

Cosmogenic Muon Induced Background Tagging in the SNO+ Detector

Lorna Jane Nolan

Submitted for the degree of Doctor of Philosophy

Queen Mary University of London

May 2023

**What doesn't kill you
makes you stronger**

KELLY CLARKSON

I, Lorna Jane Nolan, confirm that the research included within this thesis is my own work or that where it has been carried out in collaboration with, or supported by others, that this is duly acknowledged below and my contribution indicated. Previously published material is also acknowledged below.

I attest that I have exercised reasonable care to ensure that the work is original, and does not to the best of my knowledge break any UK law, infringe any third party's copyright or other Intellectual Property Right, or contain any confidential material.

I accept that the College has the right to use plagiarism detection software to check the electronic version of the thesis. I confirm that this thesis has not been previously submitted for the award of a degree by this or any other university.

The copyright of this thesis rests with the author and no quotation from it or information derived from it may be published without the prior written consent of the author.

Signature:

Date:

Details of collaboration and publications: This work was carried out by the author in collaboration with the SNO+ collaboration.

QUEEN MARY UNIVERSITY OF LONDON

LORNA JANE NOLAN, DOCTOR OF PHILOSOPHY

COSMOGENIC MUON INDUCED BACKGROUND TAGGING IN THE SNO+ DETECTOR

ABSTRACT

This thesis presents the rate of Carbon-11 production in the SNO+ scintillator phase. Carbon-11, created by a cosmogenic muon interacting with Carbon-12, is a background to the solar neutrino studies planned for the SNO+ scintillator phase. With a half-life of 20.3 minutes, Carbon-11 cannot be removed from analyses with a simple 20 or 60 second post-muon veto as many of the muon-induced backgrounds in SNO+ are. Instead Carbon-11 can be identified using a three-fold coincidence, firstly the muon passing through the scintillator, then a coincident neutron capture, and finally an event with the expected energy of a Carbon-11 decay that is sufficiently close spatially to the neutron capture within several half-lives. This thesis details the steps taken to tune the pre-existing muon identification from water phase to scintillator phase, the development of a coincident neutron capture identification, and the results of both these steps in partial fill and scintillator phases. This thesis then outlines the development of a likelihood algorithm that calculates a likelihood for every Carbon-11 candidate, and shows the results of this algorithm in the scintillator phase. A likelihood was chosen to allow for a strict cut to remove Carbon-11 from analyses, while also allowing a pure sample of Carbon-11 to be found in order to perform a data-MC calibration. A background rejection cut of 0.01 likelihood was chosen, removing 94% of the ^{11}C with only a 1% sacrifice of signal. Then two higher purity likelihood cuts are presented, 0.2 and 0.4. The 0.2 cut has a purity of 0.6, whilst the 0.4 cut has a purity of 0.9, but is statistically limited.

Acknowledgements

First and foremost I must profusely thank my incredible supervisor Jeanne Wilson, without who's endless patience, support and guidance, both personally and professionally, not one page of this work would have been possible.

I would also like to thank Seth Zenz for taking over as my Queen Mary supervisor after my first year.

During my time as a PhD student I have spent time at both QMUL and KCL, where I have been lucky enough to have the companionship of an excellent group of people, especially Joe Davis my now housemate, and our numerous tea breaks have helped keep us both sane over the last 18 months. I also must thank Katharine Dixon and Ewan Miller, for helping to make me feel fully welcomed into the group at KCL. On the QMUL side I would like to thank Billy Liggins and Mark Stringer for welcoming me into the group when I first started the PhD and helping me to get acquainted to all things SNO+.

I am grateful to the entire SNO+ collaboration for all of their help and feedback on my work. I was also lucky enough to spend time at SNOLAB where I met some of the most incredible people who helped me grow both personally and professionally. Most notably Ana Sofia Inacio, Charlie Mills, Steph Walton, Ben Tam, Rachel Richardson, Erica Caden, Cindy Lin, Matt Depatie, and Josie Paton. I would also like to thank Lisa Falk and the whole SNO+ group at Sussex for letting me join their group meetings when the pandemic started, it was invaluable to have a wider circle of people to discuss this work with. Although our time in the collaboration never overlapped, here seems the most appropriate place to thank Ashley Back, who has always supported and encouraged my interest in Physics, and is how I first heard of the experiment. I must thank Ana Sofia twice, as she also became my writing buddy this last year, without her support, company, and cheesecake recipes, I sincerely do not think I could have written this thesis.

I offer sincere thanks to Marty St-Amant, for being with me every step of the way for the last three years, offering only endless patience with every frustration I encountered during this work.

I would not have made it this far without my family always being by my side. My mum

and sister for always asking me to explain my thesis, offering to proofread it, and asking me to practise my talks on them, even if they do not understand it. My Dad, Judith, and Katie for moving me to London twice, and letting me stay with you at the very start of the pandemic. My grandma for letting me move in with her for over a year, keeping me company during the lockdowns, and always listening to me talk about this work.

Contents

List of Tables	xi
List of Figures	xvii
1 Introduction	1
2 A Brief History of Neutrino Physics	3
2.1 Neutrinos: The Elusive Solution to Beta Decay	3
2.1.1 The Beta Decay Energy Problem	3
2.1.2 Further Theoretical Indication of Neutrinos	4
2.1.3 The first experimental evidence for neutrinos	5
2.1.4 Evidence for antineutrinos and neutrino flavours	6
2.2 Neutrino Interactions in the Standard Model	7
2.3 The solar neutrino problem	10
2.3.1 The Homestake experiment	10
2.3.2 Neutrino oscillations	12
2.3.3 SNO and Super-K observe neutrino oscillations	16
2.4 Open Questions Regarding Neutrino Mass	17
2.4.1 Specific Neutrino Masses	17
2.4.2 Neutrino Mass Ordering	18
2.4.3 Mass Mechanism: Dirac or Majorana	19
2.4.4 Neutrinoless Double Beta Decay	21
3 The SNO+ Detector	24
3.1 The SNO+ detector	24
3.1.1 The location	24
Mitigation of Radon in the mine air	24
3.1.2 The Acrylic Vessel	25

3.1.3	The PhotoMultiplier Tubes	26
3.1.4	The SNO+ Trigger System	27
3.2	Data Processing in SNO+	29
3.2.1	Data Cleaning in SNO+	29
	Muon based cuts in Data cleaning	30
3.2.2	Hit Cleaning in SNO+	30
3.3	Water Phase	31
3.3.1	Detector Specifics	31
3.3.2	Physics Results	31
3.3.3	Backgrounds	32
3.3.4	Calibrations	32
3.4	Scintillator Phase	33
3.4.1	Detector specifics	33
3.4.2	Physics Goals	33
	Geoneutrinos	33
	Reactor neutrinos	34
	Supernovae neutrinos	35
	Solar neutrinos	35
3.4.3	Backgrounds	35
3.4.4	Calibrations	36
3.4.5	Scintillator Filling	36
	Scintillator Production and Transportation to SNOLAB	36
	Transporting Scintillator Underground	37
	PPO Loading	37
	The Scintillator Purification Plant	37
	Quality Assurance (QA)	38
3.4.6	Stable Partial Fill Phase	40
	Calibrations	40
3.5	Tellurium-Loaded Phase	40
3.5.1	Detector Specifics	40
3.5.2	Physics Goals	41
3.5.3	Backgrounds	41
4	Channel Software Status	43
4.1	Overview of Channel Software Status	43

