	Energy (MeV)	Branching Ratio (%)
1	e^+/e^- pair	6.05	8.0
2	γ	6.13	1.2

Table 6.9: The various energies of the decay products and the associated branching ratios for the ^{18}O$(\alpha, n)^{21}$Ne interaction (with 5.3 MeV α particles).

Excited State	Decay Product	Energy (MeV)	Branching Ratio (%)
1	γ	0.351	30.8
2	γ	1.75	10.7
3	γ	2.79	2.0
4	γ	2.80	2.0
5	γ	2.87	19.6

To calculate the total number of (α, n) reactions (expressed as a total number of produced neutrons) occurring in the detector, the thick target neutron yield formula is used:

$$Y = n_t \int_0^{E_\alpha} \frac{\sigma(E_\alpha, E)}{\epsilon(E_\alpha)} dE_\alpha, \tag{6.12}$$

where Y is the neutron yield, n_t is the density of the target atoms, E_α is the energy of the α-particles, E is the energy of the neutron, σ is cross-section for an α particle with energy E_α to produce a neutron with energy E, and $\epsilon(E_\alpha)$ is the stopping power of an α particle in the target (131). The cross sections for the (α, n) reaction on ^{13}C and ^{18}O are taken from the JENDL-3.3 database (132). Table 6.10 gives the results of this calculation for ^{210}Po in the acrylic and water. A tool developed for (α, n) calculations for the DEAP-3600 detector yields similar results (133).

The details for the total expected (α, n) rate are complicated, and require careful application of the leaching model and knowledge of the time and length of the water re-circulations.

Table 6.10: The neutron yields calculated using Equation 6.12 for the 5.3 MeV ^{210}Po α for acrylic (PMMA) and water.

Material	^{13}C (neutrons/α)	^{18}O (neutrons/α)
PMMA	4.7×10^{-8}	1.45×10^{-8}
H_2O		3.98×10^{-8}

The expected (α, n) rate calculations are detailed as a function of time in (4). Importantly, the rates are broken up for (α, n) interactions in the acrylic and (α, n) in the internal water. Several key assumptions are made in the prediction: the reactions along the surface of the AV interact with a 50% probability in the acrylic and 50% probability in the water, the concentrations of ^{210}Po on the inner and outer AV surfaces are the same, and the removal of ^{210}Po from the internal water from recirculation is done so with a reduction factor of 0.41 per recirculation pass. A constant temperature of 18.4°C is used for the leaching model, and uncertainty estimates are made by changing the temperature in the model and evaluating the effect on the rate of interactions. The total number of predicted α decays from the estimated ^{210}Po are multiplied by the neutron yields in Table 6.10 and scaled to the appropriate livetime. This gives the total number of interactions on the acrylic and in the internal water, shown in column two of Table 6.11. The relatively large error bars are from uncertainties in the initial AV surface activity, the ^{210}Po leaching rate, the water temperature, and assumptions about the water recirculation effectiveness.

Table 6.11 additionally details the leakage of the (α, n) events through the cuts and the resulting number of expected counts in the region of interest in the 190.3 days of livetime. These values are calculated using simulations of these events along the inner and outer AV and within the internal water. These numbers are provided after the cuts detailed in Table 6.1, and Section 6.9 gives the results after the cut on $\Delta\mathcal{L}$. Note that the correction scalings for the neutron detection efficiency and cut sacrifices have not been applied in generating this table.

Table 6.11: The expected background from (α, n) interactions on the acrylic vessel and in the internal water in the 190.3 days of livetime. These values are provided prior to the cut on the likelihood ratio. The uncertainties on the interaction rates are detailed in (4) and the uncertainty on the counts passing the cuts include an additional 100% uncertainty on the excited-state cross-sections. The efficiency scale factors presented in Section 6.6 are not included in these numbers.

Source	Nuclide	Interactions/Livetime (4)	Leakage (%)	Counts/ROI/Livetime
Internal Water	^{18}O	$98.5\pm^{75.5}_{58.3}$	0.28	$0.28\pm^{0.35}_{0.28}$
Inner AV	^{13}C	$163.5\pm^{79.4}_{55.5}$	0.07	$0.11\pm^{0.16}_{0.11}$
Inner AV	^{18}O	$189.3\pm^{90.8}_{62.2}$	0.02	$0.04\pm^{0.04}_{0.04}$
Outer AV	^{13}C	$163.5\pm^{79.4}_{55.5}$	0.05	$0.08\pm^{0.09}_{0.08}$
Outer AV	^{18}O	$189.3\pm^{90.8}_{62.2}$	0.02	$0.04\pm^{0.04}_{0.04}$
				$\Sigma\ 0.55\pm^{0.39}_{0.31}$

Importantly, the (α, n) backgrounds are expected to be quite small, and the five components sum to $0.55\pm^{0.39}_{0.31}$. The uncertainty on the cross sections is set to 100%, which dominates the uncertainty on the background estimate, even with the large uncertainties in the ^{210}Po rates. Lastly, the likelihood ratio cut is expected to remove more (α, n) backgrounds than signal events due to the radial and energy distributions of the background. This is supported by Figure 6.28, which indicates that the simulated (α, n) events are more 'accidental-like' than the signal.

In terms of rejecting this background, the FV cut effectively removes events produced along the inner and outer acrylic vessel. The prompt γs from $^{18}\mathrm{O}(\alpha, n)^{21}\mathrm{Ne}$ interactions are fairly low in energy (given in Table 6.9) and are effectively rejected by the prompt energy and nhit criteria. Still, the largest component of the (α, n) background is expected to come from $^{18}\mathrm{O}(\alpha, n)^{21}\mathrm{Ne}$ interactions in the internal water. Unfortunately, these events are both the most uncertain in terms of the expected rate, and the most difficult to independently measure. Section 6.8.2.1 details an attempt to measure the (α, n) interactions along the

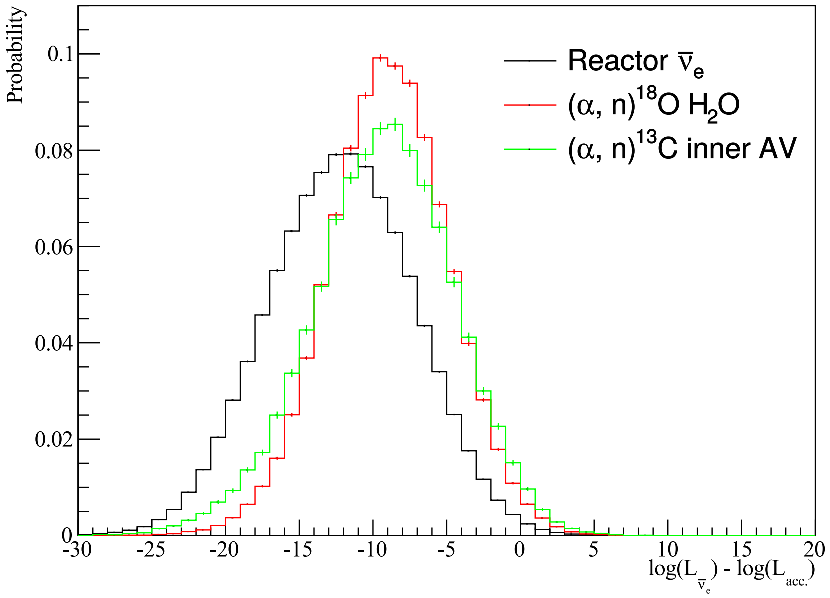

Figure 6.28: The $\Delta\mathcal{L}$ distributions for reactor IBD events, ^{18}O$(\alpha, n)^{21}$Ne interactions in the internal water, and ^{13}C$(\alpha, n)^{16}$O interactions on the inner AV. $\Delta\mathcal{L}$ is calculated according to the procedure outlined in 6.5.1 using PDFs for the signal MC and for accidental backgrounds.

acrylic.

6.8.2.1 Measurement of the (α, n) Rate on the Acrylic

Thus far the predicted (α, n) rates along the acrylic and in the internal water are based on initial ex-situ measurements of the acrylic and calculations of the time-dependence using a detailed leaching model. This section presents a measurement of the (α, n) interactions along the acrylic vessel. To perform this measurement, a sideband is chosen in radius, and the fiducial volume for the prompt events is selected as a spherical shell between 5700 and 6500 mm. This ensures that there is no overlap with the signal region where the fiducial volume of the prompt event is constrained to the internal AV volume out to 5700 mm. In this analysis, the delayed event is allowed to reconstruct between 5400 and 6600 mm to allow for the possibility the neutron has travel away from the prompt event position. The expected rates prior to the cuts are calculated in (4) and given in Table 6.12. Similarly to the analysis searching for reactor IBD events in the central volume, a likelihood ratio is constructed using PDFs for the signal (in this case the (α, n) acrylic events) and the accidental backgrounds. The PDFs generated for this calculation are identical to the ones used in the antineutrino analysis, but with the different fiducial volumes. The (α, n) likelihood ratio is defined as:

$$\Delta\mathcal{L}_{(\alpha,n)} = \log\left(\frac{\mathrm{L}_{(\alpha,\mathrm{n})}}{\mathrm{L}_{\mathrm{acc.}}}\right) = \log(\mathrm{L}_{(\alpha,\mathrm{n})}) - \log(\mathrm{L}_{\mathrm{acc.}}). \tag{6.13}$$

Figure 6.29 shows the $\Delta\mathcal{L}_{(\alpha,n)}$ for the simulated $(\alpha, n)^{13}$C events along the AV compared to the accidentals in the mixed dataset. After a cut on $\Delta\mathcal{L}_{(\alpha,n)} < -13$ a total of $3.45 \pm^{3.62}_{3.45}$ are expected.

Table 6.12: The number of (α, n) events that pass the cuts from Table 6.1 with the adjusted FV around the AV. The 100% uncertainty on the excited state cross-section is included in the expected counts/ROI/livetime. These numbers include the cut on $\Delta\mathcal{L}_{(\alpha,n)}$. Notably, the contribution from the contributions from the (α, n) interaction on oxygen are only about 5% of the total expected signal. In the sum the uncertainties on the inner and outer acrylic from each nuclide are taken as 100% correlated because they are dominated by the cross-section uncertainty.

Source	Nuclide	Interactions/Livetime (4)	Efficiency (%)	Counts/ROI/Livetime
Inner AV	^{13}C	$163.5\pm^{79.4}_{55.5}$	0.82	$1.26 \pm ^{1.42}_{1.26}$
Outer AV	^{13}C	$163.5\pm^{79.4}_{55.5}$	1.20	$1.97 \pm ^{2.19}_{1.97}$
Inner AV	^{18}O	$189.3\pm^{90.8}_{62.2}$	0.064	$0.12 \pm ^{0.13}_{0.12}$
Outer AV	^{18}O	$189.3\pm^{90.8}_{62.2}$	0.053	$0.10 \pm ^{0.11}_{0.10}$
				$\Sigma\ 3.45 \pm ^{3.62}_{3.45}$

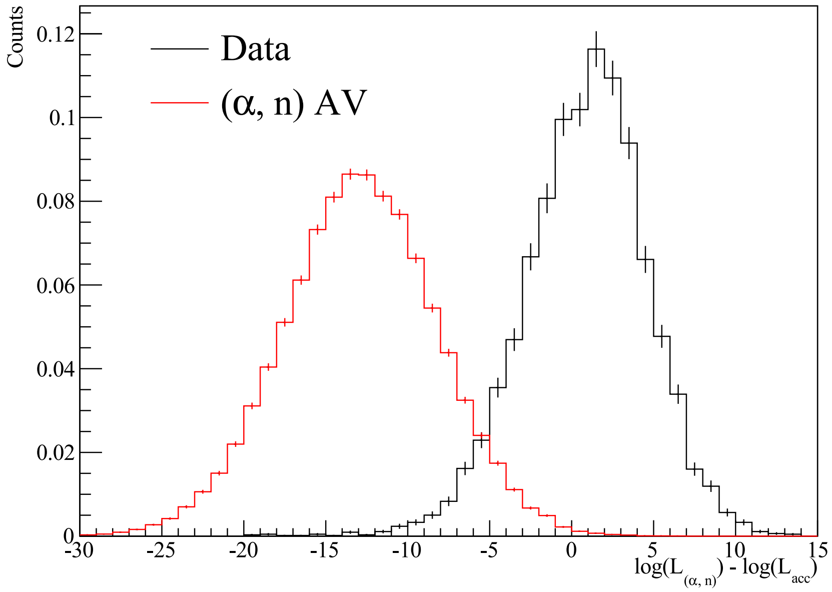

Figure 6.29: $\Delta\mathcal{L}_{(\alpha,n)}$ for the mixed dataset and for the $(\alpha,n)^{13}$C along the inner AV.

The method for generating the mixed dataset, outlined in Section 6.8.1.1, is used for this FV around the AV to make predictions for the number of accidental backgrounds in the (α, n) search (given a cut on the $\Delta\mathcal{L}_{(\alpha,n)}$). A sideband in Δt is used to check the prediction. Figure 6.30 shows the data in the signal and sideband regions compared to the mixed dataset. The left panel in Figure 6.30 shows the observation compared to the mixed dataset in the Δt sideband between 500 and 1000 μs, which shows good agreement. Similar to the reactor analysis, this indicates the mixed dataset can be used as a prediction for the number of accidental backgrounds in the signal region. Interestingly, the right panel in Figure 6.30 shows the signal region, and a clear tail that deviates from the mixed dataset at small values of $\Delta\mathcal{L}_{(\alpha,n)}$ due to (α, n) and reactor IBD events. A cut selected at $\Delta\mathcal{L}_{(\alpha,n)} < -13$ yields an expected 2.39 ± 0.23 accidental background events. After shifting by S, 1 event in the sideband region passes this cut, consistent with the accidental background expectation. The efficiency for detecting reactor antineutrinos in this limited volume is 0.71%, which yields 1.17 events in the 190.3 days of livetime (scaling from 314.6 expected interactions per year in the PSUP volume).

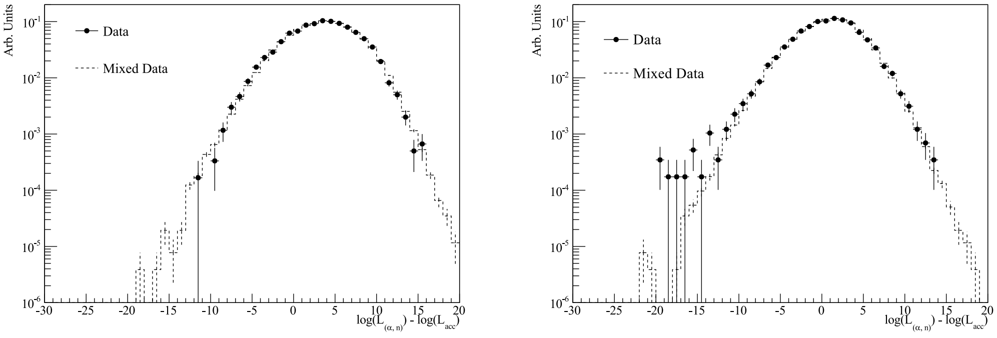

Figure 6.30: The $\Delta\mathcal{L}_{(\alpha,n)}$ for the data in the sideband (left) and signal (right) region compared to the mixed dataset.

Figure 6.31 shows the Δt distribution between 0 and 1000 μs for events that pass the cuts, including $\Delta\mathcal{L}_{(\alpha,n)} < -13$. The fit is performed using a fixed flat background, assuming 2.38 accidental backgrounds per 500 μs. The fitted time-constant is consistent with the neutron capture time. Table 6.13 summarizes the signal and background expectation and the observation. A total of 15 events are detected with values of Δt between 0 and 500 μs. Using the signal and background expectation and the number of observed events, the fraction of nominal is calculated as:

$$F = \frac{N_{\text{obs.}} - N_{\text{bkg}}}{N_{\text{sig}}}. \tag{6.14}$$

Using the results summarized in Table 6.13 gives $F = 3.32 \pm 1.12$. The prompt energy, prompt $u \cdot r$, Δr, and prompt position distributions for the 15 observed events are shown in Figure 6.32. The distributions are compared to the $(\alpha, n)^{13}$C events simulated along the inner and outer AV, which is arbitrarily scaled. Overall, the shape of each of the distributions is well-described by the simulated $(\alpha, n)^{13}$C.

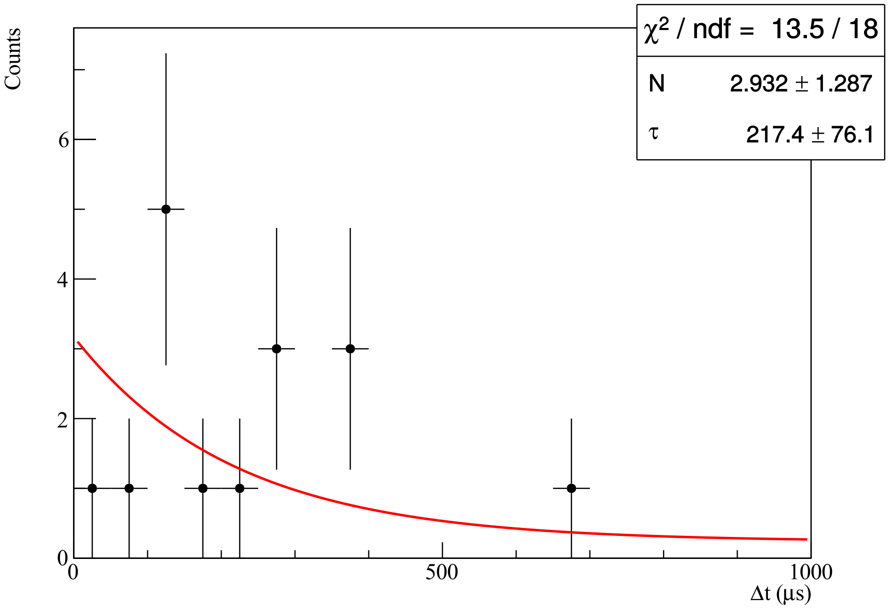

Figure 6.31: The Δt distribution for the events that pass the (α, n) cuts. The fitted time constant is consistent with the neutron capture time on hydrogen.

For the antineutrino analysis, a scale factor of 3.32 is used for the (α, n) events, and the uncertainty on that factor is propagated. Note thought that this measurement is primarily sensitive to $(\alpha, n)^{13}$C and the contribution from $(\alpha, n)^{18}$O is minimal. The measured scale factor could be different for the acrylic $(\alpha, n)^{13}$C and $(\alpha, n)^{18}$O events, particularly because the cross-section uncertainties are so large. Still, the same scale factor for $(\alpha, n)^{18}$O events is used, which would be valid in a scenario where the higher observed rate was due to a higher initial concentration of ^{210}Po on the acrylic than was measured in 2013. It turns out that the contribution to the total background from $(\alpha, n)^{18}$O interactions on the acrylic is

Table 6.13: The (α, n) signal and backgrounds for a Δt window of 0 to 500 μs, a prompt FV of 5700 to 6500 mm, and a delayed FV of 5500 to 6600 mm. The expected (α, n) events comes from Table 6.12. The accidental background is estimated using the mixed dataset in Figure 6.30. The fraction of nominal, F, is defined in equation 6.14. The factor F is used to scale the (α, n) MC as a background to the reactor antineutrino search. The uncertainty in F is dominated by the statistics in the observation.

	Expectation
Total (α, n)	$3.45 \pm {}^{3.62}_{3.45}$
Reactor IBD	1.17 ± 0.12
Accidental	2.39 ± 0.23
Observed	15
F	3.32 ± 1.12

almost negligible in the final reactor antineutrino analysis, so the precise scale factor used for these backgrounds is inconsequential.

More concerning for the antineutrino analysis is the unmeasured contribution from $(\alpha, n)^{18}$O interactions in the internal water. From Table 6.11 these events are the largest (α, n) contribution to the total background. Because they are both low in energy and distributed throughout the fiducial volume, there is no reasonable sideband to choose that would not be dominated by contributions from reactor antineutrino events. However, given an observation in the signal region, there is no method to clearly distinguish the reactor IBD interactions from the (α, n) interactions in the internal water. In other words, although it is believed, based on detailed leaching calculations, that the internal water (α, n) background is very small, there is no way to constrain the rate. In principle, it could be twenty or fifty times higher than predicted and no measurement can be performed to say otherwise. Further discussion of this point is highlighted in the summary of the analysis in Section 6.9. For now, a reasonable assumption to make it that the reason the measured concentration along the acrylic is higher than predicted is due to a higher initial concentration

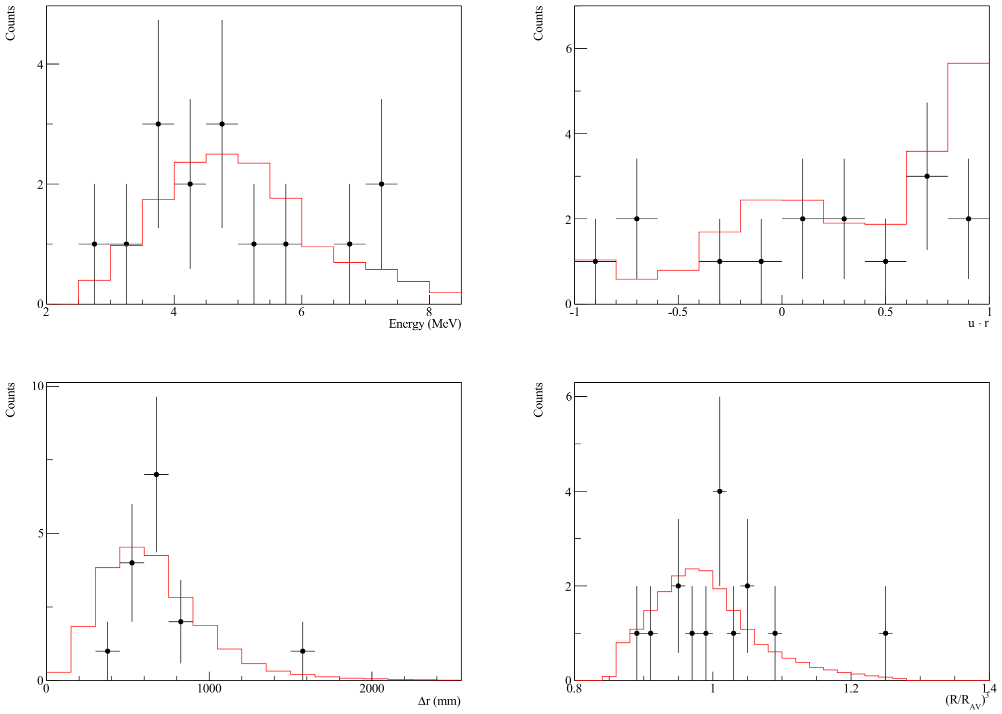

Figure 6.32: The 15 observed events in the (α, n) fiducial volume with Δt between 0 and 500 μs. The $(\alpha, n)^{13}$C simulations along the inner and outer AV is shown in red, scaled arbitrarily so that the shape can be compared to the data. From top left to bottom right is the prompt energy, prompt $u \cdot r$, Δr, and the prompt $(R/R_{AV})^3$ position.

than was measured in 2013. In this case, we would expect an increased concentration in the internal water at the same factor. Thus, the internal water $(\alpha, n)^{18}$O contribution will be scaled by the same factor in the final analysis. Lastly, I note that this measurement of the (α, n) events along the acrylic is the first ever observation of this process in a pure water Cherenkov detector.

6.8.3 Atmospherics

Atmospheric neutrinos, introduced in Section 3.2, span a broad range of energies from roughly 100 MeV to 10 GeV. This results in an array of possible interactions, discussed briefly in Section 2.5, some of which can produce a background for the reactor antineutrino search. To estimate this background, the total atmospheric production rate must be calculated, which requires accounting for several complicated effects. For example, the details of the Earth's magnetic field and the solar activity cycle are both important to understand in the atmospheric neutrino flux calculation. The values used here are the 'Bartol-04' rates calculated in (21), specifically for the SNO detector location. These rates are calculate at the solar minimum and maximum, which describe the variation in the solar activity. The solar cycle effects the total number of high energy particles bombarding the atmosphere, and thus the atmospheric neutrino production rates (134). At solar minimum, the production rates are highest. The data considered in this dissertation was collected from 2017 to 2019, during a time at which the solar cycle is near minimum. The Bartol-04 fluxes for Sudbury at solar minimum and maximum are shown in Figure 6.33.

Additionally, the atmospheric neutrinos undergo oscillations on their path to the detector. The application of these oscillations uses tools developed for and described in A. Mastbaum's thesis (22). The path length distribution is given in (135) as a function of neutrino energy and zenith angle. Neutrino energies and zenith angles, θ_z are sampled from the Bartol-04 flux distribution, and a production height is sampled using the distributions calculated in (135). The oscillation baseline is calculated as:

$$L = \sqrt{(R_E + h)^2 - (R_E - d)^2 \sin(\theta_z)} + (R_E - d)\cos(\theta_z), \tag{6.15}$$

where R_E is the radius of the earth and d is the depth of the SNO+ detector. Oscillation are applied using a three-neutrino model. The qualitative result of applying oscillations is a suppression of the $\nu_\mu/\bar{\nu}_\mu$ flux, which is converted to $\nu_\tau/\bar{\nu}_\tau$. There is only a small impact on the $\nu_e/\bar{\nu}_e$ fluxes. The resulting rates are shown in Table 6.14 for each flavor and each

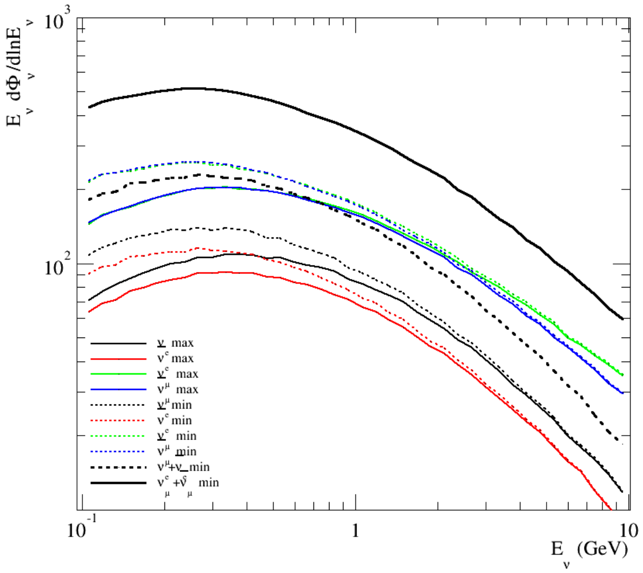

Figure 6.33: The Bartol-04 fluxes (21) averaged over zenith angle. This figure is reproduced from (22).

detector volume. The total expected interaction rate is 655.2 year^{-1} in a spherical volume of 8500 mm.

Many complicated interactions can occur as the high energy neutrinos inelastically scatter off of the target nucleons, producing various hadronic debris. This often includes the production of light mesons (π and K primarily) and baryons (Δ). These unstable particles decay and can often be identified by coincidence tagging looking for a Michel electron or multiple neutrons. CC interactions produce high energy charged leptons, making them relatively easy to tag and reject. The primary background for the reactor antineutrino search is from NC interactions on ^{16}O. An interaction such as:

$$\nu +^{16}\mathrm{O} \rightarrow {}^{15*}\mathrm{O} + \mathrm{n}, \tag{6.16}$$

Table 6.14: The expected number of atmospheric interactions (in interactions/year) at solar minimum for each flavor and detector volume after applying oscillations. The atmospheric neutrinos are generated with an energy range of 0.1 to 10 GeV. Note the outer AV extends from the AV to a radius of 8500 mm.

Flavor	Outer AV	Inner AV	AV	Total
ν_μ	142.4	79.0	2.2	223.6
$\bar{\nu}_\mu$	54.7	30.3	0.8	85.8
ν_e	133.2	73.9	2.1	209.2
$\bar{\nu}_e$	41.3	22.9	0.6	64.8
ν_τ	31.0	17.9	0.5	49.4
$\bar{\nu}_\tau$	14.2	7.9	0.2	22.3
$\sum$	416.8	231.9	6.5	655.2

leaves the oxygen nuclei in an excited state. The energy and branching ratios for the de-excitation γ-rays are calculated in (136). The primary γ is 6.18 MeV, which provides a distinct prompt signal before the neutron captures. This process has recently been used to measure the NC cross section on ^{16}O by Super-Kamiokande using atmospheric neutrinos (137) and by T2K using beam neutrinos (30).

The atmospheric neutrino interactions are simulated using the `GENIE` generator (138); more details are given in Appendix G, including a comparison to the `NEUT` generator used by the T2K and Super-Kamiokande collaborations. For each neutrino type, a little less than 100000 events are simulated, almost 1000 times the total expected flux. Specifically, for each run, 20 events are simulated for each neutrino type using the run-conditions of the detector. This is achieved by converting the output of `GENIE`, containing the final-state particle vertices, into a `ROOT` file, which is provided as input into `RAT`.

The output of the simulations are passed through the standard cuts, described in Table 6.1, as well as the follower cuts outlined in Section 6.4.3.3. Importantly, the multiplicity cut rejects prompt events with more than one detected neutron follower, which reduces the expected atmospheric background by around 30%, depending on the neutrino type. Additionally, the 100 ms veto after high nhit events is important to reduce contribution from CC events. Specifically, coincidences where both the prompt and delayed events are neutrons contribute as a background, and can be reduced by using the high nhit prompt lepton interaction to reject the following events. This has a larger effect for ν_e and ν_μ events than for ν_τ events. For ν_e interactions, adding the high nhit cut reduces the expected background by about 30%. A summary of the fractional leakage through the cuts is given in Table 6.15.

Table 6.15: About ten thousand of each neutrino flavor was simulated using the `GENIE` generator and propagated through `RAT`. The expected interactions per year are for a volume of 8500 mm, from Table 6.14. These numbers are given before the likelihood ratio cut is applied. The efficiency scale factors presented in Section 6.6 are not included in these numbers. The error bars are from cross-section and flux uncertainties, and are assumed to be 100% correlated between the various components.

Flavor	Interactions/Year	Simulated	Leakage (%)	Counts/ROI/Year	Counts/ROI/Livetime
ν_e	209.2	82760	0.572	1.20 ± 0.67	0.626 ± 0.365
ν_μ	223.6	96560	0.771	1.72 ± 0.96	0.897 ± 0.501
ν_τ	49.4	95100	1.98	0.850 ± 0.476	0.443 ± 0.248
$\bar{\nu}_e$	64.8	91920	0.676	0.438 ± 0.147	0.228 ± 0.077
$\bar{\nu}_\mu$	85.8	96700	0.967	0.830 ± 0.278	0.433 ± 0.145
$\bar{\nu}_\tau$	22.3	96720	1.92	0.428 ± 0.142	0.223 ± 0.074
				$\sum$ 5.47 ± 2.67	2.85 ± 1.41

Unlike the (α, n) background, the likelihood ratio, calculated using the accidental background PDFs, is not expected to provide additional discrimination. Indeed, Figure 6.34 shows that the atmospheric backgrounds, because the prompt event is high energy and uniformly distributed in the detector, look even less like accidental backgrounds than the signal events.

To further reject the atmospheric backgrounds, a second likelihood ratio, referred to as

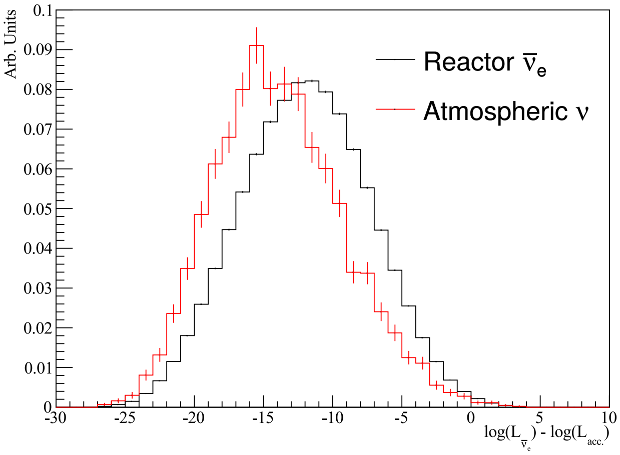

Figure 6.34: The $\Delta\mathcal{L}$ distribution for reactor IBD and atmospheric neutrino events. The atmospheric events are more 'signal-like' when using the accidental PDFs, so this parameter is not useful for rejecting the atmospheric neutrino background.

$\Delta\mathcal{L}_{\text{atm}}$, is calculated. $\Delta\mathcal{L}_{\text{atm}}$ uses Δr, prompt energy, and prompt β_{14} PDFs for the reactor IBD and atmospheric events that pass the cuts in Table 6.1. The PDFs for these variables, as well as for Δt, are shown in Figures 6.35 and 6.36. The Δr distribution shows that the prompt and delayed events in the atmospheric background are expected to be further apart, primarily because the high energy neutron produced in the atmospheric neutrino interaction travels further than the lower energy neutron from the IBD reaction. However, due to the relatively poor position reconstruction at low energy, this effect is smeared out such that it is not a particularly powerful handle to reject the atmospheric backgrounds. The prompt energy and β_{14} distributions shown in Figure 6.36 are the most powerful obervables to discriminate the atmospheric background from the signal. The $\Delta\mathcal{L}_{\text{atm}}$ for the signal and background is shown in Figure 6.37, which indicates that this parameter can be used to effectively reject a fraction of the atmospheric background without much signal sacrifice.

Directly using the T2K results provides a data-driven estimate of the cross-section uncer-

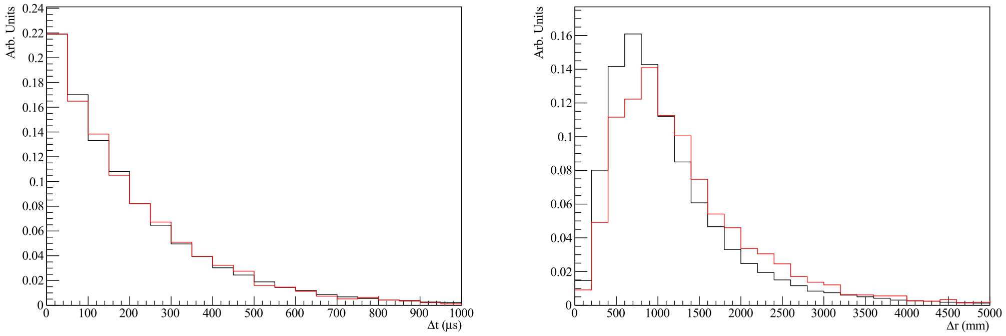

Figure 6.35: The signal (black) and atmospheric (red) Δt and Δr distributions.

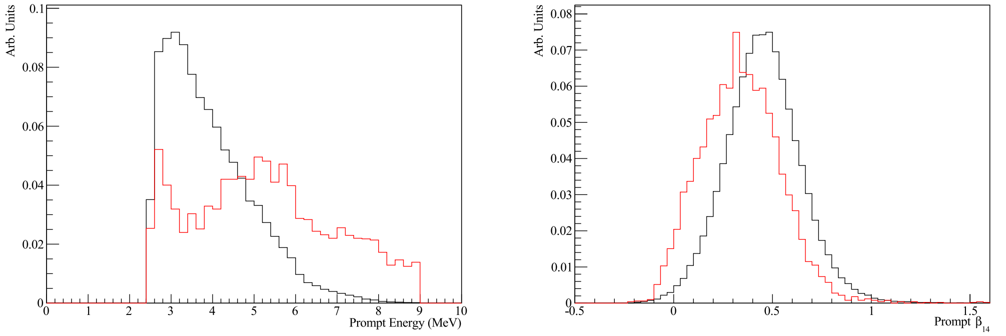

Figure 6.36: The signal (black) and atmospheric (red) prompt energy and β_{14} distributions.

tainty. Taking the quoted statistical and systematic uncertainties from (30) in quadrature, the neutrino and antineutrino cross-sections have uncertainties of 32%. Given the disagreement between the `GENIE` model (as seen in Figure G.1) and the measurement in neutrino mode, the cross-section uncertainty is increased to 55% to cover the T2K measurement. The antineutrino data agrees very well with the `GENIE` and `NEUT` models prediction, so the uncertainty is not scaled from 32%. The flux uncertainty is around 10% for atmospheric neutrinos below 10 GeV (139), which is taken in quadrature with the cross-section uncer-

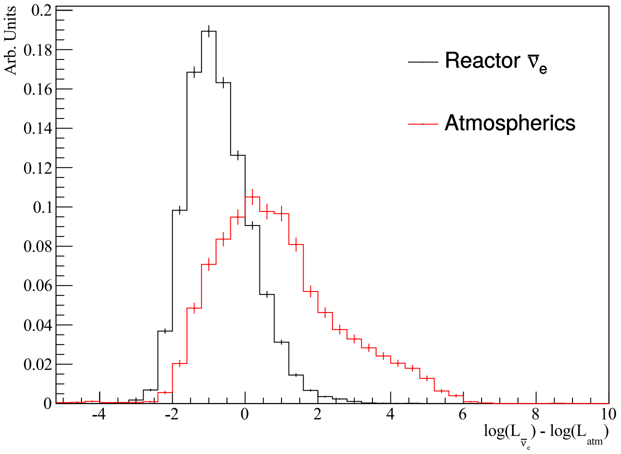

Figure 6.37: The atmospheric likelihood ratio, $\Delta\mathcal{L}_{\text{atm}}$, calculated using the Δr, prompt energy, and prompt β_{14} PDFs shown in Figures 6.35 and 6.36.

tainty. Accounting for this, the total expected number of events per year is 5.47 ± 2.67 or 2.85 ± 1.41 in 190.3 days.

6.8.3.1 Multiplicity Sideband

To constrain the atmospheric background outside of the signal region, a sideband is selected. In this case, it is important to constrain the NC background that produces prompt de-excitation γ-rays. A sideband in the delayed event multiplicity is chosen. Specifically, prompt events with two or more neutron-like events, which are rejected by the multiplicity cut in the standard analysis, are selected for this search. Figure 6.38 shows the neutrons multiplicity produced in neutral current atmospheric events, which indicates that about one-third of the final states have more than one neutron. Additional neutrons can be knocked out of oxygen nuclei by the high energy protons and neutrons.

A set of cuts for the prompt and delayed events are tuned to detect events for this

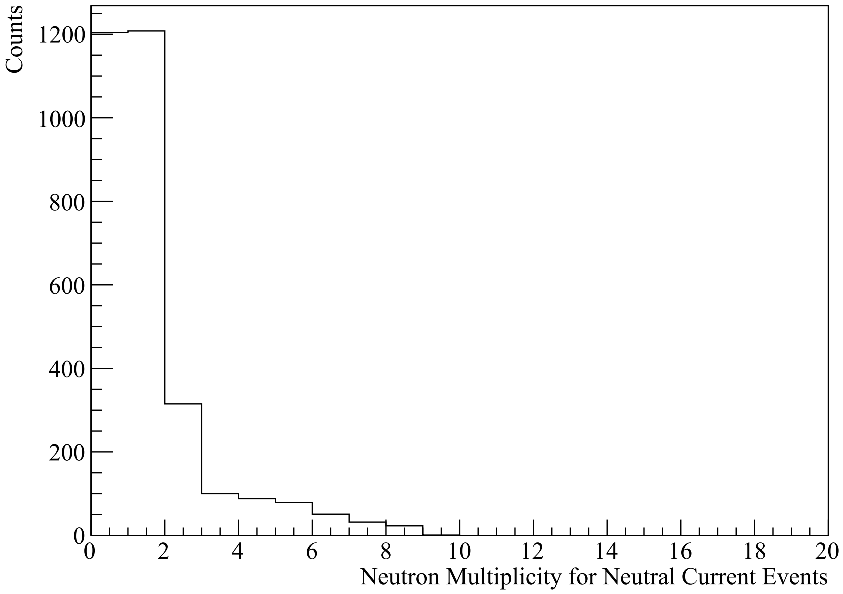

Figure 6.38: The neutron multiplicity for NC atmospheric interactions. The first bin corresponds to 0 generated neutrons. Note that this figure shows the final state particles in the neutrino interaction, as simulated by `GENIE`. More neutrons can be created if neutrons are knocked out of oxygen nuclei by the high energy neutrons and protons.

atmospheric sideband, presented in Table 6.16. The prompt selection is already tuned to have a high efficiency for the prompt γ-ray, but a lower energy cut of 4.0 MeV is added to largely reject neutron-neutron events and to drive the accidental background pile-up down.

Due to the high lower energy prompt energy cut and the multiplicity requirement, the possible backgrounds from reactor IBD interactions or accidental coincidences are reduced to effectively zero. Table 6.16 presents the total number of expected events in this sideband after the cuts. The total number after scaling to the 190.3 days of livetime is 2.45 ± 1.21. Including the efficiency corrections from Table 6.6 yields 2.16 ± 1.18, propagating the

Table 6.16: Various cut variables and values for prompt and delayed event for the atmospheric sideband search.

Variable	Cut Value Prompt	Cut Value Delayed
Data-cleaning	0xFB0000017FFE	0xFB0000017FFE
NHits Clean	> 15	
Fit Valid	True	True
Energy	$\geq$ 4.0 MeV	$\leq$ 4.0 MeV
	$\leq$ 25.0 MeV	
Position	$\leq$ 5700 mm	
ITR	> 0.5	> 0.4
β_{14}	$\geq$ -0.6	$\geq$ -0.6
	≤ 1.6	≤ 1.6
σ_R	$\leq$ 600 mm	$\leq$ 900 mm
R_{FOM}	≥ 9.9	≥ 9.9
Energy G-Test	> 0.0	> 0.0
	< 1.45	< 1.45
Energy U-Test	> 0.5	> 0.5
Energy Z-Factor	< 1.0	< 1.0
Δt		≥ 3 μs
		≤ 500 μs
Δr		$\leq$ 3000 mm
Multiplicity		> 1

uncertainty on the neutron detection efficiency scaling.

Running this analysis over the data-set produces 0 observed events in 190.3 days of livetime. Because the background expectation for this search is negligible, the Poisson 90% confidence-level (CL) upper limit, s_{max} can be found using:

$$\sum_{n=0}^{N} \frac{e^{-s_{max}} s_{max}^{n}}{n!} = 1 - \mathrm{CL} = 0.10, \tag{6.17}$$

Table 6.17: The expected number of multiple coincidence events selected using the cuts outlined in Table 6.16.

Flavor	Interaction/Year	Simulated	Efficiency (%)	Counts/ROI/Year	Counts/ROI/Livetime
ν_e	209.2	82760	0.478	1.00 ± 0.56	0.521 ± 0.292
ν_μ	223.6	96560	0.704	1.57 ± 0.88	0.821 ± 0.456
ν_τ	49.4	95100	1.65	0.815 ± 0.456	0.425 ± 0.238
$\bar{\nu}_e$	64.8	91920	0.487	0.316 ± 0.106	0.165 ± 0.055
$\bar{\nu}_\mu$	85.8	96700	0.788	0.676 ± 0.226	0.352 ± 0.118
$\bar{\nu}_\tau$	22.3	96720	1.42	0.317 ± 0.106	0.165 ± 0.055
				$\sum$ 4.69 ± 2.33	2.45 ± 1.21

where $N = 0$ is the total number of detected events (140). This yields a 90% upper limit of 2.3 counts. Because this is consistent with the total number of expected events, no additional scaling is applied to the atmospheric background in the reactor antineutrino analysis. However, this sideband search does hint that the atmospheric background might be overestimated by the `GENIE` MC.

6.8.4 Geoneutrinos

Geoneutrinos are discussed in Section 3.4 and contribute a negligible source of background for this search, particularly because the majority of the antineutrinos are emitted below the IBD threshold, shown in Figure 3.9. The unknown composition of uranium and thorium in the Earth's crust and mantle leads to uncertainties in the total flux expected. However, using a geodynamical 'high-Q' model, which predicts higher concentrations of uranium and thorium than other models (141), the total expected number of interactions per year across the entire energy spectrum in the SNO+ detector is 24 from the uranium chain and 7.6 from the thorium chain. After applying the IBD threshold of 1.8 MeV, oscillations, and the cuts detail in Table 6.1, the efficiency for detecting these events is less than 0.1%. Thus, geoneutrinos are a negligible background for the antineutrino search in water.

6.8.5 Cosmogenics

Cosmic-ray muons interacting with the nuclei in the detector can induce radioactive isotopes that can be backgrounds for the reactor antineutrino search. The muon flux as a function of underground depth is shown in Figure 6.39, which indicates the depth of various underground labs. Because of the depth of SNOLAB, the average number of muons interacting per hour is less than three (142), whereas in the Super-Kamiokande detector the rate is around 2 Hz (116).

Because the cosmic muon rate is so low, a muon-follower cut that removes 20 seconds after every identified muon is utilized. This removes the vast majority of the cosmogenically activated isotopes as well as neutrons produced in the muon interaction. However, cosmogenic isotopes that have long half-lives that undergo $\beta^- + n$ decays could produce a background for the antineutrino search. A list of these isotopes is given in Table 6.18; the values of which are reproduced from (143), (116), (144).

Table 6.18: The cosmogenically activated isotopes that undergo $\beta^- + n$ decays with end-point kinetic energies above 1 MeV. Some of the decay energies are not precisely known. All isotopes other than ^{17}N are expected to contribute less than 10^{-10} events per year after the 20 second follower cut, and are included in this table as zero.

Isotope	τ (s)	E_{kin}	SNO+ events/year
^{8}He	0.17	$9.9 + 1.0(\gamma)$	0.0
^{9}Li	0.26	$\sim$10	0.0
^{11}Li	0.0085	$\sim$16	0.0
^{16}C	0.75	$\sim$4	0.0
^{17}N	4.2	7.9	7.0×10^{-3}

Super-Kamiokande presented the measured production rates of several of these isotopes in (116), which can be straightforwardly scaled to the SNO+ detector, accounting for the different fiducial volumes and the 20 second follower cut. The details of this calculation

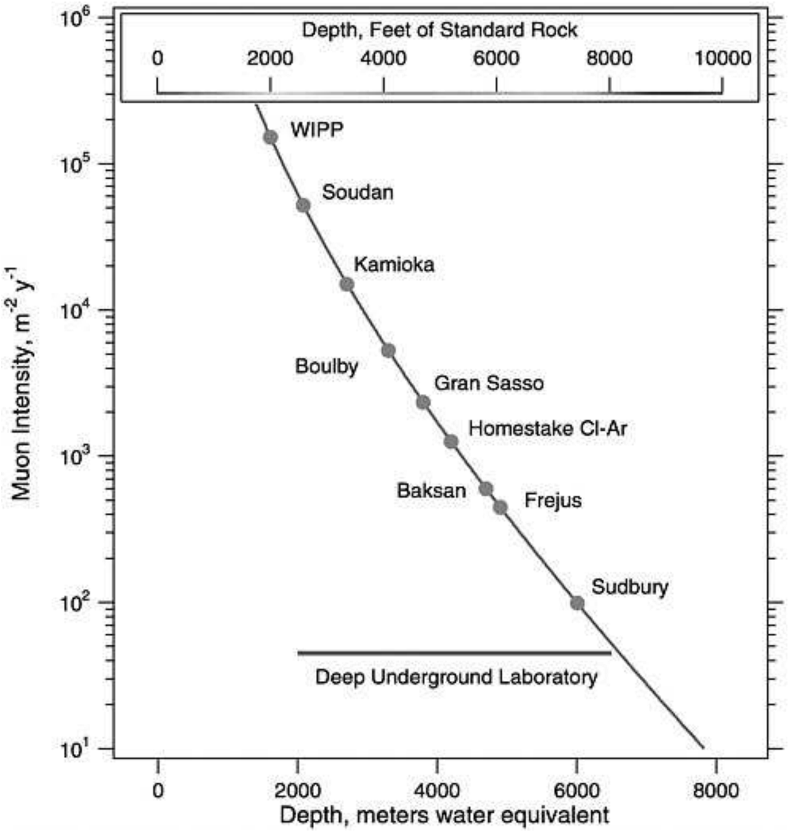

Figure 6.39: The cosmic muon flux as a function of depth, in units of meters water equivalent. SNOLAB is the second deepest underground laboratory; the Jinping underground laboratory, not shown on this plot, is the deepest.

are given in (145) and the calculated rates are reproduced in Table 6.18. It is easy to understand the very low rates – the production rates in Super-Kamiokande are already fairly low (typically less than tens of events per kton-day), and the half-lives are fairly short. The largest contribution to the background is from ^{17}N, which has the longest half-life of the $\beta^- + n$ decay isotopes, and contributes only 7.0×10^{-3} events per year. This background is negligible for the antineutrino search.

6.9 Background Summary with $\Delta\mathcal{L}$ Cut

With the total rate of signal events and accidental backgrounds all estimated, we are ready to provide a cut on the likelihood ratio. The data-cleaning sacrifice, multiplicity cut sacrifice, and neutron detection efficiency correction factors (summarized in Table 6.6) are applied to the reactor, (α, n), and atmospheric events. As discussed in Section 6.6.1, the neutron detection efficiency scaling to the simulated events is 0.874 ± 0.240 for the (α, n) events on the AV, and 0.921 ± 0.213 for the others. The accidental background is estimated from data and is not subject to the scalings that are applied to the other backgrounds.

The uncertainty on the reactor signal is dominated by the uncertainty in the neutron detection efficiency correction (which itself is dominated by the correction for the trigger efficiency). The uncertainty on the accidentals is dominated by the uncertainty in the efficiency, determined from the mixed dataset. The uncertainty on the (α, n) events are dominated by the 100% cross-section uncertainty. The uncertainty on the atmospheric events are dominated by the cross-section uncertainties, although both the flux uncertainty and the neutron detection efficiency scaling uncertainty are also important.

The cut values on $\Delta\mathcal{L}$ and $\Delta\mathcal{L}_{atm}$ are simultaneously tuned by scanning over cut values and calculating the number of signal (s) and background (b) events that pass the cuts. The optimal cuts are chosen by identifying the maximum $s/\sqrt{b}$ while scanning over the two cut values in steps of 0.05. Using this process, values of $\Delta\mathcal{L} <$ -12.35 and $\Delta\mathcal{L}_{atm} <$ -0.05 were identified as optimal cut values. Table 6.19 shows the resulting signal and background counts after cutting on the likelihood ratio distributions. After the likelihood cut a total of 2.16 ± 0.49 signal events and 1.06 ± 0.33 background events are expected. This gives a $s/\sqrt{b} = 2.1$.

Table 6.19: A breakdown of the backgrounds and signal expectation after an optimized cut of $\Delta\mathcal{L}$ < -12.35 and $\Delta\mathcal{L}_{atm}$ < -0.05. The second column gives the total number of counts per livetime for the signal and each background component, after applying the correction factors summarized in Table 6.6 (some of which are specific to the reactor neutrinos), propagating the uncertainties, and prior to applying the likelihood ratio cuts. All (α, n) backgrounds are scaled by 3.32 ± 1.12. The efficiency reduction from the likelihood cuts is given in the third column. The expected number of signal and background events after the likelihood ratio cuts is given in the fourth column.

Background	Counts/Livetime	Likelihood Efficiency (%)	Counts/Livetime after $\Delta\mathcal{L}$ cuts
Reactor IBD	5.65 ± 1.12	38.26	2.16 ± 0.49
Accidental	1991 ± 44	0.015	0.293 ± 0.053
^{18}O water	0.820 ± 0.886	16.75	0.130 ± 0.140
^{13}C inner AV	0.306 ± 0.330	16.43	0.050 ± 0.054
^{13}C outer AV	0.222 ± 0.240	20.97	0.045 ± 0.048
^{18}O inner AV	0.111 ± 0.120	16.55	0.014 ± 0.015
^{18}O outer AV	0.111 ± 0.120	16.74	0.018 ± 0.020
Total (α, n)	1.57 ± 1.17		0.258 ± 0.190
ν_e	0.552 ± 0.323	19.24	0.106 ± 0.064
ν_μ	0.791 ± 0.438	19.00	0.150 ± 0.090
ν_τ	0.391 ± 0.216	21.15	0.083 ± 0.050
$\bar{\nu}_e$	0.201 ± 0.078	21.48	0.043 ± 0.017
$\bar{\nu}_\mu$	0.382 ± 0.148	21.14	0.081 ± 0.032
$\bar{\nu}_\tau$	0.197 ± 0.075	21.71	0.043 ± 0.017
Total atmospheric	2.51 ± 1.28		0.506 ± 0.270
Total background			1.06 ± 0.33

However, as discussed in (146), $s/\sqrt{b}$ is not a good estimate of the expected discovery sensitivity when the background has non-negligible systematic uncertainty. In this case, the Gaussian significance, Z, is better estimated by $s/\sqrt{b+\sigma_b^2}$, where σ_b is the systematic uncertainty on the background. However, this estimate also fails in the limit that the background expectation is very small. A derivation of a more general formula is given in (146), which is reproduced here:

$$Z = \left[2\left((s+b)\log\left[\frac{(s+b)(b+\sigma_b^2}{b^2+(s+b)\sigma_b^2}\right] - \frac{b^2}{\sigma_b^2}\log\left[1+\frac{\sigma_b^2 s}{b(b+\sigma_b^2)}\right]\right)\right]^{1/2}. \tag{6.18}$$

Using the numbers in Table 6.19 gives $Z = 1.6$. Optimizing the likelihood ratio cut values according to Equation 6.18 gives very similar optimal cut values to those identified when optimizing against $s/\sqrt{b}$.

6.10 Partial Unblinding

A partial unblinding of the dataset is performed to ensure the analysis framework and background expectations discussed in previous sections. Runs up to 202030, corresponding to 48.9 days of livetime, are selected for unblinding. The full analysis is run over events in the signal window, defined by the cuts in Table 6.1, as well as the likelihood ratio cuts $\Delta\mathcal{L} < -12.35$ and $\Delta\mathcal{L}_{\text{atm}} < -0.05$. Scaling to 49.8 days of livetime from Table 6.19, a total of 0.56 ± 0.13 reactor IBD events and 0.27 ± 0.08 background events are expected. Applying the analysis to the data, a total of 1 event is detected, consistent with the signal plus background expectation. Figure 6.40 shows the likelihood ratio distribution for events in the signal window in linear and log scale. A total of 421 events are identified, with 3 of them having values of $\Delta\mathcal{L}$ less than -10. Information about these three events is given in Table 6.20. $\Delta\mathcal{L}$ is plotted against $\Delta\mathcal{L}_{\text{atm}}$ for the events in the signal box in Figure 6.41. This partial unblinding confirms the framework and background expectation in the signal window, but with low statistics.

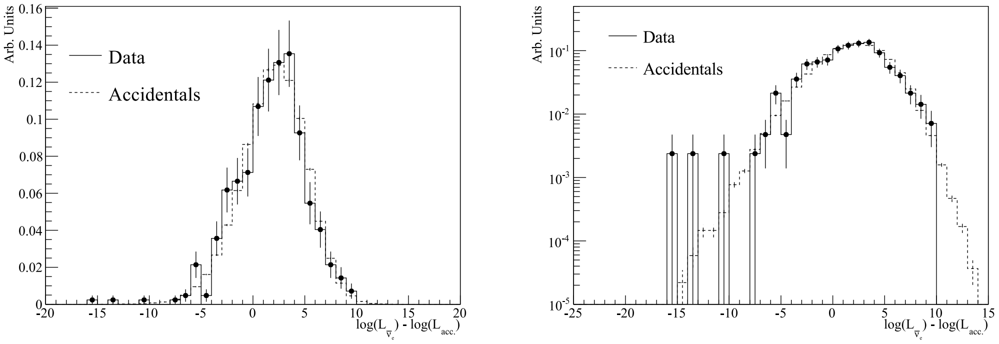

Figure 6.40: The likelihood ratio distribution for events in the signal window for the 25% of the dataset that was unblinded, in linear (left) and log (right) scale. The three events with likelihood ratio values less than -10 are detailed in Table 6.20.

Table 6.20: Information about the three events in the 25% unblinded dataset (between runs 200004 and 202030) that pass a likelihood ratio cut of $\Delta\mathcal{L} < -10$. The first event is cut by both the optimized likelihood ratio cut, $\Delta\mathcal{L} < -12.35$, and the atmospheric likelihood ratio cut, $\Delta\mathcal{L}_{\mathrm{atm}} < -0.05$. The second event passes all cuts. The third event is cut by the optimized atmospheric likelihood cut. The values of the observables are given for the prompt event unless otherwise specified.

Run	GTID	Δt (μs)	Δr (mm)	E (MeV)	E (delayed) (MeV)	β_{14}	r (mm)	$\Delta\mathcal{L}$	$\Delta\mathcal{L}_{\mathrm{atm}}$
201003	7173577	47	1448	3.2	1.9	-0.08	4553	-10.4	0.60
201902	11118965	438	1432	4.3	2.5	0.44	4447	-13.7	-0.37
201933	14743076	251	1513	4.4	2.4	-0.09	5604	-15.2	1.6

6.11 Summary

In this Chapter I presented an analysis designed to detect antineutrinos from nuclear reactors in the SNO+ detector, using 190.3 days of livetime. Because of the relatively low flux and detection efficiency, a detailed understanding of the signal and background expectations is required. The AmBe calibration source data is used to understand the relative neutron detection efficiency between the simulation and data. The background is

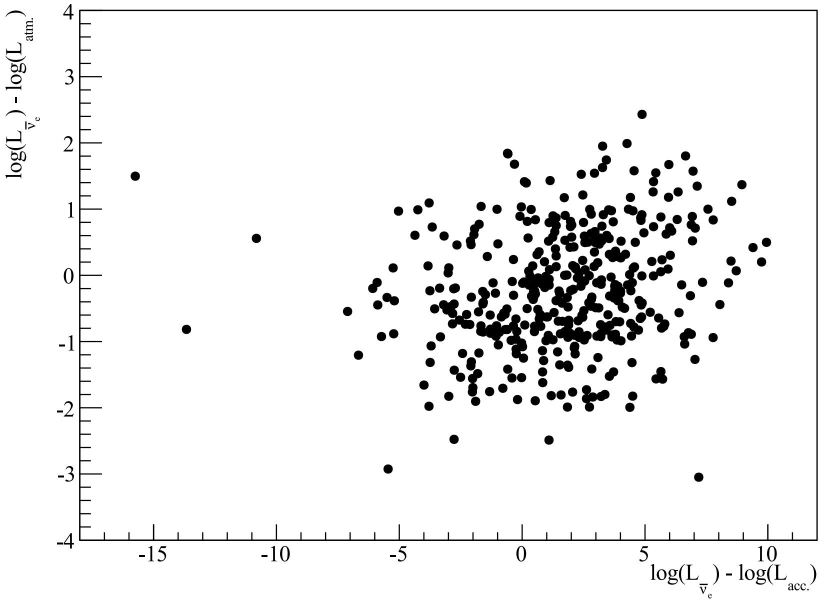

Figure 6.41: The likelihood ratio vs. the atmospheric likelihood ratio for events in the signal window for the 25% of the dataset that was unblinded.

broken up into three distinct components: accidental coincidences, (α, n) interactions, and atmospheric neutrino interactions. Each of these background is studied independently in separate sidebands. The accidentals are directly measured in the data using a window in Δt past 500 μs, while the (α, n) and atmospheric backgrounds are checked in sidebands in radius and multiplicity respectively. Lastly, a cut on the likelihood ratio parameters is tuned to maximize the sensitivity of the search.

Importantly, a particularly difficult background for this search is $(\alpha, n)^{18}$O interactions within the internal water. Using the predicted rates of ^{210}Po in the detector, this background is expected to be fairly small; indeed, from Table 6.19 the contribution is less than 0.2 events in 190.3 days of livetime. However, this is the only background that has not been measured in a sideband, primarily because it is virtually indistinguishable from the signal (other than the prompt event being slightly lower in energy). If the rate of these interactions was about 15 times larger than expected, it would be the dominant background for this search. While a measurement of the (α, n) rate along the acrylic yielded an observation about 3 times larger than the calculation, this result is not directly applicable for the internal water. The most direct way to constrain this background is through ex-situ assays of the water, which measure the radioactivity levels in a sample removed from the AV. This has been performed for SNO+, but the efficiency of the assays has yet to be completely understood and the results are not included in this thesis. In the future, ex-situ measurement might be able to directly constrain this background and improve the interpretability of the final result.

For this dissertation, the signal region for 25% of the livetime has been unblinded, in which one event is observed – consistent with the expectation. In the final analysis, additional data will be included, increasing the livetime to more than 300 days. This future work will build on the analysis presented in this thesis to analyze the full dataset. Furthermore, it is expected that the signal sensitivity can be increased by including more fiducial volume, particularly in the external volume, which was not considered in this thesis. With these improvements, it is expected that SNO+ will make the first significant detection of reactor antineutrinos in a pure water Cherenkov detector. Already established in this work is the first observation of (α, n) interactions in a water Cherenkov detector. In addition to demonstrating the feasibility of a technology, the analysis approach and background estimation presented in this thesis will be useful to future detectors searching for antineutrinos, such as WATCHMAN (121) and Super-K Gd (101).

Chapter 7

R&D for SNO+ and Future Large-Scale Water Cherenkov and Scintillator Detectors

Critical to the success of SNO+ is the performance of the PMTs and liquid scintillator. The R1408 PMTs in the SNO+ detector are reused from the SNO detector, and are identical to the ones used in the LSND experiment, which were reused by MiniBooNE (147). These PMTs are fairly low quantum efficiency ($\sim 15\%$), have a transit time spread of about 1.6 ns, and an average charge peak-to-valley ratio slightly above one. An example of the transit time distribution for the R1408 PMT, as measured for one PMT removed from the SNO detector, is shown in Figure 7.1. The R1408 provides a good reference for large-area PMT R&D, as any new PMT should outperform this PMT.

For future detectors, such as upgrades to ANNIE (122) and SNO+ (40), WATCHMAN (121), JUNO (16), `THEIA` (86), Jinping (87), and Hyper-K (148), the use of high-efficiency and fast-timing large-area PMTs with relatively low dark rates will be essential. Testing of several large-area PMTs is discussed in Section 7.1.

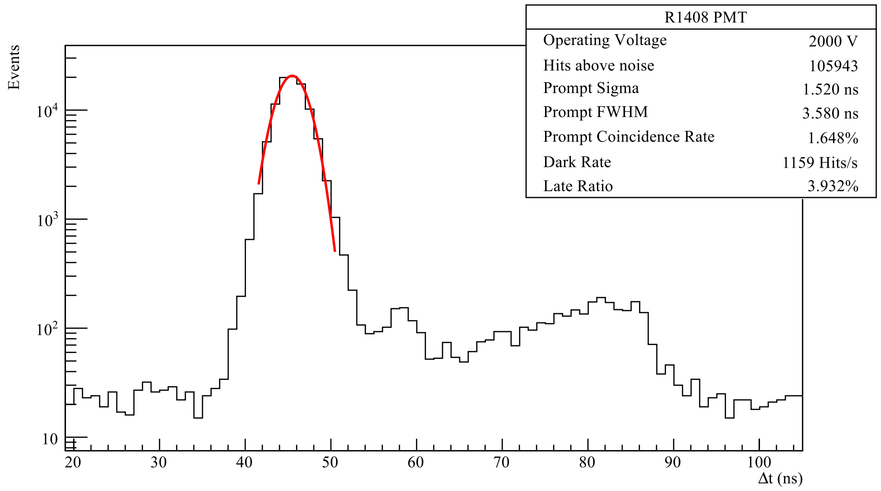

Figure 7.1: Transit time spread of a SNO PMT (removed from the detector), measured relative to a fast trigger PMT using a source of Cherenkov light. The σ of the prompt peak shows the expected TTS of approximately 1.5 ns. The operating voltage of 2000V is typical for the R1408 PMTs and the dark rate of about 1 kHz is typical at room temperature. The coincidence rate is kept below 5% in these measurements to ensure the events are primarily SPE. The late ratio is the fraction of PMT hits outside of the prompt peak.

The success of the SNO+ physics program will largely be determined by the performance of the liquid scintillator. Excellent optics and radiopurity of the scintillator would ensure a low backgrounds search for $0\nu\beta\beta$. The measurement of several important optical properties are discussed in Section 7.2. In particular the scintillator emission timing is critical for β and α discrimination which allows the rejection of backgrounds from BiPo coincidence events. The light yield of the scintillator, critical for the energy resolution of the detector, is also measured. Simulations of the bench-top setups are performed to ensure the scintillation model in `RAT` reproduces the results, which provides justification for using the model to extrapolate to large scales.

7.1 Large-area PMTs

In this section, I focus on the characterization of several large-area PMTs. Two of these PMTs are already used in current generation experiments; the HQE R5912 PMTs are used in the DEAP dark-matter experiment (149) and the 11-inch ETEL PMT are used in the ANNIE experiment (122). Other large-area PMTs of note are the 12-inch HQE Hamamatsu PMT, also characterized in our darkboxes in (150), and the 10-inch Hamamatsu PMT is characterized by the IceCube detector in (149). Perhaps the most intersting PMT characterized in this section is the R5912-MOD PMT, a prototype 8-inch Hamamtasu PMT that provides state-of-the-art timing and charge reponse.

7.1.1 Experimental Setup

The experimental setup used to characterize each of the PMTs varies little between the measurements. The basic setup uses a UV transparent acrylic block embedded with two ^{90}Sr source, which is optically coupled to an R7600-U200 1-inch square Hamamatsu PMT. The R7600-U200 PMT is used as a fast trigger; the TTS of this trigger PMT is around 250 ps so it adds negligible jitter to the measured TTS of the large-area PMTs. The acrylic block provides a source of Cherenkov light and is referred to as the 'Cherenkov source'. The Cherenkov light is produced in the β decays of ^{90}Sr and ^{90}Y. The ^{90}Sr undergoes a 0.546 MeV β^- decay to ^{90}Y with a half life of 29.1 years. The ^{90}Y undergoes a 2.28 MeV β^- decay to ^{90}Zr with a half-life of 64 hours. The β^-s from both decays enter the acrylic and create Cherenkov light. The emitted Cherenkov light has the advantage of being produced with an extremely narrow spread in timing. Additionally, the wavelength spectrum of Cherenkov light is well-known and spans the same spectrum as many common scintillators. The Cherenkov process produces few enough photons per interaction to make it easy to move the PMT far enough away from the source to primarily see single photons. The UV transparent acrylic is the same acrylic as used in the construction of the SNO acrylic vessel.

The PMT to be characterized is kept at least 30 cm away from the Cherenkov source, depending on the PMT used, in order to maintain a primarily SPE source of light. Additionally, this distance ensures that the entire front-face of the PMT is illuminated. This setup is housed in a darkbox and the high voltage is provided to each PMT by the ISEG NHS supply, a high precision, six-channel NIM module. The darkbox is lined with Finemet magnetic shielding in order to minimize interference from the Earth's magnetic field. A Lecroy WaveRunner 606Zi 600MHz oscilloscope is used to digitize the signals from both the R7600-200 and the characterization PMT. The signal from the R7600-200 PMT is used to trigger the oscilloscope readout, and the PMT is referred to as the trigger PMT in this paper. For the SPE measurements the waveforms extracted were 500.2 ns long using 50 ps samples. The scope has an 8-bit ADC with a variable dynamic range, which allowed for roughly 300 μV vertical resolution. The `LeCrunch` software (151) was used to readout the data from the scope over ethernet as well as format the data into custom `hdf5` files.

7.1.2 Analysis

The DAQ and analysis used for the various measurements in this section is consistent across the PMTs. All of the setups are chosen such that the PMT detects primarily SPEs. The shape of the SPE distribution, the width of the SPE peak, and the lack of a high or low charge tail are important components of the average PMT pulse shape. This is characterized by taking data where the coincidence rate between the trigger and a pulse on the measurement PMT is kept low (less than 5%), which ensures the data captured at the large-area 'measurement PMT' is primarily SPE. The charge of the SPE pulses is histogrammed, and should peak around 1.6 pC, corresponding to a gain of 1×10^7. In order to produce the SPE charge distribution, the analysis code integrates each digitized waveform using a 30 ns window around the arrival time of the prompt light. The first 100 ns of the waveform is used to calculate the baseline of the waveform. The prompt light comes well after the baseline window has ended. However, on occasion a PMT pulse generated by dark current ends up in the baseline window, which drags the baseline down. If there is any pulse above the electronics noise in the baseline window the entire waveform is thrown out. This ends up rejecting less than 0.1% of the total collected events. The following parameters are used in the characterization of the SPE charge distribution:

Charge FWHM: $2\sigma_{peak}\sqrt{2\log(2)}$, the FWHM as determined by the fit to the charge peak.

Peak-to-Valley (P/V): The height of the charge peak divided by the height of the minimum of the valley. The minimum of the valley is determined by a quadratic fit between the electronics noise and the charge peak. This parameter is a strong indicator of how often a PMT's pulse will cross a given discriminator threshold (higher P/V means higher efficiency).

High charge tail: The number of events above $3\sigma_{peak}$ divided by the number of events above $3\sigma_{ENW}$. This indicates the amount of multi-PE contamination into the SPE sample.

In addition to the PMT charge distribution, the transit time distribution is critically important. The transit time is the amount of time it takes for a photoelectron created at the photocathode to travel through the PMT and be detected as an output pulse from the anode. This time varies from one photoelectron to the next and the spread in the transit time distribution is one of the most important characteristics of a PMT. The transit time spread is largely determined by the electron optics of the PMT particularly between the photocathode and first dynode as well as between the first and second dynodes.

In order to extract the timing distribution, coincidence events are identified and the peaks of both the measurement PMT and the trigger PMT signals are found. Then a constant fraction discriminator is applied in analysis to each waveform and the samples corresponding to 20% of the peak height are found. The time difference, Δt, between those samples is histogrammed. To extract the full transit time distribution the waveform is stepped through in 30 ns windows, looking for PMT pulses. The 30 ns window was chosen to span the entire width of the even the largest SPE waveform. If the charge of the window is larger than 0.2 pC, than the peak of the waveform is found and Δt is histogrammed. In this scheme, a single waveform can contribute multiple times to the transit time distribution. This happens most often for double pulsing and for waveforms that have both a dark pulse and a prompt pulse. The 0.2 pC is a conservative threshold determined by the width of the electronics noise (10 × ENW) so as to be sure no electronics noise is included in the transit time histogram. The following timing distribution characteristics are extracted:

Pulses (PMT hits) above the noise: The number of PMT pulses in the transit time distribution. Unless otherwise specified, the total number of waveforms analyzed is around one million.

Prompt sigma: The prompt light is the primary contribution to the transit time distribution. The spread in time of the prompt light is characterized by fitting around 10% of the peak height on either side of the peak. The σ_{SPE} of a Gaussian fit is referred to as the transit time spread (TTS). The TTS extracted is influenced by several factors: the statistics in the peak, the systematics of the setup, the contamination of MPE hits, and uncertainties associated with the Gaussian fit to the peak. The number of waveforms in the datasets was intentionally taken in order to maintain around 1% statistical uncertainty on the prompt peak. The systematics associated with the experimental setup and the multi PE contamination, which are tied together in the distance and angle from the source, were studied by taking data at various distance and angles from the source, up to a coincidence rate of roughly 5%. The measured TTS varied by up to 3%. Finally, the fit uncertainties were determined by running the fit over various reasonable ranges, rather than 10% of the peak height on either side. This uncertainty turned out to be the largest at 5%. By fitting over smaller or larger ranges, the extracted TTS changes by around 30 ps.

Prompt FWHM: $2\sigma_{SPE}\sqrt{2\log(2)}$, the FWHM of the prompt fit.

Prompt coincidence rate: The coincidence rate of the prompt light. This parameter was intentionally kept to less than 5% by placing the R5912-MOD PMT 30 cm from the Cherenkov source. That was done to minimize MPE contamination into the SPE sample.

Dark rate: The rate of PMT pulses that fall outside the late pulsing and prompt pulsing regions. The cause of dark noise is discussed in B.0.4.

Late ratio: The rate of PMT pulses that fall within a late window, between roughly 10 ns and 60 ns after the prompt peak, depending on the PMT. These values were chosen empirically and are somewhat arbitrary.

Other than the charge and timing distribution, the relative efficiency between PMTs can be compared using the coincidence rate between the trigger and the measurement PMT. Given the distance to the source remains constant and the size of the photocathode is the same, the relative coincidence rate is indicative of a change in efficiency. In all cases, dark counts are removed when calculating the coincidence rate. The total efficiency is given by:

$$E = QE \times CE \times FE, \tag{7.1}$$

where the quantum efficiency (QE) is the efficiency a photoelectron is created given an incident photon, and is wavelength dependent, the collection efficiency (CE) is the efficiency in which a created photoelectron is focused, multiplied, and produces a signal at the anode, and the front-end efficiency (FE) is the efficiency at which the PMT pulse crosses the discriminator and analysis thresholds. Between various PMTs each one of these factors can be different, and the relative coincidence rate only measures the relative total efficiency, E.

7.1.3 R5912-MOD PMT

The R5912-MOD is an 8-inch PMT developed by Hamamtasu Photonics with 10 linearly focused dynode stages. Shown in Figure 7.2 is the PMT specifications showing the dimensions, photocathode area, and basing diagram. Three R5912-MOD PMTs were tested and compared to similar PMTs. These prototype PMTs are expected to be less efficient than Hamamatsu's R5912-200 PMTs, which peak around 35% quantum efficiency (QE); however Hamamatsu could incorporate the super bialkali (HQE) photocathode on these PMTs (23).

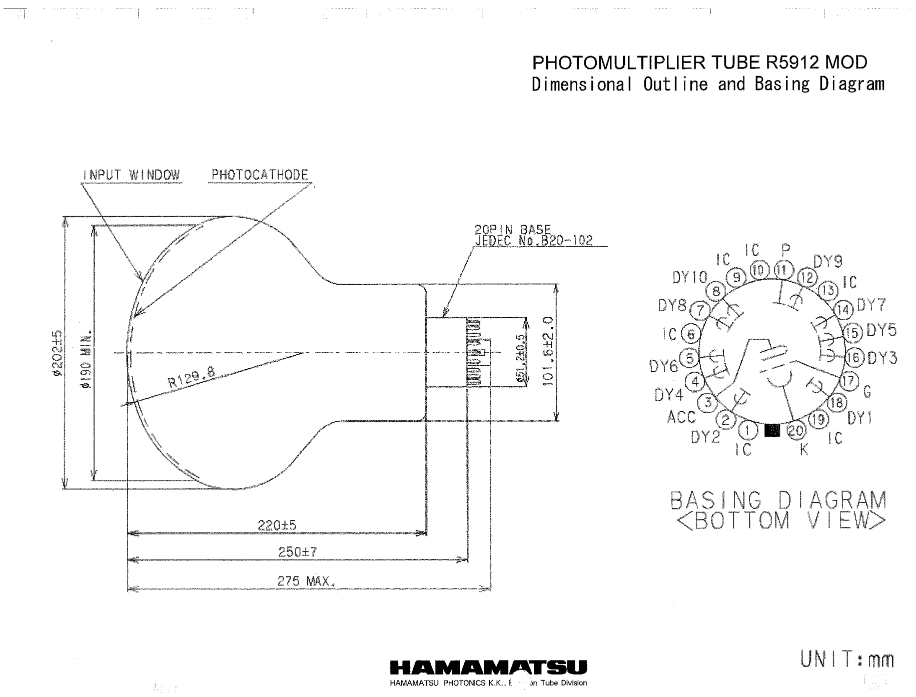
PHOTOMULTIPLIER TUBE R5912 MOD
Dimensional Outline and Basing Diagram
INPUT WINDOW
PHOTOCATHODE
20PIN BASE
JEDEC No.B20-102
φ202±5
φ190 MIN.
R129.8
φ51.2±0.5
101.6±2.0
220±5
250±7
275 MAX.
IC
IC
P
DY9
DY10
IC
DY8
DY7
IC
DY5
DY6
DY3
DY4
G
ACC
DY1
DY2
IC
K
IC
BASING DIAGRAM
<BOTTOM VIEW>
UNIT:mm
HAMAMATSU
HAMAMATSU PHOTONICS K.K.

Figure 7.2: The R5912-MOD specifications provided by Hamamatsu photonics (23).

Several base designs were tested for use with the R5912-MOD PMT. The voltage divider ratios for each dynode stage as implemented in the final base design is shown in Table 7.1.

The base design determines the high voltage that needs to be supplied to achieve a gain of 1×10^7, and was tuned so that the PMT operated at around 1800V.

-	K	DY1	G	DY2	DY3	DY4	DY5	DY6	DY7	ACC	DY8	DY9	DY10
Res. (ratio)	11.5	1	3.5	4	2	2	1	1	0	1	1	1	1
Cap (μF)											0.01	0.01	0.01

Table 7.1: Voltage divider ratios between each stage. For example the ratio for the drop from the cathode to the first dynode is 11.5 times larger than the voltage drop between the first dynode and the grid. The total resistance across the entire base is 16.8 MΩ.

The SPE charge and timing distributions for the R5912-MOD PMT are shown in Figures 7.3 and 7.4. These distributions are extracted using the setup described in Section 7.1.1 and the analysis described in Section 7.1.2. The charge distribution features a peak-to-valley of almost five, more than a factor of four better than the R1408 PMTs. The small two PE peak in the charge distribution can be clearly identified around 3.2 pC due to the narrow width of the charge distribution. The transit time distribution features a spread of about 640 ps, which is state of the art for large-area PMTs. The fast timing featured by this PMT makes it a particularly interesting candidate for use in future detectors.

The components of the transit time distribution can be broken up into the following categories:

Prompt pulses: The prompt peak makes up the primary response of the PMT, the spread of which is influenced heavily by the electron optics in the PMT. The prompt peak makes up 91.7% of the timing response.

Late pulses: The late light is the second largest component of the timing distribution, and is responsible for the peak at about 95 ns in Figure 7.5. The late light is caused by an elastic scatter off of the first dynode in which the photoelectron that travels back toward the photocathode before returning to the first dynode and causing the emission of secondary

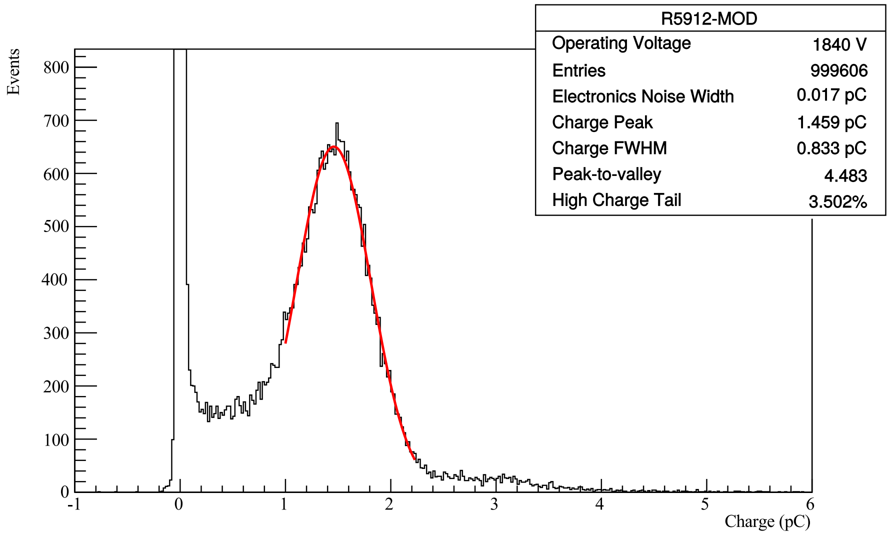

Figure 7.3: The charge distribution of the R5912-MOD PMT. Shown in the statistics box is some important characteristics of the SPE charge response. The Guassian fit to the SPE peak is shown in red. The χ^2/NDF is 70.68/59.

electrons. The late pulsing for these PMTs makes up about 6.1% of the PMTs response, which can be compared directly to the late ratio statistic in Figure 7.4, which does not correct for double pulsing.

Double pulses: The double pulsing has a similar time structure to the late light; however, in addition to the late pulse there is also a prompt pulse in the waveform. The time structure of both the initial prompt pulse and the later pulse is shown in Figure 7.5. The double pulsing is caused by an inelastic scatter off of the first dynode. There is enough energy transfer for secondary emission to take place; however, the photoelectron also recoils back toward the photocathode. Because the initial photoelectron does not maintain its full energy, it recoils over a shorter distance and thus it is expected that second of the double pulses to come slightly earlier than the late pulses. Figure 7.5 shows that is indeed the

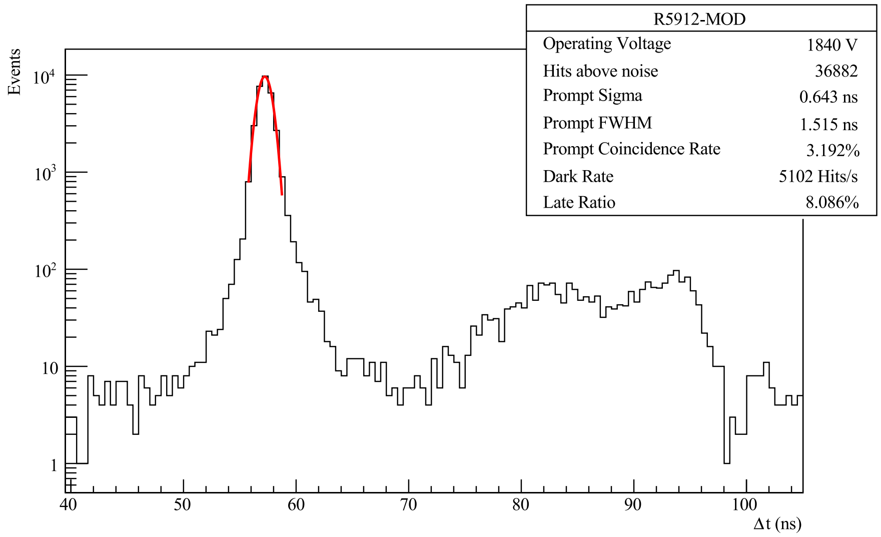

Figure 7.4: The transit time profile of the R5912-MOD. Shown in the statistics box is some important characteristic of the SPE time response. The Gaussian fit to the prompt light peak is shown in red.

case; the second pulse in the double pulsing arrives early on average than the late pulsing. Double pulsing makes up about 2.2% of the total timing structure.

Pre pulses: Pre pulsing is caused when a photon is transmitted, rather than absorbed or reflected, by the PMT glass and photocathode. The photon travels through the PMT vacuum and can strike the first dynode, causing the creation of a photoelectron at the first dynode rather than at the photocathode. One would expect to see this pre pulsing peak in the transit time distribution about 10ns before the prompt peak; however we cannot resolve any pre pulsing above the dark rate.

Dark rate: Dark pulses are caused primarily by thermionic emission of an electron at the photocathode and are not caused by incident light. There are several other ways to get dark current at the photocathode, including Cherenkov light produce by muons passing through the glass. Dark pulses are not a part of the transit time distribution, but provide a flat background that is accounted for when calculating the various percent contributions.

Figure 7.5 shows the R5912-MOD transit time distribution broken down into the various components. For the three PMTs tested the dark rate was between 2 to 5 kHz. The PMTs sit in the dark box for eight hours before the dark rate is extracted. With the statistics collected for this PMT, pre pulsing is not evident and must make up less than 0.1% of the total transit time distribution.

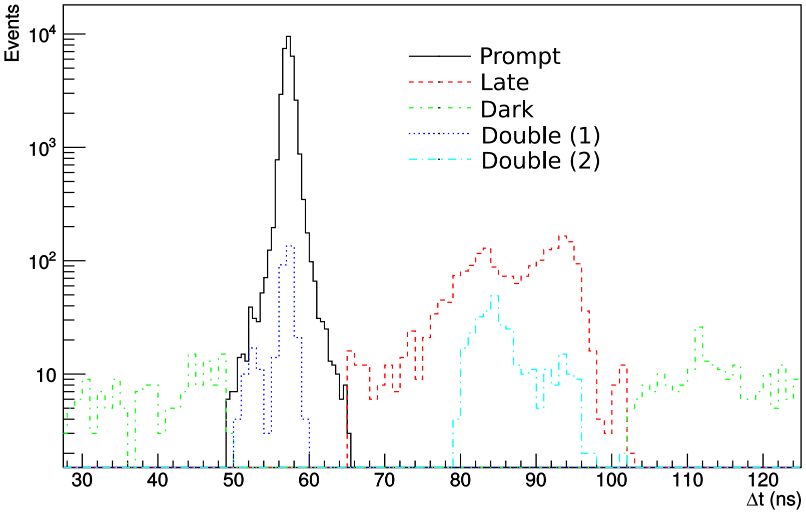

Figure 7.5: The transit timing profile of the R5912-MOD broken down into the various components that make up the structure.

Both the SPE charge and time response were characterized as a function of gain. The high voltage was changed from 1700 to 2100V in 50V steps, providing gains that range from about 0.45 to 1.7 $\times 10^7$. Various important SPE characteristics are tabulated against the high voltage supplied in Table 7.2. Notably, the relative efficiency, characterized by the coincidence rate, increases by 9% across the range of gains. The TTS remains constant across the high voltage values.

High Voltage (V)	Coincidence Rate (%)	Peak (pC)	TTS (ns)
1700	3.75	0.73	0.67
1760	3.82	0.91	0.64
1800	3.91	1.06	0.64
1880	4.00	1.36	0.65
1920	4.10	1.55	0.65
1960	4.12	1.78	0.62
2000	4.11	2.01	0.62
2100	4.13	2.69	0.65

Table 7.2: The SPE performance for the ZC2722 R5912-MOD PMT across gains spanning roughly 0.45×10^7 to 1.7×10^7. The coincidence rate changes by about 9% across the gain change and the TTS is constant within uncertainties.

In addition to measuring the charge and time response of the PMTs, the shape of the SPE waveform is characterized. The PMT waveforms are fit to a sum of three lognormal functions, given by:

$$f = \sum_{i=0}^{2} \frac{N_i}{t\sigma_i\sqrt{2\pi}} e^{-\frac{1}{2\sigma_i^2}\log(\frac{t}{\tau_i})^2}, \tag{7.2}$$

where σ_i, N_i, and τ_i for all three lognormals are floated in the fit. A example of this fit to a PMT waveform is shown in Figure 7.6. The PMT waveform is clearly non-Gaussian and is well characterized by the triple lognormal fit. Additionally, this function allows one to easily characterize common PMT features, such as the overshoot on the falling edge. This model is useful for building an accurate pulse shape model in simulations for correctly modeling the PMT response.

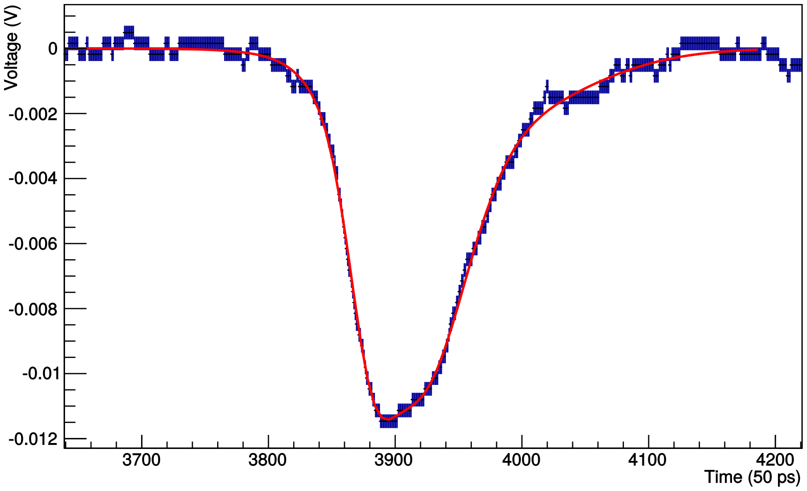

Figure 7.6: An example of an R5912-MOD PMT waveform with a triple lognormal fit shown in red.

The efficiency of the R5912-MOD PMT is compared to the R5912-200 and R1408 PMTs, both 8-inch PMTs developed by Hamamatsu, by measuring the coincidence rate for each PMT in the same setup. The relative efficiencies and other SPE parameters are compared between these PMTs in Table 7.3. It is clear that the R5912-200 PMT is the most efficient PMT, as expected, about a factor of 2 more than the R1408 PMT and a factor of 1.5 more than the R5912-MOD PMTs. The TTS and P/V is best in the R5912-MOD PMTs, which have a TTS around 0.65 ns and a P/V around 4.0.

Th afterpulsing of the R5912-MOD PMTs was also measured using a slightly different setup. Afterpulsing is caused by ionization of residual gases in the PMT, and is described in Section C.0.2. The setup used to probe afterpulsing consists of a 390 nm LED is collimated and passed through a 10nm wide optical filter, directed at the center of the R5912-MOD PMT. The beam spot of the LED is directed at the center of the PMT and is tuned to be

PMT	Coinc. Rate (%)	Rel. Eff.	TTS (ns)	Late Ratio (%)	P/V	HV
R5912-HQE	5.32	1.0	0.87	7.58	2.96	1740V
R1408 PBUT	2.47	0.46	1.51	7.18	1.08	2000V
R5912-MOD ZC2722	3.47	0.65	0.72	8.34	4.28	1840V
R5912-MOD ZC2723	3.44	0.65	0.63	8.09	4.46	1940V
R5912-MOD ZC2728	3.39	0.64	0.69	8.94	3.96	1740V

Table 7.3: The table of coincidence rates, TTS, late ratios, and peak to valleys of the three R5912-MODs tested as well as an R1408 and R5912 HQE for comparison.

about 2 cm in diameter. The signal to the LED is provided by a Agilent 33503A waveform generator, which pulses the LED with roughly 30 ns wide square pulse. The intensity of the LED is tuned using the pulse amplitude, and measurements have been made at several different intensities. The frequency of the pulses is set to 1 kHz so that there is no pile-up.

The signal from the rising edge of the signal provided by the waveform generator is used as a trigger. This ensures that the prompt light at the PMT comes after the trigger at a fixed amount of time. The scope settings are also adjusted so that the waveforms extracted are much longer in length, 50 μs, and the sampling time used is 0.1 ns which provides far better resolution than necessary. There is an intrinsic jitter on the prompt signal at the PMT associated with the width of the pulse driving the LED, which is dealt with in analysis. Because the afterpulsing distribution is very broad in time this 30 ns is a negligible jitter.

Figure 7.7 shows a 2D histogram of the Δt against charge of the afterpulses identified. There are two broad peaks around 1 μs and 6 μs with some additional structure early in time (around 400-600ns) and late in time (around 10 μs). The PMT for this measurement was run with a gain such that the peak of the charge distribution was around 2 pC. Figure C.2 shows that the majority of the afterpulsing are SPE, with a fairly large tail out to about 10 PE. The flat distribution after 20 μs is primarily due to dark pulses. Lastly, the afterpulsing rate per prompt photoelectron is identified as 16%, about an order of magnitude higher than the R1408 PMTs.

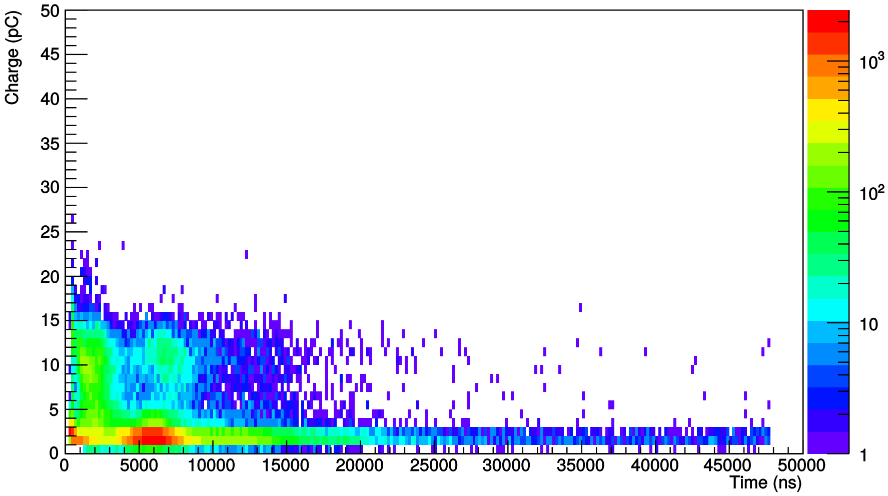

Figure 7.7: A 2D histogram showing the time between the prompt light and the afterpulse plotted against the charge of the afterpulse.

Overall, the R5912-MOD PMT shows excellent SPE characteristics, particularly an extremely narrow TTS for a PMT of its size. However, in conversations with Hamamatsu, the 100 mm neck shown in Figure 7.2 leads to pressure stability issues and is not feasible to be used in a large scale detector. New versions of the R5912 PMTs with similarly excellent timing, improved efficiency, and reduce after-pulsing are possible for the future (23).

7.1.4 11-inch ETEL PMT

The D784KFLB PMT developed by ET Enterprises has an 11-inch diameter photocathode and twelve linear-focused dynode stages. A known issue with these PMTs is a roughly 0.35 inch in diameter ‘swirly’ spot in the center of the front face, where the efficiency of the tube is anticipated to be impacted. The expected efficiency of these PMTs peaks at 30%.

Fifteen different ETEL PMTs were analyzed in the setup described in Section 7.1.1

and the digitized waveforms are analyzed using the procedure outline in Section 7.1.2. A summary of the SPE results for these PMTs are shown in Table 7.4. The details of the characterization can be found in (152) and the PMTs are now deployed in the ANNIE detector (122). The ETEL PMTs are particularly interesting because they provide a US-based competitor to Hamamatsu for the production of large-area PMTs.

11-Inch ETEL	Average	Standard Deviation	Minimum	Maximum
Charge FWHM (pC)	1.44	0.40	1.11	2.73
Peak/Valley	2.32	0.67	1.15	3.68
High Charge Tail (%)	3.86	1.28	0.921	5.71
TTS (σ_{prompt})(ns)	1.98	0.17	1.79	2.47
Late Ratio	4.51	0.74	3.0	5.76
Operating Voltage (V)	1330	117	1183	1575

Table 7.4: Summary of the SPE results for the D784KFLB ETEL PMTs. The operating voltage for all PMTs was adjusted for a gain of 1×10^7, corresponding to a SPE peak of 1.6pC. This table includes the results for 19 PMTs.

7.1.5 R5912-200 PMT

The R5912-200 high quantum efficiency PMTs are a version of the Hamamatsu R5912 8-inch PMT with a higher efficiency photocathode (improved from a quantum efficiency peaking at about 20% to 35%). A detailed characterization of these PMTs is performed in-situ by the DEAP collaboration, which deploy 255 of these PMTs (153).

Four R5912-200 PMTs were added to the SNO+ detector in 2017 while the detector was filling with water. To install them in SNO+, a custom base based on the DEAP designed was fabricated, shown on the left in Figure 7.8. Additionally, the wider neck (relative to the R1408 PMTs) required a new, custom waterproof potting. The finished product is shown on the right in Figure 7.8. These PMTs operate at high voltages around 1500V, which required careful tuning of the feed resistors in the PMTICs corresponding to the channels for the four HQE PMTs.

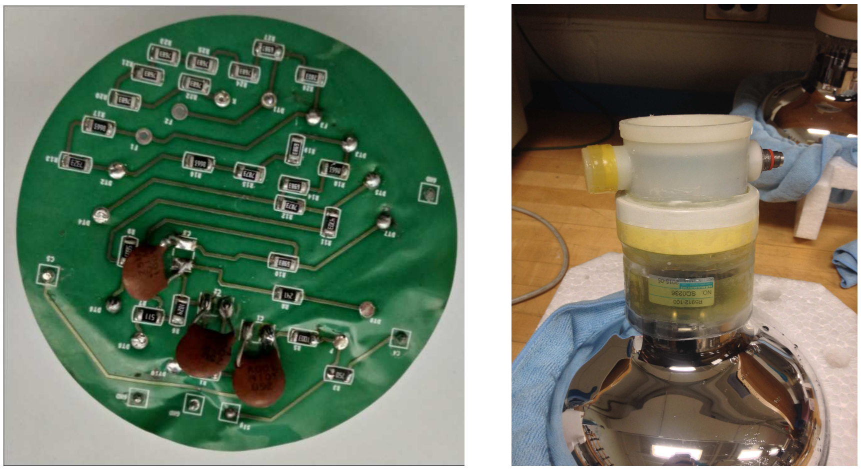

Figure 7.8: The base used for the R5912-200 HQE PMTs (left). An R5912-200 HQE after custom potting, before going into the SNO+ detector (right).

The goal in installing these PMTs is to make in-site measurements of the PMT response and compare directly to the R1408 response. A measure of the charge, timing, and efficiencies of these PMTs using SNO+ calibration data is useful for understanding how SNO+ would perform if the collaboration was able to replace all of the R1408 PMTs with these PMTS. The R5912-200 PMTs are installed in the same concentrators as the R1408 PMTs.

Before entering the detector, the four PMTs were characterized at Penn using a source of SPEs in a darkbox. The quantum efficiency of the R5912 compared to two SNO PMTs is shown in Figure 7.9 (left). At 400 nm, where the peak of the scintillation light lies, the R5912 HQE PMTs is about twice as efficienct. The SPE charge distribution for the R5912 compared to the SNO PMTs is shown in Figure 7.10. The charge distribution of the R5912 HQE's have a much higher peak-to-valley ratio and a smaller high-charge tail. Thes features improve the efficiency at which a PMT pulse crosses a discriminator threshold at one-quarter

PE (the 'front-end efficiency') because more there are fewer small SPE PMT pulses that fail to cross threshold. Thus, these PMTs would improve both the detection efficiency and the front-end efficiency of the detector. The transit time distribution for these HQE PMTs is shown in Figure 7.11. A TTS of 0.8 to 0.9 ns is measured, compared to approximately 1.5 ns for the R1408 PMTs. Additionally, the late-pulsing makes up about 5% of the total distribution, which is slightly lower than the SNO PMTs. The details regarding these measurements can be found in much more detail in (154). Overall, if SNO+ were able to replace the PMTs in the detector with these HQE PMTs it would drastically improve the energy and position resolution of the detector, allowing SNO+ to lower backgrounds.

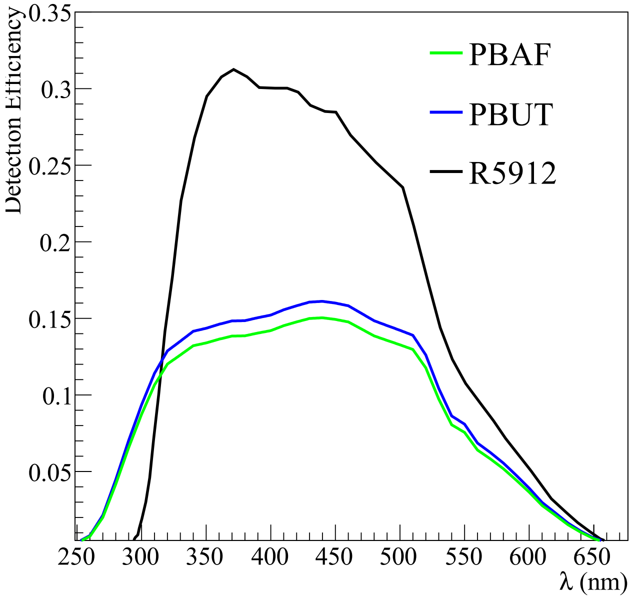

Figure 7.9: The R5912-200 HQE detection efficiency curve compared to two SNO PMTs. Relative scaling between PMT efficiencies is from bench-top data at Penn.