4.1.1	Standard Run	44
4.2	Changes to CSS for Scintillator Phase	44
4.3	Channel Software Status Tests	45
4.3.1	High Occupancy	45
	High Occupancy in scintillator phase	46
4.3.2	Low Occupancy	47
	Low Occupancy in Stable Partial Fill Phase	47
	Low Occupancy in Scintillator Phase	48
4.3.3	Short Window Charge Cuts	50
5	What did Carbon-11 ever do to us?	54
5.1	Introduction	54
5.2	Cosmogenic Muons	54
5.2.1	Primary Cosmic Rays	54
5.2.2	Secondary Cosmic rays and Cosmic muon production	55
5.2.3	Muons in the SNO+ detector	57
	Muon Reconstruction in SNO+	59
5.3	Muon Induced Backgrounds in Scintillator	60
5.3.1	Production method	60
5.3.2	Long Lived Isotopes	60
	^{14}C	60
	^7Be	61
	^{39}Ar	61
5.3.3	Short Lived Isotopes	61
5.3.4	Muon Induced Isotopes in SNO+	62
	^{11}C	62
	^{16}N	63
	^8Li and ^8B	64
	^6He	64
	^{10}C	64
5.4	Three-fold coincidence Tagging	64
6	Partial Fill	65
6.1	Water Phase Muon Tag in Partial Fill Data	65
6.2	Altering the Muon Tag for Partial Fill Phase	68

	Nhits threshold adjustment	70
6.3	Results of these modifications	72
6.3.1	In stable partial Fill	72
	OWL hit distribution in partial fill	73
6.4	Neutron Tag	75
6.5	Integrating these changes into the DC framework	77
7	Tagging Muons and Coincident Neutron Captures in Scintillator Phase	78
	Livetime and run list	78
7.1	Muon Tagging	79
7.2	Neutron Tagging	79
7.2.1	Data Cleaning Level Cuts	80
	Secondary Cuts to Address High Volume of Early Time Neutron Capture Candidates	80
8	Carbon-11 Likelihood Classifier	84
8.1	The Framework	84
8.1.1	The PDFs	84
	Nhits	85
	Time	85
	Spatial	85
8.2	Tuning the cuts	87
8.2.1	10 percent of sample	87
	Tuning the alpha cut value	87
8.3	Systematic considerations	90
8.4	The results	91
9	Conclusion	95
	Bibliography	97
A	AppendixA	104

List of Tables

2.1	The best-fit values and 3σ ranges of Δm_{21}^2 and Δm_{32}^2 neutrino mass parameters. From [34]. For details on normal and inverted mass ordering see section 2.4.2.	17
2.2	An overview of some current and predicted half-life sensitivities to $0\nu\beta\beta$ [43]. The (*) denotes a predicted sensitivity. A longer half life sensitivity shows an experiment is more sensitive.	23
4.1	A table showing the CSS charge cut criteria. A channel is flagged by the calibration if the percentage of charge in any of these regions meets the criteria compared to the amount of charge across the entire channel.	52
5.1	Cosmogenically induced backgrounds found in Super-K [76] and Borexino [65].	62
6.1	A comparison of the criteria for a muon between partial fill and water phase.	67

List of Figures

2.1	Measured Electron Energy for Tritium beta Decay [1].	4
2.2	Pion decays as observed by Powell, [a] shows just pion decay, [b] includes the subsequent muon decay [1].	5
2.3	A visualisation of the two parity states from Wu's cobalt decay experiment [14].The curved arrow indicates the direction of parity, with the second arrow demonstrating the direction the electron was released in.	10
2.4	The neutrino energies from solar decay chains [17].	11
2.5	Kaon decay into pions, showing a strange quark changing into a down quark, a transition not thought possible before quark mixing.	13
2.6	Feynman diagrams depicting neutral current scattering off electrons (a) and nucleons (c) and charged current scattering off electrons (b) and nucleons (d)	16
2.7	The two possible mass orderings (hierarchies) for neutrino mass, inverted and normal [37].	18
2.8	Double beta Decay (top) and Neutrinoless Double beta Decay (bottom) Feynmann diagrams [42]	22
2.9	The sum of electron energies, assuming an energy resolution of 2% [42]. The dotted line shows the energy spectrum for $2\nu\beta\beta$, and the solid line peak shows the energy spectrum for $0\nu\beta\beta$	22
3.1	The SNO+ Detector	25
3.2	A depiction of the electronic set up of SNO+ [51]	26
3.3	The flow of data through the SNO+ Data Acquisition System [52].	27
3.4	Predicted antineutrino spectrum in SNO+ scintillator phase [61]	34
3.5	Predicted backgrounds for the solar neutrino searches at SNO+, assuming that Borexino levels of backgrounds [65] are achieved.	36

3.6	UV-Vis spectra for two samples of LAB-PPO. [a] shows a good sample with excellent optical clarity, whilst [b] demonstrates a slight build up of contaminants as indicated by the peak appearing at 380 nm. The baseline shift seen in [b] characterised by the curve beginning at 360 nm instead of 330 nm is due to the higher PPO concentration in that sample and is no indication of the optical quality.	39
3.7	A plot showing the normal and inverted hierarchy, as well as some of the current experimental limits [66].	41
3.8	Predicted backgrounds for the SNO+ Tellurium phase in the neutrinoless double beta ROI.	42
4.1	High Occupancy Plot for run 207417, a Standard Run in the water phase. High Occupancy Statistic refers to the number of Pulsed Global Trigger hits registered on that channel divided by all Pulsed Global Trigger hits registered across the entire detector. Both of these quantities are over a whole run, and this statistic is plotted for every online channel in the run. The shaded area above the peak shows the region marked as high occupancy	46
4.2	The High Occupancy statistic distribution for water phase, partial fill phase, and scintillator phase. It shows no significant variation between the phases that would require the cut to be retuned.	47
4.3	Low Occupancy Plot for run 207417, a Standard Run in water phase. Low Occupancy Statistic refers to the number of hits registered of any trigger type on that channel divided by all hits registered across the entire detector of any trigger type. Both of these quantities are over a whole run, and this statistic is plotted for every online channel in the run. The shaded area below the peak shows the region marked as low occupancy.	48
4.4	Flat Map Plot for low occupancy water phase	49
4.5	Flat Map Plot for low occupancy stable partial fill phase	49
4.6	The Low Occupancy statistic distribution for water phase, partial fill phase, and scintillator phase. It shows the progression of the single peak into two, back into one higher, thinner peak as fill was finished.	50
4.7	Average number of channels flagged as low occupancy at a variety of FWHMs over multiple runs. Errors are statistical from averaging over multiple runs.	51

4.8	Low Occupancy statistic plotted for a run with 3.5 FWHMs (blue line) and 5 FWHMs (green line) below the fitted mean shown to compare the water phase and scintillator phase thresholds. The red line is a gaussian fit of the first half of the peak.	51
4.9	This shows the expected distribution of QHS charge per channel over an entire run (black histogram), and the possible deviations from this. The purple histogram represents a channel with a too high pedestal, registering negative charge. Red histogram in conjunction with the expected black peak means a double pedestal has been recorded for this channel. The blue peak represents a railed channel, which has become saturated.	52
5.1	Primary cosmic ray nuclei fluxes in particles per energy-per-nucleus [70]. . .	55
5.2	Cosmic ray energy distribution. Measurements taken by the Auger observatory [71]. The green line represents the approximate integral flux. . . .	56
5.3	A depiction of a hadronic shower, showing both the hadronic component and the electromagnetic component. Figure from [72].	56
5.4	The variation of muon flux with depth of several underground laboratories [73].	58
5.5	A power-law fit of the cosmogenic neutron yield in liquid scintillator (by the Daya Bay Collaboration). There are from experiments at different depths resulting in different average muon energies. Results from SNO Phase 1 are overlaid but not included in the fit, as the measurement is in water not liquid scintillator [74].	59
5.6	^{11}C Spectrum shown alongside the expected low energy solar signals.[61] . .	63
6.1	OWL PMT hits distribution of events identified as muons by running water phase criteria on water phase data (blue), and partial fill data (black). Area normalised for comparison.	66
6.2	Nhits distribution of events identified as muons by running water phase criteria on water phase data (blue), and partial fill data (black). Area normalised for comparison.	66
6.3	Nhits against OWL PMT hits of events identified as muons by running water phase criteria on partial fill data. The red line is at 8 OWL PMT hits, to demonstrate the potential impact that increasing the OWL PMT threshold could have.	67