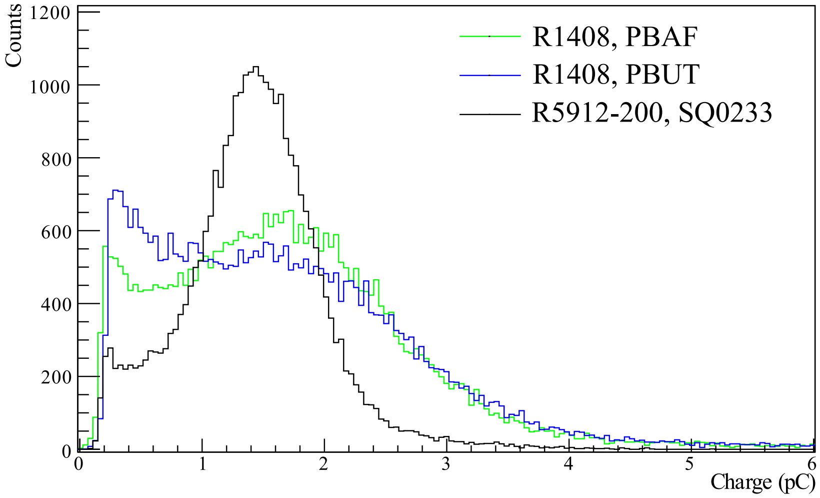

Figure 7.10: The R5912-200 HQE SPE charge distribution compared to two SNO R1408 PMTs.

7.2 Liquid Scintillator

The optical performance and radiopurity of the SNO+ liquid scintillator, LAB+PPO, are cruicial to the success of the experiment. In addition to LAB+PPO, enormous R&D effort has gone into developing and understanding the optics and stability of tellurium loaded liquid scintillator. The measurements presented in the section focus on unloaded LA+PPO (2 g/L), LAB+PPO+bisMSB (15 mg/L), and Te-loaded LAB+PPO. The Te is loaded using 1,2 butanediol and N,N-Dimethyldodecylamine (DDA) as surfactants. Notably, the light output of the Te-loaded LAB+PPO is quenched relative to LAB+PPO, leading to a lower light yield. Additionally, the α and β emission time profiles are altered by the additional quenching. Interestingly, the DDA reduces the overall quenching and recovers about 15% of the light yield lost by loading the tellurium with 1,2 butanediol.

Prior to filling and running the detector with liquid scintillator, predictions regarding the

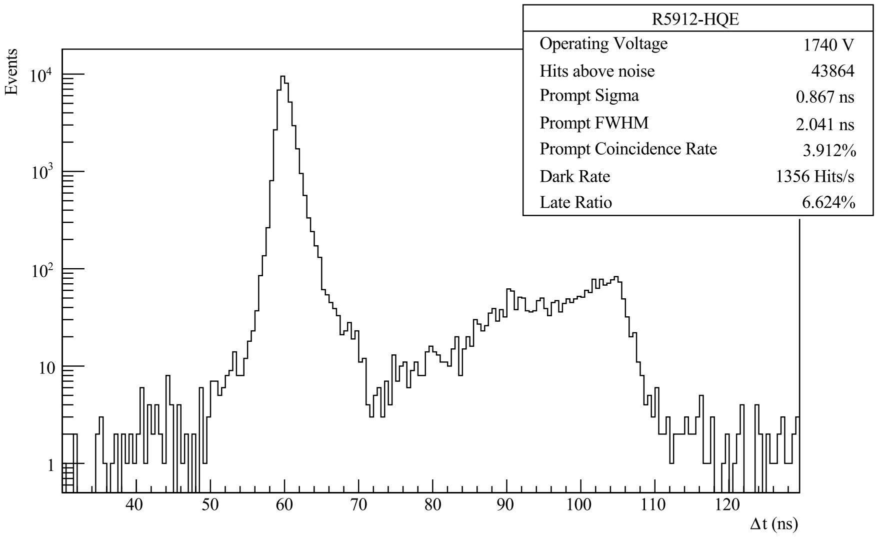

Figure 7.11: The R5912-200 HQE SPE transit time distribution.

detector performance and sensitivity to various physics ($0\nu\beta\beta$, solar neutrino interactions, etc) are critically important to guide the collaborations effort in maximizing the potential of the detector. In order to perform these predictions, developing a detailed scintillator model that accurately predicts the behavior of the scintillator is crucial. This is not an easy task – liquid scintillator behaves observably different at different scales due to the absorption and re-emission of the scintillator. The average absorption length is wavelength dependennt and can range from less than a mm to several meters, depending on the concentration of the various fluors and wavelength shifters and in the production of the LAB. Measurements on the bench-top presented in this chapter are critical components of the full scintillator model. Additionally, significant effort is expended to model each of the bench-top setups in the `RAT` Monte Carlo software, used for the full detector simulation, in order to make predictions with the scintillator model that can be tested on the bench. This at least gives small-scale verification of the detailed model. A larger scale test of the model is presented in Section 7.2.5.

Liquid scintillators are a commonly used target for counting of radioactive samples. The ionizing radiation excited the solvent molecules, which transfers its energy non-radiatively to the fluor. As the excited molecules relax to the ground state, de-excitations photons are emitted. There are several critical characteristics of the liquid scintillator and the emitted photons: the decay time of the scintillator, the total number of emitted photons (light yield), and the wavelength emission and absorption spectra of the fluor.

Upon excitation from a charged particle, the scintillator can enter a spin 0 singlet excited state or spin 1 triplet excited state. The singlet de-excites quickly and is often referred to as the fast scintillation component. The triple state de-excited over a longer period off time and results in a long-time scintillation component. Empirically, the scintillation emission time is often described by summing several decay exponentials with different time constants of the form:

$$\sum_{i=1}^{n} \frac{N_i}{\tau_i} \exp(-t/\tau_i), \tag{7.3}$$

where τ_i is the decay time constant associated with the ith exponential component weighted by N_i. This model is often chosen to include three or four components, depending on the scintillator used. For the measurements presented in Section 7.2.3 an empirical model using an $n = 4$ decay exponential is selected. Similar methods for measuring the LAB+PPO time profile are discussed in (155), (156), (157), (158).

7.2.1 Experimental Setup

The experimental setup for the scintillator light yield and emission timing are similar. In both cases, the source of scintillation light is a ^{90}Sr or ^{210}Po source deployed above a hollowed-out UVT acrylic block filled with scintillator. The acrylic block is 3 x 3 x 3.5 cm and the cylindrical volume hollowed-out is 2 cm in diameter. The source is similar to the Cherenkov source described in Section 7.1.1. Nitrogen is bubbled through the 15 mL of liquid scintillator for fifteen minutes in order to remove dissolved oxygen, which can cause some quenching.

A picture of the setup used to characterize the scintillator emission timing is shown in Figure 7.12. This setup utilizes a 1-inch square Hamamatsu R7600-U200 PMT optically coupled to the scintillation source, which is used to trigger on the scintillation signal. A second R7600-U200 PMT is located about 30 cm away from the scintillation source, which detects primarily SPEs.

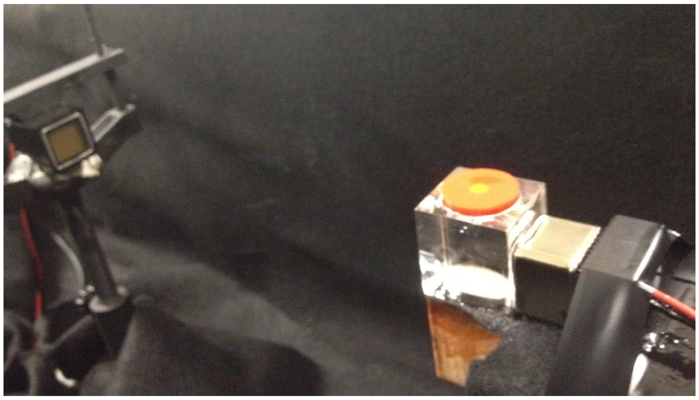

Figure 7.12: The scintillator emission timing setup, which includes an R7600-U200 PMT optically coupled to the scintillation source, which is used as a fast trigger. In this picture, a red ^{210}Po disk source is deployed above a sample of liquid scintillator, held inside a UVT acrylic block. A second R7600-U200 PMT located about 30 cm away from the scintillation source.

The light yield setup is similar to the timing setup, with the second R7600-U200 PMT replaced by an 8-inch Hamamatsu R1408 PMT, the same type as those deployed in the SNO+ detector. This PMT was chosen because the existing PMT model in the SNO+ Monte Carlo is well understood, which allows us to build a detailed model of our setup into `RAT`. The DAQ and electronics are identical to the one described in Section 7.1.1. In both setups the digitized waveforms are analyzed using offline C++ code.

7.2.2 `RAT` Model

The simulation of the darkboxes utilizes `RAT`, the same software used to simulate the SNO+ detector. The goal of including a detailed model of the timing and light yield setups is to confirm the scintillator model against bench-top data at various scales in order to confidently extrapolate to the scale of the full detector.

The `RAT` model includes the PMTs and scintillation source locations, as shown in Figure 7.13 for the emission timing and light yield setups. The ^{90}Sr or ^{210} Po sources are modeled as a point source of isotropically emitted β or α particles according to the appropriate spectrum.

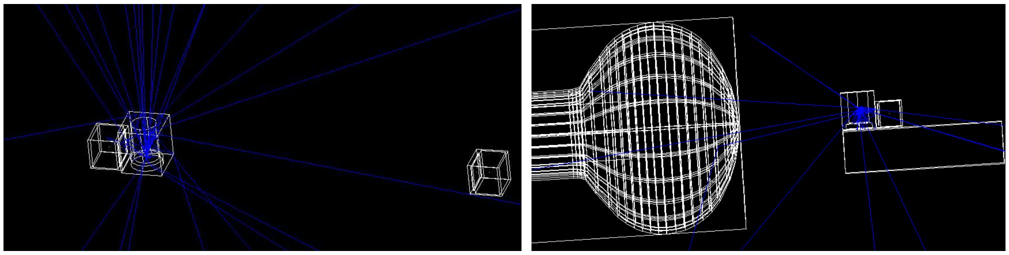

Figure 7.13: The `RAT` model of the scintillation emission timing setup (left), corresponding to Figure 7.12, and the light yield setup (right). The blue lines show the paths of various optical photons. The 1-inch Hamamatsu trigger and measurement PMTs are shown for the timing setup and the R1408 PMT is shown in the light yield setup.

The simulated efficiency of the PMTs is taken from the Hamamatsu datasheet, and the overall normalization to the efficiency curve, which impacts the light yield measurement, is calibrated. The calibration method is discussed in Section 7.2.4. For each detected photon, a simulated PE is created. The PMT pulses for each PE are modeled using a lognormal distribution, shown in Equation 7.2, that draws from parameters taken from fit to the SPE data. An example of a fit is shown in Figure 7.14, which includes the list of parameters for that particular fit. After creating a pulse for each PE, the pulses are summed in-time for every PE detected by a particular PMT. Lastly, electronic noise is added to the pulses, the magnitude of which comes directly from in-situ data.

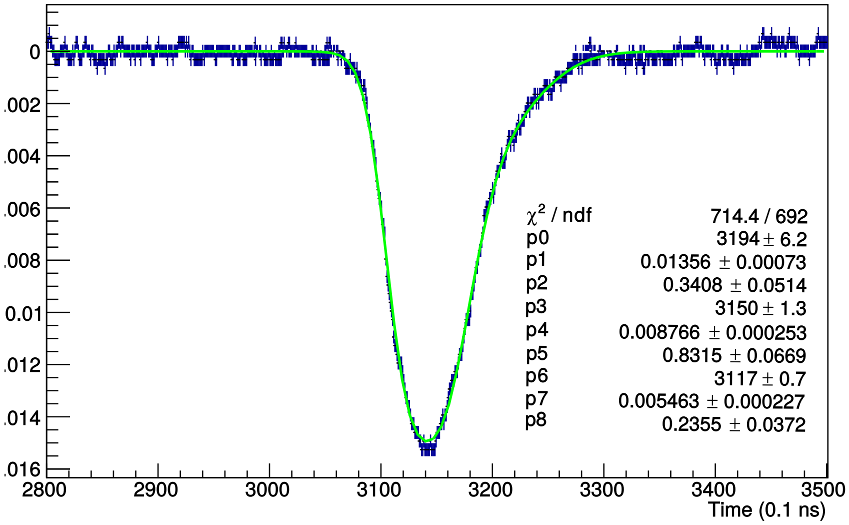

Figure 7.14: A R7600-U200 SPE waveform fit to a triple lognormal distribution (Equation 7.2). The corresponding best fit parameters are shown.

After generating the total pulse shapes for both the trigger and measurement PMT, the waveforms are digitized using a model of the oscilloscope. The oscilloscope model samples the waveform at the appropriate rate and saves the output voltage for each sample and channel. The waveform for the channel chosen as the trigger is compared against a tunable threshold and if it crosses threshold, the waveform is written to disk. The files are formatted identically to the data, in `hdf5` files that can be processed by the same code used to analyze the data. Figure 7.15 shows an example simulated event that crosses trigger threshold for a setup with an R7600-U200 trigger PMT and an R1408 measurement PMT.

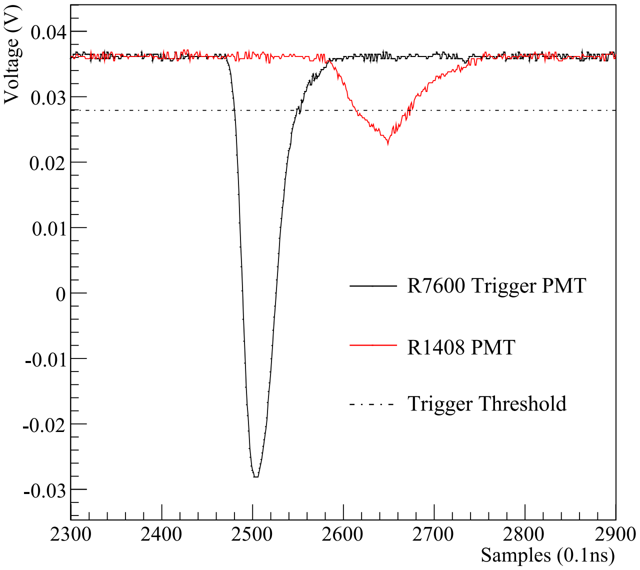

Figure 7.15: Simulated digitized waveforms for a R7600-U200 and R1408 PMTs. The R7600-U200 waveform is compared agianst a trigger threshold. For triggered events, the simulated waveforms for up to four channels are written to disk in .hdf5 format.

7.2.3 Emission Timing

The emission time profiles for LAB+PPO and Te-loaded LAB+PPO are characterized in this section. The LAB+PPO emission time is important for background rejection during the unloaded phase, which would improve the solar neutrino and reactor antineutrino measurements. Similarly, the emission profile of the Te-loaded LAB+PPO scintillator is critical to understand for background rejection in the $0\nu\beta\beta$ phase.

In order to select a pure sample of either β or α particles interaction in the scintillator, a cut on the charge of the trigger PMT is used. For events passing the charge cut, the analysis finds time of the 40% crossing for both the trigger PMT and the measurement PMT and histograms the time difference, Δt. A measurement taken with no source indicated a negligible background rate from non-source related events. The LAB+PPO and Te-loaded LAB+PPO time profiles for β and α particle excitation is shown in Figure 7.16. Notably, the triplet state excitation in the Te-loaded LAB+PPO is quenched relative to LAB+PPO, which causes less light to be emitted at later times. For both scintillators the α particle excitation yields a significantly longer tail, which can be used for pulse shape discrimination between βs and αs.

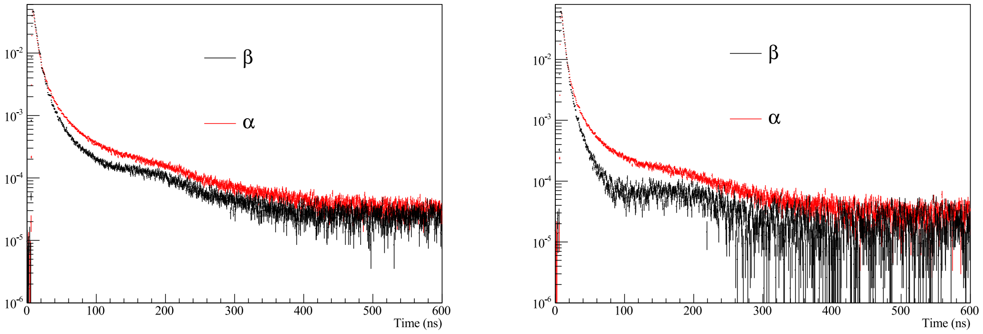

Figure 7.16: The β and α timing profiles compared for LAB+PPO (left) and Te-loaded LAB+PPO (right). For both scintillators the α particles causes a time profile with a longer tail, which can be used to discriminate between β and α particle excitation. The Te-loaded LAB+PPO has a quenched triplet state excitation relative to LAB+PPO, which reduces the number of photons emitted at larger times.

The Δt spectra are fit using the `RooFit` package (159), where the scintillation emission time profile is modeled as the sum of four decay exponentials, as shown in Equation 7.3. The full fit equation, which includes several detector related components given in Equation 7.4.

$$F = A_1 \times \left[\frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(t-\mu)^2}{2\sigma^2}} \times \sum_{i=1}^{4} \frac{N_i}{\tau_i} e^{-t/\tau_i} \right] + A_2 \times \left[\frac{1}{\sqrt{2\pi\sigma_{ap}^2}} e^{-\frac{(t-\mu_{ap})^2}{2\sigma_{ap}^2}} \right] + A_3 \qquad (7.4)$$

The fit to the time-profile can be separated into three distinct components. The first component describes the emission time-profile of the scintillation light convolved with the PMT response function. This is modeled as a Gaussian convolved with the sum of four exponentials. The exponentials are convolved with a Gaussian, used to model the transit time spread of the PMTs. The PMT response function has a width, σ, of 600 ps and mean, μ, constrained around the prompt peak. The width of the PMT response is determined in a separate measurement using a source of Cherenkov light, and is allowed to float in the fit in the range of the uncertainty provided by that measurement

The second component is a small Gaussian after-pulsing peak that occurs around 160 ns after the initial prompt light. The details of this after-pulsing feature are measured separately and used to constrain the parameters, μ_{ap} and σ_{ap}, in the fit. This after-pulsing peak is can be identified in the Δt histograms as a small peak around 200 ns. It is more clear in the β particle data, due to the fact that there is less scintillation light emitted with the longest time-constants.

The final component is used to model the dark-rate, which contributes a flat background in the Δt histogram. The relative normalizations of the three components of the fit are handled using A_i. The fit to the β and α particle LAB+PPO data is shown in Figure 7.17. Note that the Δt offset from zero is arbitrary and included various cable delays and DAQ timing offset.

The results for the normalizations, N_i, and time-constants, τ_i, from the fit to the LAB+PPO time profiles for β and α particle excitation are given in Table 7.5 and 7.6.

A simple and common method for quantifying the discrimination is using the fraction of prompt light relative to the total amount of light, called f_{prompt} and discussed in (160).

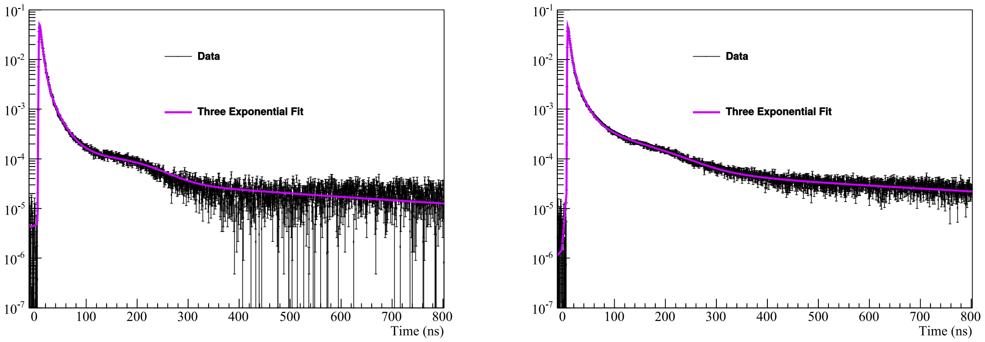

Figure 7.17: Example fits for LAB+PPO for β (left) and α (right) particles. The fit equation used is shown in Equation 7.4.

Cocktail	Particle	N_1	N_2	N_3	N_4
LAB+PPO	α	0.47 ± 0.02	0.32 ± 0.01	0.14 ± 0.01	0.07 ± 0.02
Te-loaded LAB+PPO	α	0.63 ± 0.02	0.23 ± 0.02	0.07 ± 0.02	0.07 ± 0.02
LAB+PPO	β	0.66 ± 0.02	0.19 ± 0.01	0.08 ± 0.01	0.05 ± 0.02
Te-loaded LAB+PPO	β	0.72 ± 0.02	0.23 ± 0.02	0.02 ± 0.02	0.03 ± 0.02

Table 7.5: The normalizations, N_i, from the fit to the LAB+PPO time profiles given in Equation 7.4.

Cocktail	Particle	τ_1	τ_2	τ_3	τ_4
LAB+PPO	α	4.28 ± 0.18	14.3 ± 1.9	71.1 ± 8.5	792 ± 371
Te-loaded LAB+PPO	α	3.69 ± 0.10	15.5 ± 1.3	79.3 ± 10.0	489 ± 164
LAB+PPO	β	5.10 ± 0.20	15.3 ± 2.0	48.4 ± 6.8	499 ± 150
Te-loaded LAB+PPO	β	3.70 ± 0.26	10.0 ± 2.2	52.0 ± 12.0	500 ± 176

Table 7.6: The time constants, τ_i, from the fit to the LAB+PPO time profiles given in Equation 7.4.

Here f_{prompt} is defined as:

$$f_{\text{prompt}} = \frac{\int_0^{10} \Delta t dt}{\int_0^{800} \Delta t dt}, \tag{7.5}$$

where the Δt spectrum is integrated in a prompt window and compared to total number of events in the entire distribution. The integrated Δt spectra are created by randomly drawing 1200 PEs according to the distributions shown in Figure 7.4. The 1200 PEs were selected as roughly corresponding to approximately the number of detected photoelectrons expected in the SNO+ detector for a 2 MeV β. After these spectra are generated for both β and α particles, f_{prompt} is calculated. The f_{prompt} distributions for 200000 events is shown in Figure 7.18. The distributions demonstrate excellent separation for β and α particles, which is actually improved for the Te-loaded LAB+PPO over the pure LAB+PPO. This level of separation is encouraging for background rejection in the SNO+ detector.

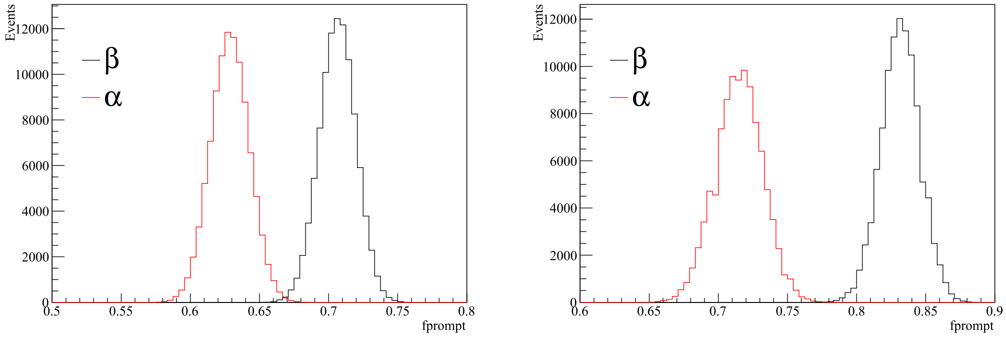

Figure 7.18: The f_{prompt} distribution for LAB+PPO and Te-loaded LAB+PPO, which show excellent discrimination between β and α particles.

In `RAT`, the scintillator emission time profile is modeled according to Equation 7.3. As a test of the model, a `RAT` geometry that includes the R7600-U200 PMTs and the scintillation source is built and shown in Figure 7.13 and described in Section 7.2.2.

The results of the simulation for the ^{210}Po, using the scintillation time constants and normalizations from Tables 7.5 and 7.6, is shown compared to the data in Figure 7.19. The χ^2/NDF between the data and simulation is around 2, which improves to around 1.2 when including a rise-time of 800 ps in the simulation. This nice match between data and `RAT` indicates that the scintillator model is accurately reproducing the bench-top data, which is encouraging for extrapolating to the full-scale detector.

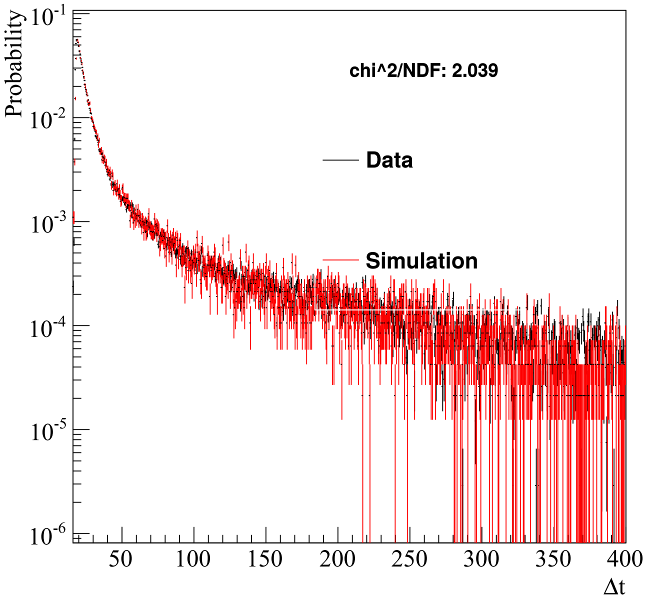

Figure 7.19: A comparison between the data and `RAT` simulation for LAB+PPO with the ^{210}Po α source.

Additionally, measurements of the re-emission time are made with the bisMSB wavelength shifter dissolved in the LAB+PPO were made using the full `RAT` model. The bisMSB absorbs the emitted light from the PPO and re-emits it at longer wavelengths around 440 nm. The absorption and re-emission processes add an additional delay to the time profile. The re-emission time-constant in `RAT` is set to 1.4 ns. Additionally, the bisMSB absorption length is critical for this measurement because the path length of the typical optical photon through the scintillator in our experimental setup is only a couple of centimeters. Because the bisMSB abosrption length is on the same scale as our setup in order to measure the time-constant one must actually run the full `RAT` model and compare against data. Figure 7.20 shows the result of the simulation compared to data for bisMSB loaded scintillator.

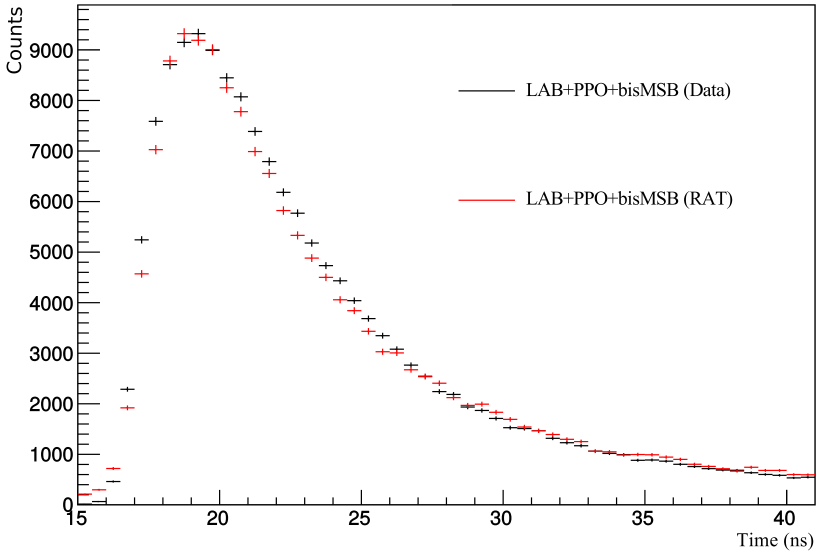

Figure 7.20: The LAB+PPO+bisMSB emission timing for data and simulation using a 1.4 ns re-emission constant in `RAT`. The fast-time component of the emission time is slightly slowed down due to additional absortion and re-emission, and this effect is well modeled in `RAT`.

7.2.4 Light Yield

The total number of photons emitted by the liquid scintillator is described empirically by Birk's law, given in Equation in 4.3. Birk's constant for LAB has been measured in

(161) as 0.0079 cm/MeV. Measurements of S or LAB+PPO and Te-loaded LAB+PPO are performed by generating a full `RAT` model of our darkbox setup described in Section 7.2.2. Using the simulation, a measurement of S can be performed by tuning the value in the simulation until it matches data. This technique is different than a typical method, which compares the light yield to a different scintillator with a known light yield. The method presented in the section has the advantages of simultaneously performing a check of the model, which can be used to extrapolate to larger scales, as discussed in Section 7.2.5. The light yield measurement is performed with the ^{90}Sr β source, and separate measurements of the α quenching are made.

In the light yield setup, the R7600-U200 PMT provides the trigger signal; however only the waveform from the R1408 is digitized. The R1408 is placed within 10 cm of the source and the waveforms are multi PE. The baseline is identified using a window before the trigger and accounted for when integrating the R1408 pulse. The integral is converted into a total charge (in pC) using Ohm's law and the total integration time.

In order to model the setup in `RAT`, the gain and overall efficiency of the PMT must be calibrated. This is performed using the Cherenkov source in the same setup, which provides the R1408 PMT with a source of primarily SPE light. The pulse size and shapes for a typical SPE pulse are extracted using lognormal fits and input into the simulation. The efficiency is extracted by scaling the quantum efficiency by an overall tuning factor, referred to as the collection efficiency (it does not reflect the value of the true collection efficiency). Figure 7.21 illustrates the process of tuning the collection efficiency to match the Cherenkov source data, which is calibrated to 1.3.

The collection efficiency is held fixed at 1.3 for the light yield measurement, and no additional tuning of parameters other than that of the light yield occurs. The LAB+PPO scintillation efficiency paramater S is tuned in `RAT` and compared against the data. The light yield that minimizes the residuals between data and simulation is selected, correspond to an $S = 11900$ photons/MeV. The data and `RAT` simulation with this light yield is shown in

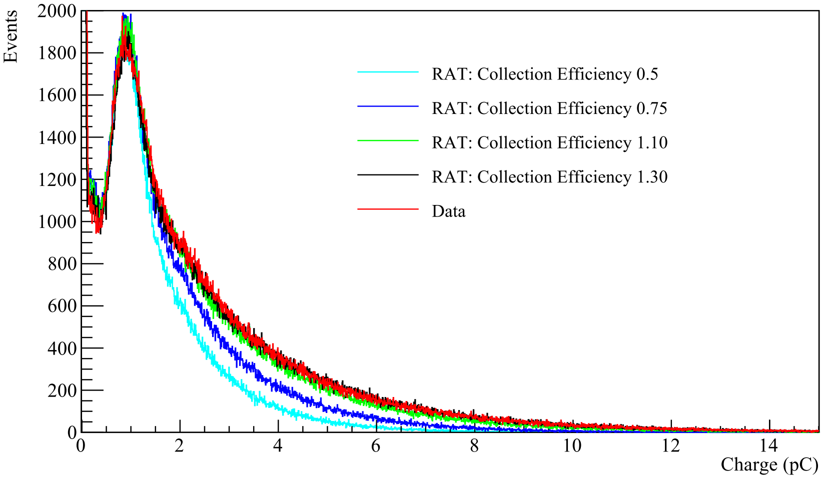

Figure 7.21: The Cherenkov source data for the R1408 PMT used for calibration prior to the light yield measurements. The overall efficiency of the PMT in `RAT` is scaled by a factor called the collection efficiency. The parameters are tuned in `RAT` until the Cherenkov source data and simulation agree, and is held fixed in the light yield simulations. For this measurement, a collection efficiency of 1.3 is selected based on the agreement with the data.

Figure 7.22. The same procedure is carried out with the Te-loaded LAB+PPO and $S = 7850$ photons/MeV is found. As expected from the scintillation emission timing measurements, some of the light is quenched by the additional components in the Te-loaded scintillator.

7.2.5 Scaling to Larger Volumes

The scintillator model in `RAT` contains the light yield and timing measurements presented, as well as measurement of the wavelength-dependent refractive indices (162), emission spectra, and absorption lengths from other sources. In order to be successful, this model should be predictive at larger scales. To test this, an approximately 30 L volume of liquid scintillator is added to a spherical acrylic vessel, shown in Figure 7.23. In this 30L

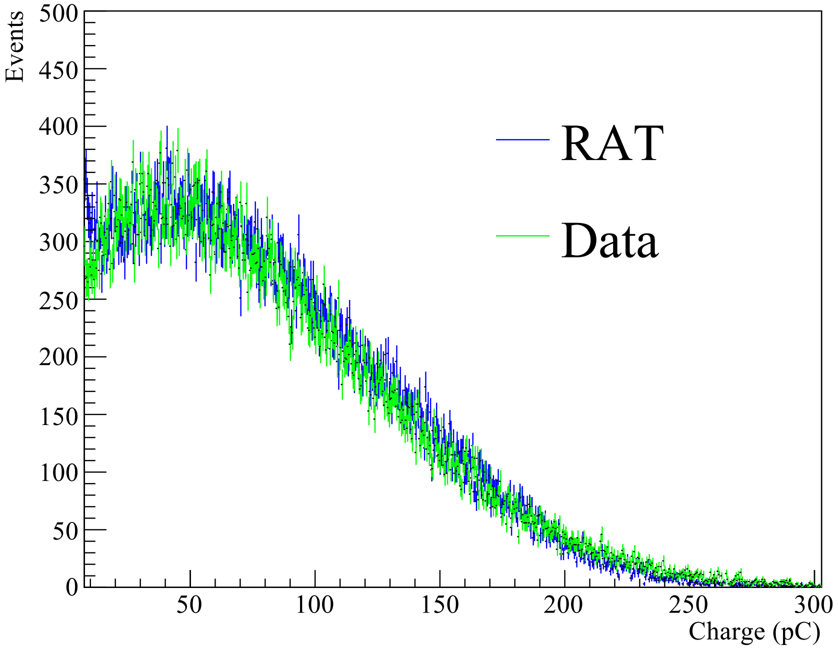

Figure 7.22: The R1408 PMT in data and simulation for the light yield setup. The scintillation efficiency, S, is tuned to 11900 photons/MeV.

volume, the ^{90}Sr β source is deployed at the center. Four PMTs are placed around the volume, one of which is the same R1408 PMT used for the light yield measurement. The PMT efficiencies are calibrated according the procedure in Section 7.2.4, and the R1408 is kept at the same high voltage and the collection efficiency is set to 1.3.

With no tuning of the model, the output of `RAT` matches data nicely, as seen in Figure 7.24, which are results for the same R1408 PMT used to perform the light yield measurement. The sensitivity projections for the SNO+ rely on the accuracy of the scintillator model, which is shown to be predictive at two scales, where the optics behaves differently due to additional absorption and re-emission across longer path lengths.

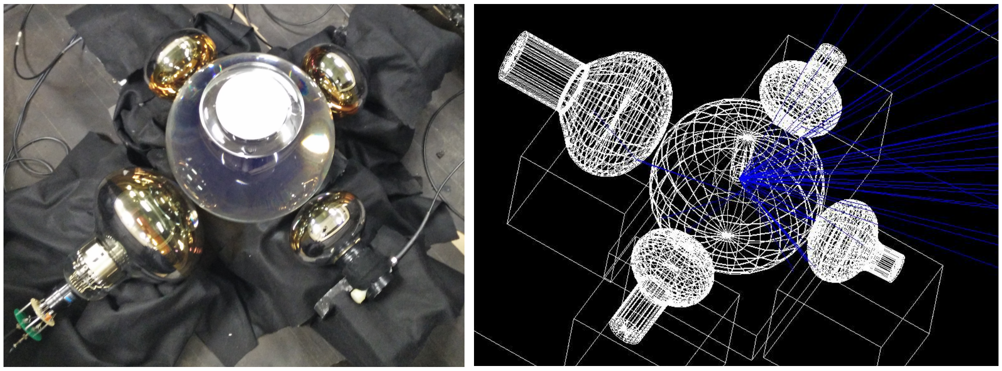

Figure 7.23: The 30 L liquid scintillator setup with four PMTs. This setup is used as a large-scale test of the `RAT` scintillator model. In the `RAT` model the blue lines show optical photons from a ^{90}Sr β decay at the center of the 30 L volume.

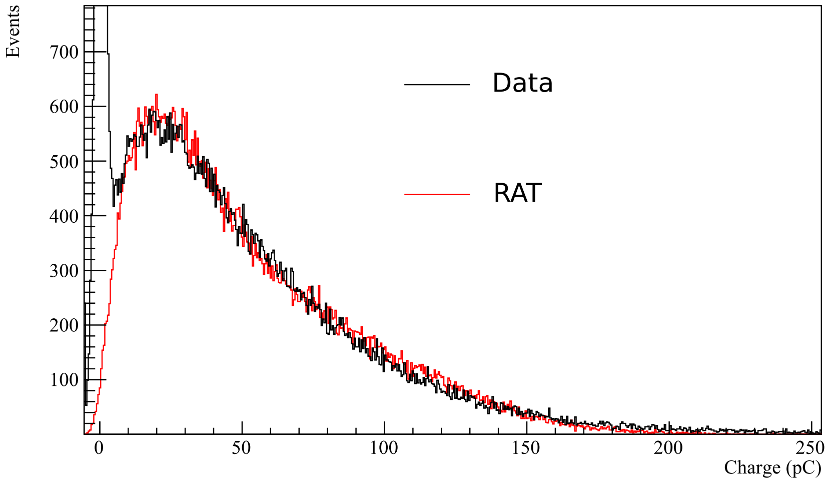

Figure 7.24: The R1408 PMT data and MC for the 30L setup. The collection efficiency is set to 1.3 and the scintillation efficiency is set to 11900.

The predicitive power of the optical model becomes even more important for the tellurium-loaded scintillator. As discussed in Section 8.2, the $0\nu\beta\beta$ sensitivity estimates depend strongly on the light yield of the scintillator, which is extracted from the measurements described in this section. Figure 7.25 shows the total number of PMTs that detect light (primarily SPE) for a simulated 1 MeV electron in the center of the SNO+ detector filled with Te-loaded scintillator without DDA. Adding the DDA improves the expected PE per MeV to about 450.

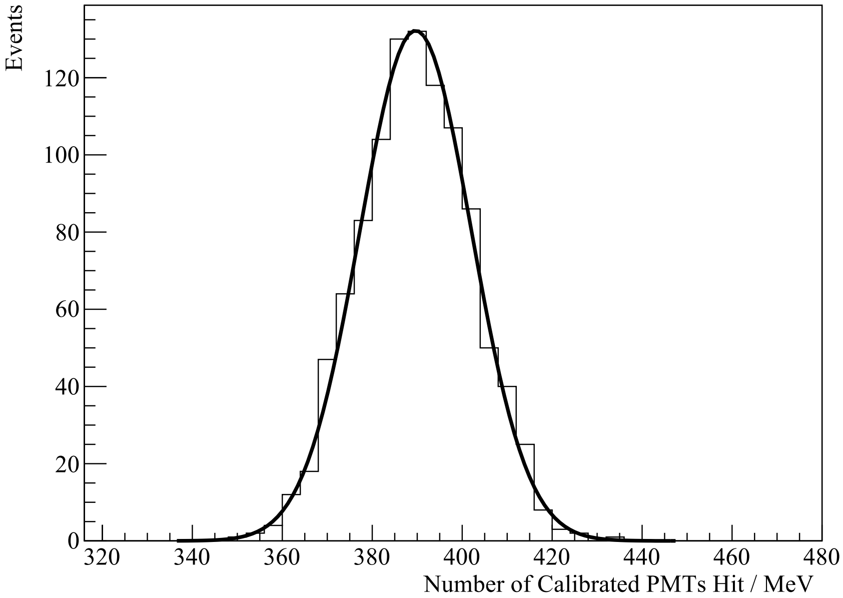

Figure 7.25: The number of expected PMTs that detect light (primarily SPE) per MeV, determined by simulating 1 MeV electrons at the center of the SNO+ detector using the optical model for Te-loaded LAB+PPO, without DDA. Adding the DDA improves the PE per MeV to about 450.

Chapter 8

The Dichroicon

The dichroicon is an instrument I co-invented, built, and tested at the University of Pennsylvania. The motivation for the dichroicon comes from years of work of detailed background modeling for large-scale liquid scintillator and water-based liquid scintillator detectors, such as SNO+ and `THEIA` (40), (86). These efforts have made clear the importance of direction reconstruction in scintillation-based detectors for background reduction. I give an explicit example of the impact of direction reconstruction in my discussion of the $0\nu\beta\beta$ decay sensitivity with the SNO+ detector in Section 8.1. In Sections 8.3 and 8.4 the Dichroicon is introduced and bench-top measurements of a prototype device are discussed. These sections largely follow two papers that I wrote: (163), (164).

8.1 Neutrinoless Double Beta Decay

The nature of the neutrino mass remains the most pressing open question in the field, and the clear consensus to determine whether the neutrino is Majorana or Dirac is through the search for neutrinoless double beta decay ($0\nu\beta\beta$). Double-beta decay ($2\nu\beta\beta$) is a rare type of radioactive decay in which two neutrons are simultaneously converted into two protons in an atomic nucleus. Each of these decays comes with the emission of an electron and an electron antineutrino:

$$N(A,Z) \to N(A,Z+2) + 2e^- + 2\bar{\nu}_e. \tag{8.1}$$

This process occurs in nuclei in which the $N(A,Z)$ nucleus has a ground state energy larger than the $N(A,Z+2)$ nucleus and less than the energy of the $N(A,Z+1)$ nucleus. This is shown schematically in Figure 8.1. The two neutrino transition is a second order weak process, leading to extremely slow decay rates. This process was first proposed by Goeppert-Mayer (165) in 1935 and first measured in ^{82}Se in 1987 (166). Several isotopes known to undergo double beta decay are shown in Table 8.1 with the corresponding Q-value, isotopic abundance, and half-life for each isotope.

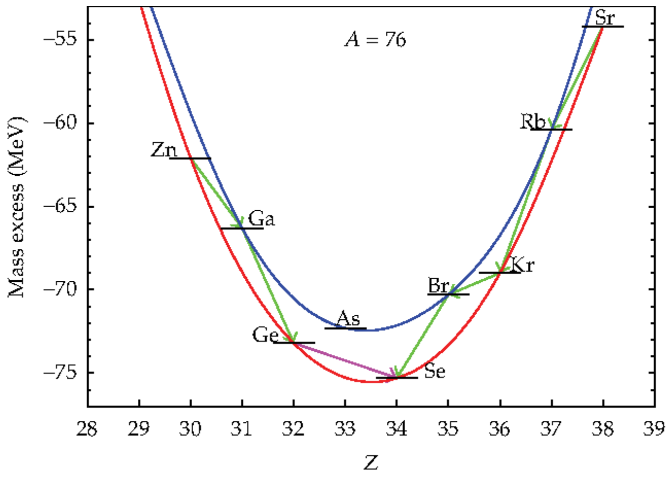

Figure 8.1: A schematic of the energies of the A = 76 isobars. The single-beta decay (shown in each case by the green arrow) is energetically forbidden between ^{76}Ge and ^{76}As. However, the double beta transition between ^{76}Ge and ^{76}Se (shown by the purple arrow) is allowed and proceeds with a very long half-life given in Table 8.1. This figure is from (5).

If the neutrino is Majorana is nature, double beta decay can occur without the emission of neutrinos. In the simplest scenario, $0\nu\beta\beta$ decay is mediated by the exchange of light

Nuclide	Q-value (MeV)	Isotopic Abundance (%)	Half-life (10^{21} years)
^{48}Ca	4.27	0.187	0.064
^{76}Ge	2.04	7.8	1.93
^{82}Se	3.00	9.2	0.096
^{100}Mo	3.03	9.6	0.69
^{116}Cd	2.81	7.5	0.028
^{130}Te	2.53	34.1	0.82
^{136}Xe	2.46	8.9	2.17
^{150}Nd	3.37	5.6	0.11

Table 8.1: List of some isotopes known to undergo double beta decay. Typically isotopes with higher Q-value and longer $2\nu\beta\beta$ half-life are desirable as they allow detectors to achieve lower backgrounds. Additionally, higher isotopic abundance means larger isotope mass can be easily be deployed without requiring expensive enrichment. Information for this table was partially taken from (5).

Majorana neutrinos, as shown in Figure 8.2. However, the process can be mediated by other mechanisms, which might introduce new particles or interactions. Discussion of a possible mechanism through the introduction of right-handed charged current and associated heavy right-handed neutrinos can be found in (167). Regardless of the mechanism, $0\nu\beta\beta$ is a lepton number (L) violating process where $\Delta L = 2$, and is thus forbidden in the standard model.

Under the assumption that $0\nu\beta\beta$ is realized through the exchange of light Majorana neutrinos, the half-life of the decay can be expressed as:

$$(\tau_{1/2}^{0\nu})^{-1} = g_A^4 G_{0\nu} \left|M_{0\nu}\right|^2 \frac{\left|m_{\beta\beta}\right|^2}{m_e^2}, \tag{8.2}$$

where $G_{0\nu}$ is the phase space factor, $M_{0\nu}$ is the matrix element unique for each isotope, m_e is the electron mass, g_A is the axial vector coupling, and $m_{\beta\beta}$ is the effective Majorana

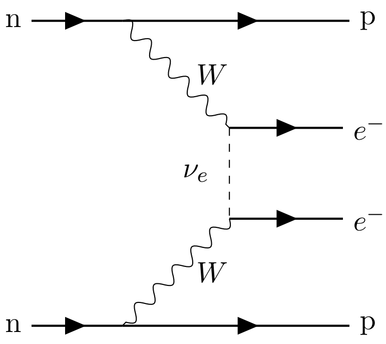

Figure 8.2: The Feynman diagrams for $0\nu\beta\beta$.

mass given by:

$$|m_{\beta\beta}| = \left|\sum_i U_{ei}^2 m_i\right|. \tag{8.3}$$

The decay rate is sensitive to not only the values of the Majorana neutrino masses, but also the elements of the mixing matrix connected to the electron flavor neutrino. A measured event rate in the detector can be converted into a half-life limit, which is then interpreted as a limit on the Majorana mass, $m_{\beta\beta}$.

The phase space factor $G_{0\nu}$ can be calculated with small uncertainties (168, 169); however, the nuclear matrix element is very difficult to calculate and requires a accurate nuclear model of these isotopes with dozens of nucleons. In (170) the results for $M_{0\nu}$ are shown as a function neutron number for various models, and, depending on the isotope, can vary by a factor of three. For example, the values for ^{130}Te range between two to four depending on the model. Additionally, current models cannot provide uncertainties on the calculated values. This presents an experimental challenge for interpreting a measurement of $\tau_{1/2}^{0\nu}$ in terms of $m_{\beta\beta}$. Typically experiments take the minimum and maximum of the calculated $M_{0\nu}$ as an uncertainty when converting to $m_{\beta\beta}$.

The search for neutrinoless double beta decay is experimentally rich and includes over a dozen experiments, many of which approach the measurement with very different experimental techniques. The SNO+ experiment hopes to search for $0\nu\beta\beta$ by loading tellurium directly into the liquid scintillator. This is a similar technique to the KamLAND-Zen experiment (24), which loads Xenon gas into their liquid scintillator. In addition to SNO+ and KamLAND-Zen, experiments such as CUORE (171), EXO (172), GERDA (173), MAJORANA (174), NEXT (175), and SuperNEMO (176) deploy a variety of alternative experimental techniques to search for $0\nu\beta\beta$ decay. In general, $0\nu\beta\beta$ experiments have excellent energy resolution, large amounts of isotopic mass, and very low levels of radioactivity. The experimental signature is a small peak at the Q-value of the $2\nu\beta\beta$ decay, as shown in Figure 8.3. Thus, good energy resolution is critical for these detectors in order to avoid leakage from the $2\nu\beta\beta$ tail.

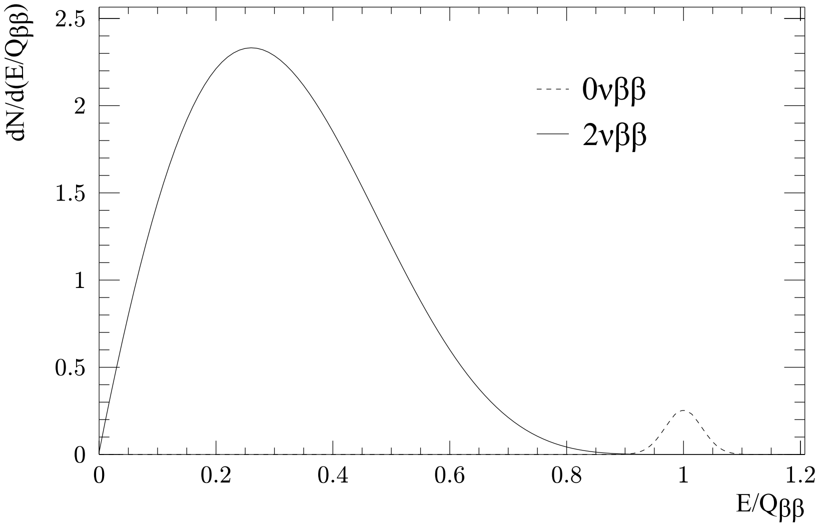

Figure 8.3: The energy spectra for $0\nu\beta\beta$ and $2\nu\beta\beta$ decay. The signature of $0\nu\beta\beta$ is a peak at the end-point of the $2\nu\beta\beta$ decay spectrum, which is shown spread by an arbitrary detector resolution.