6.4	An example of an event falsely tagged as a muon as viewed in the SNO+ event display.	68
6.5	An example of an event correctly tagged as a muon as viewed in the SNO+ event display.	69
6.6	Mean of the position in Z of OWL PMT hits. Red line shows events believed to be muons, blue line shows events believed to have been falsely tagged whilst the black line has not been identified either way.	71
6.7	RMS of the TAC of OWL PMT hits. TAC is a measure of the timing of PMT hits within an event, for more detail see Chapter 3. Red line shows events believed to be muons, blue line shows events believed to have been falsely tagged whilst the black line has not been identified either way. . . .	71
6.8	RMS of the position in Z of OWL PMT hits, with [b] zoomed in to show only the relevant section of [a]. Red line shows events believed to be muons, blue line shows events believed to have been falsely tagged whilst the black line has not been identified either way	72
6.9	Average hourly muon tagging rate over a sample of partial fill data with varying minimum nhits. Error bars too small to be seen, error is statistical from averaging over multiple runs.	73
6.10	The muon rate per hour in water phase [a] compared to the muon rate per hour in partial fill with the adapted cuts [b]. Data from the stable partial fill period of summer 2020.	73
6.11	The time between muons in water phase [a] compared to the time between muons in partial fill with the adapted cuts [b]. This is only muons within a run, the time between muons from other runs is excluded. Data from the stable partial fill period of summer 2020.	74
6.12	The time between muons in water phase [a] compared to the time between muons in partial fill with the adapted cuts [b]. This includes the time between muons even if they're from a different run. Data from the stable partial fill period of summer 2020.	74
6.13	Comparison of water phase muons and partial fill muons with retuned cuts for nhits [a] and OWL PMT hits [b].	75

6.14	An example of an event that saturated the OWLs viewed in xsnoed. The charge in crate 4 indicates this is a breakdown rather than a Physics event, as almost the whole crate has a very high charge, much higher than the majority of the rest of the PMTs. In a true Physics event we see channels which are close spatially have similar charge distributions, rather than just channels that are in the same crate. The number of OWL PMT hits are shown in the textbox in the event viewer as well as the small dots on the very edge of the sphere.	76
6.15	The results of the neutronTag processor over stable partial fill data. [a] shows the time in nanoseconds between the muon and the neutron capture candidate, whilst [b] shows the nhits of the neutron capture candidates compared to the simulation.	76
7.1	Time between muons (in seconds) for muons in scintillator phase of SNO+.	79
7.2	Time between muons and coincident neutron captures in scintillator phase of SNO+. The red line shows a fit using an exponential plus a constant. Tau represents the neutron capture livetime.	80
7.3	Distance a muon travels through the AV (m) plotted against the number of PMTs triggered in the event. Made from MC simulations in RAT-6.18.12.	82
7.4	NhitsCleaned divided by nhits for all coincident neutron capture candidates in a 3 week dataset. Events which score low are ‘confetti’ events, events which are purely electronic noise. These events have no correlation in time unlike Physics events which makes them appear multi-coloured in the SNO+ event display, resulting in their nickname of confetti events.	82
7.5	Time between muon and coincident neutron capture for a 3 week dataset in scintillator phase of SNO+. This plot contains only events with a minimum nhitsCleaned of 70. The red line shows a fit using an exponential plus a constant. Tau represents the neutron capture livetime.	83
8.1	The nhits PDF of ^{11}C candidates. This PDF contains simulated events, made using RAT-6.18.12. A fiducial volume of 5.5 m was set on these simulations to match the fiducial volume used by the likelihood algorithm. . . .	85
8.2	The time PDF of ^{11}C candidates. Made from exponential function with τ of 20.3 minutes.	86

8.3	The spatial separation PDF of ^{11}C candidates and coincident neutron captures. This was made from MC simulations of coincident neutron captures, and was smeared with the uncertainties of the ^{11}C and coincident neutron capture reconstructions.	86
8.4	A sample of 10% of the data sample for ^{11}C candidates, showing the likelihood distribution for this sample.	87
8.5	A sample of 10% of the data sample for ^{11}C candidates, showing the timing distribution for events with a likelihood of 0.2 or above in black, and events with a likelihood of 0.4 or above in red.	88
8.6	The black line shows the theoretical likelihood distribution of ^{11}C , made from randomly sampling values for nhits, spatial separation, and time since coincident neutron capture from simulations and then running those values through the likelihood algorithm. The red line shows the random coincidence distribution, found by selecting events below the ^{11}C nhits region, and running these events through the likelihood algorithm with an nhits value randomly sampled from the ^{11}C simulations.	89
8.7	The red line is the efficiency and the black line is the purity. As this is to remove a background for the solar neutrino studies, it was also important to ensure that any cuts do not also remove a significant fraction of the solar neutrinos. The fraction of the solar neutrinos removed by each cut value can also be seen in the green line.	90
8.8	The nhits for all ^{11}C candidates with a probability of 0.4 or greater in black, in green is the MC simulation nhits distribution for ^{11}C . Both are normalised to area under the histogram to be one for comparison.	92
8.9	The nhits for all ^{11}C candidates with a probability of 0.2 or greater in black, in green is the MC simulation nhits distribution for ^{11}C . Both are normalised to area under the histogram to be one for comparison.	93
8.10	The timing distribution for the time between muon and ^{11}C candidate for events with a likelihood of 0.2 or higher.	93
8.11	The timing distribution for the time between muon and ^{11}C candidate for events with a likelihood of 0.4 or higher.	94

Chapter 1

Introduction

The SNO+ detector is a multipurpose neutrino detector located in Sudbury, Ontario. Consisting of an Acrylic Vessel housing the detector medium surrounded by 9300 photomultiplier tubes (PMTs) all suspended in a cavity of ultra pure water 2 km underground, SNO+ is the successor of the Nobel Prize winning SNO experiment, and will ultimately search for neutrino less double beta decay using natural Tellurium-130 dissolved in 780 tonnes of liquid scintillator. Currently, SNO+ is in its scintillator phase with a goal of understanding and categorising detector response and backgrounds before adding Te as well as carrying out physics analyses such as solar neutrino studies. Carbon-11, created by a cosmogenic muon interacting with Carbon-12, is a background to these solar neutrino studies. With a half-life of 20.3 minutes, Carbon-11 cannot be removed by the existing simple post-muon vetos. Instead Carbon-11 is identified using a three-fold coincidence; the muon, a coincident neutron capture, then an event in the energy region expected from a Carbon-11 decay close in time and space to the neutron capture. This thesis outlines the steps taken to tune and develop each of these three steps to develop a likelihood algorithm for Carbon-11 decays.

This thesis is presented as follows. Chapter 2 presents the relevant history and neutrino theory, then Chapter 3 describes the SNO+ detector construction, operation, and Physics goals. This chapter includes a discussion of the data cleaning framework, which is relevant for the muon and neutron tagging described later, and a detailed description of scintillator filling operations and the quality assurance process that the author worked on whilst at SNOLAB. In Chapter 4, the channel software status processor is described, including tuning and new charge based cuts implemented by the author. In Chapter 5, the production of ^{11}C and other cosmogenically induced backgrounds are discussed, then in Chapter 6 the steps to tag muons and neutrons in the partial fill phase are presented.

Chapter 7 then describes muon and neutron tags in the full scintillator phase. Chapter 8 presents the full likelihood classifier developed to identify the ^{11}C events in scintillator phase and the main results achieved from the limited data available at the time of writing. Chapter 9 presents the conclusions of this thesis work.

Chapter 2

A Brief History of Neutrino Physics

2.1 Neutrinos: The Elusive Solution to Beta Decay

2.1.1 The Beta Decay Energy Problem

The theory of neutrinos begins with a discrepancy between the theoretical prediction and the experimental observations of the energy of electrons emitted by beta decay in the 1930s. Beta decay at the time was thought to be a two-body decay, $A \rightarrow B + e^-$, and thus the energy of the resulting electron should be fixed, specifically at the value showed in [1]:

$$E = \left(\frac{m_A^2 - m_B^2 + m_e^2}{2m_A} \right) c^2 \quad (2.1)$$

where m_A and m_B are the masses of the parent and daughter nucleus respectively. During experimental observations of the energy of electrons emitted via beta decay, the electron energy was found to vary. Equation 1 was found to be the maximum possible energy not the fixed energy. An example of this variation from an experiment measuring the energy of electrons emitted via beta decay of tritium can be seen in Figure 2.1.

This result seemed to imply that energy conservation was not being obeyed, leading Niels Bohr to question the validity of the principle [1]. Others however looked into alternate possibilities of the missing energy, with Meitner ruling out neutral gamma rays as a possibility [2].

This led Pauli to his theory of an electrically neutral, spin-1/2 particle with mass on the order of that of an electron or lower. However he did not believe the particle to be detectable, which led him to be hesitant in his announcement, as he did not think




Figure 2.1: Measured Electron Energy for Tritium beta Decay [1].

it was right for a theorist to suggest a particle that cannot be found. He proposed the new particle in his now famous ‘Dear Radioactive Ladies and Gentlemen’ letter to an experimental meeting in Germany. He suggested it should be called a ‘neutron’, meaning neutral one in Italian (a translation of this letter can be found in [3]).

2.1.2 Further Theoretical Indication of Neutrinos

When Chadwick discovered a neutral particle [4], [5], in 1932, it was named the neutron before it was found to be too heavy to match with Pauli’s initial predictions. Fermi began working on a theory of beta decay which would incorporate Pauli’s predicted particle, which he then named ‘neutrino’, for small neutral one in Italian. Fermi’s theory made beta decay a three-body decay, including this massless or near massless particle which accounts for the variations found in the energy of the electron. This theory proved widely successful, which led to far wider support for the neutrino.

In the 1940’s, Powell [6] was studying pion decay, by exposing photographic plates to cosmic rays. This resulted in seeing the tracks from the charged pion and the charged decay products, as can be seen in Figure 2.2. Small deviations can be seen in the path of the pion, as a result of collisions with other particles in the plate. However, when the pion decays as can be seen in Figure 2.2[a], it decays with an almost 90° angle. This is not the result of collisions, and implies that another particle is released in the decay, which cannot be seen on the slide, meaning it has no charge. The natural assumption was that this unseen particle was the same neutrino released in beta decay. As can be seen in Figure 2.2[b], an image released shortly after, the same angle occurs when the muon decays. This observation again implies that there is another neutral particle released at




Figure 2.2: Pion decays as observed by Powell, [a] shows just pion decay, [b] includes the subsequent muon decay [1].

the same time, but measurements of the electron it decays into show the same variations in energy as was observed in beta decay, meaning the decay was a three-body problem. The subsequent conclusion drawn was that two electrically neutral particles were released, or two neutrinos. Whilst this was a good assumption, it later proved to only be partially correct as three different flavours of neutrinos were discovered corresponding to the three flavours of charged leptons, meaning the neutrino released in beta decay was not the same as released in pion decay.

2.1.3 The first experimental evidence for neutrinos

After these discoveries the theoretical evidence for the existence of neutrinos was very strong, but there had still been no experimental verification of these predictions, due to the neutrinos interacting weakly with matter. In the mid-fifties, an experiment took place with Cowan and Reines setting up a water tank by the Savannah River nuclear reactor [7]. The nuclear reactor would give off a huge flux of antineutrinos, 5×10^{13} anti-neutrinos per square centimetre per second. A very small amount of these would trigger inverse beta decay


$$\bar{\nu} + p^+ \rightarrow n + e^+ \quad (2.2)$$

as they passed through the water. They planned to detect the positrons being released from this reaction as they annihilated with electrons on their journey through the water tank. To achieve this they placed PhotoMultiplier Tubes (PMTs) around the tank and scintillator material in the tank to identify the gamma rays.

Unfortunately, the results were not conclusive enough so an additional layer of verification was given by adding Cadmium Chlorine into the water, the material used in nuclear reactors due to its powerful neutron absorption properties. When it absorbs a neutron, the Cadmium then releases a gamma ray. By timing coincidence between the gamma ray from the positron annihilation and the gamma ray from the neutron absorption, they observed the first experimental evidence of the existence of the neutrino. They measured around three neutrinos per hour interacting with their detector. As a secondary check, they shut off the reactor for a short time to prove that these reactions were indeed a result of the neutrino flux from the reactor.

2.1.4 Evidence for antineutrinos and neutrino flavours

This was evidence of antineutrinos, which led many to question the difference between the neutrino and the antineutrino: were they in fact distinct particles? This questions the nature of neutrinos: if the antineutrino is indistinguishable from the neutrino, they would be Majorana, if they are two unique particles they would be Dirac (more on this in section 2.4.2). Davis and Harmer [8] designed an experiment designed to test whether the two particles were distinct. This was based on the knowledge that the crossed reaction of Cowan and Reines (see Equation 2.2), given in Equation 2.3 would also be possible, and in fact occur at a very similar rate.



Davis and Harmer set out to search for this crossed reaction, but occurring from an antineutrino instead.



They concluded the reaction does not occur with an antineutrino, establishing that there is a difference between neutrinos and antineutrinos.

This outcome was expected by Konopinski and Mahmoud, after their paper in 1953 [9] where they introduced the early concept of Lepton numbers, albeit by a different terminology. They assigned the leptons to be ‘normal’ or ‘antiparticles’. They assumed that neutrinos were Dirac particles, stating they believed the possibility of them being

Majorana was slim. They then assigned the neutrino and electron as ‘normal’ particles, along with the positive muon after some consideration. Their theory was that the amount of normal particles on each side of the interaction had to balance or the interaction was not possible. Lepton numbers now are assigned as a +1 for the ‘normal’ particle and -1 for the antiparticle with them still having to balance either side in a valid interaction. The muon assignments have also been switched, with the positive muon now being considered the antiparticle. Using this logic it can be seen that while Equation 2.3 is fine, Equation 2.4 violates lepton number conservation. This property is one of the distinguishing features of neutrinos and antineutrinos, with the other being helicity. The neutrino is left-handed while the antineutrino is right-handed.

These observations reshuffle the neutrinos/antineutrinos around slightly in earlier theories involving them, but left one key question in the theory, why a muon is never observed to decay into an electron and a photon. This would appear to be a valid decay, as it does not violate charge or lepton number conservation. The totalitarian saying ‘whatever is not forbidden is mandatory’ was co-opted by some Physicists of the time, implying there is a rule preventing this decay that has not yet been discovered. In the early sixties this rule was revealed to be flavour conservation. It was predicted that there was a muon neutrino and an electron neutrino, leading to lepton number to be separated for muons and electrons. A test of this new theory was devised at Brookhaven [10] using pion decay, identified $\bar{\nu}_\mu + p^+ \rightarrow \mu^+ + n$ decay and none of $\bar{\nu}_\mu + p^+ \rightarrow e^+ + n$ decay which, if there was only one neutrino flavour should have been equally as likely. In 1975, the Tau lepton was discovered. It also has its own lepton flavour number and was immediately predicted to have its own neutrino as well (a prediction ultimately confirmed by Fermilab in 2000 [11]).

2.2 Neutrino Interactions in the Standard Model

The Standard Model (SM) describes all elementary particles and their interactions via the strong, weak, and electromagnetic forces. Neutrinos are unique within this as being the only particles to interact mainly through only one of the forces, the weak force. The SM can be written as $SU(3)_C \times SU(2)_L \times U(1)_Y$ where C is colour, L is left-handed chirality and the Y is weak hypercharge. Eight massless bosons in the SM that mediate the strong force correspond to the eight generators of the $SU(3)_C$ component. There are three generators of $SU(2)_L$ and one of $SU(1)_Y$. These are the three massive bosons, W^\pm, Z^0 , and the massless photon, γ . The symmetry of the colour group, $SU(3)_C$ is unbroken and there

is no mixing between this and the other two components of the SM, meaning that there is no need to study the strong force connected with the electromagnetic and weak forces. An overview of the electroweak sector ($SU(2)_L \times U(1)_Y$) of the SM is presented here, for a more detailed account see [12].