A region of interest (ROI), often a narrow energy window around $Q_{\beta\beta}$ and a small fiducial volume (FV) away from external components, is chosen to maximize sensitivity. The number of observed events in the ROI is related to the half-life by:

$$N = \ln(2)\frac{N_A}{W}\left(\frac{\eta\epsilon M t}{\tau^{0\nu}_{1/2}}\right), \tag{8.4}$$

where N_A is Avogadro's number, W is the molar mass of the source, η is the isotopic abundance of the double-beta isotope, ϵ is the detection efficiency, M is the total isotope mass, and t is the measurement time. Given that each of these detectors expects backgrounds, the sensitivity to the half-life depends on N as:

$$\tau^{0\nu}_{1/2} \propto \eta\epsilon\sqrt{\frac{Mt}{B\Delta E}}, \tag{8.5}$$

where B is the background index in units of (keV kg y)$^{-1}$ and ΔE is the energy resolution. Different detector technologies are designed with distinct priorities in mind. For example, detectors such as GERDA, CUORE, and MAJORANA have excellent energy resolution, far exceeding what is possible in a liquid scintillator detector; however, in these detectors it is difficult to deploy large quantities of isotopes. In SNO+ and KamLAND-Zen an enormous amount of isotope can be deployed in these large liquid scintillator detectors, and their size allows them to fiducialize away from the external backgrounds. Currently, the best limit on $m_{\beta\beta}$ comes from the KamLAND-Zen experiment, shown in Figure 8.4. A full review of $0\nu\beta\beta$ and the various experiments is given in (50).

8.2 Neutrinoless Double Beta Decay with SNO+

The primary goal of SNO+ is to perform a measurement of neutrinoless double beta decay using tellurium loaded scintillator. At full capacity, the SNO+ plan is to load 0.5% Te by weight in the liquid scintillator. The primary advantages of using ^{130}Te (as opposed to other possible isotopes) is its high natural abundance of 34.1% and the ability to load tellurium into the liquid scintillator without optical absorption lines (which would limit

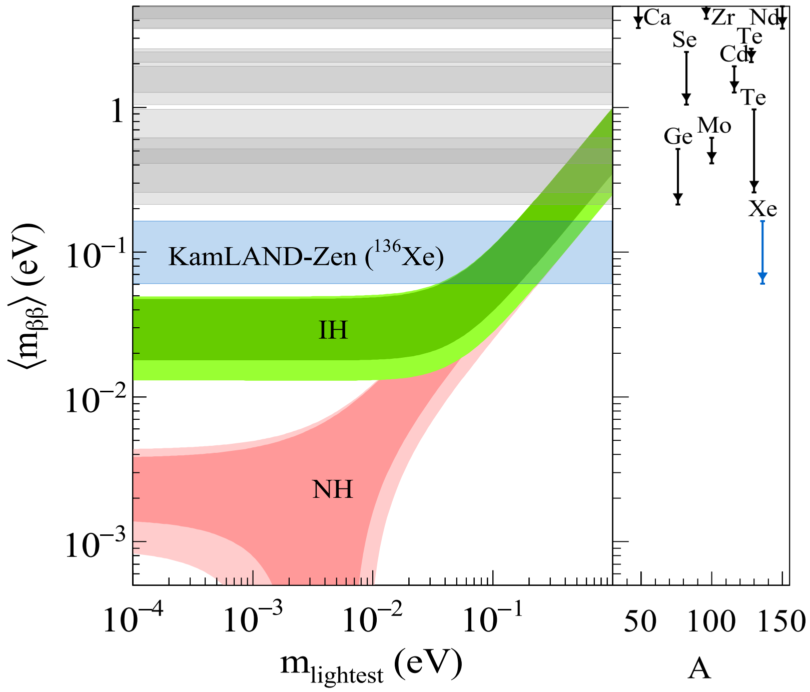

Figure 8.4: The effective Majorana mass as a function of the lightest neutrino mass. The KamLAND-Zen limit is shown in the blue band. The width of the band correspond to different choices of nuclear matrix elements. The green and pink regions show the allow parameter space, which is obtained from scanning over the unknown Majorana phases for the inverted and normal hierarchies. The best limits as of 2016 for each isotope are shown in the inset. This figure is from (24).

loading to higher fractions). Additionally, the Q-value of ^{130}Te's $2\nu\beta\beta$ decay is 2.53 MeV, which is fairly high, but still coincides with several radioactive backgrounds that must be carefully characterized.

The full background budget for SNO+ for one year is shown in Figure 8.5 for an energy

ROI between 2.42 to 2.56 MeV and a fiducial volume of 3.3 m. The total backgrounds expected for five years are shown in Table 8.2. The energy distribution of the backgrounds and the signal are shown for a five year data-set in Figure 8.6 for a hypothetical signal with $m_{\beta\beta} = 100$meV. This illustrates the reason for an asymmetric ROI, typically chosen to be -0.5σ to 1.5σ around the mean of the signal to avoid the falling $2\nu\beta\beta$ background.

Background	Counts/ROI/5 Years
^{8}B ES	23.1
2ν	6.05
Uranium	1.99
Thorium	9.60
External	6.05
(α, n)	0.11
Cosmogenics	0.54
Total	47.4

Table 8.2: Backgrounds for 0.5% Te loaded SNO+ detector for 5 years of data taking with an ROI of 2.42 to 2.56 MeV and a FV of 3.3 m.

One important background considered is γ-rays produced in the external volumes: the AV, the external water, the hold-down and hold-up ropes, and the PMTs. The γs come primarily from the decay of ^{208}Tl, which β decays with the emission of a 2.6 MeV γ. The γ can penetrate or misreconstruct into the internal detector volume. To reject these external backgrounds, a FV, centered in the middle of the AV, with a radius of 3.3 m is chosen. The radial distribution of these backgrounds is compared against the expected signal in Figure 8.7. In addition to selecting a small FV, some of these backgrounds can be rejected by using the shape of the time-residual profile and wiith similar techniques to those described in (177).

Another radioactive background comes from the ^{214}BiPo and ^{212}BiPo events. These coincidences consist of ^{214}Bi/^{212}Bi β decays followed closely in time by ^{214}Po/^{212}Po α decays.

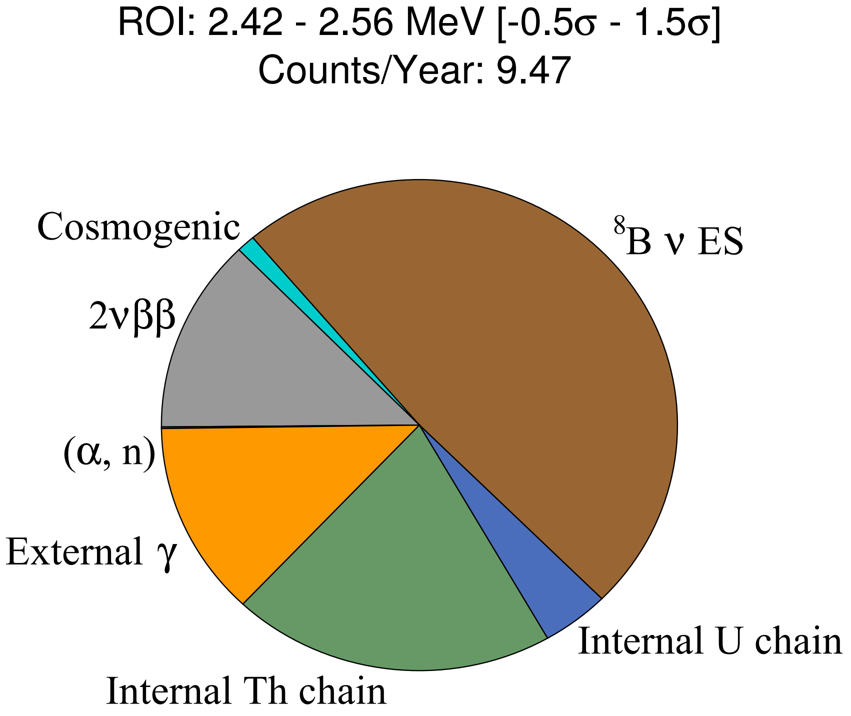

Figure 8.5: The number of background counts per year expected for the SNO+ detector in the energy ROI from 2.42 to 2.56 MeV and a FV of 3.3 m. The backgrounds are dominated by ^{8}B solar neutrino interactions. There are also significant contributions from external γ-rays, internal Thorium chain (dominated by ^{212}BiPo decays), and $2\nu\beta\beta$ leakage into the ROI.

These BiPo events can broadly be classified as either in-window or out-of-window which correspond to whether the α light is contained within the same or a different trigger window (there is a grey area where the Po decays right at the end of the trigger window). The out-of-window events can be effectively rejected using the Δt, Δr, and α/β pulse-shape

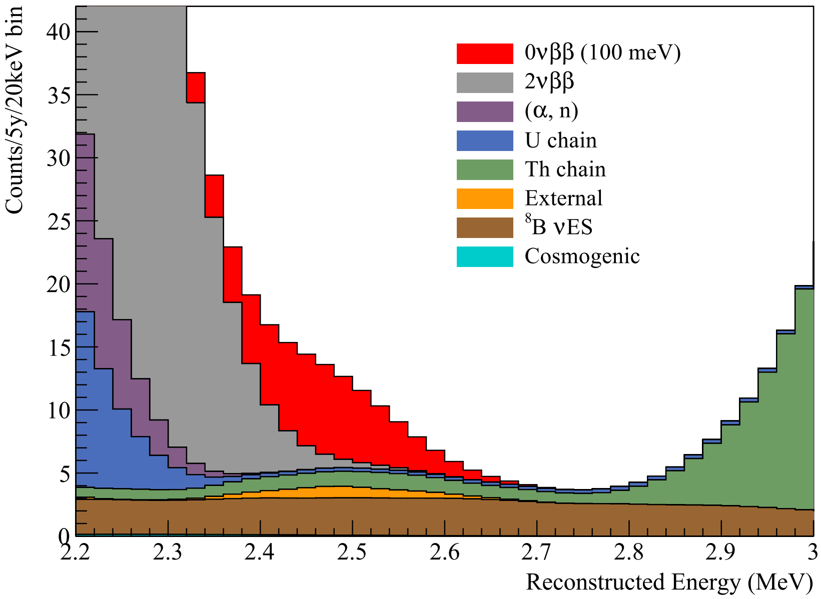

Figure 8.6: The expected reconstructed energy distribution for the backgrounds and a hypothetical $0\nu\beta\beta$ signal where $m_{\beta\beta}$ = 100 meV. The FV used is 3.3 m. An asymmetric ROI of -0.5σ to 1.5σ around the mean of the signal is chosen to avoid the falling $2\nu\beta\beta$ background.

discrimination between the events, similarly to the analysis described for the antineutrino analysis in Section 6.5. The in-window coincidences are occasionally difficult to reject when the Po decays very quickly after the Bi decay (the two prompt peaks in the time-residuals pile-up), which occurs more often for ^{212}BiPo, as ^{212}Po has a short half-life of 299 ns.

Most importantly, the expected dominant background for the SNO+ $0\nu\beta\beta$ search is ^{8}B solar neutrino elastic scattering. The only possible method for rejecting these events is to reconstruct their direction. The scattered electron's direction is strongly correlated with the direction of the incident neutrino, so recoil electrons will produce Cherenkov radiation directed away from the sun. The $0\nu\beta\beta$ signal will be isotropic with respect the solar

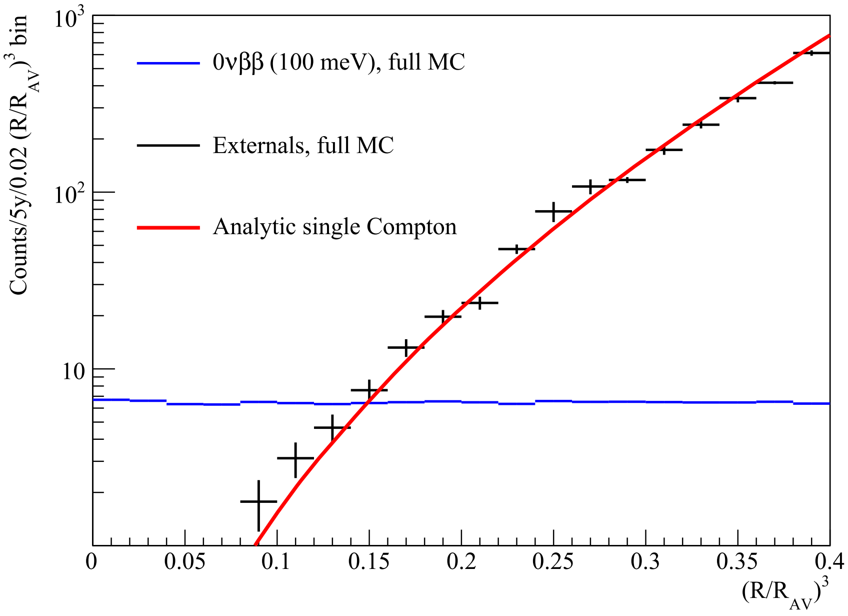

Figure 8.7: The radial distribution of the external backgrounds (which includes the ^{208}Tl from the AV, external water, ropes, and PMTs) compared to the $0\nu\beta\beta$ signal. The fiducial volume of 3.3 m corresponds to a value of $(R/R_{AV})^3 = 0.17$. The counts are normalized to 5-years of data-taking.

direction. However, reconstructing direction in scintillator is extremely difficult due to the isotropic nature of the emitted scintillation light. This consideration was the motivation for designing the dichroicon.

8.3 The Dichroicon: Concept

Despite their great success to date, water Cherenkov and scintillator neutrino detectors use only a small amount of the information available in the photons they detect. A typical large-scale photon-based detector records at most the number of detected photons and their arrival times. But photons may also carry information about physics events in their direction (178), their polarization, and their wavelength.

The focus of this chapter is on the development of a device that is capable of providing information on photon wavelength in a large-scale detector. In a Cherenkov detector—whether in water, ice, or oil—photon wavelength carries information about the propagation time from the source vertex to the photon sensor. Across 50 m of water, for example, a 550 nm photon will arrive nearly 2 ns earlier than a 400 nm photon, easily resolvable by modern photomultiplier tubes (PMTs) (154). Thus, measuring the difference in time between many long-wavelength and short-wavelength photons that lie along a Cherenkov ring provides information about event position, independent from the overall timing and angular information usually used in reconstruction. The resolution of dispersion in such a detector also allows improved timing, as both the approximately 2 ns spread from dispersion and the differential effects of Rayleigh scattering broaden the prompt time window used for reconstruction.

In a scintillation or water-based scintillation detector, photon wavelength can be used to detect Cherenkov light independently from scintillation light. For these detectors, the scintillation light typically lies in a narrow band at short optical wavelengths, while Cherenkov light is naturally broadband. Future large-scale scintillation experiments like `THEIA` (86, 179) plan to detect both scintillation and Cherenkov light as a way of providing a very broad range of physics measurement capabilities with a single detector.

The scintillation light provides a high light yield that is critical for good energy resolution and position reconstruction. The time profile of the scintillation light is also important, because it affects position reconstruction and provides ways of discriminating βs from α particles. Importantly, independent detection of Cherenkov light allows reconstruction of event direction. In addition to classification of neutrinoless double beta decay candidates against the solar neutrino background, the importance of which is discussed in Section 8.1 and in (180, 181, 182), this can aid in identifying solar neutrino events (183, 184) or discrimination of high-energy ν_e events from π_0's, which is important for studying long-baseline neutrino oscillations.

There are several possible techniques for measuring the Cherenkov light in liquid scintillator detectors. The timing of the detected photons is a powerful handle, as the Cherenkov light is produced promptly, whereas it may take as long as a nanosecond for the scintillation light to be emitted (185). Additionally, the angular distribution of the Cherenkov light around the event direction distinguishes it from isotropic scintillation light. Benchtop scale experimental setups, such as in (185, 186, 187, 188, 189), use the timing and directionality to identify the Cherenkov light.

Separating the two components in current generation large-scale scintillation-based neutrino detectors is nevertheless very difficult. The transit time spread (TTS) of PMTs is generally around 1.5 ns or larger, making it difficult to resolve the early Cherenkov light. An illustration of the typical timing spectra of the detected light for a SNO+-like detector using liquid scintillator is shown in Figure 8.8, generated using `RAT`. Even with using both the timing and spatial distributions of the hits, no current generation large-scale scintillator detector has been able to demonstrate the detection of Cherenkov light.

The challenge in discriminating Cherenkov and scintillation light by photon wavelength in large-scale detectors is doing so while maintaining the high detected light yield needed for a low-energy physics program or precision reconstruction and particle ID. Using water-based liquid scintillator (190), for example, increases the ratio of Cherenkov to scintillation light by reducing the total scintillation light. Using a scintillator like linear alkyl benzene (LAB) with only a small amount of fluor can also be done to slow down the scintillation time profile (87) and then timing can be used to identify Cherenkov light; however, this again comes with a consequent reduction in scintillation light yield. Adopting a simple filtering scheme, or using sets of photon sensors of different wavelength sensitivities (191), also reduces total light yield because the detection area taken up by filtered photon sensors can only be used for one photon wavelength band. What is needed is a way to *sort* photons

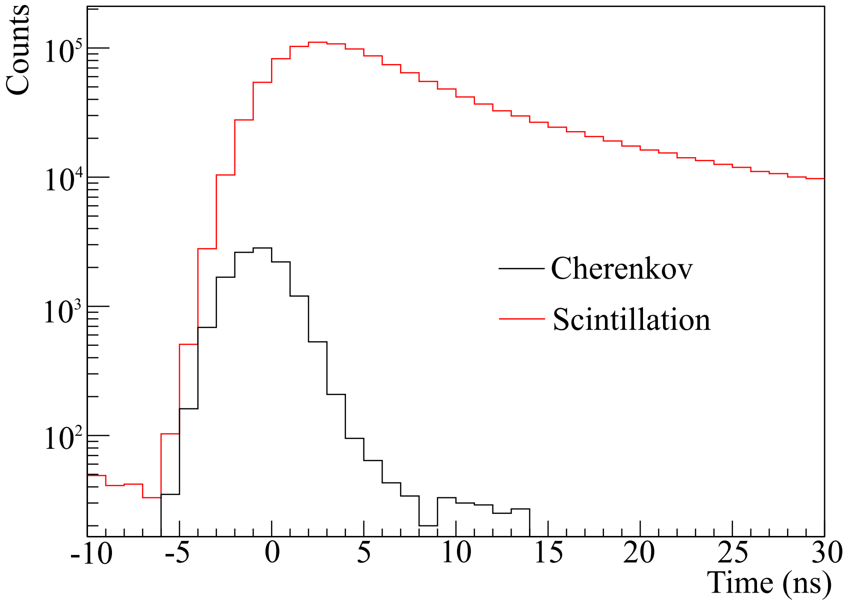

Figure 8.8: The time profile of the detected light for simulated 2.5 MeV electrons at the center of a SNO+-like detector, consisting of about 9000 PMTs with transit time spreads of 1.4 ns, a 6 meter radius acrylic vessel, and about 50% coverage. The Cherenkov light arrives promptly, but is difficult to identify due to the intrinsic resolution of the photodetectors and high light yield of the scintillator.

by wavelength, directing different wavelength bands toward relevant photon sensors, and doing this in a way that loses as little timing or position information as possible.

In a first set of tests (163), I show that sorting by wavelength can be done using dichroic reflectors, and that broadband (falling as $1/\lambda^2$) Cherenkov light can be distinguished from narrow-band scintillation light, in LAB scintillator doped with 2,5-Diphenyloxazole (PPO). To turn this approach into something that could be used in a large-scale detector, the dichroic filters are configured into a Winston-style light concentrator, the "dichroicon." As is well known, Winston cones provide optimal light collection for non-imaging detectors

(90), and have been used in other large-scale neutrino detectors (79, 192). An additional advantage to a large-scale detector using long-wavelength photons to identify Cherenkov light is that the long-wavelength photons travel faster and are scattered and absorbed far less than short-wavelength photons (193), thus preserving more of the directional information of the Cherenkov light.

In this context, photons with wavelengths between around 450 to 900 nm are referred to as ‘long-wavelength’ and photons between 350 to 450 nm are referred to as ‘short-wavelength’. The shortest wavelength photons between 300 to 350 nm are absorbed by scintillator and re-emitted at longer wavelengths. The emission spectra of common fluors such as PPO, as shown in Figure 8.9, peak around 360 to 380 nm and tail off by 450 nm, leaving primarily Cherenkov light emission above this wavelength. Red-sensitive PMTs can be used to detect this long-wavelength light and have quantum efficiencies that can extend to around 800 to 900 nm, also shown in Figure 8.9.

The dichroicon follows the off-axis parabolic design of an ideal Winston cone but is built as a tiled set of dichroic filters and thus does not achieve the idealized shape. The filters are used to direct long-wavelength light towards a central red-sensitive PMT, while transmitting the shorter wavelength light through the ‘barrel’ of the Winston cone to secondary photodetectors. This is possible because of the remarkable property of the dichroic reflectors, which reflect one passband of light (below or above a ‘cut-on’ wavelength) while transmitting its complement, with very little absorption. As shown schematically for two possible designs in Figure 8.10, the ‘barrel’ of the dichroicon is built from shortpass dichroic filters and a longpass dichroic filter is placed at the aperture of the dichroicon. The shortpass filter passes short-wavelength light while reflecting long-wavelength light; the longpass has the complementary response.

The parameter space for optimization of the dichroicon is large, and only a small fraction of it has been explored. The features of the dichroicon that can be varied include: length and geometric field of view, cut-on wavelength for different dichroic filters (they need not

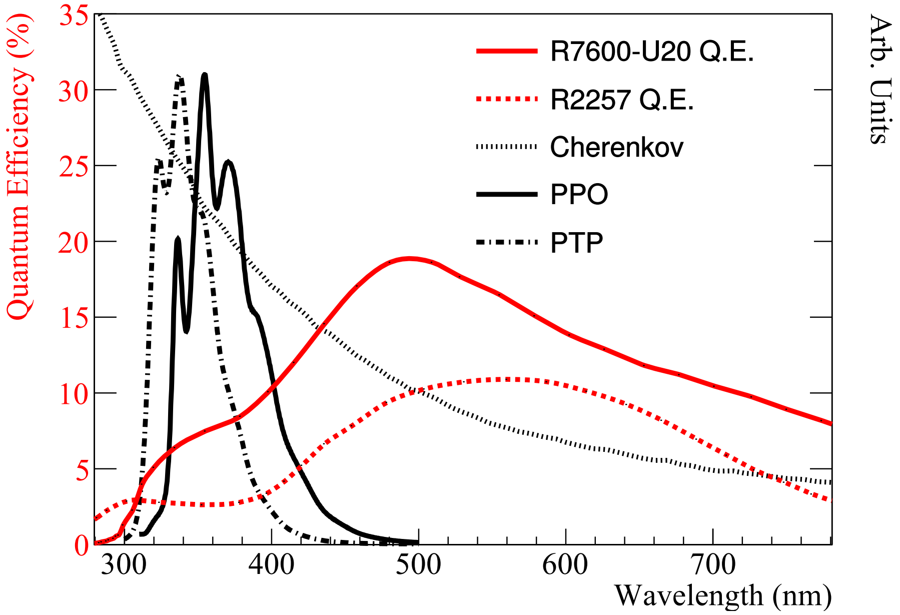

Figure 8.9: The quantum efficiency of the PMTs used in the various measurements compared to the Cherenkov emission spectrum and the emission spectra of the fluors PPO and PTP. The three emission spectra are arbitrarily scaled and show shape only. The quantum efficiency curves for the R7600-U20, R2257, and R1408 are taken from (25, 26, 27) respectively. The PPO and PTP emission spectra were taken from the PhotochemCAD database (28).

all be the same across the cone), photon sensor type and response for both short- and long-wavelength bands, presence or absence of additional filtering at the aperture, shape and reflectivity of the lightguide, and size and configuration of the photon sensors used to detect both sets of light. In fact, a multi-band dichroicon could be designed, which simply nests various dichroicons within each other, if more than two passbands were needed. The design here has been constrained by available PMT sizes and sensitivities, and sizes and shapes of available dichroic filters. For a dichroicon that would be deployed in a detector, the optimization would depend upon the physics goals, the target material and any added

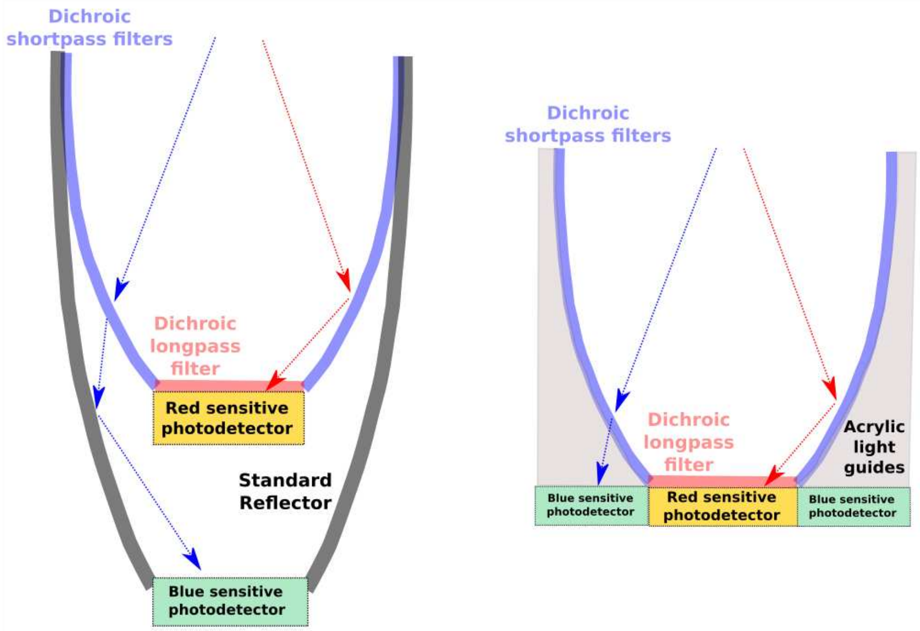
Dichroic shortpass filters
Dichroic longpass filter
Red sensitive photodetector
Standard Reflector
Blue sensitive photodetector
Dichroic shortpass filters
Dichroic longpass filter
Acrylic light guides
Blue sensitive photodetector
Red sensitive photodetector
Blue sensitive photodetector

Figure 8.10: Simple schematics of two potential options for a system to detect light sorted by the dichroicon. The schematic on the left shows an option where a parabolic reflector is built around the dichroicon to detect the short-wavelength light. The schematic on the right an option using acrylic light guides to direct the short-wavelength light to one or more photodetectors. In this case, a pixelated light detector such as a large area picosecond photodetector (LAPPD) might be an ideal sensor for the dichroicon. In both designs the long-wavelength light is detected at the aperture of the dichroicon. Neither of these full designs are constructed, but are included to show potential detection schemes for the short- and long- wavelength light. The blue and red lines show possible photon tracks for short- and long-wavelength light respectively.

fluors, and the fiducial volume of the detector.

In the prototype design, a 3D printed structure is tiled with dichroic filters, as shown in Figure 8.11. Around the barrel of the Winston cone are eighteen shortpass filters from Edmund Optics and eighteen from Knight Optical, the latter of which are custom cut to trapezoidal shapes to fill the surface area. A longpass filter from Knight Optical, 50 mm in diameter, sits at the center of the dichroicon. A second dichroicon is also constructed using several different filters. The specific filters used for the two dichroicons are presented in Tables 8.3 and 8.4 respectively. The dichroicon pictured in Figure 8.11 and presented in Table 8.3 is used for the majority of the measurements, and is referred to simply as ‘the dichroicon’. The measurements with the dichroicon were done with both a Cherenkov source and with a LAB-PPO source. The second dichroicon is used primarily for measurements in LAB with p-Terphenyl (PTP) as the fluor, which has a shorter wavelength emission spectrum than PPO. That device is referred to as dichroicon-2, and the only difference from the dichroicon is that the cut-on wavelength of the rectangular shaped short-pass filters used in the barrel of the dichroicon-2 are slightly reduced and the cut-on wavelength of the central longpass filter used at the aperture is slightly reduced.

Cut-On (nm)	Pass	Shape	Dimensions (mm)	Quantity	Manufacturer
500	Short	Rectangular	25.2 × 35.6	18	Edmund Optics
453	Short	Trapezoidal	35 × 35 × 35 × 25	6	Knight Optical
453	Short	Trapezoidal	25 × 35 × 35 × 14	6	Knight Optical
453	Short	Triangular	14 × 35 × 35	6	Knight Optical
480	Long	Circular	⌀ 50	1	Knight Optical

Table 8.3: The details for the filters used for the dichroicon shown in Figure 8.11. The cut-on wavelength is given for an average incidence angle of 45°.

Cut-On (nm)	Pass	Shape	Dimensions (mm)	Quantity	Manufacturer
450	Short	Rectangular	25.2 × 35.6	18	Edmund Optics
453	Short	Trapezoidal	35 × 35 × 35 × 25	6	Knight Optical
453	Short	Trapezoidal	35 × 35 × 35 × 14	6	Knight Optical
453	Short	Triangular	14 × 35 × 35	6	Knight Optical
462	Long	Circular	⌀ 50	1	Knight Optical

Table 8.4: The details for the filters used for the dichroicon-2. The cut-on wavelength is given for an average incidence angle of 45°.

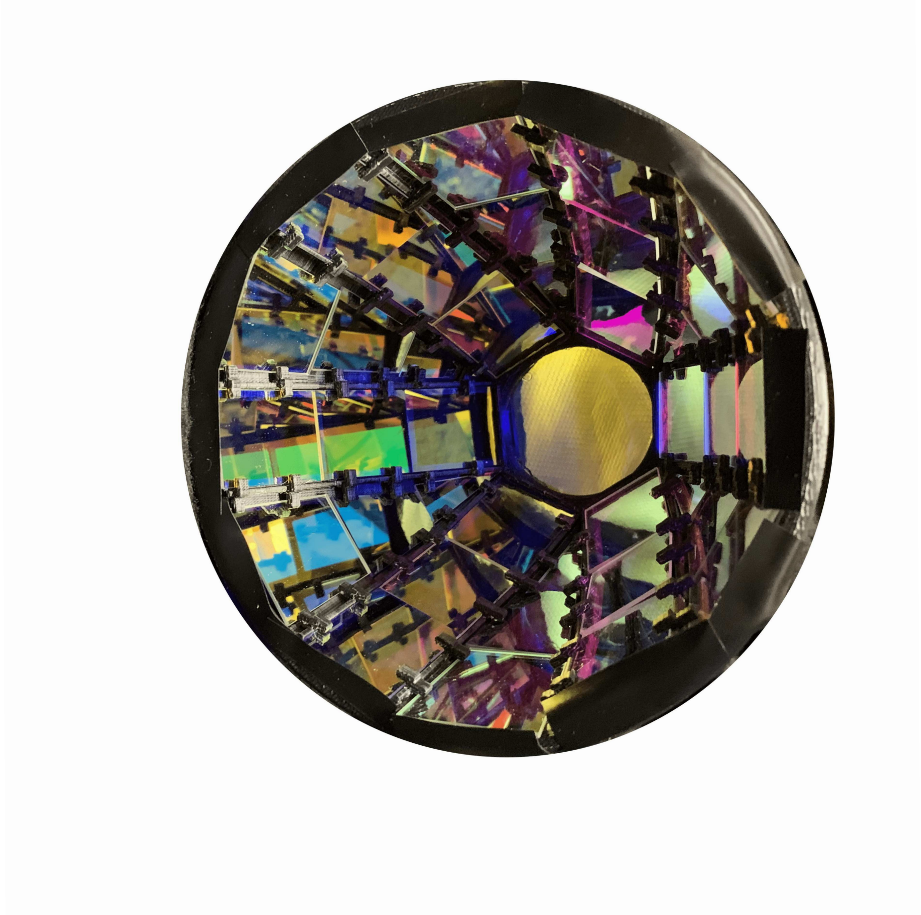

Figure 8.11: A head-on view of the dichroicon. The filters are held in a custom 3D printed plastic holder and can be easily swapped. The shortpass filters tile the barrel of the Winston cone and a central longpass filter is placed at the aperture. A small amount of black electrical tape is used to block a small gap between the filters and the holder at the top of the dichroicon. The outer diameter of the dichroicon is about 150 mm and the inner radius, where the long-pass filter is located, is about 50 mm.

The cut-on wavelength specified in Tables 8.3 and 8.4 refers to wavelength corresponding to the 50% transmission crossing, as the filter either goes from transmitting to reflecting or vice versa. The choice for the cut-on, manufacturer, size, and shape of the filters was motivated primarily by availability, cost, and design limitations. The choice to use cut-on wavelengths between 450 to 500 nm for the dichroicon was motivated largely by the result in (163), where measurements made with LAB+PPO and a single 500 nm dichroic filter showed excellent performance. Shorter pass filters were used for the dichroicon-2.

The measurements focus on detecting the light at both the aperture and barrel of the Winston cone using PMTs. There are, however, many possible options for photon sensing. One interesting option is LAPPDs, which provide excellent time resolution of around 50 ps (194) over their active areas, which consist of 350 square centimeter pixels. By coupling an LAPPD to a dichroicon one could use the central pixels to detect the long-wavelength light and the outer pixels for the short-wavelength light. Also possible instead of PMTs is an array of silicon photomultipliers (SiPMs), similar to the device described in (195), to perform pixelated detection of the light sorted by the dichroicon.

The method for detecting the short-wavelength light is complicated by the fact that the entire barrel of the Winston should be instrumented to provide maximal collection efficiency. Placing PMTs directly behind the barrel of the cone would be an expensive and inefficient way to collect the light, although it would preserve the timing of the photons. There are many ideas for better ways to collect these photons – one could use acrylic light guides to direct the light back towards one or several blue sensitive PMTs. Another idea is to use a second parabolic reflector, built from reflecting material instead of filters, that wraps around the dichroicon, more efficiently directing the short-wavelength light toward the photon sensor. A simpler option consists of a cylinder that is lined with reflective material or paint, which reflects the short-wavelength photons back to a single PMT. This final option was chosen for the experimental setup.

All of these options will slightly degrade the timing of the short-wavelength photons by reflecting the photons one or more times before they are detected. Particularly for scintillation detectors, where the timing of the photons is already spread out by the intrinsic emission spectrum of the scintillator, this will be a small effect, and it is part of the reason for choosing to tile the barrel of the dichroicon with shortpass filters, rather than building the complementary design.

Although I have chosen to include a longpass dichroic filter at the dichroicon aperture, its presence ultimately depends on the physics goals of an experiment and the configuration of the photon sensors. The advantage of a longpass dichroic filter at the aperture rather than a simple longpass absorbing filter, is that short-wavelength light that hits the aperture is reflected rather than absorbed, and thus total short-wavelength light yield is not affected. The majority of this light, in the configuration, would have to be detected by another device—for example, by PMTs on the other side of a large detector. A different optimization of the dichroicon, however, might lead to a different choice for the aperture filter.

Nevertheless, even with a longpass dichroic filter at the aperture of the dichroicon, some short-wavelength light does leak through. For Cherenkov/scintillation separation, even this small amount of leakage is noticeable, because there is so much more scintillation than Cherenkov light generated at the source. Thus a longpass absorbing filter is included *behind* the longpass dichroic filter at the dichroicon aperture in some of the measurements, which leads to improved purity of Cherenkov photons detected there. While this means that some of the short-wavelength light is lost from the system, it is a tiny amount, as it is only the small number that leak through the longpass dichroic filter that are absorbed. A photon sensor with better timing at the aperture would likely make the longpass absorbing filter unnecessary for Cherenkov/scintillation separation, as the small amount of scintillation light that leaks through could be distinguished by timing.

The primary goal in the measurements of this dichroicon is to demonstrate the ability to sort photons for both a Cherenkov and scintillation source. The measurements with a Cherenkov source are designed as a demonstration that the dichroicon works as intended, in an easy to study system, with application for large-scale neutrino detectors such as Hyper-Kamiokande. The measurements with scintillation sources will further demonstrate the photon sorting technique in addition to providing a way to separate Cherenkov and scintillation light.

In Section 8.4.1 I discuss the experimental setup and results of the dichroic filter characterization, which provides critical input into the simulation software. In Section 8.4.4.1 the calibration of the PMTs used in the dichroicon setup is discussed, which is necessary for a quantitative understanding of the dichroicon results. In Section 8.4 the benchtop setup, data analysis, and results for the dichroicon measurements are presented. These results include measurements using a Cherenkov source and two different scintillation sources, a variety of different dichroic filters and absorbing longpass filters, and two different red-sensitive PMTs at the aperture of the dichroicon.

Finally, it should be noted that dichroic filters have appeared recently in several other potential photon detection devices for large-scale neutrino detectors. These filters are starting to be studied in more detail for use as a photon-trap device for Hyper-Kamiokande (196) and for the ARAPUCA and X-ARAPUCA light trap designs for ProtoDUNE and DUNE (197).

8.4 The Dichroicon: Measurements

The characterization of the dichroic filters used in the dichroicon is presented in Section 8.4.1. These measurements are critical for the understanding and simulation of the full dichroicon response. Preliminary measurements using single bandpass filters or single dichroic filters were performed, presented in Sections 8.4.2 and 8.4.3, to understand if Cherenkov and scintillation separation was possible using a simple setup. The success of these measurements led to the construction of the first full dichroicon and the measurements performed with this prototype dichroicon are presented in Section 8.4.4.

8.4.1 Filter Characterization

A schematic of the dichroic filter characterization experimental setup using PMTs is shown in Figure 8.12. In this setup, a collimated LED is directed toward a 50/50 beamsplitter. One output of that beamsplitter goes toward an R7600-U200 PMT, referred to as the normalization PMT, which provides a measure of the LED intensity, which changes slightly across datasets. The other output goes to a dichroic filter, which both transmits and reflects the incoming photons. The transmitted and reflected light is detected by R7600-U200 PMTs, referred to as the transmission and reflection PMTs respectively. The three PMTs are operated at -800 V.

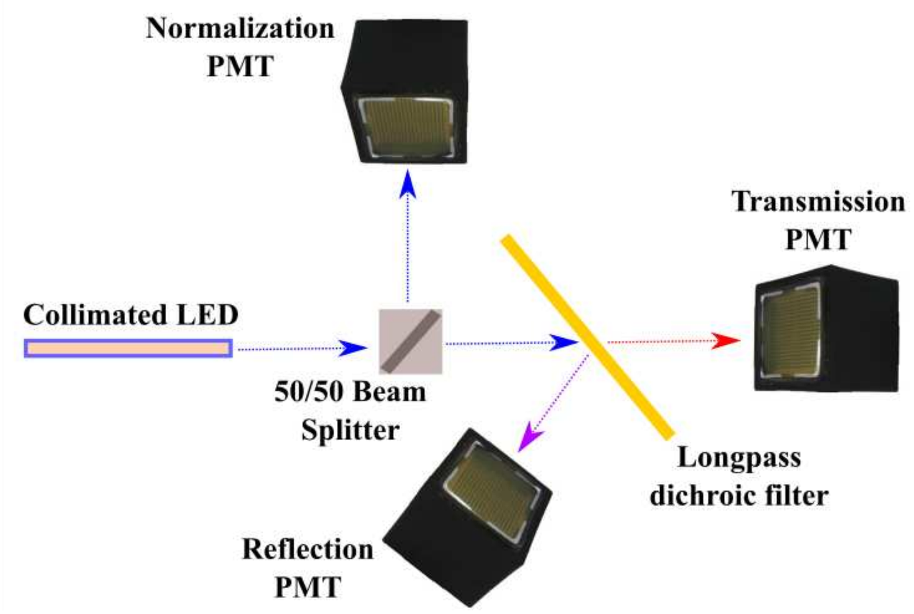

Figure 8.12: The dichroic filter characterization setup. The normalization, transmission, and reflection PMTs are R7600-U200 PMTs. The dichroic filter is held on a rotating stage.

The dichroic filter is located on a rotating stage and the angle can be chosen in 1° increments. The response at incidence angles from 0 to 60° is measured, where 0° indicates light impinging normal to the surface. Given the geometry of the setup, especially the opening angle of the collimated LED beam, very large incidence angles $> 60°$ were not possible to measure. Additionally, measurements of reflection were not made at incident angles less than 15° because of shadowing by parts of the setup of the reflected light.

LEDs at wavelength of 385, 405, 450, 505, 555, 590, and 630 nm from Thorlabs were used to probe the filter response across wavelength. The data sheets are available online. The spectral FWHM of the LEDs range from 12 to 30 nm and no filters were used to narrow the wavelength range of the beam.

The LEDs are pulsed with 40 ns wide 3 V square pulses at 1 kHz. This output is split, one side is used to trigger the oscilloscope acquisition and the other goes to the LED. At these settings the LED output provides a relatively high intensity source, resulting in the collection of around 100 photoelectrons (PEs) per triggered event at the normalization PMT. In general, due to the nature of the dichroic filter, either the reflection or transmission PMT views a similar number of PE, while the other detects very few photons over the entire dataset.

The data acquisition (DAQ) system is a Lecroy WaveRunner 606Zi 600 MHz oscilloscope which digitizes the analog signals from the PMTs. The data is sampled every 100 ps in 200 ns long waveforms. The oscilloscope has an 8-bit ADC with a variable dynamic range, which allows for roughly 100 μV resolution. The `LeCrunch` software (151) is used to read out the data, formatted in custom `hdf5` files, over ethernet connection.

`C++`-based analysis code runs over the `hdf5` files to calculate the amount of light collected by the normalization, transmission, and reflection PMTs. Each PMT signal is integrated to produce a charge. The gain of the PMTs is set such that if a single photon is detected the peak of the charge distribution sits around 1 pC. For each triggered event, the charge is converted to number of photons detected, which is summed over the total data set.

For each LED, two calibration datasets are taken with no dichroic filter – one with the LED directed at the transmission PMT and the other with it directed toward the reflection

PMT (still including the beamsplitter and normalization PMT). These datasets are used to measure both PMT responses under the condition where no dichroic filter is blocking the LED output and are used to normalize to an expected intensity for 100% transmission or reflection.

The calculated transmission T through the dichroic filter is given below in Equation 8.6:

$$T = \frac{T_F}{T_{NF}} \times \frac{N_{NF}}{N_F} \tag{8.6}$$

The first term, T_F, is the total amount of light detected at the transmission PMT and is divided by the total amount of light detected by the transmission PMT when no dichroic filter was present, T_{NF}. This gives the fractional transmittance of the filter, under the assumption that the intensity of the LED did not change. To correct for realistic variations in the LED intensity, a second term is included. This term is the normalization PMT measurement of the LED intensity during the data taking with no filter, N_{NF}, divided by the normalization PMT measurement during data taking with the filter, N_F. This provides the relative change in intensity of the LED between the two measurements. The same equation is used to calculate the reflectance, R, where the data for the reflection PMT is used instead:

$$R = \frac{R_F}{R_{NF}} \times \frac{N_{NF}}{N_F}. \tag{8.7}$$

The results for the transmissivity and reflectivity of the 500 nm short-pass filters used in the barrel of the dichroicon are shown in Figure 8.13. and 8.14.

In addition to measuring the response using PMTs, the wavelength and incident angle dependence of the transmission is measured using a spectrometer. The experimental setup to measure the transmission as a function of wavelength and incidence angle with an Ocean Optics USB-UV-VIS Spectrometer is a simplified version of the setup shown in Figure 8.12. The PMTs and beam-splitter are removed and the transmitted light is detected with the spectrometer. While this simple setup does not provide reflectivity values, the transmitted data captured is far more detailed in terms of the behavior as a function of wavelength. To

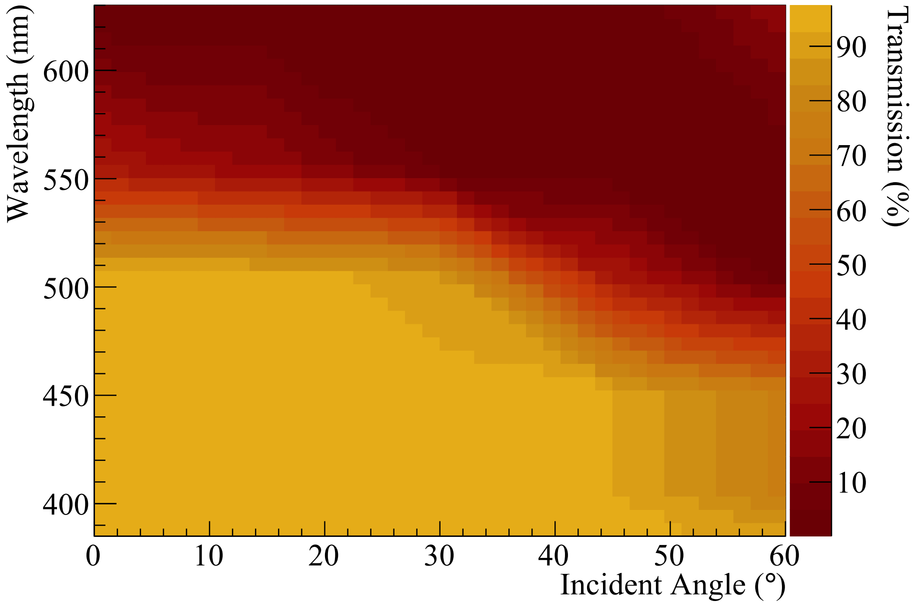

Figure 8.13: The transmission data for the 500 nm shortpass filter for incidence angles between 0 to 60°. The percent transmission was calculated using Equation 8.6.

span relevant wavelengths between 350 to 750 nm, three different light sources are used: a 365 nm LED, a 405 nm LED, and a white LED which spans 420 to 750 nm.

Data is taken with no filter to understand the spectrum and intensity of each of the LEDs. The dichroic filter is added between the collimated LED and the spectrometer at varying incidence angles. This data is normalized to the no filter data to calculate the absolute transmission. The resulting transmission is shown in Figure 8.15 for the 480 nm longpass filter that is used at the aperture of the dichroicon.

This data is consistent with the PMT data that was taken for this filter. Perhaps most interestingly, the small amount of short-wavelength leakage through this filter at high incidence angles becomes clear and will be important in understanding the results for the

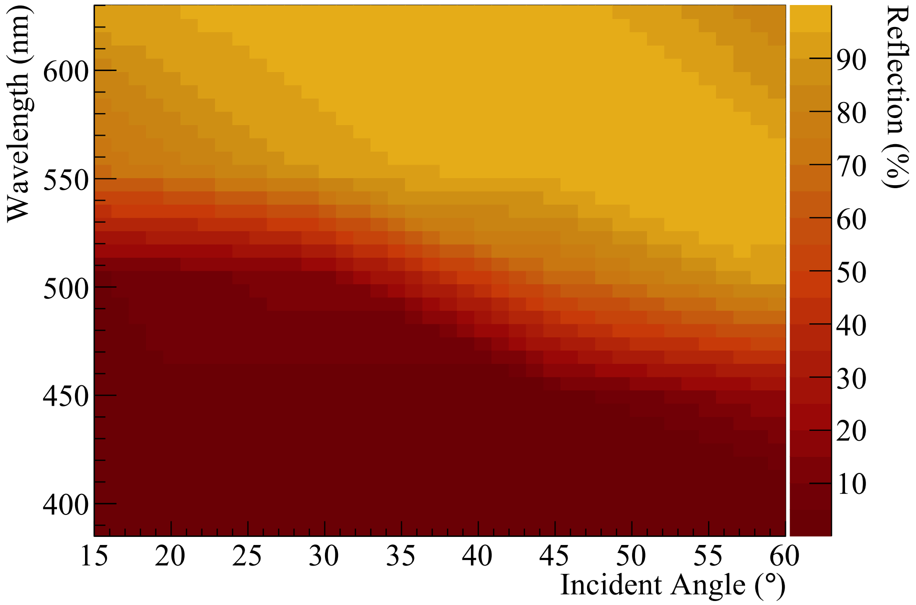

Figure 8.14: The reflection data for the 500 nm shortpass filter for incidence angles between 15 to 60°. Note that the reflected data only extends to incident angles of 15° due to shadowing effects in the setup. The percent reflection was calculated using Equation 8.7.

dichroicon. This measurement is performed for all of the different types of filters deployed in the dichroicon, specified in Table 8.3.