The symmetry group $SU(2)_L$ is the weak isospin and only acts on the left-handed chiral components of the fermion fields. It has three generators,

$$I_a \quad (a = 1, 2, 3). \quad (2.5)$$

The hypercharge symmetry group $U(1)_Y$ is generated by hypercharge operator Y . This is linked to I_3 and charge operator Q through the Gell-Mann-Nishijima relation, which implies the unification of electroweak interactions:

$$Q = I_3 + \frac{Y}{2}.$$

In order to have local gauge invariance, three vector gauge bosons need to be introduced. These are $A_a^\mu (a = 1, 2, 3)$ and are connected with the generators in Equation 2.5, for the $SU(2)_L$ group and a single vector gauge boson field B^μ connected with the generator Y from the $U(1)_Y$ group.

In gauge theories, the covariant derivative D_μ replaces the normal derivative, δ_μ . This covariant derivative can be written as:

$$D_\mu = \partial_\mu + ig \underline{A}_\mu \cdot \underline{I} + ig' B_\mu \frac{Y}{2}, \quad (2.6)$$

where coupling constant g is associated with $SU(2)_L$ and coupling constant g' is associated with $U(1)_Y$.

Choosing the representations of the fermion fields is the next step, which is usually based on the V-A theory of weak interactions. To simplify the derivation, consider just the first generation of leptons and quarks (for the three-generation expansion, see [12]), with the left-handed components grouped into weak isospin doublets:

$$L_L = \begin{pmatrix} \nu_{eL} \\ e_L \end{pmatrix}, \quad Q_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix}.$$

This fixes the generators of the $SU(2)_L$ group to be $I_a = \tau_a/2$,

$$\underline{I}L_L = \frac{\tau}{2}L_L, \quad \underline{I}Q_L = \frac{\tau}{2}Q_L,$$

where $\underline{\tau} = (\tau_1, \tau_2, \tau_3)$, while Y is still fixed by the Gell-Mann-Nishijima relation as:

$$YL_L = -L_L, \quad YQ_L = \frac{1}{3}Q_L.$$

This means the left-handed lepton doublets have $Y = -1$ and the left-handed quark doublets have $Y = \frac{1}{3}$.

The right-handed components of the fermions (e_R, u_R, d_R) can be assumed as singlets under the weak isospin transformations. They have hypercharge

$$Ye_R = -2e_R, \quad Yu_R = \frac{4}{3}u_R, \quad Yd_R = -\frac{2}{3}d_R.$$

To obtain the interactions between fermions and gauge bosons, expand the covariant derivative in Equation 2.6 to gain the interaction Lagrangian. Considering just the lepton terms, and splitting this into the off-diagonal terms (charged-current) and the diagonal terms (neutral-current), gives the charged current interaction Lagrangian as:

$$\mathcal{L}_{1,L}^{(CC)} = -\frac{g}{2} \{ \bar{\nu}_{eL} (A_1 - iA_2) e_L + \bar{e}_L (A_1 + iA_2) \nu_{eL} \}.$$

By introducing a field that annihilates W^+ bosons to create W^- bosons,

$$W^\mu \equiv \frac{A_1^\mu - iA_2^\mu}{\sqrt{2}},$$

this gives the charged current interaction Lagrangian as:

$$\mathcal{L}_{1,L}^{(CC)} = -\frac{g}{2\sqrt{2}} j_{W,L}^\mu W_\mu + \text{H.c.},$$

where $j_{W,L}^\mu = 2\bar{\nu}_{eL}\gamma^\mu e_L$, the leptonic charge current.

By next considering the Quantum Electrodynamics Lagrangian,

$$\mathcal{L}_{1,L}^{(\gamma)} = -e j_{\gamma,L}^\mu A_\mu,$$

where $j_{\gamma,L}^\mu = -\bar{e}\gamma^\mu e$, the leptonic electromagnetic current, and including θ_W , the Weinberg mixing angle, the vector gauge boson fields introduced earlier become:

$$\begin{aligned} A^\mu &= \sin \vartheta_W A_3^\mu + \cos \vartheta_W B^\mu \\ Z^\mu &= \cos \vartheta_W A_3^\mu - \sin \vartheta_W B^\mu \end{aligned}.$$

The weak force is mediated by massive gauge bosons, the W^\pm boson which mediates the charged current interactions, and the neutral Z boson involved in the neutral current interactions. The weak interaction's coupling constant is much stronger than the electromagnetic interaction ($\alpha_{EM} = \frac{1}{137}, \alpha_W = \frac{1}{30}$), however it is weaker than the electromagnetic




Figure 2.3: A visualisation of the two parity states from Wu’s cobalt decay experiment [14]. The curved arrow indicates the direction of parity, with the second arrow demonstrating the direction the electron was released in.

interaction due to the large masses of the gauge bosons ($m_Z = 91 \text{ GeV}$, $m_W = 80 \text{ GeV}$). It is about an equivalent strength to the electromagnetic force at energies close to the mass of the W boson. The weak interaction has a small range of around 0.1% the diameter of a proton.

The weak interaction is the only part of the Standard Model that has been observed to violate parity. This was first shown by Wu in 1957 [13] by studying the decay of polarised cobalt-60. In a parity-inverted ‘mirror world’ as shown in Figure 2.3, the only quantity that changes sign is the vector momentum of the emitted electron, meaning the expectation under parity conservation is equal numbers of electrons emitted in the two directions. Instead, most of the electrons were released in the parity-inverted state, demonstrating that parity is not a conserved quantity in the weak interaction.

2.3 The solar neutrino problem

As neutrino theory advanced, experiments were designed to measure the flux of solar neutrinos reaching the Earth, the results of which ended up being known as ‘The Solar Neutrino Problem’, and became one of the biggest questions in neutrino theory until the discovery of neutrino oscillations.

2.3.1 The Homestake experiment

The Homestake experiment was set up in 1970 by Davis and Bahcall [15], in the old Homestake gold mine in South Dakota. It set out to measure the flux of neutrinos from




Figure 2.4: The neutrino energies from solar decay chains [17].

the sun. From studies of the solar neutrino flux and its decay chains it was known that solar neutrinos are exclusively electron neutrinos. As such, the experiment was set up to detect only electron neutrinos by filling a tank with tetrachloroethylene (C_2Cl_4) which contains one atom Chlorine-37 per molecule, a specific isotope that undergoes inverse beta decay when it interacts with an electron neutrino to become Argon-37. Neutrinos are produced from multiple decay chains in the sun with different energies, and the more energetic neutrinos leave the Argon-37 in an excited state. Using the known parameters of the ground states of Argon and Chlorine, Bahcall [16] calculated the interaction cross section for solar neutrinos from the Be-7, pep and CNO chains, which do not have the energy to excite the resulting Argon particle above the ground state. The energies of the neutrinos released from the different chains can be seen in Figure 2.4.

The B-8 decay in the sun is more energetic, and the neutrinos it produces are capable of exciting the Argon past the ground state. The branching ratio for this more energetic decay was calculated by Bahcall and Barnes [18], using the mirror interaction of Ca-37 into K-37. This found that the lower energy neutrinos had a lower branching rate, and it was determined that the detector was most sensitive to neutrinos with energies higher than 5.8 MeV. Using this sensitivity, the expected neutrino cross section from the B-8 chain was predicted to be $2.0 \pm 1.2 \times 10^{-35} s^{-1}$ per ^{37}Cl atom [15]. The detector was built 1478 m below the surface, in an unused portion of the Homestake gold mine in South Dakota. This location aided in shielding the detector from muons and cosmogenically activated isotopes, and the rock had an advantage of being low in Thorium and Uranium, avoiding adding another source of background. The tank was made completely leak proof in order to avoid any contamination from the Argon in the air. The produced Argon was removed

from the tank by purging with Helium and then the Argon was purified to remove other elements that had been purged from the tank with it and placed in a proportional counter. This then allowed them to determine the number of Argon atoms that had been removed from the tank.

They also conservatively calculated their background reductions, and after collecting a suitable amount of data, they calculated the flux they had obtained. This flux was significantly lower than predicted. The observed rate was $0.3 \times 10^{-35} \text{s}^{-1}$ per ^{37}Cl atom which can be calculated as a flux of $1.14 \pm 0.037 \times 10^{-42} \text{cm}^{-2}$. This observation being much lower than the prediction was later confirmed by Super-Kamiokande [19]. In 1990, IMB [20] and Kamiokande [21] also discovered a similar anomaly, this time from studies of atmospheric neutrinos not solar neutrinos. This lower than expected flux remained an open question in particle physics until the discovery of neutrino oscillations.

2.3.2 Neutrino oscillations

Neutrino oscillations bear significant resemblance to quark mixing. The flavour eigenstates are linear combinations of the mass eigenstates. In down-type quark mixing (using just two families for simplicity), the mass eigenstates are $\begin{bmatrix} u \\ d \end{bmatrix}$ and $\begin{bmatrix} c \\ s \end{bmatrix}$. The flavour states are denoted as $\begin{bmatrix} u' \\ d' \end{bmatrix}$ and $\begin{bmatrix} c' \\ s' \end{bmatrix}$. The two families have coupling constants, g_{ud} (coupling of up-type quarks to down-type quarks) and g_{cs} (the coupling of charm-type quarks to strange-type quarks), which are approximately equal and labelled as g_w . This shows allowed quark interactions, but kaon decay, as seen in Figure 2.5, has been observed, which violates these principles, showing a strange quark changing into a down quark. This is due to quark mixing, a theory suggested by Nicola Cabbibo [22], the down and strange quarks are linked, see Equation 2.7, where θ_c is the Cabbibo angle, which can only be measured experimentally.

$$d' = d \cos \theta_c + s \sin \theta_c \quad (2.7)$$

Lepton-quark symmetry, a theory that builds on the similarities of leptons and quarks [23], is then taken to apply to $\begin{bmatrix} u \\ d' \end{bmatrix}$ and $\begin{bmatrix} c \\ s' \end{bmatrix}$. Kobayashi and Maskawa expanded on this theory to include all three generations of quarks, producing the CKM matrix [24]. A simplified




Figure 2.5: Kaon decay into pions, showing a strange quark changing into a down quark, a transition not thought possible before quark mixing.

version of this matrix can be seen in Equation 2.8.

$$\begin{bmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{bmatrix} \approx \begin{bmatrix} \cos \theta_c & \sin \theta_c & 0 \\ -\sin \theta_c & \cos \theta_c & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (2.8)$$

Quark mixing allows the neutral Kaon to oscillate with its antiparticle, due to the down-type and strange-type quark mixing, causing the strangeness of the Kaon to change. When Pontecorvo [25] first put forward his theory of neutrino oscillations, he predicted them to follow a similar oscillation. Specifically the oscillation between a right-handed antineutrino and a corresponding right-handed neutrino. This right-handed neutrino would be sterile, meaning that it would not interact through the forces in the Standard Model. The neutrino and the antineutrino would be eigenstates of two Majorana (see 2.4.3) neutrinos, as Equation 2.9 shows.

$$\begin{aligned} |\bar{\nu}_R\rangle &= \frac{1}{\sqrt{2}}(|\nu_1\rangle + |\nu_2\rangle) \\ |\nu_R\rangle &= \frac{1}{\sqrt{2}}(|\nu_1\rangle - |\nu_2\rangle) \end{aligned} \quad (2.9)$$

This theory is based on the assumption that there exists a weaker version of the weak interaction which does not preserve lepton number. A second paper he published on the subject [26] suggested instead that they oscillate between flavours. This was first put forward just for the ν_e and ν_μ flavours, the tau neutrino had not been theorised yet. The superpositions of the neutrinos made from the mass eigenstates ν_1 and ν_2 , are the observable neutrino states, ν_e and ν_μ , which bear a strong resemblance to the quark mixing shown in Equation 2.7.

$$\begin{aligned} \nu_e &= \cos \theta \nu_1 + \sin \theta \nu_2 \\ \nu_\mu &= -\sin \theta \nu_1 + \cos \theta \nu_2 \end{aligned}$$

The flavour states of the neutrinos (ν_e, ν_μ, ν_τ) are each a superposition of the mass states, as in Equation 2.10.

$$|\nu_\alpha\rangle = \sum_k U_{\alpha k}^* |\nu_k\rangle \quad (2.10)$$

The $U_{\alpha k}^*$ is a unitary mixing matrix, α represents the three flavour states (e, μ, τ) and the $|\nu_k\rangle$ term is a neutrino with mass m_k . This mass state is an eigenstate of the Hamiltonian of a free neutrino:

$$H |\nu_k\rangle = E_k |\nu_k\rangle$$

When the dispersion relation ($E_k = \sqrt{p^2 + m_k^2}$) is taken in the ultra-relativistic limit, it gives Equation 2.11.

$$E_k \approx E + \frac{m_k^2}{2E} \quad (2.11)$$

The free, isolated mass states evolve over time, as the quark mass states do. They evolve as a plane wave, meaning that they must be a solution of the Schrödinger Equation, making Equation 2.10:

$$|\nu_\alpha\rangle = \sum_k U_{\alpha k}^* e^{-iE_k t} |\nu_k\rangle$$

This equation is valid if the differences between neutrino masses are unresolvable in weak decays. Using Heisenberg's uncertainty principle, it would be possible to resolve a mass difference between neutrinos with orders of magnitude energy difference. For a neutrino with energy uncertainty ΔE , that changes state over time interval Δt , Heisenberg's uncertainty relation becomes:

$$\Delta E \Delta t \gtrsim 1$$

Using Equation 2.11 as well as approximating $\Delta t \approx L$, this leads to:

$$|\Delta m_{kj}^2| \frac{L}{2E} \gtrsim 1$$

For just the electron and muon flavour neutrinos (for simplicity), U is a unitary 2×2 matrix:

$$U = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix}$$

Using the unitary properties of this matrix, the original relation in Equation 2.10 becomes:

$$|\nu_\alpha(t)\rangle = \Sigma(\Sigma U_{\alpha k}^* e^{-iE_k t} U_{\beta k}) |\nu_k\rangle$$

The survival probability of a neutrino refers to the probability that it is detected at Earth as the same flavour that it originated as. A short derivation of the probability of an electron

neutrino turning into a muon neutrino will be shown, for more detailed derivation, or for the three flavour neutrino survival probability see [27].

$$\begin{aligned} P_{e \rightarrow \mu} &= |\langle \nu_\mu(t) | \nu_e(0) \rangle|^2 \\ &= |\sum e^{iE_i t} U_{\mu i}^* U_{ei}|^2 \end{aligned}$$

Using the U defined for 2 flavour oscillations:

$$\begin{aligned} P_{e \rightarrow \mu} &= |-e^{iE_1 t} \sin \theta \cos \theta + e^{iE_2 t} \sin \theta \cos \theta|^2 \\ &= \sin^2 \theta \cos^2 \theta [2 - (e^{i(E_2 - E_1)t} + e^{-i(E_2 - E_1)t})] \\ &= 2 \sin^2 \theta \cos^2 \theta [1 - \cos(\Delta E_{21} t)] \\ &= \sin^2 2\theta \sin^2 \left(\frac{\Delta E_{21} t}{2} \right) \end{aligned}$$

Using the relativistic dispersion relation defined in Equation 2.11, this becomes:

$$P_{e \rightarrow \mu} = \sin^2 2\theta \sin^2 \left(1.27 \frac{\Delta m_{21}^2 L [\text{ev}^2][\text{km}]}{E [\text{GeV}]} \right)$$

This demonstrates that if $\Delta m_{12} = 0$, then there would be no possibility of an oscillation, meaning that the neutrinos must have distinct masses.