The behavior of the filters will ultimately need to be mapped for the specific fluid that the dichroicon is submerged in. The buffer fluid around the PMTs is commonly water, and measurements with the filter placed in a water bath are made to understand the expected change in performance. The setup and technique is identical to those in air, and the results show that the dichroic filters do perform differently in water. Figure 8.16 shows the results for the 480 nm longpass filter at two incidence angles compared to the data taken in air. As expected, there is a small change to the behavior of the filter, particularly where it

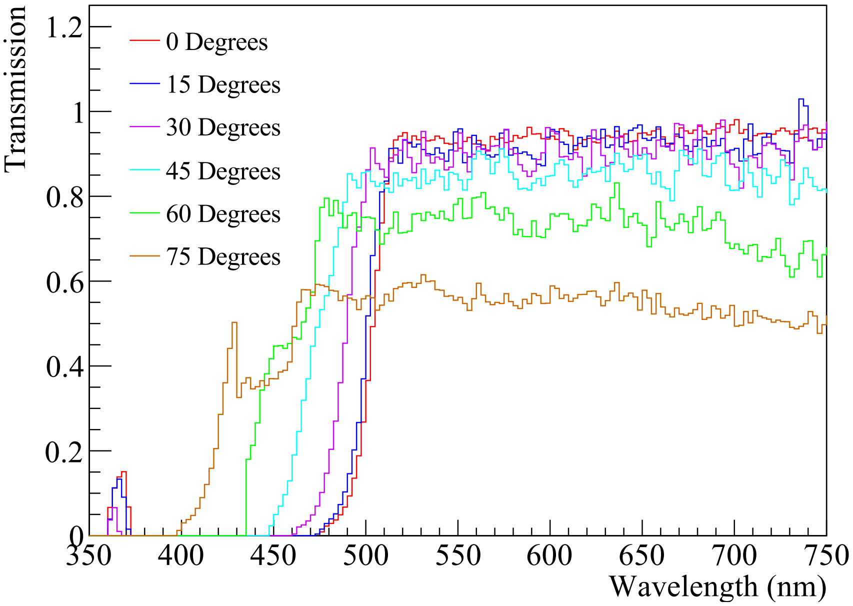

Figure 8.15: The transmission through the 480 nm longpass filter, used at the aperture of the dichroicon, as a function of wavelength, for multiple incidence angles.

transitions from reflecting to transmitting at larger incidence angles. It is anticipated that larger shifts would occur if the dichroicon were submersed in scintillator or oil.

The results described in this section are used as input into the simulation software used to model the dichroicon. More detailed studies of these dichroic filters, with data taken at more incidence angles will improve the model in the future.

8.4.2 Measurements with Single Bandpass Filters

Prior to full measurements of the dichroicon, a simple setup was used to understand the achievable levels of Cherenkov and scintillation separation using wavelength. Before deploying dichroic filters, bandpass filters were used to select specific portions of the LAB+PPO

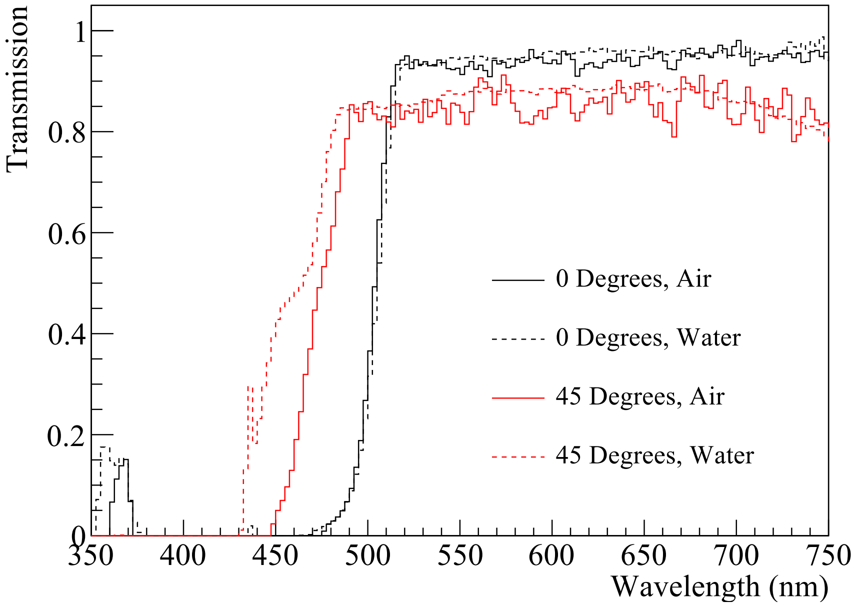

Figure 8.16: The transmission through the 480 nm longpass filter, used at the aperture of the dichroicon, as a function of wavelength, for 0° and 45° incidence angles, in air and in water.

emission spectrum. The experimental setup with bandpass filters is shown in Figure 8.17. The scintillation source and trigger PMT used is described in 7.2.1. The transmitted light through the bandpass filter is detected by a second R7600-U200 PMT, which has a broad efficiency that spans the entire wavlength range of interest. The scintillation light is collimated before the bandpass filter using a mask with a small aperture 1 cm in diameter. Both PMTs are operated at -800 V where their gains are about 5×10^6.

The bandpass filters selected for use are shown in Table 8.5, which were selected to span the LAB+PPO emission spectrum shown in Figure 8.9. The bandpass filters provide excellent blocking outside of the passband – less than 0.01% of light outside the pass is transmitted. This allows for a high fidelity selection of certain wavelengths to understand

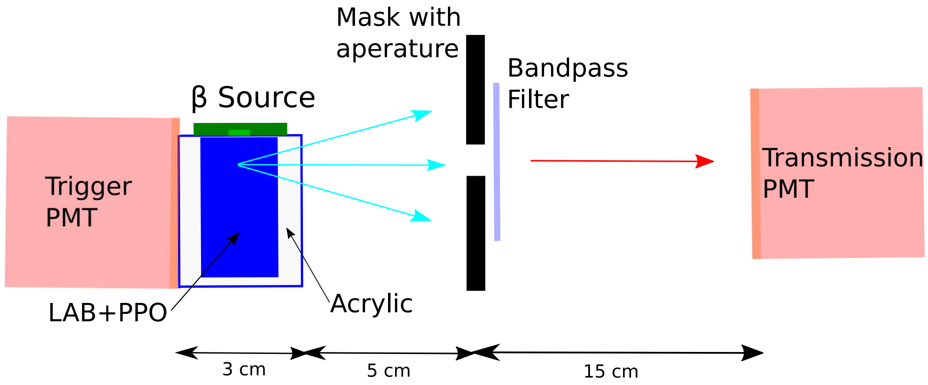
Mask with
aperature
Bandpass
Filter
β Source
Trigger
PMT
Transmission
PMT
LAB+PPO
Acrylic
3 cm
5 cm
15 cm

Figure 8.17: A schematic of the experimental setup, showing the β source deployed above the LAB+PPO target and the locations of the mask, PMTs, and bandpass filter. The colored lines indicate example optical photon paths before and after the filter. The aperture is 1 cm in diameter.

the scintillation and Cherenkov emission in specific bands. This is used eventually to help understand the best dichroic filters to select for Cherenkov and scintillation separation.

Center (nm)	FWHM (nm)	Peak Transmission (%)
355	10	95
387	11	95
405	10	96
430	10	46
450	10	98
470	10	53
494	20	95
510	10	60
530	10	54

Table 8.5: The central wavelength, FWHM, and transmission at the central wavelength for each of the bandpass filters. The tolerances on the central wavelengths of the filters are less than or equal to 3 nm.

The standard DAQ system described in 7.1.1 is used to digitize the waveforms from the trigger and transmission PMT. Given the collimation by the aperture and the narrow bandpass filter, the transmission PMT detects primarily single photoelectrons, and the standard analysis to calculate Δt is applied, described in Section 7.2.3.

Data taken for the nine bandpass filters is shown in Figures 8.18 and 8.19. Less data was necessary for the shorter wavelength bandpass filters to achieve roughly the same level of statistics as the longer wavelength filters. In general, sufficient data was taken for each filter until the statistical uncertainties were smaller than 3% in the peak of the emission spectrum. The 3% level ensures that the rise times of the time profiles can be clearly distinguished. Because very little scintillation light is detected when using the bandpass filters above 500 nm, the statistics are lower for those data sets. Data was also taken with the aperture masked off with black felt. This data set showed a coincidence rate consistent with the dark rate of the PMTs. Additionally, a data set was taken with no bandpass filter to extract the time profile integrated across the entire PPO emission spectrum.

The Δt spectra are fit using Equation 8.8 using the `RooFit` package (159).

$$F = C \times f_{PMT}(t - t') + (1 - C) \times \sum_{i=1}^{2} \frac{A_i \times (e^{-t/\tau_i} - e^{-t/\tau_R})}{(\tau_i - \tau_R)} * f_{PMT}(t - t'). \tag{8.8}$$

The scintillation light is fit to the sum of two decay exponentials with time constants τ_1 and τ_2, weights A_1 and A_2, and an exponential rise time, τ_R. The weights A_1 and A_2 are constrained to sum to one. The PMT transit time spread (TTS) is determined offline using a Cherenkov light source to be 350 ps, which is in good agreement with the Hamamatsu datasheet. Additionally, f_{PMT} accounts for the TTS of both the trigger and transmission PMTs. That distribution is convolved with the scintillation profile in the fit. An offset t' allows for cable delays and other arbitrary time offsets. The second component of the model accounts for the Cherenkov light, which is modeled simply as f_{PMT} because the Cherenkov light is emitted promptly in comparison to the TTS of the PMTs. The scintillation and Cherenkov light are weighted appropriately by C.

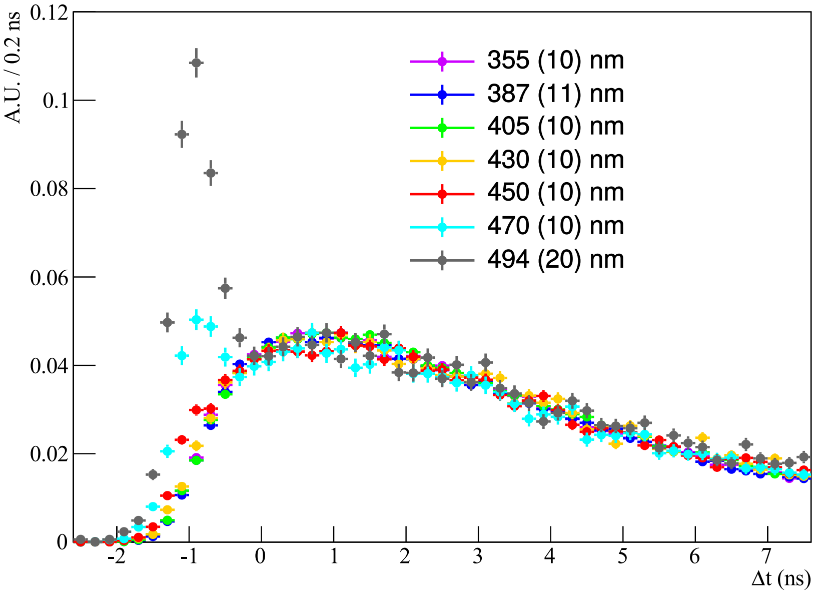

Figure 8.18: The time emission profile, zoomed into the rise time of the LAB+PPO for different wavelength regimes. The light is selected using the bandpass filters listed in Table 8.5. The central wavelength and width of the bandpass filters are specified in the legends. The histograms are normalized to the peak of the scintillation light and shown in arbitrary units (A.U.). The Cherenkov light can be clearly identified at early times in the data for the filters longer than 450 nm. These histograms, as well as additional ones, are shown separately in Figure 8.19.

As discussed in Section 7.2.3, the full scintillation spectrum is typically fit with three or four exponential decay time constants. The primary goal here, however, is to identify the Cherenkov light, and thus only the first 30 ns of the spectrum are fit. Neither the length of the waveform nor the length of data taking provides for accurate measurement of the full scintillation spectrum, which has already been measured by several experiments. Figure 8.20 shows two examples of the fit with the Cherenkov and scintillation components

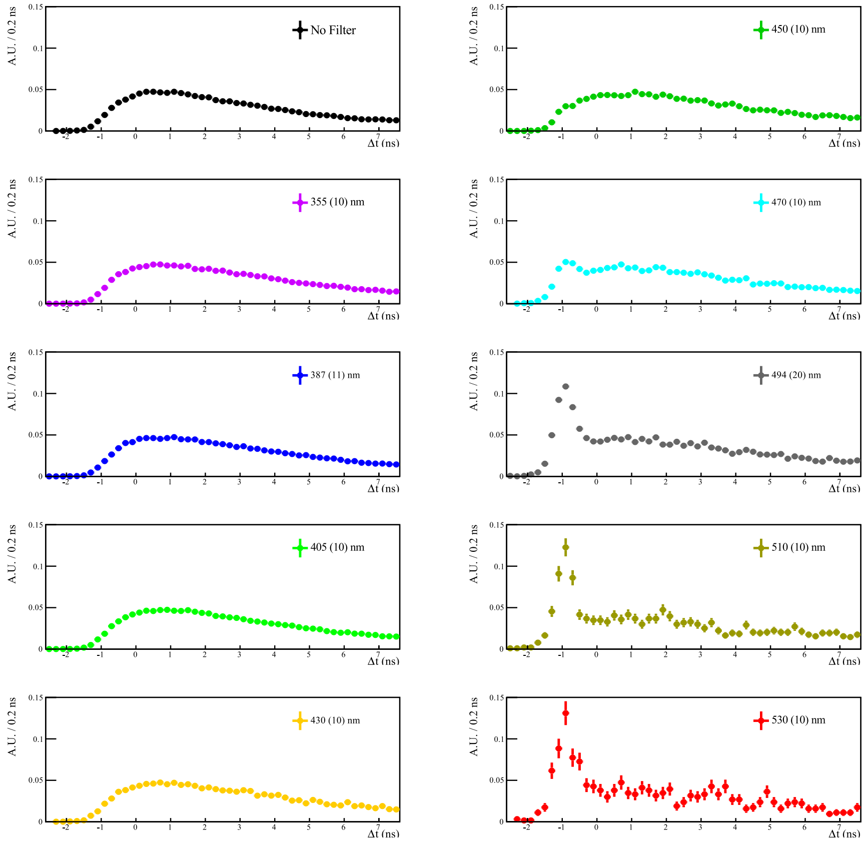

Figure 8.19: The time emission profile for each of the bandpass filters specified in Table 8.5. Each of the data sets are normalized to the peak of the scintillation light. The Cherenkov light becomes clearly separated above 450 nm.

separated.

One can select a relatively pure sample of Cherenkov light by integrating the time profile in a prompt window. Equation 8.9 defines the purity, P, of the Cherenkov light in a prompt window, where F_C is the Cherenkov component in the fit. The upper and lower bounds of the integral are taken to span the Cherenkov component (Figure 8.20) and are arbitrarily

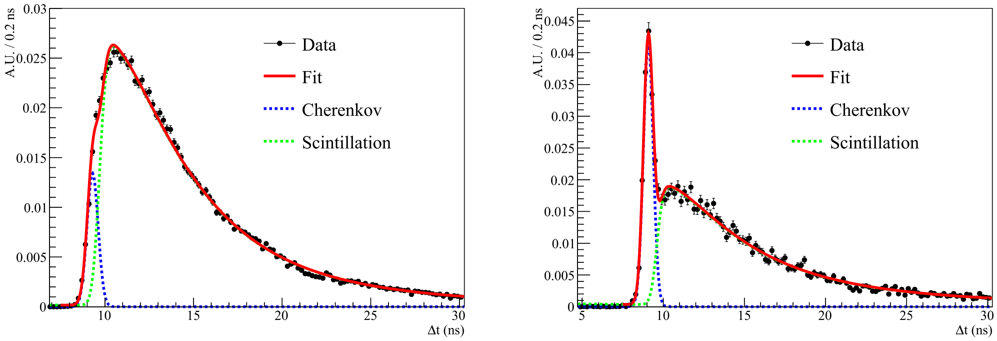

Figure 8.20: The full fit for the 355 nm bandpass filter data (left) and 494 nm bandpass filter data (right) is shown in red. The Cherenkov and scintillation components of the fit are shown explicitly. Note the arbitrary Δt offset from 0 ns, mostly due to cable delays, was not removed for this plot, as it was for Figures 8.18 and 8.19.

offset from zero based on cable delays and other time offsets.

$$P = \int_{8.0}^{9.5} \frac{F_C}{F} dt. \tag{8.9}$$

Table 8.6 shows the full fit results for the ten data sets. The time constants and their relative fractions (τ_1, τ_2, and A_1) are consistent across a wide range of wavelengths. This suggests that the scintillation light in each wavelength regime is being produced via the same mechanism. The purity, P, of the Cherenkov light selection is high even for the short wavelength bandpass filters due to the excellent timing characteristics of the PMTs used. The fits to the data for filters above 450 nm show the Cherenkov light can be selected with a purity of over 90%. This is apparent in Figure 8.20 (right) where the fit to the 494 nm bandpass filter shows that the prompt window is dominated by the Cherenkov component.

The amount of scintillation light detected for each bandpass filter can be compared to the expectation from the PPO emission spectrum. In order to do this, the PPO emission spectrum output from the set-up must be measured. The emission spectrum of the

Filter	$\boldsymbol{A_1}$ **(%)**	$\boldsymbol{\tau_R}$ **(ns)**	$\boldsymbol{\tau_1}$ **(ns)**	$\boldsymbol{\tau_2}$ **(ns)**	$\boldsymbol{1-C}$ **(%)**	**P (%)**
None	68 ± 3	1.6 ± 0.1	3.6 ± 0.1	11.2 ± 1.2	94 ± 2	73 ± 4
355	71 ± 4	1.7 ± 0.1	3.6 ± 0.1	10.5 ± 1.8	94 ± 2	70 ± 4
387	73 ± 4	1.6 ± 0.1	3.7 ± 0.1	10.8 ± 1.8	95 ± 2	69 ± 4
405	68 ± 3	1.7 ± 0.1	3.6 ± 0.1	10.1 ± 2.0	96 ± 2	68 ± 4
430	70 ± 3	1.7 ± 0.1	3.6 ± 0.1	11.4 ± 1.8	94 ± 2	78 ± 4
450	68 ± 3	1.7 ± 0.1	4.0 ± 0.2	11.6 ± 2.2	93 ± 2	76 ± 4
470	68 ± 4	1.7 ± 0.2	4.0 ± 0.2	11.9 ± 2.4	89 ± 2	90 ± 4
494	68 ± 3	1.7 ± 0.1	4.0 ± 0.2	12.0 ± 2.0	82 ± 2	95 ± 4
510	68 ± 4	1.7 ± 0.3	4.0 ± 0.3	12.0 ± 3.2	80 ± 4	$98 \pm {}^{2}_{5}$
530	68 ± 6	1.8 ± 0.4	4.0 ± 0.4	11.9 ± 4.2	77 ± 7	$98 \pm {}^{2}_{7}$

Table 8.6: The fit results for the data with each of the bandpass filters. The fit parameters are defined in Equation 8.8. P is the purity of the Cherenkov light selected in a prompt window and is defined in Equation 8.9.

LAB+PPO is measured using a UV-Vis Ocean Optics Spectrometer by directly exciting the scintillator with a 335 nm LED. The measurement is made by placing the spectrometer at the same relative location as the aperture to probe the spectrum of light that the transmission PMT will view. The scintillator for this measurement is held in the same acrylic block. The PPO absorbs strongly on its own emission spectrum, which causes the lack of the short wavelength peak when comparing to Figure 8.9. The self-absorption of PPO is described in more detail in (198).

In the bandpass data the total amount of scintillation light per triggered event is scaled based on the expected transmission of the filter and efficiency of the PMT. The result is shown in Figure 8.21, compared directly to the measured emission spectrum. The bandpass filter data is consistent with the measured emission spectrum, including the long wavelength emission beyond 450 nm.

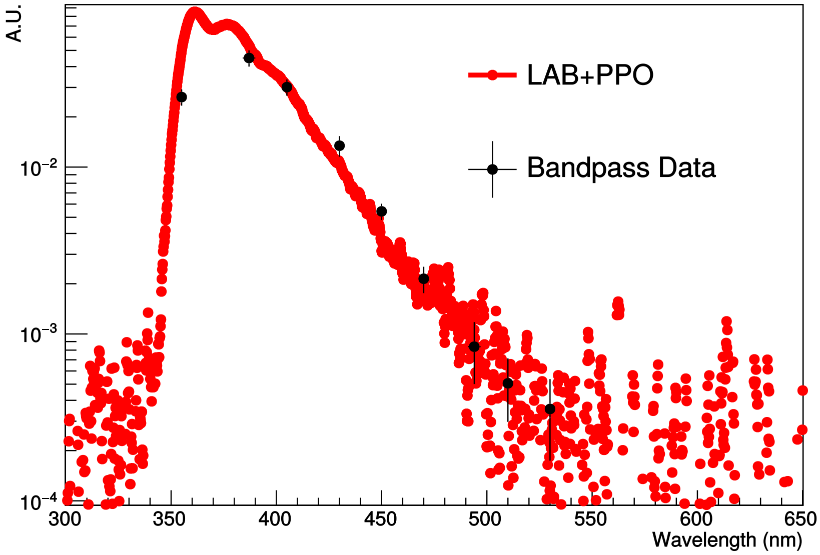

Figure 8.21: The measured emission spectrum compared to the amount of scintillation light detected for each bandpass filter. The amount of scintillation light is scaled based on the total number of triggered events, the efficiency of the PMT, and the transmission of the filter. The shape extracted from the bandpass filter measurement agrees well with the emission spectrum. The noise above 550 nm and below 340 nm is from the dark noise of the spectrometer.

8.4.3 Measurements with Single Dichroic Filters

Similar to the bandpass setup, measurements with a single dichroic filter are useful for understanding the measurements performed with the more complicated dichroicon setup. The experimental setup using a single dichroic filter is similar to Figure 8.17. Instead of a bandpass filter, a dichroic filter placed at an incident angle of 45° is placed behind the aperture. Additionally, a third PMT, called the reflection PMT, is placed to view the reflected light from the dichroic filter. The reflection PMT is the same type as the other two PMTs and is run at the same high voltage. This setup is shown schematically in Figure

8.22. The dichroic filters are both designed for optimal performance at 45°, so a rotating stage with 1° accuracy was used to hold the filter.

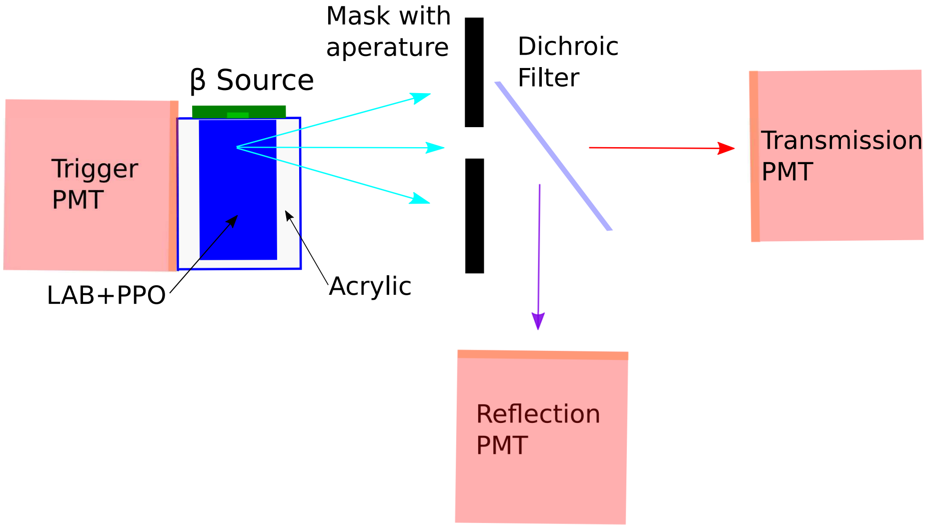

Figure 8.22: A schematic of the experimental setup with the dichroic filter. The setup is the similar to Figure 8.17, with the replacement of the bandpass filter with a dichroic filter and the addition of the reflection PMT.

In this setup, measurements are made with both a 506 nm longpass dichroic filter and 500 nm shortpass dichroic filter. Shown in Figures 8.23 and 8.24 are the fits for the longpass and shortpass filter respectively using LAB+PPO. The data for both the transmission PMT and the reflection PMT are shown with the corresponding fit. For the longpass filter, only the long wavelength light is transmitted, and a clear Cherenkov peak can be identified at the transmission PMT, similar to the bandpass data presented. The reflected light is primarily the LAB+PPO scintillation light, and the data for the reflection PMT data shows the typical LAB+PPO scintillation time-profile. For the 500 nm shortpass filter the reflected light shows more modest Cherenkov separation at prompt times. The shortpass filter's performance was not quite as good (in terms of the separation) as the longpass filter, which is primarily due to the non-negligible reflection (about 8% for wavelengths between 350 - 480 nm) at the LAB+PPO emission wavelengths.

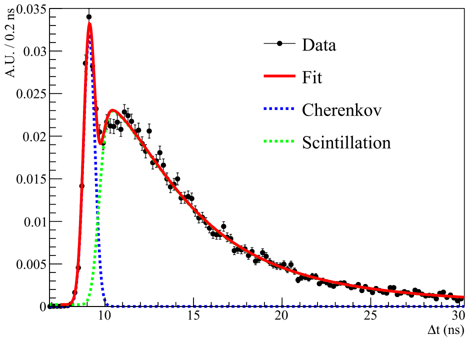

Figure 8.23: The fit results for the longpass 506 nm dichroic filter using LAB+PPO. The transmitted light (left) shows clear Cherenkov and scintillation separation, while the reflected light (right) shows the LAB+PPO emission time-profile.

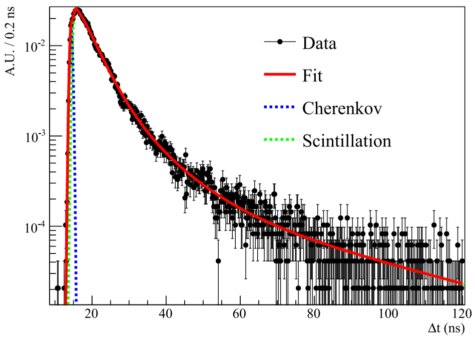

Figure 8.24: The fit results for the shortpass 500 nm dichroic filter using LAB+PPO. The transmitted light (left) shows the LAB+PPO emission time-profile, while the reflected light (right) shows modest Cherenkov and scintillation separation.

The fraction of lost photons can be estimated by comparing the sum of the total number of detected photons in the reflection and transmission PMTs to the total number of detected photons in the measurement with no filter. This was found to be 97.1 ± 1.5% and 98.7 ± 1.5% for the longpass and shortpass filters respectively. The uncertainties are primarily systematic and were estimated by swapping the transmission and reflection PMTs and taking the difference as a two-sided uncertainty.

8.4.3.1 Results with a Large-area PMT

The PMTs used in the measurement presented thus far have extremely good timing properties, and are probably unrealistically expensive and small for a large-scale detector. The transit-time spread of modern large area tubes has been improving, however. Recently Hamamatsu has developed a prototype version of the R5912 8-inch PMT with a TTS of around 700 ps (154), called the R5912-MOD, detailed in Section 7.1.3. The PMT was deployed in an identical setup to Figure 8.22, with the transmission PMT replaced with the R5912-MOD PMT. For this measurement, a 506 nm dichroic longpass filter is used to identify the long-wavelength Cherenkov light. The results of this measurement are shown in Figure 8.25 with the corresponding fit. The only change to the fit is to allow for a wider PMT TTS, which fit to 820 ± 150 ps, consistent with the results in (154). The fit parameters are consistent with those presented in Table 8.6. With an adjusted prompt window, a purity, P, of 52 ± 4 % is found. The lower purity is expected due to the broadening of the PMT TTS.

8.4.4 Dichroicon Characterization

8.4.4.1 PMT Calibration

In the dichroicon measurements presented in Section 8.4, light is detected using three different Hamamatsu PMTs: an R2257, an R7600-U20, and an R1408. To compare various dichroicon measurements, the relative detection efficiencies of the PMTs is needed. The

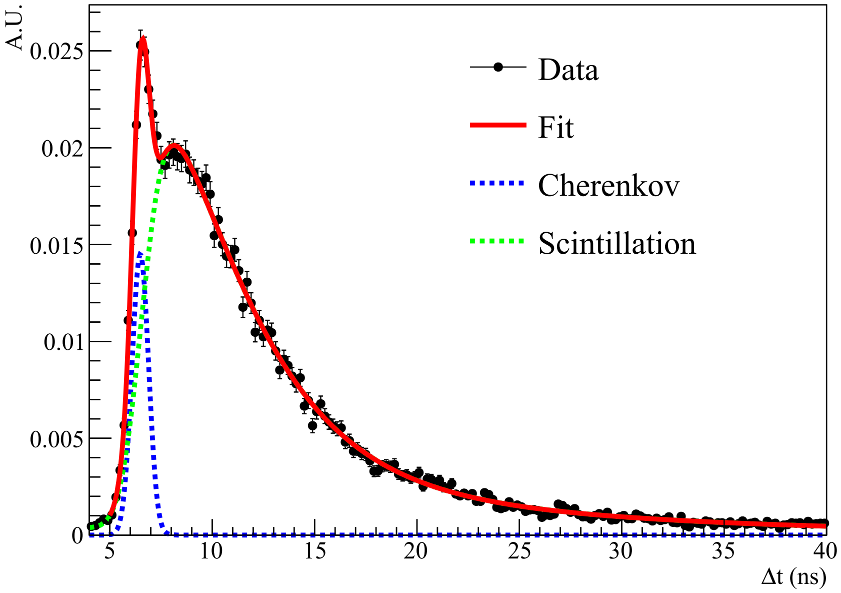

Figure 8.25: The data using the R5912-MOD PMT to detect the transmitted light through the longpass dichroic filter. The fit result shows a clear Cherenkov peak at early times.

total detection efficiency is given by:

$$\mathrm{QE} \times \mathrm{CE} \times \mathrm{F}, \tag{8.10}$$

where QE is the quantum efficiency of the PMT, CE is the collection efficiency of the PMT, and F is the front-end efficiency of the PMT. The QE measures the likelihood a photoelectron is created given an incident photon, and depends on the wavelength of the photon. The CE is the probability that a created photoelectron is guided to the dynode stack and multiplied, creating a signal at the anode. The F measures the efficiency for detecting the signal, given that the PMTs pulses might be lost in the noise or fall below the analysis threshold. In principle an additional efficiency factor is needed for absorption of the PMT glass and photocathode, however this factor should be similar between the three PMTs and thus not impact the relative detection efficiencies.

The QE curves are provided by Hamamatsu for the R2257 and R7600 PMTs (25, 26), and the QE of the R1408 has been measured for the SNO collaboration (27) but the factor CE × F must be measured for the setup. For this, an experimental setup very similar to the one described in Section 8.4.1 is used, including identical DAQ and analysis software.

The measurement is perform for two LEDs, at peak wavelengths of 505 nm and 590 nm. The LED is collimated and directed toward the 50/50 beamsplitter, one output of which goes towards an R7600-U200 normalization PMT. The other output is directed toward either a 494 nm or 587 nm bandpass filter, which narrows the wavelength spectrum of the LED so that it only spans a small portion of the QE curve. The output of the bandpass filter is detected by one of the three measurement PMTs. The ratio of the number of photoelectrons detected by the measurement PMT and the normalization PMT is calculated and compared between the three PMTs. The difference in this ratio between the PMTs measures the difference in detection efficiencies. The known QE for each PMT at 494 nm and 587 nm can be factored out, giving a measure of the relative CE × F factors, referred to as $\mathrm{R}_{\mathrm{CEF}}$. The values of $\mathrm{R}_{\mathrm{CEF}}$ calculated relative to the R2257 PMT (the least efficient) at both wavelengths are shown in Table 8.7. Note the agreement between the R_{CEF} values at both wavelengths, as expected based on the fact that neither the collection or front-end efficiency should be wavelength dependent.

8.4.4.2 Experimental Setup

The dichroicon experimental setup consists of either a pure source of Cherenkov light or a scintillation source. The Cherenkov source is described in Section 7.1.1 and the scintillation source is described in Section 7.2.1.

The scintillator of choice is LAB, which has become a popular liquid scintillator due to its ease of handling, high light yield, and compatibility with acrylic. The fluors used in the

PMT	$R_{\mathbf{CEF}}$ at 494 nm	$R_{\mathbf{CEF}}$ at 587 nm
R7600-U20	1.81 ± 0.04	1.86 ± 0.04
R1408	2.02 ± 0.04	1.97 ± 0.04

Table 8.7: The relative efficiencies of the R7600-U20 and R1408 PMTs. The $R_{\mathbf{CEF}}$ measures the efficiency ratio between the R2257 PMT and the other two PMTs, after factoring out the expected difference in the quantum efficiencies. This effectively provides the value for the collection and front-end efficiencies of the other two PMTs. This factor is used when comparing expected numbers of detected photons for each PMT in the dichroicon measurements.

various measurements are PPO and PTP, both added at a concentration of 2 g/L to the LAB. The primary fluors drastically increase the light output and shift the emission from the UV into the visible, where the PMTs operate most efficiently. Because of its popularity, the LAB+PPO properties have been characterized on the benchtop as discussed in Chapter 7 and discussed in (155, 156, 157, 158, 199). The properties of the fluor PTP are also well-studied (200), but for slightly different applications, such as for the X-ARAPUCA devices (201), and it is not a particularly common fluor to dissolve in LAB. Thus, the majority of the measurements are performed with LAB+PPO.

A Hamamatsu ultra-bialkali R7600-U200 PMT is optically coupled to the acrylic block using Eljen Technology EJ 550 optical grease. The PMT acts as a high efficiency fast trigger and provides the time-zero, and is referred to as the trigger PMT. On the other side of the cube is the dichroicon. The filter at the central aperture of the dichroicon is coupled to a PMT using the EJ 550 optical grease.

Measurements are made separately with two different Hamamatsu PMTs at the aperture: the R7600-U20 and the R2257, which are referred to as the aperture PMTs. The former is a 1-inch (25.4 mm) square PMTs operated at -900V, while the R2257 is a 2-inch (50.8 mm) diameter cylindrical PMT operated at 1500V. These high voltage values were chosen based on recommendations from the Hamamatsu datasheets, and the PMTs were not operating at the same gain. The QE of these PMTs, as well as the emission spectra of the fluors is shown compared to the Cherenkov spectrum in Figure 8.9.

The R7600-U20 has the advantage of a very high efficiency for long-wavelength light, peaking around 20% at 500 nm, and still at 10% by 700 nm. Additionally, this PMT has very fast timing, with a measured TTS of around 350 ps. In comparison, the R2257, while relatively efficient at long wavelengths, only peaks around 10% efficiency close to 600 nm. Additionally, the 900 ps TTS, while still very good, is not as impressive as the R7600-U20. The photocathode area, however, is about five times larger than the R7600-U20, and its cylindrical shape makes it match very well at the center of the dichroicon design.

The full setup is shown schematically in Figure 8.26. The front face of the R2257 or R7600-U20 is placed 215 mm away from the light source. In the complete configuration, a cylinder with reflective Mylar coating is used to direct the short-wavelength light back to an R1408 PMT, operated at 2000 V with a gain of 10^7. Additionally, the R2257 or R7600-U20 are wrapped in reflective foil to ensure photons are not lost when they hit the back of central PMT and its base. The reflective cylinder is about 152 mm in diameter, to fit tightly around the dichroicon and any small gaps were closed with black tape. The R1408 PMT is an 8-inch (203 mm) diameter Hamamatsu PMT, so the outside of the PMT is masked off using felt to ensure only the central area, viewing inside the reflective cylinder, is used. Various pictures of the setup can be found in Figures 8.27 and 8.28.

As discussed, this particular design for the detection of the scintillation light is not expected to be optimal. The length of the R2257 requires a long lightguide, and the R1408 PMT is larger than necessary for the narrow-view dichroicon. The design is directed only toward achieving the primary goal here: demonstrating the sorting of photons in a way that preserves as much of both wavelength bands as possible. A more robust and integrated design would be needed for a realistic large-scale detector.

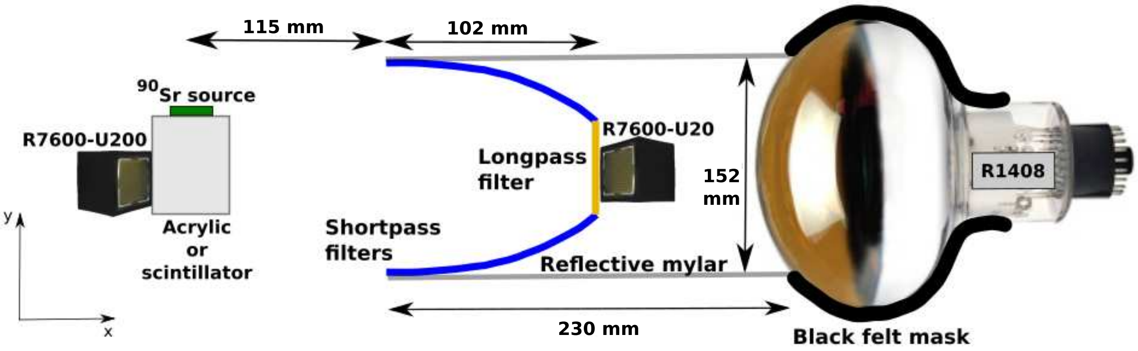

Figure 8.26: A schematic showing the setup with the dichroicon and reflective cylinder. The R7600-U200 PMT is optically coupled to the acrylic or scintillator source and used as a fast trigger. The long-wavelength light is detected at the aperture of the dichroicon. The short-wavelength light is transmitted through the dichroicon, reflected off of the Mylar lining the cylinder, and detected by an R1408 PMT. The setup with the R2257 aperture PMT is identical to the one shown, except due to the length of the aperture PMT, the reflective cylinder is extended 150 mm. The back of the aperture PMT is covered in reflective foil. The area of the R1408 PMT outside of the reflective cylinder is masked off using felt.

8.4.4.3 DAQ and Data Analysis

The same DAQ system as the one described in Section 8.4.1 is used. The rate of coincidences where both the trigger and aperture PMTs detect light is kept low, and thus two million triggered events were recorded for all data sets to maintain reasonable statistics. This low coincidence rate (about 1% for the Cherenkov source) ensures that the detected events at the aperture PMT are single photoelectron (SPE). This simplifies calculating the photon arrival times where no correction for multiple photon hits is needed.

`C++`-based analysis code runs over the `hdf5` files to identify interesting events, in which light is detected by the PMTs around the dichroicon. A software-based constant fraction discriminator is used to find the time difference between the trigger PMT and the dichroicon PMTs. This is done by scanning the digitized waveform of the dichroicon PMTs looking for

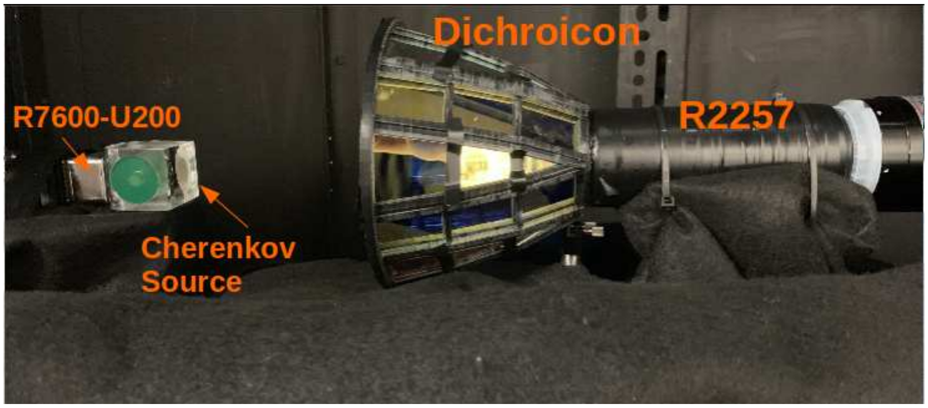

Figure 8.27: A side view of the dark-box setup with the Cherenkov source. The dichroicon is shown with the R2257 PMT at the aperture. In front of the R2257 PMT is a 480 nm longpass dichroic filter. The barrel of the dichroicon consists of shortpass dichroic filters. The reflecting cylinder and R1408 PMT are not shown in this setup. The distances and size of the various important components is shown in Figure 8.26.

a threshold crossing where the voltage is three times larger than the width of the electronics noise. After each threshold crossing the number of consecutive samples above threshold are counted in a 15 ns window. If the waveform stays above threshold for longer than 3 ns, the analysis flags the threshold crossing as associated with a true PMT pulse, rather than a spike in the electronics noise. The peak of the PMT pulse is identified, and the sample associated with the 20% peak-height crossing is found. The time of the trigger PMT is identified similarly and the large signal at the PMT makes the threshold crossing easy to find.

In general, the dark-rate of each PMT is estimated by looking for PMT pulses in a window before the prompt light, and in all cases a correction is applied when making quantitative comparisons. This turns out to be a small correction given the relatively low dark-rates of these PMTs.

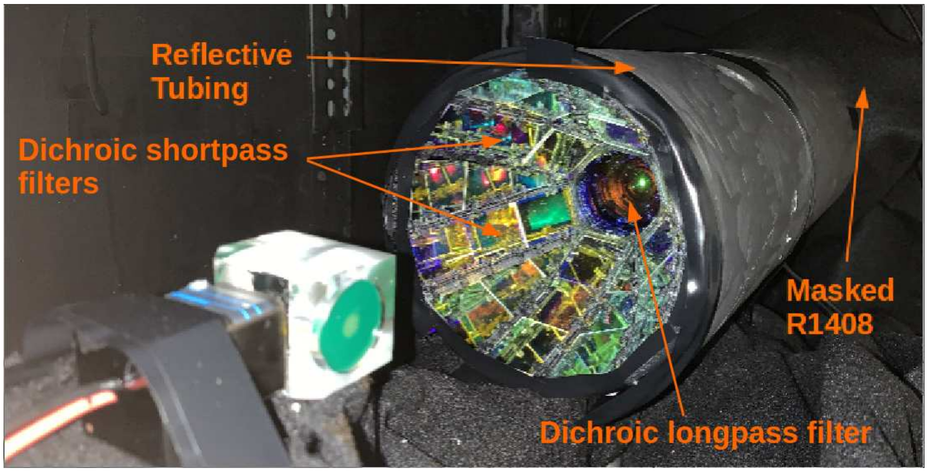

Figure 8.28: The full setup with the short-wavelength light detection system, which consists of a cylinder with a Mylar-lined reflecting interior that ends at an R1408 PMT. The part of the R1408 PMT outside the cylinder is masked off using black felt.

8.4.4.4 Simulation Models

A simulation of the bench-top setup with the dichroicon was developed in the `Chroma` software package (202). `Chroma` is open-source and can be found on Github (203). `Chroma` provides a fast real-time ray-tracer, as well as a full optical Monte Carlo that can be nearly 400 times faster than `GEANT4`. The detector geometries are defined by triangulated surfaces rather than constructive solid geometry. This provides a reasonable alternative to the standard `GEANT4`-based Monte Carlo software, particularly for very large geometries such as `THEIA` and Hyper-Kamiokande with tens of thousands of PMTs. A model of the dichroicon and bench-top setup is implement in `Chroma` and compared to the Cherenkov source data in Section 8.4.4.5.

A detailed optical model of the dichroicon is implemented in `Chroma` and shown in Figures 8.29 and 8.30. `Chroma` allows the triangular mesh defined by CAD software to be directly used in the simulation, so the CAD drawing for the dichroic filter holder is used to accurately reproduce the positions and orientations of the dichroic filters. For other geometry components, a triangular mesh is constructed at runtime according to measurements taken of the dichroicon. The simulation defines the surface properties of each triangle in the mesh along with the bulk properties of the material between triangles. Photons are initially produced in a `GEANT4` simulation and transferred to a GPU where they are propagated from triangle to triangle. This propagation is done on a GPU using `CUDA` ray tracing code where each `CUDA` core handles a single photon, allowing many photons to be propagated in parallel much faster than a single thread could achieve.

`Chroma` implements several surface models that govern the behavior of photons on the triangular mesh. By default, the Fresnel equations are used to refract or reflect photons between materials of different refractive indices. A surface model that defines absorption, diffuse reflection, and specular reflection probabilities is used for the surfaces of opaque materials. Finally, a model that uses wavelength and angle-of-incidence dependent reflection and transmission probabilities is used to model the behavior of the dichroic filters. Between mesh triangles, Rayleigh scatter, attenuation, and reemission of photons is simulated according to the defined bulk material properties.

The measured transmission properties of the various filters described in Section 8.4.1 were used to create the dichroic surface models in the simulation. The filter holder mesh was set to perfectly absorbing. For the PMTs, Hamamatsu specifications were used to set the QE of the photocathodes and construct the overall geometry. A simple model of the acrylic Cherenkov source implemented in `GEANT4` was used to generate Cherenkov photons from the energetic electrons in Y^{90} decays, which are propagated in the `Chroma` geometry.

After photons are propagated and absorbed on the photocathode, a simple DAQ simulation is performed whereby the hit times of the photons are smeared by a Gaussian distribution with a width matching the Hamamatsu provided TTS of the PMT. The earliest hit time, if any photons were detected, is taken as the hit time for the PMT in that event. Analyzing the time differences between the hit time on the tag PMT coupled to

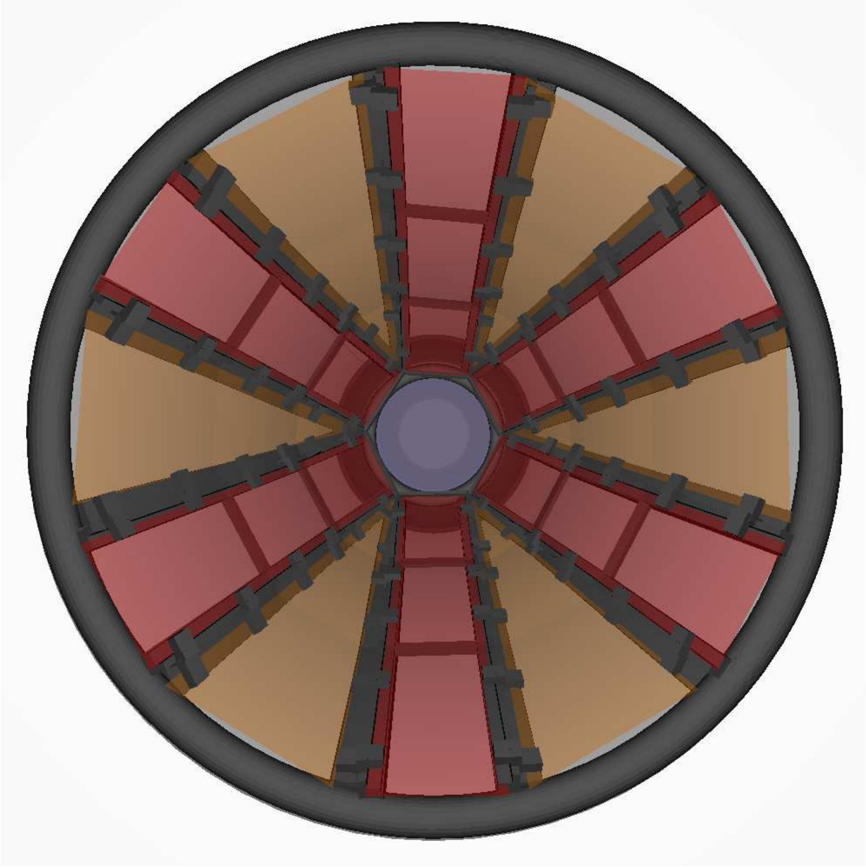

Figure 8.29: A direct view of the `Chroma` dichroicon model with the R2257 PMT at the center of the dichroicon. The outer diameter of the dichroicon is about 150 mm and the inner radius, where the long-pass filter is located, is about 50 mm. The two different colors in the barrel of the dichroicon indicate the two different types of short-pass filters, as detailed in Table 8.3. This figure is created by B. Land.

the Cherenkov source and the PMTs in the dichroicon can proceed in the same way as the measured data.

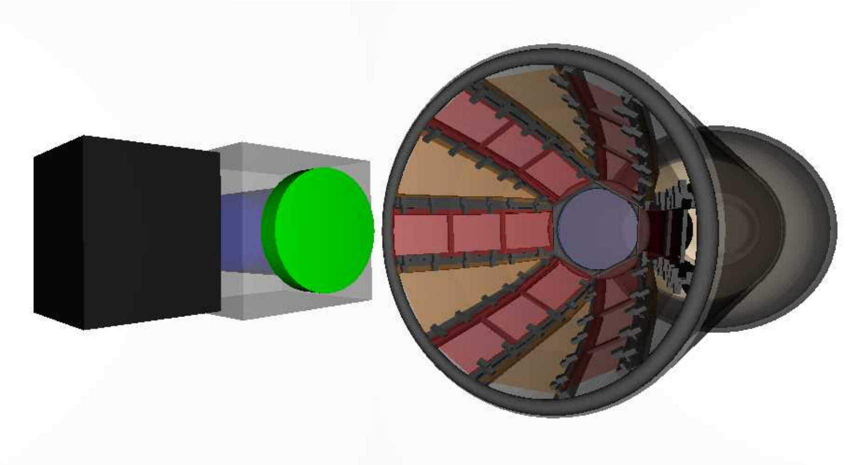

Figure 8.30: The full `Chroma` simulation setup for reproducing the Cherenkov source results. Geometry components with the same optical properties are colored similarly, however the colors are arbitrary. The sizes of the components and distances between the objects are identical to those used for the data, shown in Figure 8.26. This figure is created by B. Land.

8.4.4.5 Cherenkov Source Results

The primary goals in the measurements of the Cherenkov source are to determine the effectiveness of the dichroicon for spectral sorting of Cherenkov light, to understand how each component of the dichroicon and reflecting cylinder affects the overall performance of the full device, and to develop and test the `Chroma` model of the dichroicon. The Δt profiles presented are generated using the analysis described in Section 8.4.4.3.