Pontecorvo, Maki, Nakagawa and Sakata expanded on this to include all three neutrino flavours, resulting in three mixing angles governing these oscillations, with the U in Equation 2.10 being a 3×3 unitary matrix. This results in the PMNS matrix as can be seen in Equation 2.12. The 3 matrices, in order, represent oscillations of atmospheric neutrinos, reactor neutrinos and solar neutrinos.

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (2.12)$$

c_{12} and s_{12} are $\cos \theta_{12}$ and $\sin \theta_{12}$ etc.

The middle term, as seen from reactor neutrinos, can also be observed through long baseline experiments such as T2K [28], Hyper-Kamiokande [29] and DUNE [30], and is also where there is sensitivity to search for CP violations.

The MSW matter effects, a theory developed by Mikheyev Smirnov and Wolfenstein [31], are a term that explains the effects on neutrino oscillation as a result of passing through matter. All neutrino flavours can interact with protons and neutrons through the Z boson of the weak interaction (as Figure 2.6(c)), and as all flavours can do this equally, this causes no net change in the oscillation probabilities. However, electron neutrinos can also interact through exchange of the W boson with electrons (as Figure 2.6(b)),




Figure 2.6: Feynman diagrams depicting neutral current scattering off electrons (a) and nucleons (c) and charged current scattering off electrons (b) and nucleons (d)

something muon and tau neutrinos are unable to do due to the lack of muons and taus in matter. This happens to solar neutrinos as they pass through the Sun before they reach the Earth, but has not been observed for atmospheric neutrinos. This additional effect on the electron neutrinos is explained by this matter effect term:

$$\sin^2 2\theta_{mat} = \frac{\sin^2 2\theta}{\sin^2 2\theta + (\cos 2\theta - \frac{2V}{\Delta m^2/E})}$$

2.3.3 SNO and Super-K observe neutrino oscillations

Proving neutrino oscillations became the focus of neutrino physicists and in 1999 SNO [32], a large neutrino detector filled with heavy water located in a nickel mine in Canada, became operational and able to detect solar neutrinos. This experiment, using a large array of PMTs, was able to detect neutrino interactions from both the neutral current (Equation 2.13) and charged current (Equation 2.14) interactions, as well as the elastic scattering off electrons with a lower cross section than the other two interactions.

$$\nu_x + d \rightarrow p + n + \nu_x \quad (2.13)$$

$$\nu_e + d \rightarrow p + p + e^- \quad (2.14)$$

The neutral current interaction is open to all flavours of neutrinos, whilst the charged current interaction is only allowed for the electron neutrinos. By comparing these two channels the ratio of electron neutrinos compared to total neutrinos can be found. Super-K [19][33] became operational in 1996, using ultra-pure water it was able to detect solar neutrinos, atmospheric neutrinos and man-made neutrinos, such as those from nuclear reactors. This experiment detected the elastic scattering of neutrinos off electrons in the

water, as can be seen in Equation 2.15. Whilst all types of neutrinos can interact in this way, the cross section for electron neutrinos is around 6 times higher.

$$\nu_x + e^- \rightarrow \nu_x + e^- \quad (2.15)$$

Between these SNO result and the Super-K atmospheric neutrino result, neutrino oscillations were discovered, proving that neutrinos do have distinct masses. In 2015 the two experiments shared the Nobel Prize in Physics for this discovery. The mass differences are related such that only two of the three need to be directly measured in order to calculate the other:

$$\Delta m_{21}^2 + \Delta m_{32}^2 - \Delta m_{31}^2 = 0 \quad (2.16)$$

The Δm_{21}^2 term has limits placed on it by measurements of solar neutrinos. Δm_{32}^2 is calculated from measurements of atmospheric oscillations as well as long baseline experiments such as T2K and NOvA. Table 2.1 shows the current limits on these parameters. KamLAND's measurements help to constrain the Δm_{12}^2 parameter, and long baseline ν_μ disappearance measurements help to constrain the Δm_{23}^2 parameter.

Table 2.1: The best-fit values and 3σ ranges of Δm_{21}^2 and Δm_{32}^2 neutrino mass parameters. From [34]. For details on normal and inverted mass ordering see section 2.4.2.

Parameter (mass ordering)	best-fit value	3σ region
Δm_{21}^2	$7.37 \times 10^{-5} \text{eV}^2$	$6.39 - 7.96 \times 10^{-5} \text{eV}^2$
Δm_{32}^2 (normal)	$2.56 \times 10^{-3} \text{eV}^2$	$2.45 - 2.69 \times 10^{-3} \text{eV}^2$
Δm_{32}^2 (inverted)	$2.54 \times 10^{-3} \text{eV}^2$	$2.42 - 2.66 \times 10^{-3} \text{eV}^2$
$\sin^2 \theta_{12}$	0.297	0.250 - 0.354

2.4 Open Questions Regarding Neutrino Mass

2.4.1 Specific Neutrino Masses

Whilst the discovery of neutrino oscillations showed that neutrinos have distinct masses, it did not define the actual masses, other than to say that at least two of the masses must be non-zero. Neutrino masses are known to be small, which makes accurate measurements of them challenging. Cosmological models have constrained these masses through comparison of the abundance of photons to the abundance of neutrinos, and how that fits into their models. Surveys such as WMAP, Plank and Hubble space telescope Key Project are

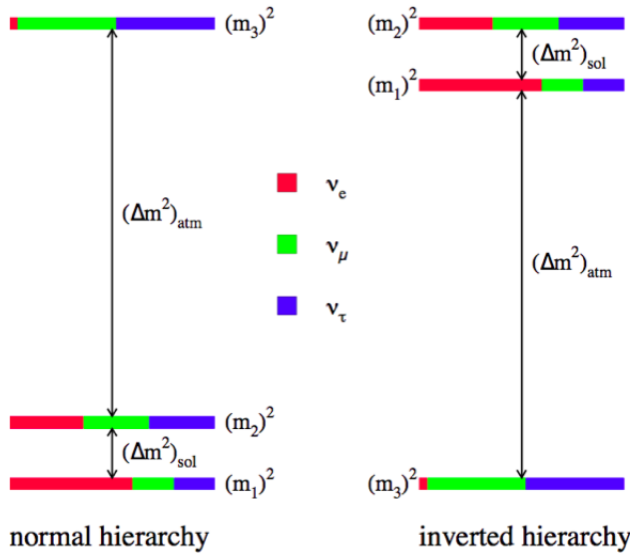


Figure 2.7: The two possible mass orderings (hierarchies) for neutrino mass, inverted and normal [37].

still ongoing with ever improving sensitivities, currently the best constraint from their combined results is <0.48 eV [35].

Modern experiments are attempting to constrain the neutrino mass through observations of beta decay. With very precise measurements of the resulting electron energy, the energy of the neutrino can be inferred which can then be converted to a mass. KATRIN [36] is an experiment to measure this with more precision. This experiment is sensitive to masses as low as 0.2 eV.

2.4.2 Neutrino Mass Ordering

Another thing that the measured mass differences do not answer is the mass ordering. Neutrinos and antineutrinos behave differently under the MSW effects, resulting in the sign for Δm_{12} . However Δm_{32} in an absolute value, with no sign. To obtain the sign, the matter effects also have to be observed in conjunction with the oscillations. This happens in solar neutrinos (Δm_{12}) due to the passage through the sun, but is not observed with previous atmospheric neutrino measurements. The two options for the masses are normal mass ordering ($m_1 < m_2 < m_3$), or the inverted mass ordering ($m_3 < m_1 < m_2$). Figure 2.7 shows the two energy options, as well as the composition of the mass eigenstates for each flavour, and Figure 3.7 shows the mass orderings of the two mass orderings as well as some current experimental limits.