Measurements are done in several staged configurations to understand the device under a variety of conditions. First, data was taken with the aperture PMT placed 215 mm away from the source, with no filters and no reflecting cylinder. This gives a baseline measurement for those to follow and is referred to as the 'no filters' dataset. Second, a 480 nm dichroic longpass filter was coupled to the front-face of the aperture PMT, which acts as it would at the center of the dichroicon. In this configuration, only the dichroic longpass filter is present, and the ‘barrel’ of the dichroicon is not included. Third, a ‘standard’ absorbing longpass filter is added behind the dichroic longpass filter. This filter is included to absorb possible short-wavelength leakage through the dichroic filter. Finally, the barrel of the dichroicon filled with the shortpass filters was deployed. Data was taken with the dichroicon both with and without the absorbing longpass filter behind the dichroic longpass filter.

In order to compare across configurations and aperture PMTs, a factor C^{*}_{NORM} is defined in Equation 8.11 that integrates the histograms, H, 5σ around the mean, μ, of the distributions, where σ and μ come from a Gaussian fit to the no filters data.

$$C_{\mathrm{NORM}} = \frac{\int_{\mu-5\sigma}^{\mu+5\sigma} H dt}{N \times A \times \mathrm{R}_{\mathrm{CEF}}} \tag{8.11}$$

The integral is then normalized by the number of triggered events, N, the photocathode area of the PMT, A (taken from the Hamamatsu data sheets (25, 26)), and the relative efficiencies of the PMT, $\mathrm{R}_{\mathrm{CEF}}$. The $\mathrm{R}_{\mathrm{CEF}}$ factor contains the collection and front-end efficiencies, which are measured relative to the R2257 PMT, as presented in Section 8.4.4.1. C_{NORM} is used to compare the amount of Cherenkov light detected across different configurations and between the two aperture PMTs. By including the photocathode area, the collection efficiency, and the front-end efficiency in the C_{NORM} factor, the difference in performance between the two PMTs comes primarily from the different QE curves.

The results of these tests for the aperture PMT are presented in Table 8.8. In order to directly compare the `Chroma` results to the data, a scaling factor in the simulation that is used to adjust the overall efficiency is tuned so that the simulation and data agree exactly for the no filters configuration. This scaling factor is kept constant for the subsequent results. As is clear from Table 8.8 the overall agreement between data and simulation is quite good. The largest discrepancies occur for data with the dichroicon and absorbing filter, likely due to small mis-modelling of the very complicated nature of the three different types of dichroic filters. Figure 8.31 shows the results for the R2257 PMT under several of the configurations compared directly to the `Chroma` results.

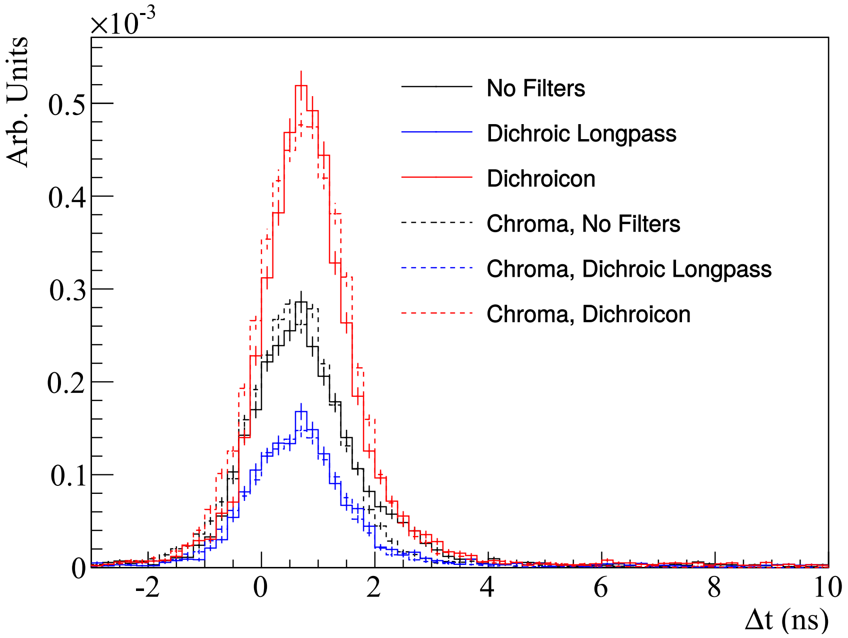

Figure 8.31: The results for the R2257 central aperture PMT and the acrylic Cherenkov source. In black is data for the configuration with no filters or dichroicon. The blue shows the data with the longpass dichroic filter optically coupled to the R2257. The data with the full dichroicon added is shown in red. The corresponding `Chroma` results are shown in the dashed lines. These results are summarized in Table 8.8. The value of Δt is determined by cable delays and transit times through the PMTs and has no impact on the analysis. The results for the R7600-U20 look similar, but with a narrower TTS.

The results presented in Table 8.8 indicate that regardless of the configuration, the the R7600-U20 PMT detects more Cherenkov light per photocathode area than the R2257 PMT, as one would expect based on the quantum efficiencies of the PMTs. The relative changes to C_{NORM} depend on the shape of the QE curve of each PMT, but the qualitative features are consistent in both setups.

Adding the dichroic longpass filter in front of these aperture PMTs reduces the Cherenkov

PMT	Dichroic	Absorbing	Dichroicon	C_{NORM} (1/m^2) [Data]	C_{NORM} (1/m^2) [`Chroma`]
R2257	-	-	-	1.73 ± 0.03	1.73 ± 0.02
R2257	✓	-	-	0.95 ± 0.02	0.91 ± 0.01
R2257	✓	✓	-	0.87 ± 0.02	0.78 ± 0.01
R2257	✓	-	✓	2.75 ± 0.03	2.94 ± 0.03
R2257	✓	✓	✓	1.80 ± 0.03	1.65 ± 0.02
R7600-U20	-	-	-	4.18 ± 0.06	4.18 ± 0.09
R7600-U20	✓	-	-	1.76 ± 0.04	1.85 ± 0.06
R7600-U20	✓	✓	-	1.60 ± 0.04	1.64 ± 0.05
R7600-U20	✓	-	✓	8.90 ± 0.09	7.31 ± 0.11
R7600-U20	✓	✓	✓	5.80 ± 0.07	3.74 ± 0.08

Table 8.8: The Cherenkov source results for the R2257 and R7600-U20 aperture PMTs, with different configurations of the filters. C_{NORM} is defined in Equation 8.11. The errors are statistical only. The ✓indicates whether a given part of the setup was used. The first column corresponds to the central longpass dichroic filter, the second column to the absorbing longpass filter behind the dichroic filter, and the third column to the barrel of the dichroicon equipped with the shortpass dichroic filters. The results from data and the Chroma simulation are shown in the next two columns. To account for unmodeled inefficiencies in the simulation, the results for each PMT are scaled such that the case with no filters has the same C_{NORM} as data.

light by about 50%, as expected given the reflection of the shorter-wavelength photons. The reduction is slightly larger for the R7600-U20 because it has a higher sensitivity to short-wavelength light, which is largely being filtered out. Introducing the absorbing longpass filter behind the dichroic filter, an additional 10% of the Cherenkov light is lost. This filter absorbs short-wavelength leakage through the dichroic filter, which is not designed to have perfect blocking below the cut-on wavelength, particularly for photons at large incidence angles. This becomes important for the measurements presented in Section 8.4.4.6, where the absorbing filter is used to remove scintillation light leakage through the dichroic filter.

By incorporating the dichroicon barrel, filled with shortpass dichroic filters, the total amount of Cherenkov light collected at the aperture (through only the dichroic longpass filter) is increased by a factor of 2.9 and 5.1 for the R2257 and R7600-U20 respectively. The larger factor for the R7600-U20 comes from the relatively larger short-wavelength sensitivity—the short-wavelength light, which can be reflected by the barrel of the dichroicon, now strikes the central dichroic filter at higher incidence angles, where the filter is more likely to leak the photon.

In addition to detecting the light at the aperture, as with a standard Winston cone, the design requires the detection of the light transmitted through the short-pass dichroic filters. The short-wavelength system, consisting of the reflective cylinder and R1408 PMT, was tested with both aperture PMTs. The inclusion of this system did not affect the amount of light detected by either aperture PMT. The R1408 detects 3.3 times more light than either aperture PMT, after correcting for the measured relative differences in the front-end and collection efficiencies.

Simulations of this setup over-predict the total light detected by the R1408 by a factor of 3. This is likely due to the fact that the model does not include many inefficiencies in the setup of the short-wavelength detection system. These primarily include non-perfect reflectivity of many of the components, non-uniform deployment of the reflective components, and the lack of the inclusion of the PMT stands in the model. The result is nevertheless encouraging, as this simple system detects 33% of the predicted short-wavelength light. A more integrated device with better and well-understood reflective coating could easily improve on these results and should in principle be easier to model.

Overall, with both aperture PMTs, the dichroicon demonstrated excellent capabilities for spectral sorting of photons towards two different PMTs, simultaneously detecting both the short- and long-wavelength light.

8.4.4.6 LAB+PPO Results

The results presented in this section use the setup described in Section 8.4.4.2, now with the LAB+PPO scintillation source. As discussed, one of the primary goals with this source will be to separate the scintillation and Cherenkov components of the light emission. Figure 8.9, which compares the shape of the PPO emission spectrum to the shape of the Cherenkov emission spectrum, illustrates that photons with wavelengths above 500 nm should be primarily Cherenkov light. In the full setup, this light will be directed toward the aperture PMT, while the shorter wavelength scintillation light will pass through the dichroicon and be detected by the R1408. As with the Cherenkov measurements presented in Section 8.4.4.5, measurements here are made with the two different aperture PMTs.

The Δt distributions are created using the analysis techniques described in Section 8.4.4.3. The scintillation time profile for LAB+PPO is well understood and has been measured in (155, 156, 199), which I use as a guide for the fits to the time profiles. As discussed in (155), deoxygenation of the LAB+PPO leads to a reduction in the quenching of the scintillation light, which primarily affects the late-light timing in the tail of the $\Delta \mathrm{t}$ distribution. The focus of these measurements is primarily on the prompt light and the spectrum is fit only to 20 ns past the prompt peak, so deoxygenation of the scintillator is not performed.

Following the procedure with the Cherenkov source, measurements are first made of light at the aperture PMT under varying conditions. These include incrementally adding an absorbing longpass filter and then the barrel of the full dichroicon, keeping track of the relative effect of each component. Unlike with the Cherenkov source, the prompt peak cannot simply be integrated to yield the total Cherenkov component, because there might be significant contamination from scintillation light leaking through the dichroic filter. Additionally, unlike the Cherenkov source measurement, no measurement is made without the central dichroic filter, as the Cherenkov light is completely dominated by the large scintillation light yield.

The Δt spectra for the central PMT are fit using Equation 8.8. The fit is performed to the prompt Cherenkov and scintillation light between 0 to 30 ns in the Δt histogram. The scintillation component of the fit uses the sum of two decay exponentials with time constants τ_1 and τ_2 and associated weights A_1 and A_2 and an exponential rise time τ_R. The weights are constrained to sum to one. The PMT TTS, f_{PMT}, is determined from the Cherenkov source data to be around 350 ps for the R7600-U20 and 900 ps for the R2257. The mean, μ, and width, σ, of f_{PMT} are constrained to the values found when performing Gaussian fits to the Cherenkov source data. This distribution is convolved with the scintillation time-profile in the fit. There is an arbitrary offset, t', which allows for cable delays and other time-offsets, that is different for the two PMTs. The Cherenkov component is modeled by weighting the PMT TTS distribution by an appropriate factor C. The Cherenkov and scintillation components are constrained to sum to one.

Figure 8.32 shows an example fit done to the LAB+PPO data with the R7600-U20 at the aperture of the dichroicon, with the Cherenkov and scintillation components shown separately. As is evident, the scintillation and Cherenkov components of the light are very nicely separated, and a prompt timing window can be selected to identify a pure sample of Cherenkov light. The purity of the selection of the prompt Cherenkov light, P, is defined in Equation 8.12, which calculates the fraction of the Cherenkov component of the fit in a non-symmetric window that goes from -5σ to $+1.5\sigma$ around the mean of the distribution.

$$P = \int_{\mu-5\sigma}^{\mu+1.5\sigma} \frac{C \times f_{PMT}(t-t')}{F} dt \tag{8.12}$$

In addition to the purity, the total amount of Cherenkov light detected is an interesting quantity that can be directly compared to the Cherenkov source data. Given that the indices of refraction of LAB+PPO and acrylic are almost identical (162), the total amount of Cherenkov light collected should be nearly the same. As with the Cherenkov source data, a correction is applied for the number of triggered events, the photocathode area of the PMT, and the collection and front-end efficiencies of the PMT, as defined in Equation

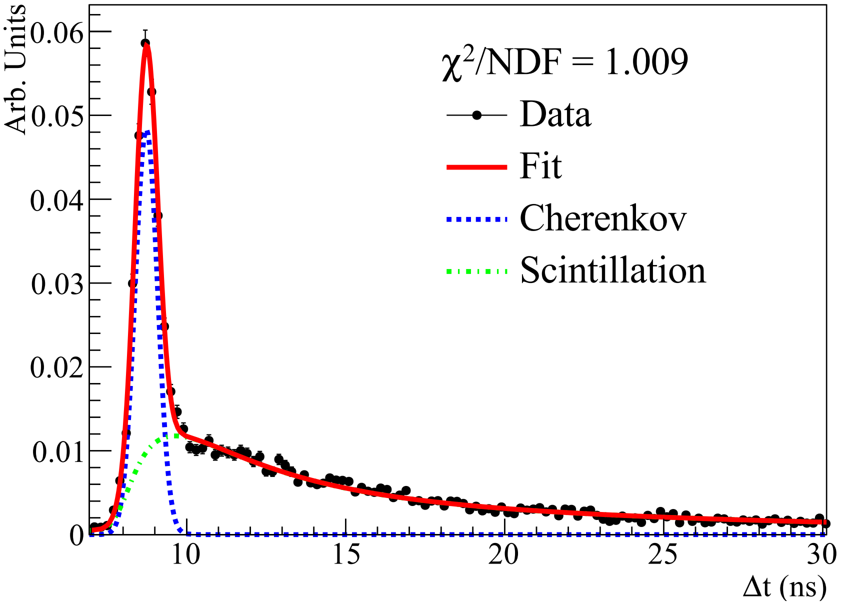

Figure 8.32: An example fit using Equation 8.8, to LAB-PPO data. The Cherenkov and scintillation components are shown separately, in addition to the total fit in red. The data were taken with the shortpass barrel of the dichroicon and the absorbing longpass filter behind the dichroic longpass filter at the aperture of the dichroicon. The data is normalized to 1.0. The Δt offset from zero is arbitrary and does not impact the fit.

8.13. Note this equation is essentially identical to Equation 8.11, but here the fit is explicitly used to remove the expected counts from scintillation leakage into the prompt window.

$$C^*_{\mathrm{NORM}} = \frac{\int_{\mu-5\sigma}^{\mu+5\sigma} C \times f_{PMT}(t-t')dt}{\mathrm{N} \times \mathrm{A} \times \mathrm{R_{CEF}}} \tag{8.13}$$

Particle identification and reconstruction in a real scintillation detector using Cherenkov light will depend on both the purity of the Cherenkov photons and their total number. One can achieve great purity by, for example, moving the cut-on wavelength toward even longer wavelengths, but this buries more of the remaining Cherenkov photons beneath the scintillation light and thus would lead to only a small number of usable Cherenkov photons. Thus, the metric selected is the total number of Cherenkov photons multiplied by the purity,

$$R = C^{*}_{\mathrm{NORM}} \times P. \tag{8.14}$$

When $P = 1$ the detector can operate identically to a Cherenkov detector, up to the total light yield given by C^{*}_{NORM}. In a real detector, additional information can be used to identify Cherenkov photons (such as an hypothesized event direction) and thus while R provides a good relative metric, it may be a pessimistic assessment of the total effective Cherenkov yield.

The results for both aperture PMTs for the various configurations are presented in Table 8.9 with the corresponding `Chroma` results in Table 8.10. The C^{*}_{NORM} results agree well with the Cherenkov source data and can be directly compared to the C_{NORM} values in Table 8.8. As with the Cherenkov source results, the `Chroma` predictions agree well with the data, particularly for the R2257. Again, the largest discrepancy is with the dichroicon with the R7600-U20, where the total amount of Cherenkov light collected is being under predicted. Additionally, the `Chroma` results systematically under predicts the amount of scintillation light leakage through the central dichroic longpass filter, leading to a higher purity selection of the Cherenkov light in the prompt window. Future efforts to characterize the dichroic filters in more detail should help improve the agreement between the data and simulation.

From these tables it is clear that an excellent purity is achieved with both the R2257 and R7600-U20 PMTs. As expected, given the faster timing of the R7600-U20 PMT, better purity is achieved with this PMT. In all setups with the absorbing longpass filter behind the longpass dichroic filter at the aperture, the purity for Cherenkov selection in the prompt window is better than 90%. It is important that this purity is achieved with a prompt window that still selects the majority of the Cherenkov light – that is, to achieve a higher purity of Cherenkov selection an extremely narrow window around the prompt peaks is not

PMT	Dichroic	Absorbing	Dichroicon	C^*_{NORM} (1/m^2)	P (%)	R
R2257	✓	-	-	0.91 ± 0.06	70.8 ± 1.4	0.64 ± 0.06
R2257	✓	✓	-	0.80 ± 0.06	90.6 ± 1.1	0.72 ± 0.06
R2257	✓	-	✓	2.77 ± 0.13	82.2 ± 1.2	2.28 ± 0.07
R2257	✓	✓	✓	1.84 ± 0.10	90.0 ± 1.1	1.66 ± 0.07
R7600-U20	✓	-	-	1.77 ± 0.13	69.9 ± 1.5	1.24 ± 0.09
R7600-U20	✓	✓	-	1.73 ± 0.13	95.5 ± 1.4	1.65 ± 0.13
R7600-U20	✓	-	✓	8.71 ± 0.64	84.4 ± 1.5	7.35 ± 0.56
R7600-U20	✓	✓	✓	5.54 ± 0.40	93.2 ± 1.3	5.16 ± 0.38

Table 8.9: Results for the LAB+PPO source with the R2257 and R7600-U20 central PMTs. C^*_{NORM} is total Cherenkov light, normalized to the number of triggers and photocathode area, defined in Equation 8.13. P is the purity of the Cherenkov light selection in a prompt window, defined in Equation 8.12. R is $C^*_{\mathrm{NORM}} \times P$. The errors come from the uncertainties on the fit parameters. The ✓indicates whether a given part of the setup was used. The first column corresponds to the central longpass dichroic filter, the second column to the absorbing longpass filter behind the dichroic filter, and the third column to the barrel of the dichroicon equipped with the shortpass filters.

selected, which would reject not only the scintillation light, but also large amounts of the Cherenkov light.

By comparing the R values for the various datasets several conclusions can be drawn. First, comparing the first two rows for each PMT in Table 8.9, the introduction of the 500 nm absorbing longpass filter without the barrel of the dichroicon increases the purity without significantly affecting the total Cherenkov light detected, effectively increasing R. However, when the full dichroicon is in place, the total Cherenkov light detected is significantly affected by the inclusion of the absorbing longpass filter, which was discussed in the Cherenkov source data and is due to the leakage of photons through the central dichroic filter at high incidence angles. This suggests that the inclusion of the absorbing longpass filter is not necessarily the optimal configuration for a large-scale detector, where not only purity but also the total number of detected Cherenkov photons is important. While the

PMT	Dichroic	Absorbing	Dichroicon	C^*_{NORM} (1/m^2)	P (%)	R
R2257	✓	-	-	0.91 ± 0.10	87.8 ± 1.1	0.80 ± 0.11
R2257	✓	✓	-	0.83 ± 0.08	97.7 ± 1.1	0.81 ± 0.08
R2257	✓	-	✓	2.60 ± 0.21	86.2 ± 1.4	2.24 ± 0.25
R2257	✓	✓	✓	1.72 ± 0.16	96.4 ± 1.1	1.66 ± 0.17
R7600-U20	✓	-	-	1.77 ± 0.22	80.3 ± 3.6	1.42 ± 0.29
R7600-U20	✓	✓	-	1.76 ± 0.21	95.0 ± 1.8	1.67 ± 0.22
R7600-U20	✓	-	✓	6.41 ± 0.49	80.0 ± 1.7	5.13 ± 0.64
R7600-U20	✓	✓	✓	3.63 ± 0.32	96.4 ± 1.9	3.50 ± 0.34

Table 8.10: Results for the `Chroma` simulations of the LAB+PPO source with the R2257 and R7600-U20 central PMTs. C^*_{NORM} is total Cherenkov light, normalized to the number of triggers and photocathode area, defined in Equation 8.13. P is the purity of the Cherenkov light selection in a prompt window, defined in Equation 8.12. R is $C^*_{\mathrm{NORM}} \times P$. The errors come from the uncertainties on the fit parameters. The ✓indicates whether a given part of the setup was simulated. The first column corresponds to the central longpass dichroic filter, the second column to the absorbing longpass filter behind the dichroic filter, and the third column to the barrel of the dichroicon equipped with the shortpass filters. To account for unmodeled inefficiencies in the simulation, the results for each PMT are scaled such that the case with only the longpass dichroic filter has the same C^*_{NORM} as data.

purity of the selection is increased (by about 10%) with the absorbing longpass filter, the total Cherenkov light lost is about 30%, so the value of R decreases.

Perhaps most notably, by including the full dichroicon the Cherenkov light collection is improved by about a factor of 5 for the R7600-U20 and a factor of 3 for the R2257, identical to the Cherenkov source results. This suggests that having a red-sensitive PMT is important at the center of the dichroicon, but having additional blue sensitivity will help with the detection of the Cherenkov light. The temporal separation with these PMTs is already quite good, and therefore a small amount of short-wavelength scintillation leakage is acceptable. Because both Cherenkov and scintillation light are leaked through the filter at the aperture, this would increase the total Cherenkov light detected while sacrificing a small amount of purity.

In addition to the Cherenkov light detection, the ability to sort scintillation light to the back R1408 PMT is a critical feature of the dichroicon design. The setup is adjusted to include the full reflective cylinder and R1408 PMT. Encouragingly, it was easy to identify individual events with coincidences of Cherenkov light at the aperture PMT and scintillation light at the back PMT. A typical waveform for the setup with the R2257 is shown in Figure 8.33. In this event, it is clear that the back PMT detected several photons, leading to a PMT pulse size corresponding to around 5 photoelectrons, while the aperture PMT detected a single Cherenkov photon. The scintillation timing profile is also evident in the pulse shape of the R1408.

The data for the aperture PMT is compared with and without the reflecting cylinder, and no difference is found. This indicates that the reflecting light guide does not lead to additional scintillation light bouncing off of the cylinder and leaking through the central dichroic filter. The total light collected at the R1408 PMT should represent the high light yield of the scintillator. To quantify this, the charge spectrum, which is heavily multi-PE, is integrated using Equation 8.15.

$$Q_{\text{TOT}} = \sum_{i=0}^{N} \frac{Q_i \times C_i}{Q_{\text{SPE}}} \tag{8.15}$$

Here Q_i is the charge of bin i in pC and C_i is the associated bin content. This sum is normalized by the charge of the SPE peak Q_{SPE}, which is 1.6 pC as the PMT is operating at a gain of 10^7. This is a good measure of the total light collected by the R1408 and can be compared to the total light collected by the aperture PMT, which remains in the SPE regime. For the aperture PMT the photon counting is done the same way as with the Cherenkov source. The R1408 detects about 550 times more light than both aperture PMTs in these setups. This highlights the traditional difficulty with separating the two components of the light – the scintillation yield overwhelms the Cherenkov yield. Using

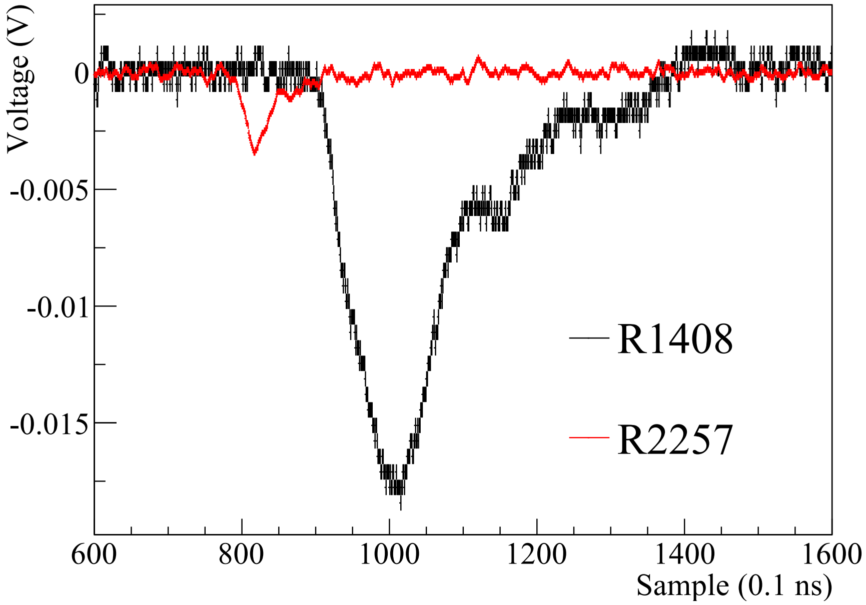

Figure 8.33: The R2257 and R1408 waveforms for a single triggered event. This event has an early-time PMT pulse, corresponding to Cherenkov light, for the R2257 PMT and an approximately 5 PE pulse at the R1408 PMT. The waveforms are specifically selected by looking at the Δt histogram and selecting an event in the Cherenkov peak. The x-axis is labelled in 0.1 ns samples, and the offset between the two waveforms is showing the additional photon travel time and transit time through the R1408 PMT.

the dichroicon, a large fraction of the scintillation light is simultaneously detected with the R1408, while detecting the Cherenkov light at the aperture with high purity. To quantify this, the results for the R1408 are compared to the `Chroma` prediction and an efficiency for detecting the short-wavelength light around 30% is found, consistent within uncertainties with the value found for the Cherenkov source. With a better engineered, integrated system, it is expected to be able to increase the efficiency for detecting the short-wavelength light.

The time distribution of the detected light for the R7600-U20 and R1408 PMTs can be seen in Figure 8.34. The scintillation light detected by the R1408 PMT swamps the Cherenkov light collected by the R7600-U20 PMT, making it very clear why this separation is so difficult without spectral photon sorting. This figure clearly illustrates the power of the dichroicon – by detecting the Cherenkov and scintillation light in separate PMTs, the high scintillator light yield is maintained while simultaneously detecting the Cherenkov light.

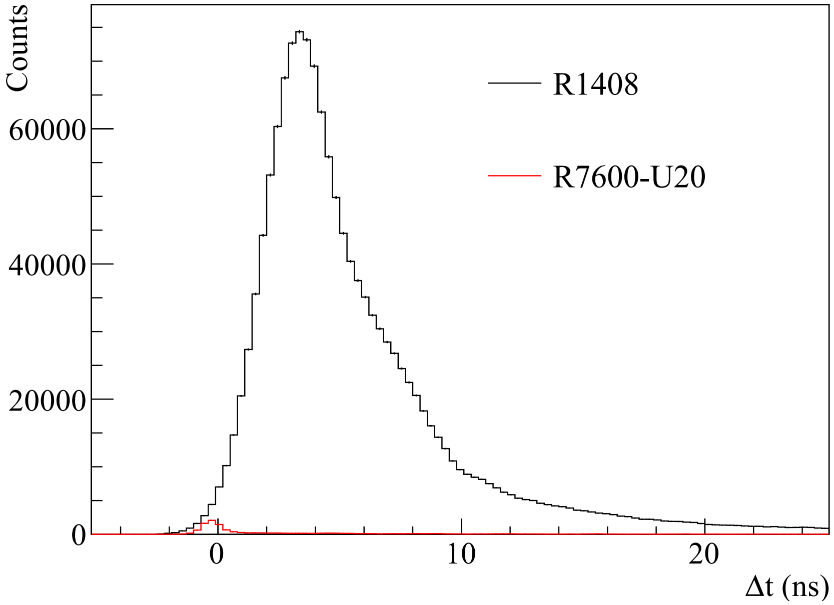

Figure 8.34: The dichroicon data for both the R7600-U20 PMT at the aperture and the R1408 PMT behind the dichroicon. The light detected by the R1408 is primarily scintillation light and the light detected by the R7600-U20 is primarily Cherenkov light.

8.4.4.7 Pulse Shape Discrimination

For liquid scintillator detectors, particle identification (PID) is critical for background rejection. LAB+PPO is well known to have good pulse-shape discrimination between β and

α excitation, a thorough discussion of which can be found in (156). In short, α particles, more than β particles, are known to excite the slow component of the scintillation emission, which is typically associated with triplet state excitation of the solute molecules. The ratio of the amount of prompt to the amount of late light can be used as a handle to separate β from α particles. A practical application of this PSD technique in the Borexino detector can be found in (204), and is discussed for SNO+ and LENA in (40, 205). Similar PSD techniques for background rejection in scintillators can be found in (160, 206).

In principle, there is another difference in the light production for β and α particle excitation in liquid scintillator. Most β particles from radioactive decays are above the Cherenkov threshold, whereas the α particles are not. Thus, if Cherenkov light could be identified at energies around a couple of MeV in a liquid scintillator detector, the absence of Cherenkov light would be a clear tag for α excitation.

This is tested in the setup using a ^{210}Po α source to irradiate the LAB+PPO, replacing the ^{90}Sr source. For this, only the R7600-U20 aperture PMT is used, as both central PMTs behave well in terms of detecting ample Cherenkov light in the setup with the dichroicon. The ^{210}Po decays 100% of the time via a 5.41 MeV α, which enters the scintillator, creating scintillation light.

Figure 8.35 shows the data for the R1408 PMT, which detects the short-wavelength scintillation light through the barrel of the dichroicon. As expected, the typical difference in the scintillation time-profiles is identified, where, under α excitation, the scintillator produces more late light. This is the typical manner in which liquid scintillator detectors perform PID.

The data for the R7600-U20 at the aperture of the dichroicon is shown in Figure 8.36. As can be seen clearly, the prompt Cherenkov light is absent in the data with the α source. Using the full dichroicon setup, both the difference in the scintillation time-profiles for the back PMT and the difference at early times for the aperture PMT can be used to discriminate between α and β particles.

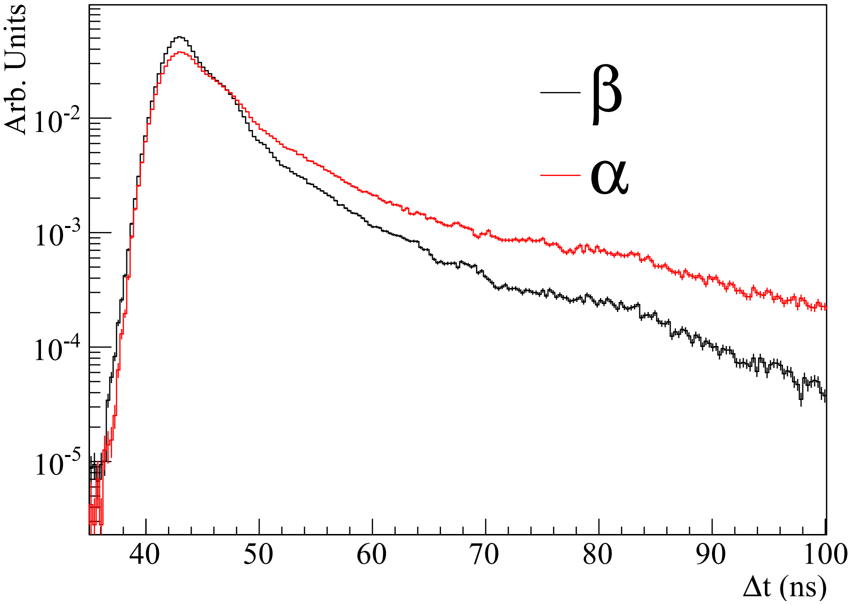

Figure 8.35: The R1408 data for a ^{90}Sr β source and a ^{210}Po α source. The small bumps in the timing spectrum are due to the complicated PMT transit time distribution, which includes two different late-pulsing peaks. The difference between these distributions is used to discriminate between β and α particle excitation in liquid scintillator detectors.

8.4.4.8 Off-Axis Source

Two measurements with the source displaced from center of the dichroicon are performed with the LAB+PPO scintillation source to understand the behavior of the dichroicon. In all cases the scintillation light is emitted isotropically, so as the source moves further from the central-axis of the dichroicon, the average incidence angle of the photons increase and the solid-angle acceptance shrinks. The coordinate system is defined in Figure 8.26 and is used to describe the movement of the source. In the first off-axis test (off-axis 1) the source is relocated to the edge of dichroicon, moving it 75 mm in the +y direction. The second

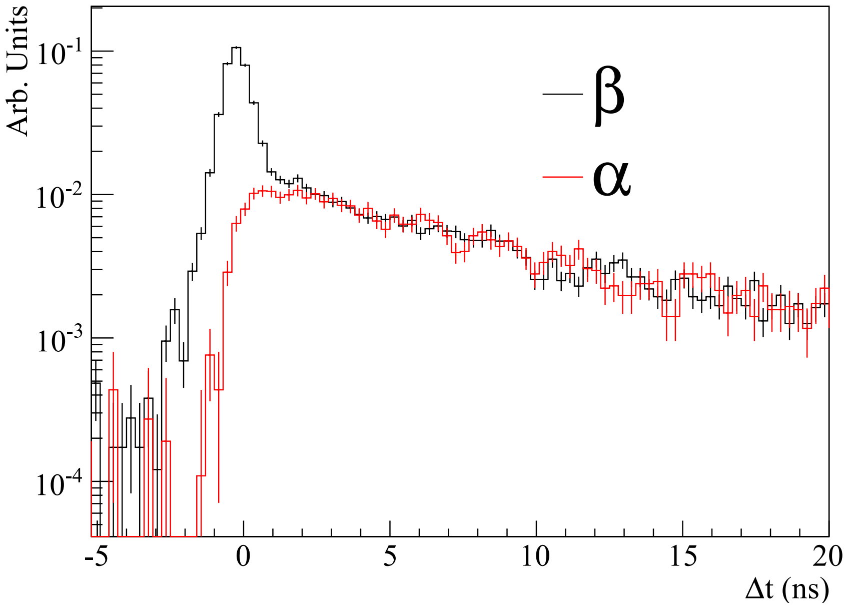

Figure 8.36: The R7600-U20 data for a ^{90}Sr β source and a ^{210}Po α source. The lack of Cherenkov light for the below-threshold α particles can be clearly identified.

off-axis test (off-axis 2) moves the source 75 mm further in the +y direction. The x-distance from the dichroicon is kept constant at 115 mm and the z-position is kept at the same level as the center of the dichroicon.

In the first off-axis measurement a reduction consistent with 50% is identified in the light collection of the Cherenkov light at the aperture PMT, as expected based on the change in solid angle. For the second off-axis measurement the source is now outside the geometric field of view, so only 6% of the Cherenkov light is collected. The only reason any light is collected at all in this setup is that the dichroicon design is not a perfect Winston cone.

For these off-axis measurements, the total light collected at the R1408 PMT is quantified using Q_{TOT}, as defined in Equation 8.15 and compare it relative to the on-axis data.

The data collected at this PMT does not require the Winston cone to work as a reflector—instead, it only relies on the short-wavelength light incident on the dichroicon to be transmitted and successfully reflected back to the PMT. For the off-axis 1 data, about 50% of the total light detected. Then for off-axis 2 data, about 29% of the light still collected, a much larger factor than found for the aperture PMT. This demonstrates the short-wavelength light collection is more robust to large angles of incidence, outside of the view of the Winston cone. Table 8.11 summarizes the results presented in this section.

y (mm)	C^*_{NORM} (1/m^2)	$Q^{\text{rel}}_{\text{TOT}}$
0	0.91 ± 0.10	1.00 ± 0.06
75	0.49 ± 0.05	0.52 ± 0.03
150	0.06 ± 0.01	0.29 ± 0.02

Table 8.11: A summary of the results for the off-axis measurements with the R2257 PMT at the aperture of the dichroicon. The value of C^*_{NORM} for the central data is the same as given in Table 8.9. The C^*_{NORM} and Q_{TOT} values for the first off-axis measurement at y of 75 mm are consistent with a 50% reduction in the collection of the scintillation and Cherenkov light. At the more extreme incident angles the Cherenkov collection efficiency is reduced by a larger factor, as the source is outside the geometric field of view of the Winston cone. The Q_{TOT} values are given relative to the central data, which is normalized to one.

In a large detector, the geometric field of view of the dichroicon will need to be chosen given the detector design – in particular, the size of the detector, the size and shape of the photodetectors, and the expected fiducial volume. In general, an inner fiducial volume is usually several meters away from the PMT array. Thus the maximum angle of incidence on a Winston cone is quite small. The behavior of the dichroicon under far-field illumination is not measured, and will be a priority for future measurements and simulation studies.

By using a scintillator with a shorter wavelength emission spectrum, the cut-on of the central dichroic filter can be decreased, thus increasing the total Cherenkov light collected at the aperture PMT. A fluor, PTP, was identified as being able to dissolve in LAB at 2g/L with a high light yield and a shorter wavelength emission spectrum than PPO, as presented in Figure 8.9. Measurements discussed in this section are done with LAB+PTP and compared directly to the LAB+PPO measurements. The setup included the full dichroicon, with the R7600-U20 at the aperture, but without any absorbing longpass filter behind the dichroic filter.

The second dichroicon, called dichroicon-2, detailed in Table 8.4 uses 462 nm longpass filter at the aperture to replace the 480 nm longpass filter. Additionally, the rectangular filters are replaced with 450 nm shortpass filters. These filters are chosen with shorter-wavelength cut-on values to reflect to and transmit through the aperture filter a larger fraction of the Cherenkov light.

The data for the R7600-U20 PMT with a LAB+PPO source, shown in Figure 8.37, shows an increase in the amount of Cherenkov and scintillation light detected when using the dichroicon-2. This is shown similarly for the LAB+PTP source in Figure 8.37; however, given the shorter wavelength emission spectrum of the PTP, there is only an increase in the Cherenkov light detected, with no change to the amount of scintillation light leaking through the dichroic filter.

These results are presented quantitatively in Table 8.12. The definitions of C^*_{NORM}, P, R, and Q_{TOT} are the same as the ones presented in Section 8.4.4.6. The first row shows the LAB+PPO results already presented in Table 8.9 for the R7600-U20 PMT at the aperture of the dichroicon. By replacing the LAB+PPO with LAB+PTP the amount of Cherenkov light, quantified by C^*_{NORM}, does not change, as expected, but the purity of the selection in a prompt window improves to about 88%. This improvement comes from the smaller amount of scintillation light leakage through the central filter, due to the shorter emission spectrum, and thus less scintillation light in the selected prompt window.

Rows three and four in Table 8.12 show the results for dichroicon-2, which show a notable increase of about 40% in the total Cherenkov light collected. Again, this is consistent between the LAB+PPO and LAB+PTP results. The purity of the selection for the

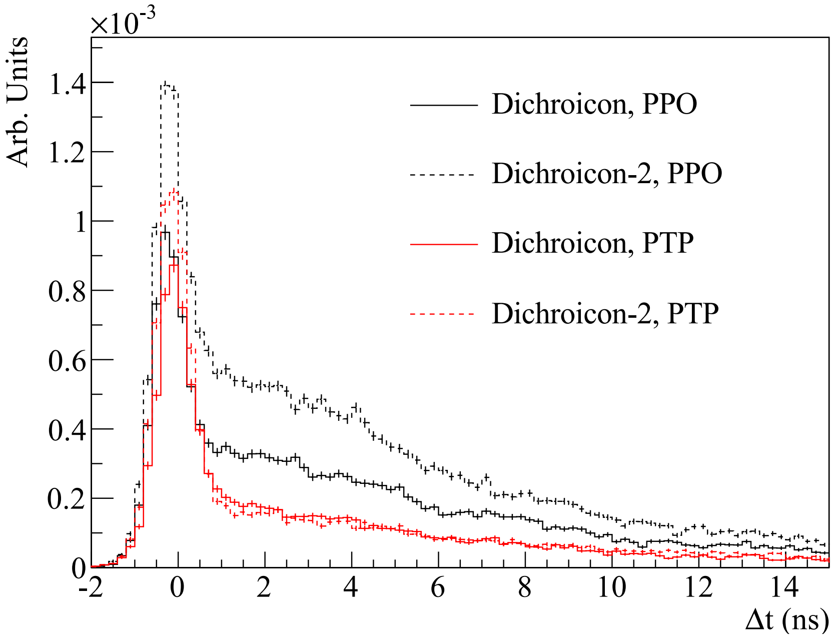

Figure 8.37: The data for the R7600-U20 at the aperture of the dichroicon using LAB+PPO and LAB+PTP targets, compared for the dichroicon and the dichroicon-2.

Scintillator	Dichroicon	C^*_{NORM} (1/m^2)	P	R	Q_{TOT} (10^6 pC)
LAB+PPO	1	8.71 ± 0.64	84.4 ± 1.5	7.35 ± 0.56	9.47 ± 0.56
LAB+PTP	1	8.94 ± 0.67	88.2 ± 1.6	7.89 ± 0.61	7.03 ± 0.30
LAB+PPO	2	12.38 ± 0.93	81.2 ± 1.9	10.05 ± 0.79	9.03 ± 0.54
LAB+PTP	2	12.54 ± 0.92	92.0 ± 1.4	11.54 ± 0.86	7.12 ± 0.33

Table 8.12: Comparison between LAB+PPO and LAB+PTP and the two different dichroicons for the R7600-U20 aperture PMT. For both scintillator cocktails, using the dichroicon-2 with shorter wavelength pass filter increased R.

LAB+PPO data decreases as there is more scintillation light leakage through the shorter wavelength central filter. However, there is no increase in the scintillation light leakage for the LAB+PTP data, so the purity of the selection increases, as there is more Cherenkov light but the same amount of scintillation light in the prompt window. Overall, the value of R for this setup with LAB+PTP and the dichroicon-2 reaches 12.54, which is the largest value for any setup tested.

In general the dichroicon-2 performs better than the dichroicon by collecting more Cherenkov light at the aperture with no significant decrease in the purity of the Cherenkov selection. This is in part due to the narrow TTS of the R7600-U20 which allows for excellent separation, regardless of the increase in the scintillation leakage for LAB+PPO. It is also not unexpected that the performance is improved by replacing some of the filters – the choice of the filters for the original dichroicon was not optimized. Overall, the dichroicon-2 measurements demonstrate that small, simple changes to the dichroicon can yield performance improvements, which will be further investigated in future studies.

The data for the R1408 PMT is compared in Table 8.12, which shows the total amount of scintillation light detected in the Q_{TOT} column. In both cases the changes to the dichroicon do not impact the total amount of light collected by the R1408 PMT. This is expected as only a very small fraction of the scintillation light is above 450 nm, so replacing half of the shortpass filters has very little impact on the total scintillation light transmitted through the dichroicon.

The disadvantage of using LAB+PTP can be understood by comparing the Q_{TOT} column between LAB+PPO and LAB+PTP. It is clear that when LAB+PTP is used the total collected scintillation light at the R1408 is about 75% of that for LAB+PPO. Part of this effect comes from the QE of the R1408, shown in Figure 8.9. However the QE of the R1408 is fairly flat across the emission spectra of LAB+PPO and LAB+PTP and does not explain the majority of the effect. The amount of light collected at the R7600-U20 trigger PMT is a good indicator of the total light yield of the scintillator. This comparison is again done by integrating the charge spectra, and shows that the LAB+PTP light yield is about 20% lower than the LAB+PPO light yield. This is consistent with the lower amount of collected scintillation light for the R1408 PMT, indicating that no additional scintillation light is lost in the setup by using LAB+PTP. Rather, it is simply the intrinsic light output of the scintillator which appears to be lower. This is not unexpected, as LAB+PPO is a very popular liquid scintillator specifically for having a very high light yield.

8.4.5 Large-Scale Detector Simulation

The `Chroma` model for the dichroicon can be used to simulate large-scale detectors and evaluate the efficacy of dichroicons. A basic model consisting of a 1-kT right cylinder active volume of LAB+PPO surrounded by 13,350 dichroicons is shown in Figure 8.38. The dichroicon model is the same as shown in Figure 8.30. A single 100 MeV electron event is shown in Figure 8.39 where dichroicons are colored blue if only a short-wavelength hit was detected, red if only a long-wavelength hit is detected, and magenta if both a short and long-wavelength hit was detected. This shows a clear Cherenkov ring on the long-wavelength PMTs despite every short-wavelength PMT being hit with scintillation photons, illustrating the usefulness of the dichroicon detection scheme. Using this simulation model, the dichroicon can be optimized for maximal physics performance in large-scale detectors, which is the focus of future studies.

The dichroicon provides a way in principle to create truly hybrid Cherenkov/scintillation detectors, in which a very broad range of physics can be done. While there are other ways to observe both Cherenkov and scintillation light in a liquid scintillator detector, the dichroicon approach has the advantage of allowing high scintillation light yield—important for low-energy physics—while retaining the fast timing of the scintillator and Cherenkov light, and with high purity of the latter. Monolithic neutrino detectors are not the only possible application: dichroicons can also be used with segmented scintillation detectors if there is an interest in observing Cherenkov light distinct from scintillation light. In any detector in which detection area is limited—either at the front-face of a detector segment or on the

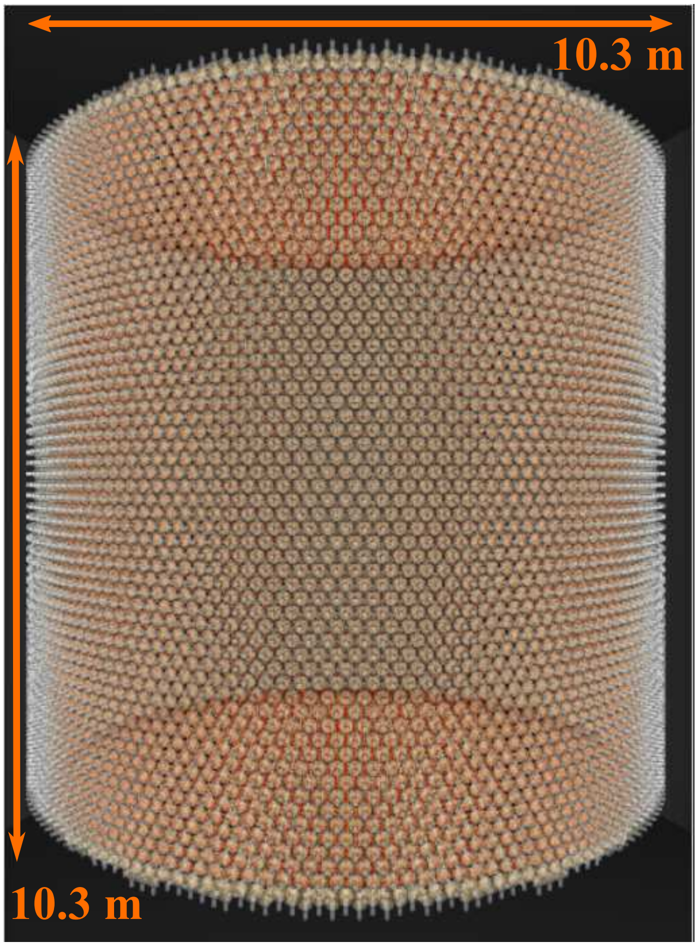

Figure 8.38: A visualization of a 1-kT right cylinder active volume of LAB+PPO instrumented with 13,350 dichroicons produced by `Chroma`. This figure was created by B. Land.

walls of a monolithic detector—the dichroicon provides an effective way of sorting photons by wavelength.

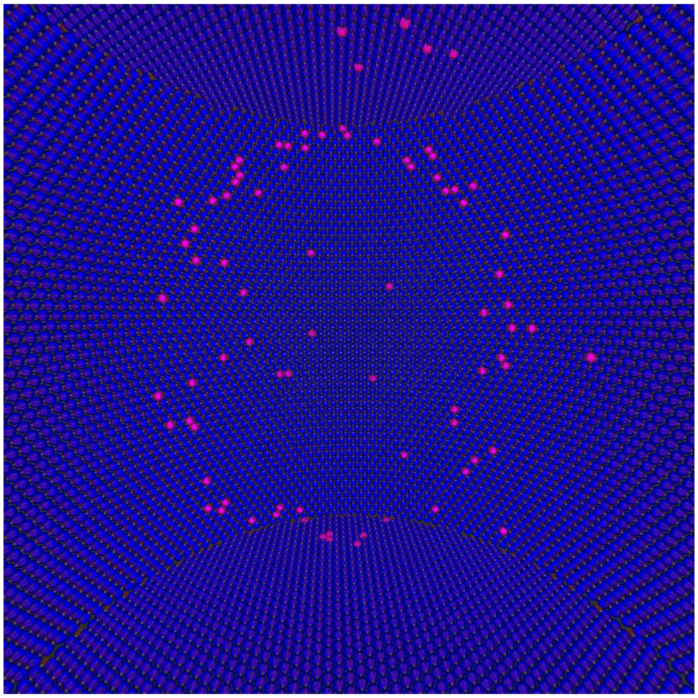

Figure 8.39: A `Chroma` event display showing a single 100 MeV electron event in LAB+PPO. Dichroicons are colored blue if a short-wavelength hit was detected, or red if a long-wavelength hit was detected. Both a long- and short-wavelength hit results in a magenta dichroicon. Despite all dichroicons detecting many short-wavelength scintillation photons, a clear long-wavelength Cherenkov ring can be seen in magenta. Every dichroicon detects at least one short-wavelength photon so there are no dichroicons colored red in this image. No selection criterion was used when choosing this event and most events appear similar during hand-scanning of the simulated data. This figure was created by B. Land.

Chapter 9

Conclusion

Historically, antineutrinos from nuclear reactors have played a crucial role in our understanding of these elusive particles. In this dissertation I presented an analysis designed to search for reactor antineutrinos using the SNO+ detector. Specically, 190.3 days of SNO+ water-fill data are considered. In the analysis, three primary backgrounds are studied: accidentals, (α, n) interactions, and atmospheric neutrino interactions. A technique for distinguishing IBD interactions from accidental backgrounds provides excellent rejection of accidental coincidences. The atmospheric neutrino and (α, n) backgrounds are carefully estimated using the simulation and the predictions for all three backgrounds are checked in carefully chosen sidebands. Furthermore, an examination of the AmBe source yielded an important calibration of the neutron detection efficiency, requiring detailed understanding of the trigger efficiency. One event is observed in the unblinded livetime, consistent with the signal plus background expectation. An expected signal sensitivity is estimated at slightly more than 2σ, which will be verified in the near future when the full livetime becomes unblinded. Additionally, future efforts to include the full $\sim$340 day livetime of the SNO+ water dataset, is expected to result in the first ever detection of reactor antineutrinos in a water Cherenkov detector. That effort will build upon the results presented in this thesis and will have implications for future water-based antineutrino detectors, such as WATCHMAN.

The second part of this thesis focuses on R&D efforts, at bench-top scales, on state-of-the-art PMTs and liquid scintillators. Future large-scale neutrino detectors searching for $0\nu\beta\beta$, low energy solar neutrinos, geoneutrinos, and the DSNB will utilize such advanced technologies. I characterized several large-area PMTs, most notably the R5912-MOD, which demonstrated excellent timing and charge responses. Bench-top measurements of these PMTs is a crucial step to provide input into simulations of future detectors. In addition, LAB-based liquid scintillators have become the target of choice for many current and future large-scale monolothic detectors. The careful understanding of the optical properties is crucial for predicting sensitivities. Measurements of the liquid scintillator's α and β time-profiles and light yield are presented in this thesis, and are valuable inputs into scintillator models used to extrapolate to large scales. Predictions of energy resolution and background rejection rely on these bench-top scale measurements.

Finally, a new instrument called the dichroicon is presented in this thesis. The dichroicon was inspired, in part, by background modeling for the SNO+ $0\nu\beta\beta$ search, which indicates that the dominant background will be ^{8}B solar neutrino interactions. To reject these events, knowledge of the particle's direction is critical, but extremely difficult to determine in detectors where the Cherenkov photons are swamped by the scintillation component. One method to seperate the Cherenkov and scintillation photons – which would provide the ability to reconstruct the particle direction – is using wavelength information about the detected photons. The dichroicon provides a method for spectrally sorting photons towards two (or more) different photodetectors (or pixels), where the short- and long- wavelength light can be detected separately. Bench-top measurements of the dichroicon show exciting promise for this device, and are used as input to generate large-scale simulations of detectors filled with these dichroic concentrators.

Appendices

Appendix A

Data-Cleaning Cuts

The following data-cleaning cuts are applied to the data and events that fail any of the cuts are removed. More details regarding these cuts is given in (207).

Zero zero: Due to an issue with the CMOS chip, all events where the bottom two digits of the GTID are 0 are discarded.

Crate isotropy: Events that are anisotropic in crate space, specifically with more than 70% of hits in a single crate, are removed.

Fitterless time spread: Identifies flasher events by using the time-spread of the hit PMTs, as many of the photons in these events travel across the entire detector.

Flasher geometry: Identifies flasher events using the typical flasher hit distribution.

In-time-channel time spread: Tags events with fewer than 60% of hits within a 93 ns sliding window.

Junk: Orphaned events, events with an ECA flag set, or events where a PMT hit shows up multiple times.

Muon: Muons are tagged using their large nhit (≥ 150) and the number of OWL PMTs that have hits (≥ 5).

Neck: Removes events originating in the neck using the neck PMTs. If two or more neck PMTs fire in any given event the event is removed.

OWL: Events with 3 or more OWL PMTs firing are removed.

Charge cluster: Events with a high charge hit surrounded by cross-talk hits are removed.

Charge/nhit Identifies shark-fin events using the large fraction of the total charge deposited in a small fraction of the total hit PMTs.

Charge/time Events with a high charge PMT hit at very early times, which removes shark-fins and flashers.

Ring of fire Removes events caused by noise with a particular hit pattern in crate space, which was more common during SNO.

Muon follower Events for 20 seconds after a muon are removed.

CAEN The peak and integral of the ESUMH CAEN trace falls within a specified range, tuned using ^{16}N data. The CAEN cut range is shown in Figure A.1. This cut is the predecessor of the SNO AMB cut.

Polling A specific, previously unused, bit in the MTCD word is used to mark events recorded during times where rate or base current polling is ongoing, which is used for data-quality monitoring. Events with this bit set are removed.

Retrigger Events occurs less than 3 μs after the previous event.

Burst Four second time-periods with more than six events with nhit greater than 40 are removed.

Missed muon follower Coincidences between a 60 nhit event followed within 1 ms by a 20 nhit event are removed, as well as all the events between the two events.

Missing CAEN data Events that do not have CAEN data, due to readout issues with the CAEN. This cut is included because the CAEN cut described above cannot be applied to these events.

Pedestal Pedestal events from the nhit monitor and ping crates are removed using the pedestal bit in the trigger word. Additionally, events within one second of any pedestal event are removed because of possible noise created by the pedestal signal. The pedestals are sent in short bursts so this cut does not remove significant amounts of livetime.

Atmospheric High nhit events (≥ 200) with less than 5 OWL hits followed by an event within 20 μs that has an nhit ≥ 100 are tagged as potentially atmospheric charged-current neutrino interactions with Michel electron followers. These events and events within 20 seconds are removed.

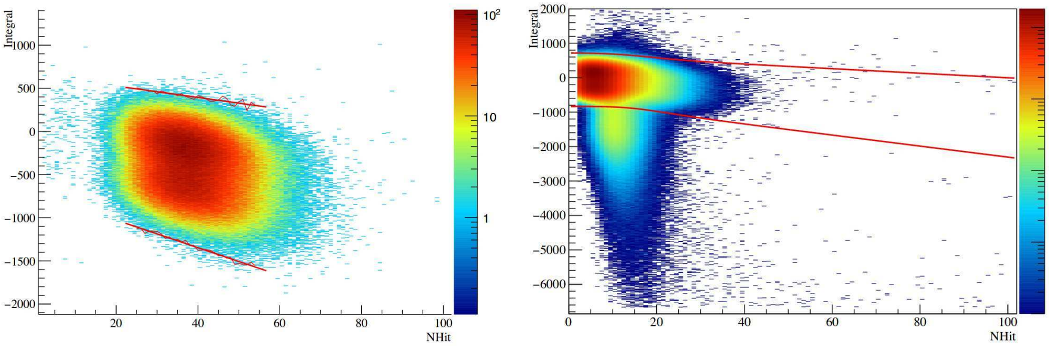

Figure A.1: (Left) The integral of the ESUM waveform for ^{16}N events as a function of nhit. The red lines show the range used by the data-cleaning cut. (Right) The integral of the ESUM waveform as a function of nhit for a physics run, which includes flashers and shark-fins that are removed by this cut. This figure is from Eric Marzec's thesis (20).

Appendix B

Electronics Calibration

There are several stages to ensuring the electronics used to read out the PMT pulses are calibrated. Specifically, the settings on the hardware must be carefully tuned in order to perform optimally. Additionally, the baseline of the charge and TAC ADC, the noise in the electronics, and the dark noise rate of the PMTs must be measured in-situ so that they can be simulated correctly. Finally, the PMT gains must be measured in order to perform a discriminator 'time-walk' correction and to understand the efficiency at which the PMT pulses cross threshold.

B.0.1 ECAL

A suite of electronics calibrations and tests (called an ECAL) is run through specialized software, called `penn_daq2`, during detector running with no high voltage on the PMTs. ECALs are run during periods of maintenance, once every couple of months. The ECAL is primarily responsible for measuring and setting the following hardware settings:

1. The location of the baseline and the size of the noise on the baseline of the discriminator, in ADC counts. The threshold is then set 1 ADC count above the noise of the baseline to ensure the discriminator is rarely triggering on electronics noise.

2. The baseline or ‘pedestal’ values of the QHS, QHL, QLX, and TAC ADCs is set by adjusting DACs on the FEC. The charge values are set around 600 ADC counts, which allows for a large dynamic range on the 12-bit ADCs.
3. The length in time of the windows for QHS, QHL, and QLX to integrate over. This is set by adjusting DAC values on the FEC.
4. The length of the GTValid window for each channel, again set by tuning a DAC on the FEC. The GTValid lengths are then measured for each channel and ensured to be less than the length of lockout.

In addition to tuning these hardware settings, the ECAL is responsible for running a suite of tests that ensure critical components of the hardware are functioning properly. The following are tested during an ECAL:

1. Each channel is counting GTs correctly and rolls over properly at the 16 bit rollover.
2. Each channel reads out the correct number of hits for a pedestal run with a fixed number of pedestals sent to each channel.
3. The FEC voltages are all at expected values.
4. The XL3 can read from and write to a set of register on the FEC.
5. Every channel has a working N20 and N100 signal.
6. The channels are responding properly to the injection of real charge, mimicking PMT pulses (this is done using a pulser on the PMTIC)

B.0.2 ECA

Electronics calibrations to measure the charge and time pedestals at high voltage are called ECAs. The ECAs consists of two types of runs: PDST and TSLP. The former uses pedestals to measure the zero-level charges on each channel in the detector, which should already be set from the ECAL. The ECA records these baseline values in a database so that the they can be removed when converting from ADC counts to charge deposited. The TSLP run also uses pedestals but delays the global trigger in order to scan the ADC-to-ns conversion for the TAC. In addition to providing critical calibrations for the charges and TACs, the ECA also identifies broken channels either not reading out hits or with bad baseline values. Both types of ECAs are run approximately weekly in order to maintain properly calibrated channels, the baseline for which can drift slowly with time.

B.0.3 PCA

The PMT calibrations (PCA) is a direct calibration of the PMT response using a light source, either the deployed laserball source (described in Section D.1.1) or using TELLIE LEDs (described in D.1.2). Using one of these two sources, light is injected, close to isotropically, into the detector. The intensity of the light is tuned carefully so that the PMTs are primarily detecting single photoelectrons. Particularly because the high voltage is provided from a single supply for the entire 512 PMTs in the crate, there is variation in the values of the PMT gains across the detector. The PCA runs provides a high statistics sample of SPE data in which the peaks of the charge distributions are measured and used to determine the gain at which each of the PMTs is running. The PCA additionally measures the channel-to-channel time delay, due to different path lengths along the backplane, and the discriminator time-walk for each channel. The time-walk effect is the dependence of the time measured on the size of the PMT pulse due to having constant discriminator thresholds, shown schematically in Figure B.1. PCAs are run fairly sporadically once every couple of months.

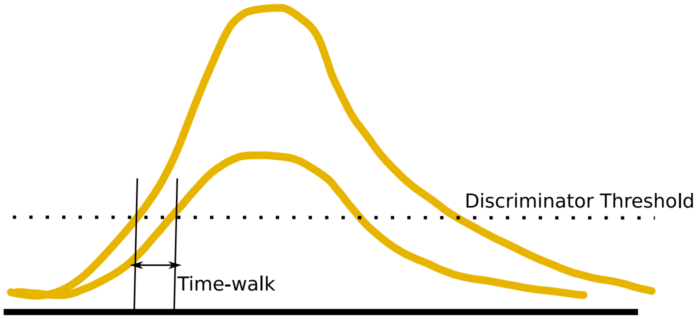

Figure B.1: A schematic demonstrating the time-walk effect for two potential PMT pulses. Due to the constant discriminator threshold, the time measured depends on the size of the pulse.

B.0.4 Dark Noise

PMTs produce a small amount of current when they are operating in a complete dark state. This output is typically referred to as 'dark current' or 'dark noise'. Dark noise is primarily caused by the thermionic emission at the photocathode. The dark noise rate of the R1408 PMTs is about 500 Hz at 16°C, the temperature that they are operating at in the cavity. Thus, a typical 400 ns event will have about two PMT hits due to dark noise. It is critical that a measurement of the PMT dark noise is performed in order to include this effect in our simulation. To do so, PGT is run at 50 Hz during physics runs and the hits in those events are assumed to originate from dark noise, as pile-up with light in the detector is rare. By counting the hits in these events a run-by-run estimate of the dark noise rate in each PMT is made and used in our simulation of the detector.

Appendix C

Hit Cleaning

Hit-cleaning refers to the rejection of PMT hits created by anything other than light generating a photoelectron at the photocathode. There are three types of PMT hits that were considered for cleaning: cross-talk, afterpulsing, and dark noise. The last of these is not discussed in this section as there is very little discriminatory power to reject these hits. Instead, we measure the noise rates for each PMT individually, which is provided as input into the run-by-run simulations. Additionally, hit-cleaning is responsible for removing and/or keeping track of channels that are either dropping data due to very high discriminator-crossing rates or channels that are out-of-sync. Hit-cleaning is important to remove these PMT hits from being considered by the event reconstruction algorithms, particualrly because effects like cross-talk are not simulated.

C.0.1 Cross-talk Cut

Due to the capacitive coupling between nearby channels, a high charge PMT signal on one channel can cause signals on one or more adjacent channels to cross threshold. The signals are bipolar in nature and thus, when integrated, have charge values around pedestal. Additionally, due to the nature of the signal, it will cross threshold slightly after the high charge hit.

The details of this cut largely follow the description in M. Dunford's thesis (208). For any hit with charge above a QHS of 50 (usually corresponding to two or more PE), the neighboring six channels on either side of the hit channel and on both neighboring slots are scanned. Any hits on these 36 channels are checked as being potentially cross-talk hits using the charge and time of the hit. A distribution showing the calibrated time-difference vs QHS for hits after a high charge hit is shown in Figure C.1 after ECA but before PCA calibration (left) and after PCA calibration (right). The data used for this plot is standard physics data. The time-walk correction and low-charge rejection from PCA calibration largely cleans up the cross-talk, but a significant contribution from cross-talk can still be identified at QHS around 5 ADC counts and Δt around 8 ns.

The cross-talk cut is applied before the PCA because the PCA makes explicit assumptions about the shape of the PMT pulse in the time-walk correction, which is not valid for the cross-talk hits. Thus, as is clear in Figure C.1, applying PCA actually pulls the cross-talk hits closer in time to the in-time hits. The adjacent hits with ECA calibrated Δt after the high-charge hit of 9 to 25 ns and QHS values between -30 to 10 ADC counts are rejected by the cross-talk cut. In the implementation of the PMT calibration selector, which selects PMTs to be used in the fit methods, the cross-talk cut is applied. Additionally, a version of the nhits, called 'nhits cleaned', contains only the hits that pass all of the various calibrations, including the cross-talk cut.

There is no simulation of the cross-talk due to the ability to cleanly cut the cross-talk without sacrificing real hits. Because the hits are adjacent to a high charge hit, out-of-time, and low in charge they can be easily distinguished from Cherenkov light, which comes in-time and should be large SPE, corresponding to a QHS of about 25 ADC counts. However, with scintillation light, where the hits are expected to come out-of-time, the cross-talk cut might need to be returned. By applying the cross-talk cut to simulated events generated in a scintillator filled detector, it was found that the sacrifice of the cross-talk cut was

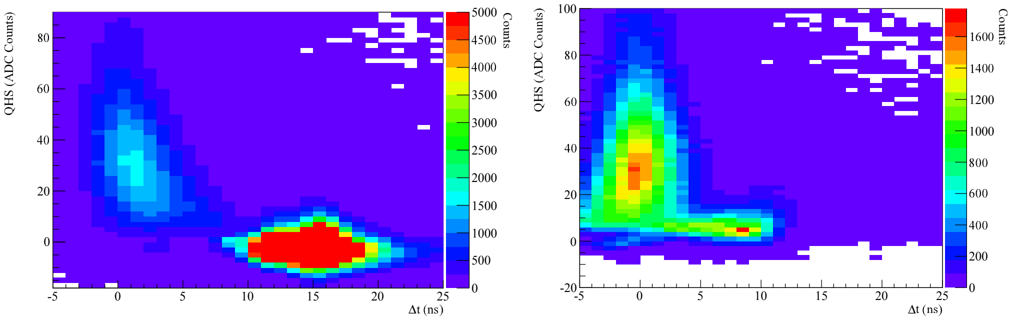

Figure C.1: The time-difference between a high charge hit (above a QHS of 50 ADC counts) and the hits in adjacent channels plotted against the QHS of the hit in the adjacent channel. There is clearly a large contribution from cross-talk at QHS around pedestal and Δt > 5 ns. The left plot shows hits after ECA calibration but before PCA and the right plots shows after PCA calibration. The cross-talk cut is applied after ECA but before PCA.

extremely small – less than 0.01% of real hits were cut. This still has no been validated by looking at the cross-talk cut in scintillator data, but the simulated results imply this cut will not need retuning.

C.0.2 PMT Afterpulsing

PMT afterpulsing, distinct from late pulsing, is caused by the ionization of residual gases in the PMT by the photoelectrons (209). Nobel gases such as Helium can diffuse through the PMT glass at a rate proportional to the pressure different inside and outside of the gas. Any photoelectron being accelerated between the photocathode and first dynode might interact and ionize the gas. These positive ions will drift back toward the photocathode in the strong electric field, where they will deposit their energy and release more electrons. The PMT pulses caused by these secondary electrons are called afterpulses.

The afterpulsing for the R1408 PMT was never characterized for the SNO PMTs due to the relatively low afterpulsing rate and overall light yield of the SNO detector. Additionally, afterpulsing is more important in time-correlated searches, where a primary and secondary event might occur within about 10 μs of one another. In this case, you might expect the secondary event to contain afterpulse hits from the first event, depending on the number of PMTs hit in the first event.

Figures C.2 and C.3 shows the results of a bench-top measurement performed in the Penn darkbox to measure the afterpulsing rate in a R1408 PMT. The PMT is exposed to a 400 nm LED at different intensities, and for each LED pulse, a 20 μs PMT waveform is digitized. Figure C.2 (left) shows the intensity of the initial pulse in terms of the charge of the PMT pulse in pC. For this PMT a SPE deposits about 1 pC of charge, making it easy to roughly convert between pC and number of PE. Figure C.2 (right) shows the corresponding time of the after-pulses for the various intensity beams. As expected, the number of after-pulses scales with number of detected photoelectrons – more electrons means a higher likelihood of ionizing the residual gas. Indeed, the relationship appears to be linear, as shown in Figure C.3. Additional tests show that the afterpulsing is affected only weakly by the high voltage the PMT is operating at.

An attempt to confirm the bench-top measurements of the afterpulsing model is made by looking at the detector response during an early portion of the water fill. This is particularly interesting because the PMTs at Penn have been sitting in a different environment than the ones in the detector and they might exhibit different afterpulsing behavior. Figure C.4 shows the time between a PMT that was hit in two events for events that occurs within 50 μs of one another. This preliminary look confirms the Δt distribution found, showing clear peaks at one and five μs. There is also an effect from what is now understood as ringing in the trigger system that causes the large spike around 1 μs, but the broader afterpulsing peak can still be identified. The primary contribution to the flat background is just from the dark-rate of the PMTs, which shows no correlation in time. Note that this data was not taken as a dedicated measurement of the afterpulsing, and it relied on random triggers occurring at the right time in order to capture the afterpulse. A dedicated run that used

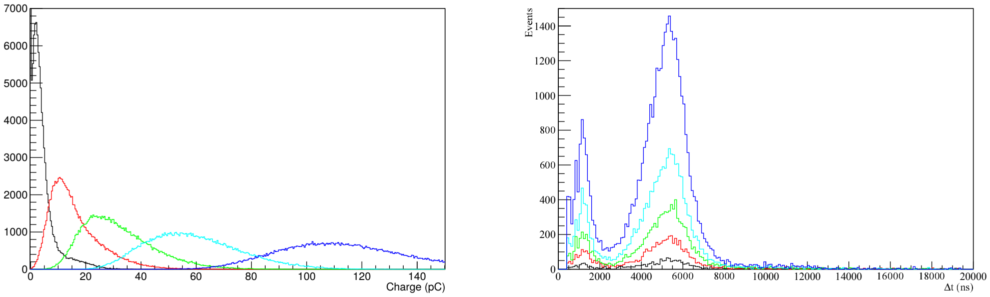

Figure C.2: (Left) The deposited charge at the R1408 PMT for an LED pointed directly at the PMT at various intensities. The LED intensity is adjusted in order to understand the afterpulsing rate as a function of number of electrons. The SPE peak for this PMT is around 1 pC, making it easy to roughly convert to number of PE. (Right) The time between the initial pulse on the PMT from the LED light and any afterpulses. There are clear afterpulsing peaks at around one and five μs. The number of afterpulses scales with number of prompt PEs.

a dedicated light source, just as TELLIE, and the MTC/A+ forced retriggering, described in Section 4.5.2, would drastically improve this measurement.

Based on these afterpulsing results, a model in the SNO+ simulation software is developed. The afterpulsing model works by, for each PMT hit, calculating a probability that an afterpulse is created based on the charge deposited b the hit and the probability of an afterpulse given in Figure C.3. The afterpulsing hit is assigned a MC time that is randomly sampled from the expected Δt distribution.

With water in the detector, there are very few retrigger events or time-correlated events, and thus we expect that these hits are, for the most part, never built into triggered events. However, for very high energy deposits, such as for cosmic muons, there is a significant afterpulsing probability so that just the afterpulses might trigger the detector. This is likely an important consideration in the search for neutron followers after muons. However, two additional facts make afterpulses – even for these events – relatively easy to reject. First, the hits must come within 5 μs and occur on the same PMT. Additionally, the hit must fall

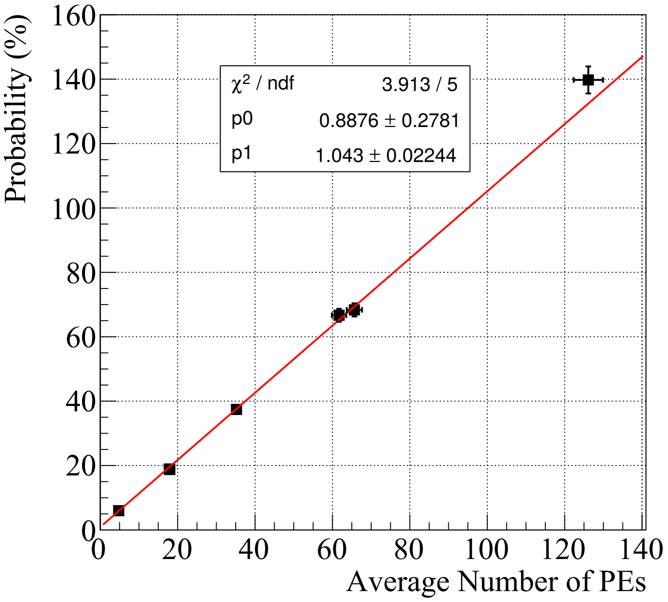

Figure C.3: The afterpulsing probability as a function of the average number of photoelectrons. The number of photoelectrons detected was tuned using a variable intensity LED, as shown in Figure C.2 (left). A probability over 100% indicates on average more than one afterpulse was detected for each initial pulse. The fit indicates that there is about a 0.9% chance of each individual PE creating an afterpulse.

within a very narrow prompt window of about 10 ns in order for it to be considered in the reconstruction of the follower event. Given the broad Δt distribution for the afterpulses, this is very unlikely to occur often. Additionally, given issues with ringing in the trigger system, often the first 10 μs or so is rejected for these types of analyses, making afterpulsing negligible.

Based on these afterpulsing results, a model in the SNO+ simulation software is developed. The afterpulsing model works by, for each PMT hit, calculating a probability that an afterpulse is created based on the charge deposited b the hit and the probability of an

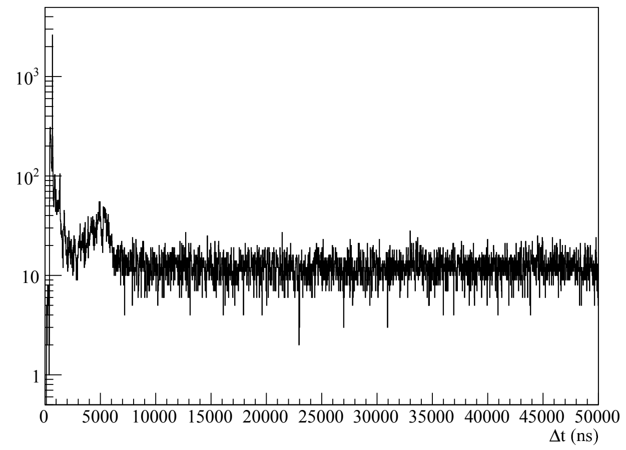

Figure C.4: The afterpulsing distribution for detector data showing the afterpulsing peaks at one and five μs, consistent with the benchtop data shown in C.2.

afterpulse given in Figure C.3. The afterpulsing hit is assigned a MC time that is randomly sampled from the expected Δt distribution.

With water in the detector, there are very few retrigger events or time-correlated events, and thus we expect that these hits are, for the most part, never built into triggered events. However, for very high energy deposits, such as for cosmic muons, there is a significant afterpulsing probability so that just the afterpulses might trigger the detector. This is likely an important consideration in the search for neutron followers after muons. However, two additional facts make afterpulses – even for these events – relatively easy to reject. First, the hits must come within 5 μs and occur on the same PMT. Additionally, the hit must fall within a very narrow prompt window of about 10 ns in order for it to be considered in the reconstruction of the follower event. Given the broad Δt distribution for the afterpulses, this is very unlikely to occur often. Additionally, given issues with ringing in the trigger system, often the first 10 μs or so is rejected for these types of analyses, making afterpulsing negligible.

In scintillator it is not so clear that afterpulsing will be a negligible effect. With the much larger light yields, higher event rate, and expected time-correlated events, such as ^{212}BiPo and ^{214}BiPo events, afterpulsing might play an important role. This was looked as using the afterpulsing generator for generated ^{212}BiPo events. The ^{212} BiPo events are characterized by a prompt β decay from the Bismuth, that deposits as much as 2 MeV of energy, corresponding to about 1200 PMT hits. This is followed by a trigger event caused by a delayed α – the half-life of ^{212}Po is about 300 ns – in which you might expect afterpulsing hits caused by the photoelectrons in the prompt event. Even in this scenario, according to the simulation, the impact of afterpulsing is very small. Given 1200 PE in the prompt event, the number of expected afterpulses is 12, which are distributed over about a 5 μs window. The trigger window is 400 ns, which means that only about 1 afterpulsing hit is expected in any given trigger window. Given that the α event will have several hundred PMT hits from scintillation light, this afterpulsing hit will have a small effect on the overall event.

C.0.3 Channel Flags

The 'channel flags' refers to flags in the PMT data bundle that flag various errors with a PMT hit. Figure C.5 shows the bit map for the PMT bundle, specifically bits 30 and 31 in word one and bits 28 and 31 in word two mark channel errors in the PMT hit. Note the CGT Sync 16 bit in word one is not used due to an issue with the hardware setting that bit in the PMT bundle.

The missed count flag gets set when the channel drops hits. This occurs when the 16 analog memory cells in the CMOS chip get filled before the channel can get read out. Channels running at very high rates (often due to a broken discriminator or a discriminator threshold set too close to the noise) often drop hits. Anytime this happens a flag in the PMT bundle is set and as long as the channel is full no new data is readout. Because no

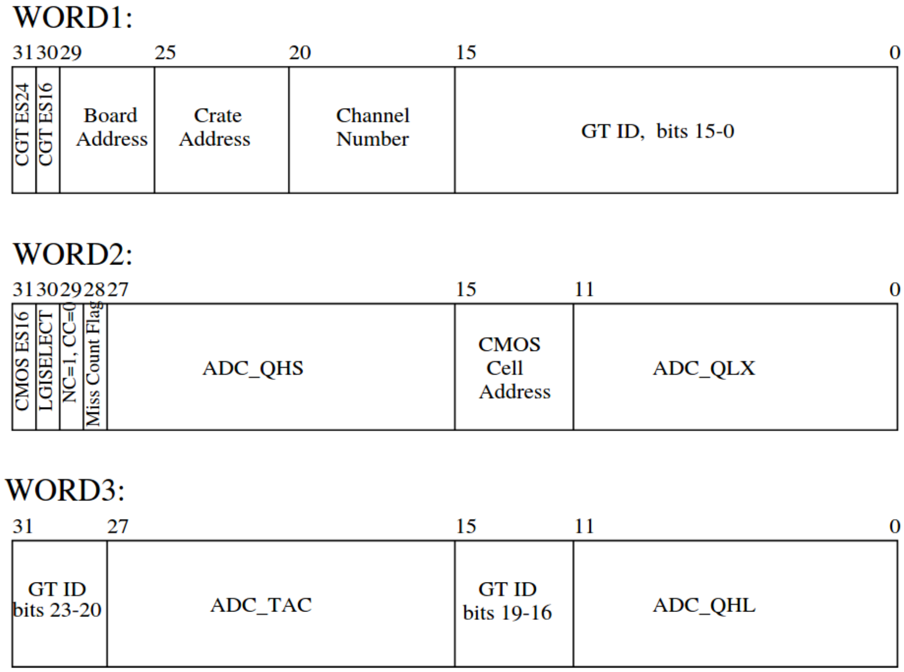
WORD1:
31 30 29 25 20 15 0
CGT ES24
CGT ES16
Board Address
Crate Address
Channel Number
GT ID, bits 15-0
WORD2:
31 30 29 28 27 15 11 0
CMOS ES16
LGISELECT
NC=1, CC=0
Miss Count Flag
ADC_QHS
CMOS Cell Address
ADC_QLX
WORD3:
31 27 15 11 0
GT ID bits 23-20
ADC_TAC
GT ID bits 19-16
ADC_QHL

Figure C.5: The PMT data bundle bitmap.

new data can be readout, there is no way to tell which events would have lost hits. Thus any channel that has a missed count flag set for the run is removed for the entire run.

The sync 16 and sync 24 flags are set when the channel is out-of-sync with GTID the trigger system. Each individual channel has a GTID counter that is incremented every time they receive a GT from the trigger system. The GTID for each channel must match the GTID recorded by the trigger in order for the channel's data to be written in the correct event. Every 2^{16} GTIDs the trigger system sends a 'SYNC_16' signal to all channels in thee detector that tells the channels to check their GTID and ensure they are in-sync with the trigger system. For channels are are not at an increment of 2^{16}, and thus out of sync with the trigger system, the sync 16 flag gets set in the PMT bundle. This happens similarly for increments of 2^{24} GTIDs, which is intended to add robustness to the system. If either the out-of-sync flags are set in a run, that channel is removed for the entire run. Note that at each of these SYNC_16 and SYNC_24 signals, the channel is forced back into sync with the trigger system – but at that time it is too late for the last 2^{16} GTIDs.

The out-of-sync check is slightly more complicated in nature than the missed count checks, because out-of-sync nature of a channel is only checked at SYNC_16 and SYNC_24s, which do not respect run-boundaries. Thus, a system running on the nearline keeps track of the time at which any channel goes out of sync and ensures the channel is discarded. This system is able to keep track of out-of-sync channels across run boundaries and throw away channels for the previous run if the channel is out of sync at the first SYNC_16 in the run.

Appendix D

Optical Calibration

Understanding the optics of the SNO+ detector is critical to properly simulate the light propogation through the various components. The optical measurements are performed using the laserball source, described in Section D.1.1.

D.1 Optical Calibration Sources

D.1.1 Laserball

The laserball is an optical calibration source, developed for SNO and used in SNO+, described in detail in (210). The laserball is a light diffusing sphere that consists of a uniform distribution of scattering spheres. A pulsed laser is connected to the source via a fiber running down the umbilical, where it enters the source and is diffused isotropically. Various laser dyes provide methods for changing the wavelength distribution of the emitted light. The laserball provides in-site calibrations of the PMT and electronics response. Specifically, the PMT calibrations described in Section B.0.3 and the trigger efficiency measurement described in Section 4.7 use laserball calibration data.

D.1.2 TELLIE

TELLIE is an optical timing calibration system that utilizes 92 optical fibers mounted on the PSUP and pointed toward the center of the detector. LEDs of various wavelengths and associated electronics are located on deck, which allows them to be easily replaced in the event of a failure. The TELLIE system is designed to provide an source of well-understood light that can be used to calibrate each PMTs time and charge response. The primary advantage over the laserball is that TELLIE does not require a source to be deployed in the detector, which is critical for scintillator phase, where internal source deployments are associated with cleanliness concerns. The trigger efficiency measurements described in 4.7 use low-intensity TELLIE runs. A full description of the TELLIE system can be found in (211), (93).

D.2 Results

Understanding the optics of the SNO+ detector is critical to properly simulate the light propogation through the various components. The optical measurements are performed using laserball data at various wavelengths, which provides a source of near-isotropic light. Using the laserball deployed at different locations, the optical response of the internal and external water, the acrylic, and the PMTs can be characterized. Specifically, the attenuation length, defined for SNO+ as the absorption length summed with the Raleigh scattering length, can be measured for the water and acrylic. Additionally, the PMT response as a function of incident angle is extracted. Full details for the SNO+ optical calibration are given in (29). The internal and external water absorption is shown in Figure D.1 and the acrylic attenutation is shown in Figure D.2. With this optical calibration, the nhit scale across the detector matches very nicely between data and MC, shown in Figure D.3 for a x-axis scan of the ^{16}N source.

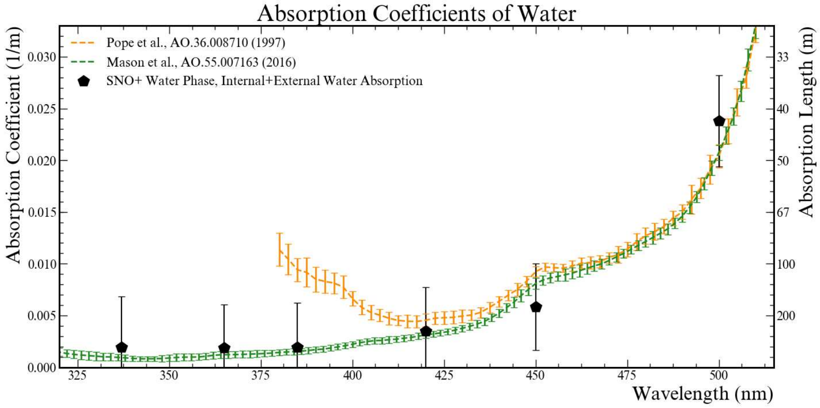

Figure D.1: Internal and external water absorption coefficients/lengths (left/right vertical axis), in the black data points (29).

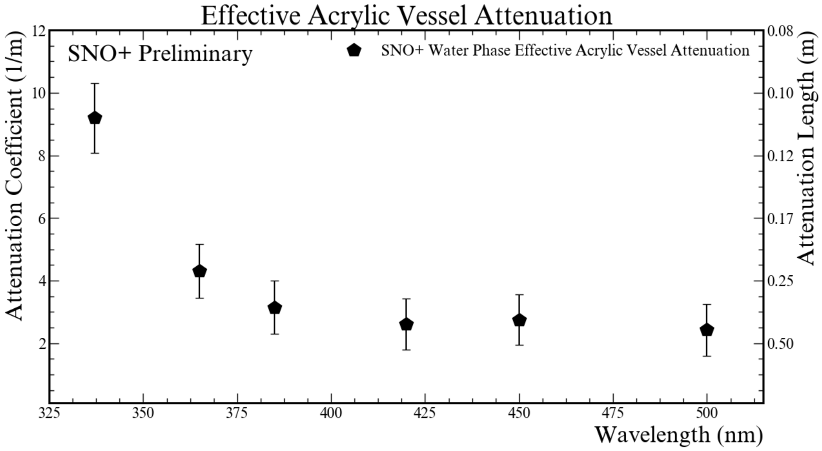

Figure D.2: The acrylic vessel attenuation coefficiency/lengths (left/right vertical axis) (29).

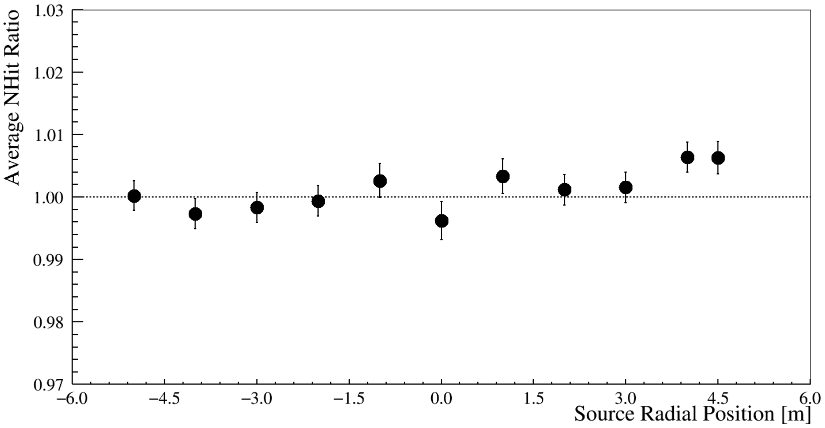

Figure D.3: The nhits ratio between data and MC for ^{16}N source positions (29). After optical calibrations, the number of PMTs hit in data and MC agree well regardless of event position in the detector.

Appendix E

Reactor Complexes

A list of all reactor complexes and the distance from the complex to the SNO+ detector, ordered by distance from the detector is given in Table E.1.

Table E.1: All reactor complexes and the distance from the complex to the SNO+ detector, ordered by distance from the detector.

Complex	Distance (km)	Complex	Distance (km)	Complex	Distance (km)
BRUCE	241.590	PICKERING	340.217	DARLINGTON	349.122
GINNA	469.504	NINE MILE POINT	500.045	FITZPATRICK	500.602
PERRY	519.236	FERMI	527.363	POINT BEACH	552.238
DAVIS BESSE	562.532	PALISADES	615.010	BEAVER VALLEY	652.707
COOK	657.757	SUSQUEHANNA	722.660	THREE MILE ISLAND	789.614
DRESDEN	799.932	BYRON	807.761	BRAIDWOOD	809.041
INDIAN POINT	819.832	LIMERICK	829.294	LASALLE	834.235
PEACH BOTTOM	846.081	QUAD CITIES	898.247	HOPE CREEK	904.001
SALEM	904.419	SEABROOK	910.046	PRAIRIE ISLAND	914.226
MILLSTONE	923.421	OYSTER CREEK	931.634	CLINTON	932.845
DUANE ARNOLD	971.719	CALVERT CLIFFS	973.767	NORTH ANNA	974.748
PILGRIM	984.738	MONTICELLO	987.475	SURRY	1097.49
POINT LEPREAU	1154.11	CALLAWAY	1215.52	HARRIS	1216.44
MCGUIRE	1224.44	WATTS BAR	1242.05	CATAWBA	1266.22
SEQUOYAH	1288.69	OCONEE	1302.30	FORT CALHOUN	1308.75
ROBINSON	1340.83	COOPER	1346.82	SUMMER	1349.45
BROWNS FERRY	1395.84	BRUNSWICK	1412.83	VOGTLE	1477.64
WOLF CREEK	1496.29	ANO	1593.28	HATCH	1612.71
FARLEY	1720.83	GRAND GULF	1807.49	RIVER BEND	1947.99
WATERFORD	1989.48	COMANCHE PEAK	2106.71	ST. LUCIE	2114.20
TURKEY POINT	2322.57	SOUTH TEXAS	2336.89	COLUMBIA	2874.37
PALO VERDE	3013.82	LAGUNA VERDE	3239.64	DIABLO CANYON	3481.20
HUNTERSTON B	5023.40	TORNESS	5145.64	HEYSHAM A	5199.33
HEYSHAM B	5199.35	HARTLEPOOL A	5266.22	HINKLEY POINT B	5312.48
FLAMANVILLE	5462.71	SIZEWELL B	5527.80	DUNGENESS B	5555.28
PALUEL	5589.53	GRAVELINES	5614.90	PENLY	5615.92
BORSSELE	5679.42	CHINON B	5697.19	DOEL	5713.67
ALMARAZ	5736.70	BLAYAIS	5750.03	ST. LAURENT B	5752.13
CIVAUX	5765.37	RINGHALS	5797.84	DAMPIERRE	5804.54
CHOOZ B	5806.39	TIHANGE	5808.96	EMSLAND	5815.19
NOGENT	5819.43	BELLEVILLE	5836.44	BROKDORF	5844.84
TRILLO	5880.36	FORSMARK	5890.46	BILIBINO	5892.02
GOLFECH	5902.92	CATTENOM	5919.16	GROHNDE	5945.21
KOLA	5969.34	OLKILUOTO	5973.47	OSKARSHAMN	5993.01
PHILIPPSBURG	6046.41	ASCO	6047.71	ST. ALBAN	6057.74
COFRENTES	6061.00	BUGEY	6064.80	FESSENHEIM	6074.63
VANDELLOS	6079.82	NECKARWESTHEIM	6097.51	CRUAS	6100.90
MUEHLEBERG	6110.47	TRICASTIN	6116.57	LEIBSTADT	6124.95
GOESGEN	6126.51	BEZNAU	6130.25	GUNDREMMINGEN-B	6191.59
GUNDREMMINGEN-C	6191.61	LOVIISA	6210.85	ISAR	6284.72
LENINGRAD	6343.34	LENINGRAD-2	6344.99	TEMELIN	6357.17
DUKOVANY	6450.86	BOHUNICE	6562.43	KRSKO	6608.10
MOCHOVCE	6613.99	KALININ	6687.48	PAKS	6737.21
ROVNO	6764.72	KHMELNITSKI	6863.28	SMOLENSK	6875.83
KURSK	7136.62	KOZLODUY	7154.30	NOVOVORONEZH	7220.24
SOUTH UKRAINE	7223.32	NOVOVORONEZH 2	7299.19	CERNAVODA	7316.69
ZAPOROZHYE	7384.60	BELOYARSK	7516.92	BALAKOVO	7533.42
ROSTOV	7673.73	ANGRA	7908.32	EMBALSE	8169.71
ARMENIAN	8261.83	ATUCHA	8397.55	TOMARI	8478.79
HIGASHI DORI-1 (TOHOKU)	8588.80	ONAGAWA	8785.98	FUKUSHIMA-DAINI	8877.08
KASHIWAZAKI KARIWA	8938.59	TOKAI	8949.08	SHIKA	9015.93
TSURUGA	9131.47	MIHAMA	9132.43	HONGYANHE	9140.95
HAMAOKA	9149.18	OHI	9151.90	TAKAHAMA	9157.28
HANUL	9190.88	SHIMANE	9219.21	BUSHEHR	9272.40
SHIN-WOLSONG	9288.24	WOLSONG	9288.29	SHIN-KORI	9318.43
KORI	9319.84	HAIYANG	9370.21	HANBIT	9373.93
IKATA	9379.42	GENKAI	9434.97	SENDAI	9543.44
TIANWAN	9545.61	CHASNUPP	9591.44	QINSHAN	9819.13
FANGJIASHAN	9819.16	QINSHAN 3	9819.21	QINSHAN 2	9820.69
SANMEN	9898.43	KANUPP	9992.28	NINGDE	10057.6
CHINSHAN	10150.7	KUOSHENG	10154.9	RAJASTHAN	10170.7
FUQING	10175.7	KAKRAPAR	10358.0	MAANSHAN	10375.1
LING AO	10422.1	DAYA BAY	10422.6	TARAPUR	10430.3
TAISHAN	10482.6	YANGJIANG	10502.9	FANGCHENGGANG	10537.8
CHANGJIANG	10669.9	KAIGA	10753.5	MADRAS	10973.6
KOEBERG	11013.3	KUDANKULAM	11168.5		

Appendix F

The AmBe Source Run List

Table F.1 lists the AmBe calibration source runs that were selected for use in the analysis presented in Chapter 6. Figure F.1 shows the reconstructed promp event position across the entire run list. The source positions are clearly identified by the hot-spots in this figure.

Table F.1: The full run list of all AmBe deployed source runs used. The source positions are determined using the rope tensions by the source manipulator system.

Run number	x (mm)	y (mm)	z (mm)
109133	0	-0.406	-0.998
109134	0	-0.406	-0.998
109135	0	-0.407	-0.929
109137	0	4406.30	-8.089
109140	0	2978.49	-3.094
109144	0	1497.13	-0.305
109147	0	-1490.18	-2.624
109150	0	-2991.33	-22.492
109153	0	-3997.97	-23.24
109156	0	-2610.05	-2608.06
109159	0	2599.22	-2605.06
109162	0	-0.688	-5498.6
109165	0	-1.048	-4999.05
109168	0	-0.395	-4503.19
109171	0	-3.568	-4503.19
109174	0	0.193	3005.12
109178	0	0.189	4511.1
109181	0	-7.678	5502.69
109208	0	-0.357	1499.05
109211	0	-0.588	-1501.07
109214	0	1483.19	-1501.28
109217	0	-1482.68	-1500.81
109220	0	-1496.85	1500.91

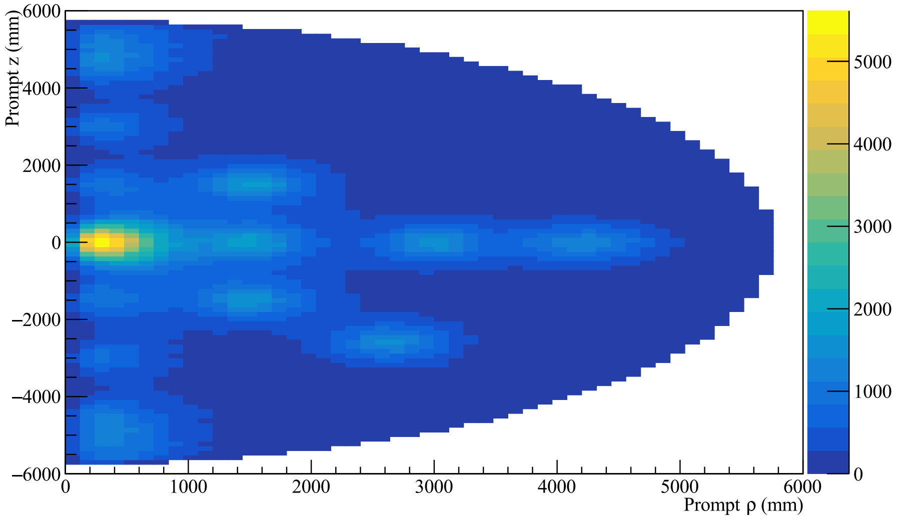

Figure F.1: The reconstructed ρ and z positions of the tagged prompt event, summing events over all AmBe runs shown in Table F.1. The z-axis (color scale) is showing the total number of events in each bin.

Appendix G

Atmospheric Neutrino Event Simulation with the `GENIE` Generator

The `GENIE` (Genertes Events for Neutrino Interaction Experiments) generator simulates high-energy (above 100 MeV) neutrino interactions with a specified target in the input detector model (138), (212). The output of the `GENIE` are a set of final state particles created in the neutrino interaction, which are importanted into `RAT`, which performs the full detector simulation. The nuclear physics model used in all interactions in `GENIE` is the relativistic Fermi gas model developed in (213). The full accounting of the paramatrizations and models used by `GENIE` for the different interactions is given in (212).

Importantly, the interacting nucleus is often left in an excited state and de-excitation γs for oxygen, important as a background for the analysis presented in Chapter 6, are handled using the branching ratios given in (136). The cross section used by `GENIE` for $\nu/\bar{\nu}$ + ^{16}O is given in Figure G.1. It is compared against the cross section used in `NEUT`, a separate atmospheric neutrino interaction generator, which is calculated in (214). The data points come from a T2K measurement of this process (30). This measurement is also performed by Super-Kamikande using atmospheric neutrinos (137).

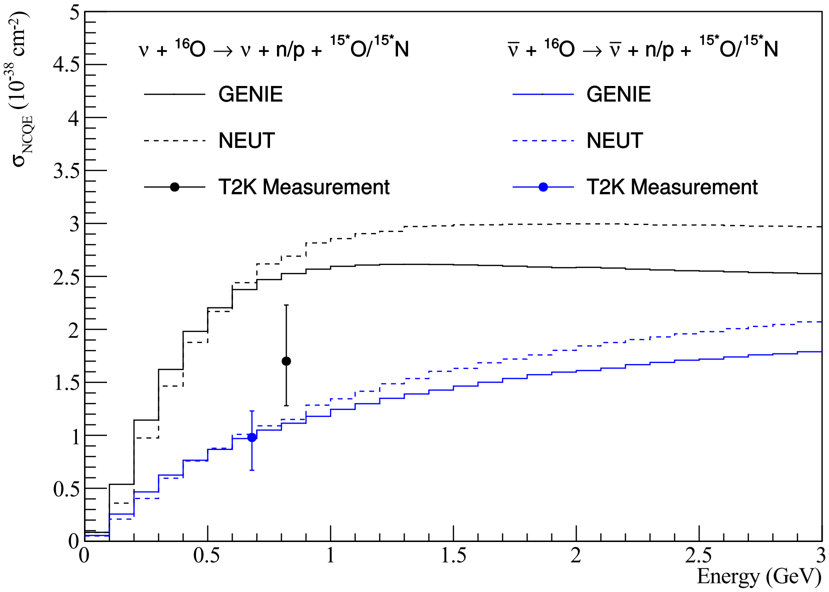

Figure G.1: The $\nu/\bar{\nu}$ + ^{16}O cross sections used in the `GENIE` and `NEUT` generators, compared to data taken by T2K (30).

A simplified version of the SNO+ detector is used for the `GENIE` simulation. The model includes the AV, 50 mm thick and with a radius of 6000 mm, created using the PMMA material ($C_5H_8O_2$). The external water is contained within an 8500 mm shell and the PMT and steel support structure are not included. The inner AV volume is filled with water. A visualization of the geometry is shown in Figure G.2.

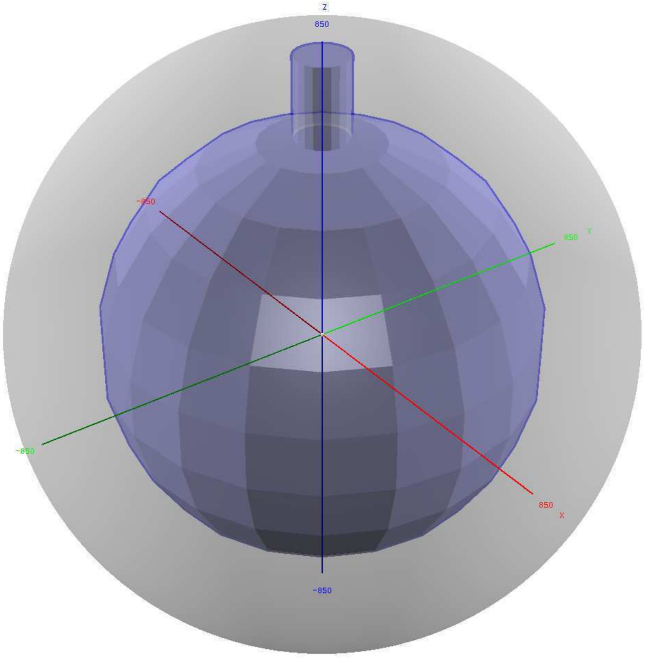

Figure G.2: A visualization of the simplified SNO+ geometry used in the **`GENIE`** primary vertex generation. The detector model includes the acrylic vessesl with its neck (blue), and the light water (grey). The facets are an artifact of the visualization. This figure and caption are reproduced from A. Mastbaum's thesis (22).

G.1 Verifying the `GENIE` MC

In order to verify the `GENIE` MC, a source of well-understood events is used to compare to data. For this comparison, a data-set ranging from run 100000 to 108416, consisting of 114.68 days of data, is used. The simulation uses `RAT` 6.5.4. This analysis was used as verification of `GENIE` for an estimate of the atmospheric background for the invisible nucleon decay paper, (39), but is also important to present in the context of the atmospheric background for the antineutrino search (Section 6.8.3).

Michel electrons are particularly easy to tag due to their high energy and tight temporal coincidence with the associated muon event. Coincidence events with prompt nhits larger than 200 and delayed nhits larger than 100 are selected. The data-cleaning and classifier cuts are loosened to maintain high signal efficiency, and a fiducial volume of 7500 mm is used for both events. A total of 20 events pass these cuts, which can be compared to an expectation of 16, calculated using the atmospheric neutrino interaction rates in Table 6.14. The Michel electron energy spectrum and the Δt between the prompt and delayed events are consistent between data and MC, and are shown in Figures G.3 and G.4.

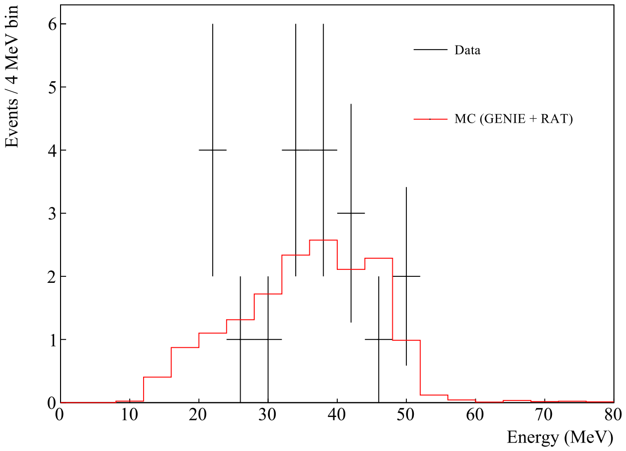

Figure G.3: The energy of the delayed events, compared for data and MC, shows the expected Michel electron energy spectrum.

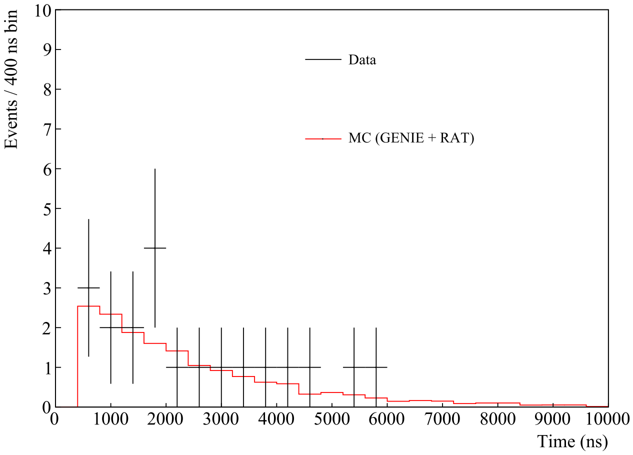

Figure G.4: The Δt between the prompt and delayed events shows the expected muon decay time.

References

[1] Ivan Esteban, M. C. Gonzalez-Garcia, Alvaro Hernandez-Cabezudo, Michele Maltoni, and Thomas Schwetz. **Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23}, δ_{CP}, and the mass ordering**. *JHEP*, **01**:106, 2019. xiii, xxiii, 9, 10, 17, 33, 34

[2] Patrick Huber. **Determination of antineutrino spectra from nuclear reactors**. *Phys. Rev. C*, **84**:024617, Aug 2011. xiii, 73, 86

[3] Th. A. Mueller, D. Lhuillier, M. Fallot, A. Letourneau, S. Cormon, M. Fechner, L. Giot, T. Lasserre, J. Martino, G. Mention, A. Porta, and F. Yermia. **Improved predictions of reactor antineutrino spectra**. *Phys. Rev. C*, **83**:054615, May 2011. xiii, xxvi, 73, 89

[4] V. Lozza. **AlphaN in Water PHase (2017 - 2019)**. *SNO+ Internal Document Database*, **6205**, 2019. xv, 129, 131, 132, 134

[5] Andrea Giuliani and Alfredo Poves. **Neutrinoless Double-Beta Decay**. *Adv. High Energy Phys.*, **2012**:857016, 2012. xviii, xxxiv, 196, 197

[6] Olga Mena and Stephen J. Parke. **Unified graphical summary of neutrino mixing parameters**. *Phys. Rev. D*, **69**:117301, 2004. xxii, 16

[7] P.F. De Salas, S. Gariazzo, O. Mena, C.A. Ternes, and M. Trtola. **Neutrino Mass Ordering from Oscillations and Beyond: 2018 Status and Future Prospects**. *Front. Astron. Space Sci.*, **5**:36, 2018. xxii, 16

[8] C. H. Llewellyn Smith. **Neutrino Reactions at Accelerator Energies**. *Phys. Rept.*, **3**:261–379, 1972. xxii, 21

[9] P. Vogel and John F. Beacom. **Angular distribution of neutron inverse beta decay, anti-neutrino(e) + p —¿ e+ + n**. *Phys. Rev.*, **D60**:053003, 1999. xxii, 19, 20, 21

[10] T. Leitner. *Neutrino Interactions with Nucleons and Nuclei.* PhD thesis, University of Giessen, 2005. xxii, 21, 22

[11] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics.* 4 2007. xxii, 6, 24, 25

[12] W. C. Haxton, R. G. Hamish Robertson, and Aldo M. Serenelli. **Solar Neutrinos: Status and Prospects**. *Ann. Rev. Astron. Astrophys.*, **51**:21–61, 2013. xxii, 26

[13] Q. R. Ahmad et al. **Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory**. *Phys. Rev. Lett.*, **89**:011301, 2002. xxiii, 1, 5, 25, 26, 27, 43

[14] D.A. Dwyer. **Antineutrinos from nuclear reactors: recent oscillation measurements**. *New J. Phys.*, **17**(2):025003, 2015. xxiii, 31

[15] S. Abe et al. **Precision Measurement of Neutrino Oscillation Parameters with KamLAND**. *Phys. Rev. Lett.*, **100**:221803, 2008. xxiii, 32

[16] Fengpeng An et al. **Neutrino Physics with JUNO**. *J. Phys.*, **G43**(3):030401, 2016. xxiii, 2, 35, 36, 45, 158

[17] M. Agostini et al. **Comprehensive geoneutrino analysis with Borexino**. *Phys. Rev. D*, **101**(1):012009, 2020. xxiv, 37, 38

[18] Stan Seibert. *A Low Energy Measurement of the ^8B Solar Neutrino Spectrum at the Sudbury Neutrino Observatory.* PhD thesis, University of Texas at Austin, 2006. xxiv, 46

[19] M. Anderson et al. **Measurement of the ^{8}B solar neutrino flux in SNO+ with very low backgrounds**. *Phys. Rev.*, **D99**(1):012012, 2019. xxiv, 2, 40, 63, 67, 93

[20] Eric Marzec. *Measurement of the ^8B solar neutrinos in the SNO+ water phase and a model of vacuum-enhanced neutrino mixing.* PhD thesis, Pennsylvania U., 2019. xxv, xlii, 33, 65, 66, 272

[21] G. D. Barr, T. K. Gaisser, P. Lipari, S. Robbins, and T. Stanev. **Three-dimensional calculation of atmospheric neutrinos**. *Phys. Rev. D*, **70**:023006, Jul 2004. xxx, 140, 141

[22] Andy Mastbaum. *Constraining the HEP Solar Neutrino and Diffuse Supernova Neutrino Background Fluxes with the Subdury Neutrino Observatory.* PhD thesis, Pennsylvania U., 2016. xxx, xliii, 24, 140, 141, 298

[23] Hamamatsu Photonics. **Private Communication**, 2017. xxxi, 165, 173

[24] A. Gando et al. **Search for Majorana Neutrinos near the Inverted Mass Hierarchy Region with KamLAND-Zen**. *Phys. Rev. Lett.*, **117**(8):082503, 2016. [Addendum: Phys. Rev. Lett.117,no.10,109903(2016)]. xxxiv, 199, 201

[25] Hamamatsu Photonics. **Hamamatsu R7600 Datasheet**. `https://www.hamamatsu.com/resources/pdf/etd/R7600U$_$TPMH1317E.pdf`, 2019. [Accessed Aug. 26, 2019]. xxxv, 210, 235, 244

[26] Hamamatsu Photonics. **Hamamatsu R2257 Datasheet**. `https://www.hamamatsu.com/resources/pdf/etd/R2257$_$TPMH1141E.pdf`, 2019. [Accessed Aug. 26, 2019]. xxxv, 210, 235, 244

[27] S. D. Biller, N. A. Jelley, M. D. Thorman, N. P. Fox, and T. H. Ward. **Measurements of photomultiplier single photon counting efficiency for the Sudbury Neutrino Observatory**. *Nucl. Instrum. Meth.*, **A432**:364–373, 1999. xxxv, 210, 235

[28] Masahiko Taniguchi, Hai Du, and Jonathan S. Lindsey. **PhotochemCAD 3: Diverse Modules for Photophysical Calculations with Multiple Spectral Databases**. *Photochemistry and Photobiology*, **94**(2):277–289, 2018. xxxv, 210

[29] Ana Sofia Inacio. **Optical Calibration of the SNO+ Detector with the Laserball in the Water Phase** . *SNO+ Internal Document Database*, **5862**, 2019. xliii, 288, 289, 290

[30] K. Abe et al. **Measurement of neutrino and antineutrino neutral-current quasielasticlike interactions on oxygen by detecting nuclear deexcitation γ rays**. *Phys. Rev. D*, **100**:112009, Dec 2019. xliii, 142, 145, 297

[31] K. Eguchi et al. **First results from KamLAND: Evidence for reactor anti-neutrino disappearance**. *Phys. Rev. Lett.*, **90**:021802, 2003. 2, 45

[32] T. Araki et al. **Measurement of neutrino oscillation with KamLAND: Evidence of spectral distortion**. *Phys. Rev. Lett.*, **94**:081801, 2005. 2, 30, 38

[33] Y. Abe et al. **Indication of Reactor $\bar{\nu}_e$ Disappearance in the Double Chooz Experiment**. *Phys. Rev. Lett.*, **108**:131801, 2012. 2, 34

[34] F. P. An et al. **Observation of electron-antineutrino disappearance at Daya Bay**. *Phys. Rev. Lett.*, **108**:171803, 2012. 2, 34, 45

[35] J. K. Ahn et al. **Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment**. *Phys. Rev. Lett.*, **108**:191802, 2012. 2, 35, 45

[36] J. Ashenfelter et al. **The PROSPECT Physics Program**. *J. Phys.*, **G43**(11):113001, 2016. 2, 37, 46

[37] I. Alekseev et al. **DANSS: Detector of the reactor AntiNeutrino based on Solid Scintillator**. *JINST*, **11**(11):P11011, 2016. 2, 37

[38] M. R. Anderson et al. **Measurement of neutron-proton capture in the SNO+ water phase**. 2020. 2, 40, 64, 67, 91

[39] M. Anderson et al. **Search for invisible modes of nucleon decay in water with the SNO+ detector**. *Phys. Rev.*, **D99**(3):032008, 2019. 2, 40, 67, 93, 298

[40] S. Andringa et al. **Current Status and Future Prospects of the SNO+ Experiment**. *Adv. High Energy Phys.*, **2016**:6194250, 2016. 3, 158, 195, 257

[41] Frederick Reines and Clyde L. Cowan. **The neutrino**. *Nature*, **178**:446–449, 1956. 4, 29, 45

[42] B. Pontecorvo. **Inverse beta processes and nonconservation of lepton charge**. *Sov. Phys. JETP*, **7**:172–173, 1958. [Zh. Eksp. Teor. Fiz.34,247(1957)]. 4

[43] V. N. Gribov and B. Pontecorvo. **Neutrino astronomy and lepton charge**. *Phys. Lett.*, **28B**:493, 1969. 5

[44] Y. Fukuda et al. **Evidence for oscillation of atmospheric neutrinos**. *Phys. Rev. Lett.*, **81**:1562–1567, 1998. 5, 28, 43

[45] P.F. de Salas, D.V. Forero, C.A. Ternes, M. Tortola, and J.W.F. Valle. **Status of neutrino oscillations 2018: 3σ hint for normal mass ordering and improved CP sensitivity**. *Phys. Lett. B*, **782**:633–640, 2018. 9, 17

[46] F. Capozzi, E. Lisi, A. Marrone, and A. Palazzo. **Current unknowns in the three neutrino framework**. *Prog. Part. Nucl. Phys.*, **102**:48–72, 2018. 9

[47] L. Wolfenstein. **Neutrino oscillations in matter**. *Phys. Rev. D*, **17**:2369–2374, May 1978. 10

[48] S.P. Mikheev and A.Yu. Smirnov. **Resonant amplification of neutrino oscillations in matter and solar neutrino spectroscopy**. *Nuovo Cim. C*, **9**:17–26, 1986. 10

[49] Vedran Brdar, Alexander J. Helmboldt, Sho Iwamoto, and Kai Schmitz. **Type I seesaw mechanism as the common origin of neutrino mass, baryon asymmetry, and the electroweak scale**. *Phys. Rev. D*, **100**:075029, Oct 2019. 15

[50] Michelle J. Dolinski, Alan W. P. Poon, and Werner Rodejohann. **Neutrinoless Double-Beta Decay: Status and Prospects**. *Ann. Rev. Nucl. Part. Sci.*, **69**:219–251, 2019. 15, 200

[51] Julien Lesgourgues and Sergio Pastor. **Massive neutrinos and cosmology**. *Physics Reports*, **429**(6):307 – 379, 2006. 16

[52] Francesco Capozzi, Eleonora Di Valentino, Eligio Lisi, Antonio Marrone, Alessandro Melchiorri, and Antonio Palazzo. **Global constraints on absolute neutrino masses and their ordering**. *Phys. Rev.*, **D95**(9):096014, 2017. 17

[53] John N. Bahcall, M.H. Pinsonneault, and Sarbani Basu. **Solar models: Current epoch and time dependences, neutrinos, and helioseismological properties**. *Astrophys. J.*, **555**:990–1012, 2001. 23

[54] John N. Bahcall and Roger K. Ulrich. **Solar Models, Neutrino Experiments and Helioseismology**. *Rev. Mod. Phys.*, **60**:297–372, 1988. 23

[55] Raymond Davis, Jr., Don S. Harmer, and Kenneth C. Hoffman. **Search for neutrinos from the sun**. *Phys. Rev. Lett.*, **20**:1205–1209, 1968. 25

[56] B. T. Cleveland, Timothy Daily, Raymond Davis, Jr., James R. Distel, Kenneth Lande, C. K. Lee, Paul S. Wildenhain, and Jack Ullman. **Measurement of the solar electron neutrino flux with the Homestake chlorine detector**. *Astrophys. J.*, **496**:505–526, 1998. 25

[57] J. N. Abdurashitov et al. **Solar neutrino flux measurements by the Soviet-American Gallium Experiment (SAGE) for half the 22 year solar cycle**. *J. Exp. Theor. Phys.*, **95**:181–193, 2002. [Zh. Eksp. Teor. Fiz.122,211(2002)]. 25

[58] W. Hampel et al. **GALLEX solar neutrino observations: Results for GALLEX IV**. *Phys. Lett.*, **B447**:127–133, 1999. 25

[59] M. Altmann et al. **GNO solar neutrino observations: Results for GNO I**. *Phys. Lett.*, **B490**:16–26, 2000. 25

[60] Y. Fukuda et al. **Atmospheric muon-neutrino / electron-neutrino ratio in the multiGeV energy range**. *Phys. Lett.*, **B335**:237–245, 1994. 28

[61] K. Abe et al. **Solar Neutrino Measurements in Super-Kamiokande-IV**. *Phys. Rev. D*, **94**(5):052010, 2016. 32

[62] B. Aharmim et al. **Combined Analysis of all Three Phases of Solar Neutrino Data from the Sudbury Neutrino Observatory**. *Phys. Rev. C*, **88**:025501, 2013. 32

[63] B. Aharmim et al. **Low Energy Threshold Analysis of the Phase I and Phase II Data Sets of the Sudbury Neutrino Observatory**. *Phys. Rev. C*, **81**:055504, 2010. 32

[64] R. Bonventre, A. LaTorre, J.R. Klein, G.D. Orebi Gann, S. Seibert, and O. Wasalski. **Non-Standard Models, Solar Neutrinos, and Large θ_13**. *Phys. Rev. D*, **88**(5):053010, 2013. 33

[65] Alexander Friedland, Cecilia Lunardini, and Carlos Pea-Garay. **Solar neutrinos as probes of neutrinomatter interactions**. *Physics Letters B*, **594**(3):347 – 354, 2004. 33

[66] M.M. Guzzo, P.C. [de Holanda], M. Maltoni, H. Nunokawa, M.A. Trtola, and J.W.F. Valle. **Status of a hybrid three--neutrino interpretation of neutrino data**. *Nuclear Physics B*, **629**(1):479 – 490, 2002. 33

[67] F. Boehm et al. **Final results from the Palo Verde neutrino oscillation experiment**. *Phys. Rev. D*, **64**:112001, 2001. 33

[68] C. Bemporad. **Results from CHOOZ**. *Nucl. Phys. B Proc. Suppl.*, **77**:159–165, 1999. 33

[69] P.F. Harrison, D.H. Perkins, and W.G. Scott. **Tri-bimaximal mixing and the neutrino oscillation data**. *Physics Letters B*, **530**(1):167 – 173, 2002. 33

[70] C.D. Shin et al. **Observation of reactor antineutrino disappearance using delayed neutron capture on hydrogen at RENO**. *JHEP*, **04**:029, 2020. 35

[71] Y. Abe et al. **First Measurement of θ_{13} from Delayed Neutron Capture on Hydrogen in the Double Chooz Experiment**. *Phys. Lett. B*, **723**:66–70, 2013. 35

[72] F.P. An et al. **New measurement of θ_{13} via neutron capture on hydrogen at Daya Bay**. *Phys. Rev. D*, **93**(7):072011, 2016. 35

[73] Janet M. Conrad, William C. Louis, and Michael H. Shaevitz. **The LSND and MiniBooNE Oscillation Searches at High Δm^2**. *Ann. Rev. Nucl. Part. Sci.*, **63**:45–67, 2013. 36

[74] Y.J. Ko et al. **Sterile Neutrino Search at the NEOS Experiment**. *Phys. Rev. Lett.*, **118**(12):121802, 2017. 37

[75] A. P. Serebrov, V. G. Ivochkin, R. M. Samoilov, A. K. Fomin, A. O. Polyushkin, V. G. Zinoviev, P. V. Neustroev, V. L. Golovtsov, A. V. Chernyj, O. M. Zherebtsov, and et al. **First Observation of the Oscillation Effect in the Neutrino-4 Experiment on the Search for the Sterile Neutrino**. *JETP Letters*, **109**(4):213221, Feb 2019. 37

[76] Y. Abreu et al. **A novel segmented-scintillator antineutrino detector**. *JINST*, **12**(04):P04024, 2017. 37

[77] N. Allemandou et al. **The STEREO Experiment**. *JINST*, **13**(07):P07009, 2018. 37

[78] A. Gando et al. **Partial radiogenic heat model for Earth revealed by geoneutrino measurements**. *Nature Geo.*, **4**:647–651, 2011. 38

[79] J. Boger et al. **The Sudbury neutrino observatory**. *Nucl. Instrum. Meth.*, **A449**:172–207, 2000. 39, 209

[80] J. Benziger et al. **The Scintillator Purification System for the Borexino Solar Neutrino Detector**. *Nucl. Instrum. Meth.*, **A587**:277–291, 2008. 40, 47

[81] S. Fukuda et al. **Solar B-8 and hep neutrino measurements from 1258 days of Super-Kamiokande data**. *Phys. Rev. Lett.*, **86**:5651–5655, 2001. 43

[82] Y. Ashie et al. **A Measurement of atmospheric neutrino oscillation parameters by SUPER-KAMIOKANDE I**. *Phys. Rev.*, **D71**:112005, 2005. 43

[83] H. H. Chen. **Direct Approach to Resolve the Solar Neutrino Problem**. *Phys. Rev. Lett.*, **55**:1534–1536, 1985. [,120(1985)]. 43

[84] A. Aguilar-Arevalo et al. **Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam**. *Phys. Rev.*, **D64**:112007, 2001. 45

[85] G Alimonti et al. **Science and technology of BOREXINO: A Real time detector for low-energy solar neutrinos**. *Astropart. Phys.*, **16**:205–234, 2002. 45

[86] M. Askins et al. **Theia: An advanced optical neutrino detector**. 2019. 45, 158, 195, 206

[87] John F. Beacom et al. **Physics prospects of the Jinping neutrino experiment**. *Chin. Phys.*, **C41**(2):023002, 2017. 45, 158, 207

[88] P. Adamson et al. **First measurement of electron neutrino appearance in NOvA**. *Phys. Rev. Lett.*, **116**(15):151806, 2016. 46

[89] John B. Birks. *The Theory and practice of scintillation counting.* 1964. 47

[90] Roland Winston and Jay M. Enoch. **Retinal Cone Receptor as an Ideal Light Collector**. *J. Opt. Soc. Am.*, **61**(8):1120–1121, Aug 1971. 49, 209

[91] C. J. Jillings, R. J. Ford, A. L. Hallin, P. Harvey, R. W. MacLeod, H. B. Mak, P. Skensved, and R. L. Stevenson. **The photomultiplier tube testing facility for the Sudbury Neutrino Observatory**. *Nucl. Instrum. Meth.*, **A373**:421–429, 1996. 51

[92] M. R. Dragowsky et al. **The N-16 calibration source for the Sudbury Neutrino Observatory**. *Nucl. Instrum. Meth.*, **A481**:284–296, 2002. 63

[93] M. Anderson et al. **The SNO+ Experiment**. *JINST*. [in preparation]. 63, 64, 288

[94] Tony LaTorre. **Measurement of the Trigger Efficiency in SNO+**. *SNO+ Internal Document Database*, **4723**, 2017. 64, 70

[95] S. Agostinelli et al. **GEANT4: A Simulation toolkit**. *Nucl. Instrum. Meth.*, **A506**:250–303, 2003. 72

[96] S. Seibert. **RAT Documentation**, 2014. 72

[97] R. Brun and F. Rademakers. **ROOT: An object oriented data analysis framework**. *Nucl. Instrum. Meth.*, **A389**:81–86, 1997. 72

[98] J. R. Wilson. *Measurement of the ^8B solar neutrino energy spectrum at the Sudbury Neutrino Observatory.* PhD thesis, Oxford, 2004. 79

[99] Mark R. Vagins. **Selective Filtration of Gadolinium Trichloride for Use in Neutron Detection in Large Water Cherenkov Detectors**. Technical report, United States, Apr 2013. DOE-FG–02-07ER41494. 82

[100] John F. Beacom. **The Diffuse Supernova Neutrino Background**. *Ann. Rev. Nucl. Part. Sci.*, **60**:439–462, 2010. 83

[101] Charles Simpson. **Physics Potential of Super-K Gd**. *PoS*, **ICHEP2018**:008, 2019. 83, 157

[102] Xin Qian and Jen-Chieh Peng. **Physics with Reactor Neutrinos**. *Rept. Prog. Phys.*, **82**(3):036201, 2019. 84

[103] X. B. Ma, W. L. Zhong, L. Z. Wang, Y. X. Chen, and J. Cao. **Improved calculation of the energy release in neutron-induced fission**. *Phys. Rev.*, **C88**(1):014605, 2013. 84

[104] Patrick Huber. **On the determination of anti-neutrino spectra from nuclear reactors**. *Phys. Rev.*, **C84**:024617, 2011. [Erratum: Phys. Rev.C85,029901(2012)]. 84

[105] Th. A. Mueller et al. **Improved Predictions of Reactor Antineutrino Spectra**. *Phys. Rev.*, **C83**:054615, 2011. 84

[106] Feng Peng An et al. **Improved Measurement of the Reactor Antineutrino Flux and Spectrum at Daya Bay**. *Chin. Phys.*, **C41**(1):013002, 2017. 84

[107] *Operating Experience with Nuclear Power Stations in Member States.* Operating Experience with Nuclear Power Stations in Member States (CD-ROM). INTERNATIONAL ATOMIC ENERGY AGENCY, Vienna, 2019. 85

[108] INFN Laboratory for Nuclear Technologies. **A reference worldwide model for antineutrinos from reactors**. `https://www.fe.infn.it/antineutrino/`, 2020. [Accessed May 7. 15, 2020]. 85

[109] Marica Baldoncini, Ivan Callegari, Giovanni Fiorentini, Fabio Mantovani, Barbara Ricci, Virginia Strati, and Gerti Xhixha. **Reference worldwide model for antineutrinos from reactors**. *Phys. Rev. D*, **91**(6):065002, 2015. 85

[110] IESO. **IESO Power Data, Data Directory**. `http://www.ieso.ca/en/Power-Data/Data-Directory`, 2020. [Accessed May 7. 15, 2020]. 85

[111] **Private Communication with Atomic Energy of Canada Limited (AECL)**, 2013. 86

[112] M. Tanabashi et al. **Neutrino Masses, Mixing, and Oscillations**. *Phys. Rev. D.*, **98**:030001, 2018. 87

[113] SNO+ Collaboration. **SNO+ Water Phase Unidoc**, 2020. 88

[114] D. Adey et al. **Improved Measurement of the Reactor Antineutrino Flux at Daya Bay**. *Phys. Rev. D*, **100**(5):052004, 2019. 88, 117

[115] D. Cokinos and E. Melkonian. **Measurement of the 2200 m/sec neutron-proton capture cross section**. *Phys. Rev. C*, **15**:1636–1643, May 1977. 91

[116] Y. Zhang et al. **First measurement of radioactive isotope production through cosmic-ray muon spallation in Super-Kamiokande IV**. *Phys. Rev.*, **D93**(1):012004, 2016. 91, 150

[117] B. Aharmim et al. **Electron energy spectra, fluxes, and day-night asymmetries of B-8 solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory**. *Phys. Rev.*, **C72**:055502, 2005. 92

[118] Shunsaku Horiuchi, John F. Beacom, and Eli Dwek. **The Diffuse Supernova Neutrino Background is detectable in Super-Kamiokande**. *Phys. Rev.*, **D79**:083013, 2009. 92

[119] John F. Beacom and Mark R. Vagins. **GADZOOKS! Antineutrino spectroscopy with large water Cherenkov detectors**. *Phys. Rev. Lett.*, **93**:171101, 2004. 92

[120] Chenyuan Xu. **Current status of SK-Gd project and EGADS**. *J. Phys. Conf. Ser.*, **718**(6):062070, 2016. 92

[121] M. Askins et al. **The Physics and Nuclear Nonproliferation Goals of WATCHMAN: A WAter CHerenkov Monitor for ANtineutrinos**. 2015. 92, 157, 158

[122] A. R. Back et al. **Accelerator Neutrino Neutron Interaction Experiment (ANNIE): Preliminary Results and Physics Phase Proposal**. 2017. 92, 158, 160, 174

[123] Freija. Descamps et al. **Water phase: list of criteria for a run to pass the Run Selection** . *SNO+ Internal TWiki*, 2018. 94

[124] L. Lebanowski. **Water Reconstruction Errors and FOMs**. *SNO+ Internal Document Database*, **6190**, 2020. 97

[125] Oleg Chkvoret. **Leaching Studies**. *SNO+ Internal Document Database*, **2384**, 2014. 128

[126] L. Anselmo et al. **Measurement of Pb-210 Activity From Acrylic Vessel (AV) LightWater Assay**. *SNO+ Internal Document Database*, **3134**, 2015. 129

[127] S. Harissopulos, H.W. Becker, J.W. Hammer, A. Lagoyannis, C. Rolfs, and F. Strieder. **Cross section of the C-13(alpha, n)O-16 reaction: A Background for the measurement of geo-neutrinos**. *Phys. Rev. C*, **72**:062801, 2005. 129

[128] B. Braizinha, J.H. Esterline, H.J. Karwowski, and W. Tornow. **Determination of the proton and alpha-particle light-response functions for the KamLAND, BC-501A and BC-517H liquid scintillators**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **623**(3):1046 – 1049, 2010. 129

[129] David W. McKee, Jerome K. Busenitz, and Igor Ostrovskiy. **A 13C(,n)16O calibration source for KamLAND**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **587**(2):272 – 276, 2008. 129

[130] W. A. Peters. **Comment on "Cross section of the ^{13}C$(\alpha,n)^{16}$O reaction: A background for the measurement of geo-neutrinos"**. *Phys. Rev. C*, **96**:029801, Aug 2017. 129

[131] G.N. Vlaskin, Yu. S. Khomyakov, and V.I. Bulanenko. **Neutron Yield of the Reaction (α, n) on Thick Targets Comprised of Light Elements**. *Atmoic Energy*, **117**:357–365, 2014. 130

[132] Keiichi. Shibata et al. **Japanese Evaluated Nuclear Data Library Version 3 Revision-3: JENDL-3.3**. *Journal of Nuclear Science and Technology*, **39**(11):1125–1136, 2002. 130

[133] S. Westerdale. **A new tool for (α,n) yield calculations and its implications for DEAP-3600**. *AIP Conf. Proc.*, **1921**(1):060002, 2018. 130

[134] G. Barr, T.K. Gaisser, and T. Stanev. **Flux of Atmospheric Neutrinos**. *Phys. Rev. D*, **39**:3532–3534, 1989. 140

[135] T. K. Gaisser and Todor Stanev. **Path length distributions of atmospheric neutrinos**. *Phys. Rev. D*, **57**:1977–1982, Feb 1998. 140

[136] H. Ejiri. **Nuclear deexcitations of nucleon holes associated with nucleon decays in nuclei**. *Phys. Rev. C*, **48**:1442–1444, Sep 1993. 142, 296

[137] L. Wan et al. **Measurement of the neutrino-oxygen neutral-current quasielastic cross section using atmospheric neutrinos at Super-Kamiokande**. *Phys. Rev. D*, **99**(3):032005, 2019. 142, 297

[138] C. Andreopoulos et al. **The GENIE Neutrino Monte Carlo Generator**. *Nucl. Instrum. Meth. A*, **614**:87–104, 2010. 142, 296

[139] T. Kajita. **Atmospheric neutrinos**. *New J. Phys.*, **6**:194, 2004. 145

[140] Ralph P. Kraft, David N. Burrows, and John A. Nousek. **Determination of Confidence Limits for Experiments with Low Numbers of Counts**. *apj*, **374**:344, 1991. 149

[141] Oleg Smirnov. **Experimental Aspects of Geoneutrino Detection: Status and Perspectives**. *Prog. Part. Nucl. Phys.*, **109**:103712, 2019. 149

[142] Billy Liggins. *Cosmic muon-induced neutrons in the SNO+ water phase*. PhD thesis, Queen Mary University of London, 2020. 150

[143] Shirley Weishi Li and John F. Beacom. **First calculation of cosmic-ray muon spallation backgrounds for MeV astrophysical neutrino signals in Super-Kamiokande**. *Phys. Rev. C*, **89**:045801, 2014. 150

[144] S. Dazeley, M. Askins, M. Bergevin, A. Bernstein, N.S. Bowden, T.M. Shokair, P. Jaffke, S.D. Rountree, and M. Sweany. **A search for cosmogenic production of -neutron emitting radionuclides in water**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **821**:151 – 159, 2016. 150

[145] C. Lin. **Cosmic-ray muon spallation backgrounds estimation for antineutrino search in SNO+ .** *SNO+ Internal Document Database*, **5716**, 2019. 151

[146] Glen Cowan. **Discovery sensitivity for a counting experiment with background uncertainty**. *en. In*, page 8, 2012. 153, 154

[147] S. J. Brice et al. **Photomultiplier tubes in the MiniBooNE experiment**. *Nucl. Instrum. Meth.*, **A562**:97–109, 2006. 158

[148] K. Abe et al. **Letter of Intent: The Hyper-Kamiokande Experiment — Detector Design and Physics Potential —**. 2011. 158

[149] R. Abbasi et al. **Calibration and Characterization of the IceCube Photomultiplier Tube**. *Nucl. Instrum. Meth.*, **A618**:139–152, 2010. 160

[150] J. Brack et al. **Characterization of the Hamamatsu R11780 12 inch Photomultiplier Tube**. *Nucl. Instrum. Meth.*, **A712**:162–173, 2013. 160

[151] Tony LaTorre. **Lecrunch Open Source Software**. `https://bitbucket.org/tlatorre/lecrunch/src/default/`, 2013. [Accessed Nov. 19, 2018]. 161, 218

[152] N. Barros, T. Kaptanoglu, B. Kimelman, J. R. Klein, E. Moore, J. Nguyen, K. Stavreva, and R. Svoboda. **Characterization of the ETEL D784UKFLB 11 in. photomultiplier tube**. *Nucl. Instrum. Meth.*, **A852**:15–19, 2017. 174

[153] P. A. Amaudruz et al. **In-situ characterization of the Hamamatsu R5912-HQE photomultiplier tubes used in the DEAP-3600 experiment**. *Nucl. Instrum. Meth.*, **A922**:373–384, 2019. 174

[154] Tanner Kaptanoglu. **Characterization of the Hamamatsu 8” R5912-MOD Photomultiplier Tube**. *Nucl. Instrum. Meth.*, **A889**:69–77, 2018. 176, 206, 233

[155] H. M. O’Keeffe, E. O’Sullivan, and M. C. Chen. **Scintillation decay time and pulse shape discrimination in oxygenated and deoxygenated solutions of linear alkylbenzene for the SNO+ experiment**. *Nucl. Instrum. Meth.*, **A640**:119–122, 2011. 179, 236, 248

[156] Paolo Lombardi, Fausto Ortica, Gioacchino Ranucci, and Aldo Romani. **Decay time and pulse shape discrimination of liquid scintillators based on novel solvents**. *Nucl. Instrum. Meth.*, **A701**:133–144, 2013. 179, 236, 248, 257

[157] E. O’Sullivan, H. S. Wan Chan Tseung, N. Tolich, H. M. O’Keeffe, and M. Chen. **SNO+ Liquid Scintillator Characterization: Timing, Quenching, and Energy Scale**. *Nucl. Phys. Proc. Suppl.*, **229-232**:549–549, 2012. 179, 236

[158] Xiao-Bo Li, Hua-Lin Xiao, Jun Cao, Jin Li, Xi-Chao Ruan, and Yue-Kun Heng. **Timing properties and pulse shape discrimination of LAB-based liquid scintillator**. *Chinese Physics C*, **35**(11):1026–1032, nov 2011. 179, 236

[159] Wouter Verkerke and David P. Kirkby. **The RooFit toolkit for data modeling**. *eConf*, **C0303241**:MOLT007, 2003. [,186(2003)]. 184, 225

[160] P. A. Amaudruz et al. **Measurement of the scintillation time spectra and pulse-shape discrimination of low-energy β and nuclear recoils in liquid argon with DEAP-1**. *Astropart. Phys.*, **85**:1–23, 2016. 185, 257

[161] Belina Von Krosigk. *Measurement of proton and α particle quenching in LAB based scintillators and determination of spectral sensitivities to supernova neutrinos in the SNO+ detector*. PhD thesis, Dresden, Tech. U., Dept. Math., 2015. 190

[162] I. S. Yeo et al. **Measurement of the refractive index of the LAB-based liquid scintillator and acrylic at RENO**. *Phys. Scripta*, **82**:065706, 2010. 191, 249

[163] Tanner Kaptanoglu, Meng Luo, and Josh Klein. **Cherenkov and Scintillation Light Separation Using Wavelength in LAB Based Liquid Scintillator**. *JINST*, **14**(05):T05001, 2019. 195, 208, 214

[164] Tanner Kaptanoglu, Meng Luo, Ben Land, Amanda Bacon, and Josh Klein. **Spectral Photon Sorting For Large-Scale Cherenkov and Scintillation Detectors**. *Phys. Rev. D*, **101**(7):072002, 2020. 195

[165] M. Goeppert-Mayer. **Double beta-disintegration**. *Phys. Rev.*, **48**:512–516, 1935. 196

[166] S. R. Elliott, A. A. Hahn, and M. K. Moe. **Direct Evidence for Two Neutrino Double Beta Decay in ^{82}Se**. *Phys. Rev. Lett.*, **59**:2020–2023, 1987. [,989(1987)]. 196

[167] Vladimir Tello, Miha Nemevsek, Fabrizio Nesti, Goran Senjanovic, and Francesco Vissani. **Left-Right Symmetry: from LHC to Neutrinoless Double Beta Decay**. *Phys. Rev. Lett.*, **106**:151801, 2011. 197

[168] J. Kotila and F. Iachello. **Phase space factors for $\beta^+\beta^+$ decay and competing modes of double-β decay**. *Phys. Rev.*, **C87**(2):024313, 2013. 198

[169] Sabin Stoica and Mihail Mirea. **New calculations for phase space factors involved in double-β decay**. *Phys. Rev.*, **C88**(3):037303, 2013. 198

[170] J. Barea, J. Kotila, and F. Iachello. **Nuclear matrix elements for double-β decay**. *Phys. Rev.*, **C87**(1):014315, 2013. 198

[171] C. Alduino et al. **First Results from CUORE: A Search for Lepton Number Violation via $0\nu\beta\beta$ Decay of ^{130}Te**. *Phys. Rev. Lett.*, **120**(13):132501, 2018. 199

[172] M. Auger et al. **Search for Neutrinoless Double-Beta Decay in ^{136}Xe with EXO-200**. *Phys. Rev. Lett.*, **109**:032505, 2012. 199

[173] M. Agostini et al. **Results on Neutrinoless Double-β Decay of ^{76}Ge from Phase I of the GERDA Experiment**. *Phys. Rev. Lett.*, **111**(12):122503, 2013. 199

[174] N. Abgrall et al. **The Majorana Demonstrator Neutrinoless Double-Beta Decay Experiment**. *Adv. High Energy Phys.*, **2014**:365432, 2014. 199

[175] J. Martn-Albo et al. **Sensitivity of NEXT-100 to Neutrinoless Double Beta Decay**. *JHEP*, **05**:159, 2016. 199

[176] R. Arnold et al. **Probing New Physics Models of Neutrinoless Double Beta Decay with SuperNEMO**. *Eur. Phys. J.*, **C70**:927–943, 2010. 199

[177] Jack Dunger and Steven D. Biller. **Multi-site Event Discrimination in Large Liquid Scintillation Detectors**. *Nucl. Instrum. Meth.*, **A943**:162420, 2019. 202

[178] Jacopo Dalmasson, Giorgio Gratta, Ako Jamil, Scott Kravitz, Milad Malek, Kevin Wells, Julie Bentley, Samuel Steven, and Jiani Su. **Distributed Imaging for Liquid Scintillation Detectors**. *Phys. Rev.*, **D97**(5):052006, 2018. 206

[179] Gabriel D. Orebi Gann. **Physics Potential of an Advanced Scintillation Detector: Introducing THEIA**. *arXiv:1504.08284 [physics.ins-det]*, 2015. 206

[180] Andrey Elagin, Henry Frisch, Brian Naranjo, Jonathan Ouellet, Lindley Winslow, and Taritree Wongjirad. **Separating Double-Beta Decay Events from Solar Neutrino Interactions in a Kiloton-Scale Liquid Scintillator Detector By Fast Timing**. *Nucl. Instrum. Meth.*, **A849**:102–111, 2017. 206

[181] Runyu Jiang and Andrey Elagin. **Space-Time Discriminant to Separate Double-Beta Decay from ^{8}B Solar Neutrinos in Liquid Scintillator**. *arXiv:1902.06912 [physics.ins-det]*, 2019. 206

[182] Steven D Biller. **Probing Majorana neutrinos in the regime of the normal mass hierarchy**. *Phys. Rev.*, **D87**(7):071301, 2013. 206

[183] R. Bonventre and G. D. Orebi Gann. **Sensitivity of a low threshold directional detector to CNO-cycle solar neutrinos**. *Eur. Phys. J.*, **C78**(6):435, 2018. 206

[184] G. D. Orebi Gann. **Solar neutrinos with Theia**. In *Proceedings, 5th International Solar Neutrino Conference: Dresden, Germany, June 11-14, 2018*, pages 345–361, 2019. 206

[185] J. Caravaca, F. B. Descamps, B. J. Land, J. Wallig, M. Yeh, and G. D. Orebi Gann. **Experiment to demonstrate separation of Cherenkov and scintillation signals**. *Phys. Rev.*, **C95**(5):055801, 2017. 207

[186] J. Caravaca, F. B. Descamps, B. J. Land, M. Yeh, and G. D. Orebi Gann. **Cherenkov and Scintillation Light Separation in Organic Liquid Scintillators**. *Eur. Phys. J.*, **C77**(12):811, 2017. 207

[187] Julieta Gruszko, Brian Naranjo, Byron Daniel, Andrey Elagin, Diana Gooding, Chris Grant, Jonathan Ouellet, and Lindley Winslow. **Detecting Cherenkov light from 1 - 2 MeV electrons in linear alkylbenzene**. *JINST*, **14**(02):P02005, 2019. 207

[188] Ziyi Guo and Zhe Wang. **Slow Liquid Scintillator for Scintillation and Cherenkov Light Separation**. *Springer Proc. Phys.*, **213**:173–177, 2018. 207

[189] Mohan Li, Ziyi Guo, Minfang Yeh, Zhe Wang, and Shaomin Chen. **Separation of Scintillation and Cherenkov Lights in Linear Alkyl Benzene**. *Nucl. Instrum. Meth.*, **A830**:303–308, 2016. 207

[190] M. Yeh, S. Hans, W. Beriguete, R. Rosero, L. Hu, R.L. Hahn, M.V. Diwan, D.E. Jaffe, S.H. Kettell, and L. Littenberg. **A new water-based liquid scintillator and potential applications**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **660**(1):51 – 56, 2011. 207

[191] C. Aberle, A. Elagin, H. J. Frisch, M. Wetstein, and L. Winslow. **Measuring Directionality in Double-Beta Decay and Neutrino Interactions with Kiloton-Scale Scintillation Detectors**. *JINST*, **9**:P06012, 2014. 207

[192] G. Alimonti et al. **The Borexino detector at the Laboratori Nazionali del Gran Sasso**. *Nucl. Instrum. Meth.*, **A600**:568–593, 2009. 209

[193] Michael Wurm et al. **Optical Scattering Lengths in Large Liquid-Scintillator Neutrino Detectors**. *Rev. Sci. Instrum.*, **81**:053301, 2010. 209

[194] Matthew J. Wetstein et al. **Systems-Level Characterization of Microchannel Plate Detector Assemblies, using a Pulsed sub-Picosecond Laser**. *Phys. Procedia*, **37**:748–756, 2012. 214

[195] Viacheslav A. Li, Timothy M. Classen, Steven A. Dazeley, Mark J. Duvall, Igor Jovanovic, Andrew N. Mabe, Edward T.E. Reedy, and Felicia Sutanto. **A prototype for SANDD: A highly-segmented pulse-shape-sensitive plastic scintillator detector incorporating silicon photomultiplier arrays**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **942**:162334, 2019. 214

[196] Carsten Rott, Seongjin In, Fabrice Retire, and Peter Gumplinger. **Enhanced Photon Traps for Hyper-Kamiokande**. *JINST*, **12**(11):P11021, 2017. 216

[197] A. A. Machado and E. Segreto. **ARAPUCA a new device for liquid argon scintillation light detection**. *JINST*, **11**(02):C02004, 2016. 216

[198] Teresa Marrodan Undagoitia. *Measurement of light emission in organic liquid scintillators and studies towards the search for proton decay in the future large-scale detector LENA*. PhD thesis, Munich, Tech. U., 2007. 229

[199] T. Marrodan Undagoitia, F. von Feilitzsch, L. Oberauer, W. Potzel, A. Ulrich, J. Winter, and M. Wurm. **Fluorescence decay-time constants in organic liquid scintillators**. *Rev. Sci. Instrum.*, **80**:043301, 2009. 236, 248

[200] T.A. DeVol, D.K. Wehe, and G.F. Knoll. **Evaluation of p-terphenyl and 2,2 dimethyl-p-terphenyl as wavelength shifters for barium fluoride**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **327**(2):354 – 362, 1993. 236

[201] A. A. Machado, E. Segreto, D. Warner, A. Fauth, B. Gelli, R. Maximo, A. Pizolatti, L. Paulucci, and F. Marinho. **The X-ARAPUCA: An improvement of the ARAPUCA device**. *arXiv:1804.01407 [physics.ins-det].*, 2018. 236

[202] Stan Seiber and Tony LaTorre. **Fast Optical Monte Carlo Simulation With Surface-Based Geometries Using Chroma**, 2011. 240

[203] Benjamin Land. **Chroma**. `https://github.com/BenLand100/chroma`, 2019. 240

[204] C. Galbiati, M. Misiaszek, and N. Rossi. **α/β discrimination in Borexino**. *Eur. Phys. J.*, **A52**(4):86, 2016. 257

[205] Michael Wurm et al. **The next-generation liquid-scintillator neutrino observatory LENA**. *Astropart. Phys.*, **35**:685–732, 2012. 257

[206] J. Ashenfelter et al. **Light Collection and Pulse-Shape Discrimination in Elongated Scintillator Cells for the PROSPECT Reactor Antineutrino Experiment**. *JINST*, **10**(11):P11004, 2015. 257

[207] M. Askins et al. **Status of Data Cleaning During the SNO+ Water Phase**. *SNO+ Internal Document Database*, **5430**, 2017. 270

[208] Monica Dunford. *Measurement of the ^8B solar neutrino energy spectrum at the Sudbury Neutrino Observatory*. PhD thesis, Pennsylvania U., 2006. 278

[209] R. Mirzoian, E. Lorenz, D. Petry, and C. Prosch. **On the influence of afterpulsing in PMTs on the trigger threshold of multichannel light detectors in selftrigger mode**. *Nucl. Instrum. Meth.*, **A387**:74–78, 1997. 279

[210] B. A. Moffat, R. J. Ford, F. A. Duncan, K. Graham, A. L. Hallin, C. A. W. Hearns, J. Maneira, P. Skensved, and D. R. Grant. **Optical calibration hardware for the sudbury neutrino observatory**. *Nucl. Instrum. Meth.*, **A554**:255–265, 2005. 287

[211] R. Alves et al. **The calibration system for the photomultiplier array of the SNO+ experiment**. *JINST*, **10**(03):P03002, 2015. 288

[212] Costas Andreopoulos, Christopher Barry, Steve Dytman, Hugh Gallagher, Tomasz Golan, Robert Hatcher, Gabriel Perdue, and Julia Yarba. **The GENIE Neutrino Monte Carlo Generator: Physics and User Manual**. 10 2015. 296

[213] A. Bodek and J.L. Ritchie. **Further Studies of Fermi Motion Effects in Lepton Scattering from Nuclear Targets**. *Phys. Rev. D*, **24**:1400, 1981. 296

[214] Artur M. Ankowski, Omar Benhar, Takaaki Mori, Ryuta Yamaguchi, and Makoto Sakuda. **Analysis of γ-Ray Production in Neutral-Current Neutrino-Oxygen Interactions at Energies above 200 MeV**. *Phys. Rev. Lett.*, **108**:052505, Feb 2012. 296