Super-K[19] and IceCube [38] are both attempting to investigate the mass ordering

through observing atmospheric neutrinos that have passed through the Earth before reaching the detector in order to study the matter effects. DUNE will also be sensitive to the mass ordering due to the long-baseline of the experiment.

2.4.3 Mass Mechanism: Dirac or Majorana

In the Standard Model there currently exists no mechanism for neutrinos to gain mass, but as explained above they have been experimentally shown to have mass. There are two possible ways they could get mass. The first is if they are Dirac particles, this is how the other particles in the Standard Model get their mass, involving a coupling to the Higgs boson. However all neutrinos are left-handed, meaning that to enable a Dirac mass, a right-handed neutrino would have to be introduced to the Standard Model. This neutrino would have to be sterile, meaning that it does not interact through any force in the Standard model but would allow the neutrinos to couple with the Higgs.

This would be done by defining a right-handed field for each of the three neutrino flavours such that the field annihilates a left-handed antineutrino and produces a right-handed neutrino, which could then interact with the Higgs boson in this interaction (shown for the electron neutrino (ν_e)):

$$\lambda_\nu \nu_e^c (\nu h^0 - e h^+) \quad (2.17)$$

where λ_ν is the Yukawa coupling constant to the neutrino, h^0 is the spin-zero Higgs field and h^+ is the second component of a weak isospin doublet with h^0 that annihilates positively charged particles.

Adding a right-handed field would result in a mixing term in the Lagrangian of

$$\mathcal{L}^D = -m_D (\bar{\psi}_L \psi_R + \bar{\psi}_R \psi_L)$$

where ψ_R and ψ_L are the right and left handed components respectively and m_D is the Dirac mass term.

This corresponds to a Dirac mass for the neutrino (where ν can be any neutrino flavour) of:

$$m_\nu = \frac{\lambda_\nu \nu}{\sqrt{2}} \quad (2.18)$$

Whilst nothing forbids adding one or more sterile neutrinos to the Standard Model, this mechanism of neutrino mass is not favoured due in part to a lack of explanation why neutrinos have such small mass if they gain mass in the same manner as the other Standard Model particles.

The other possibility for their mass is if they are Majorana particles. This means that

the particle and the antiparticle are indistinguishable from each other. This requires that the particle must be neutral, which makes neutrinos a candidate. While neutrinos and anti-neutrinos have been identified, their only difference is helicity, which is not always a conserved quantity. This mechanism is favoured as it does not require any new particles to be added to the Standard Model and also is a source of CP violation. Majorana fields give mass in the Standard Model Lagrangian through the following term of left-handed components:

$$\mathcal{L}^M = -m_L \left(\overline{\psi_L^{(c)}} \psi_L + \text{H.c.} \right).$$

Majorana neutrinos can be detected through a very rare decay called neutrinoless double beta decay. The Majorana mass term (M) has the form

$$\frac{1}{2} M \nu_e^c \nu_e^c.$$

The see-saw mechanism [39] of mass means that neutrinos being Majorana does not rule out a Dirac mass term as well. The see-saw mechanism proposes combining a Dirac mass term with two Majorana mass terms.

Assuming that $\lambda_\nu \approx \lambda_e$, the relevant terms of both the Dirac and Majorana mass terms can be written into the mass Lagrangian as

$$L_{mass} = \lambda_\nu (h^0 \nu_e - h^+ e) \nu_e^c + \frac{1}{2} M \nu_e^c \nu_e^c.$$

Adding this into the Lagrangian results in the following mass terms

$$L_{\nu_e mass} = m_{\nu_e} \nu_e \nu_e^c + \frac{1}{2} M \nu_e^c \nu_e^c.$$

where $m_{\nu_e} \approx \frac{\lambda_e \nu}{\sqrt{2}}$, which is the same as defined above but with the approximation of the Yukawa constants being equal. Writing these terms in the mass matrix and diagonalising them yields two masses, a light mass and a heavy mass.

$$\mu_{light} \approx \frac{m_{\nu_e}^2}{M}$$

$$\mu_{heavy} \approx M$$

The Dirac mass is reduced by this factor shown in the light mass, leading to the term ‘see-saw’. If M is sufficiently large then this theory could result in neutrino masses low enough to fit current experimental data.

This gives two fields, one for the heavy and for the light mass.

$$\nu_{light} \approx \nu_e + \left(\frac{m_{\nu_e}}{M} \right) \nu_e^c \approx \nu_e$$

$$\nu_{heavy} \approx \nu_e^c - \left(\frac{m_{\nu_e}}{M}\right) \nu_e \approx \nu_e^c$$

As can be seen, the light neutrino field corresponds to the three flavours of neutrinos we have observed, while the field from the heavy mass term is the right-handed neutrino field, which it is thought would be very difficult to directly observe. Both of these fields describe Majorana neutrinos.

2.4.4 Neutrinoless Double Beta Decay

Even-even nuclei, nuclei that have an even number of both protons and neutrons, are the most stable atomic configuration, however some of them can decay further into more stable states. For some isotopes, beta decay from this even-even state is forbidden as it would require energy rather than release it.

Some of these isotopes can decay through double beta decay [40] ($2\nu\beta\beta$), where two beta (or inverse beta) decays happen simultaneously in the nucleus. This by-passes the forbidden transition while allowing the isotope to decay to a more stable state. To be allowed to decay in this manner the atom must, at minimum, have the combined mass of the decay nucleus and two electrons. This is most commonly a decay from the ground state to the ground state, however in some instances the decay can go to the second excited state of the decay atom. This is a higher order process, and is rare. If an isotope can theoretically decay through this process, but also through single beta decay, the isotope would undergo single beta decay first as this would have a much shorter half-life. This means that $2\nu\beta\beta$ decay can only happen in an isotope where single beta decay is energetically forbidden, or heavily suppressed. Of the 35 isotopes theoretically able to decay via $2\nu\beta\beta$, 14 have been observed to decay in this manner, including Tellurium-130 [41].

Neutrinoless double beta decay ($0\nu\beta\beta$) is only possible if neutrinos are Majorana rather than Dirac in nature. This is predicted to be rarer than the observed $2\nu\beta\beta$. It is predicted that it can happen in some of the same isotopes that $2\nu\beta\beta$ is observed in. Figure 2.8 shows double beta decay on the top and neutrinoless double beta decay on the bottom. If neutrinos are Majorana particles, it can be considered that the neutrinos are simply virtual particles exchanged in the decay, leading the electrons to have a larger energy. This is how they can be distinguished from $2\nu\beta\beta$, by the higher energy of the electrons, as they carry the full Q-value of the decay. An example of this spectrum can be seen in Figure 2.9. There are a number of experiments [43] searching for $0\nu\beta\beta$, which use a variety of isotopes thought to decay in this manner. Table 2.2 shows a summary of some of these searches with their predicted (*) or achieved half-life sensitivities. Whilst the SuperNEMO [44]




Figure 2.8: Double beta Decay (top) and Neutrinoless Double beta Decay (bottom) Feynman diagrams [42]




Figure 2.9: The sum of electron energies, assuming an energy resolution of 2% [42]. The dotted line shows the energy spectrum for $2\nu\beta\beta$, and the solid line peak shows the energy spectrum for $0\nu\beta\beta$

experiment is predicted to be less sensitive, the approach they use is different to the other current searches, meaning that if there was a discovery they would be able to reconstruct the individual electron tracks in a way unique to that of other experiments.

Table 2.2: An overview of some current and predicted half-life sensitivities to $0\nu\beta\beta$ [43]. The (*) denotes a predicted sensitivity. A longer half life sensitivity shows an experiment is more sensitive.

Experiment	Isotope	Half-life Sensitivity
SNO+ [45]	^{130}Te	$T_{1/2}^{0\nu} > 2.0 \times 10^{26}\text{yrs}^*$
NEXT [46]	^{136}Xe	$T_{1/2}^{0\nu} > 6.0 \times 10^{25}\text{yrs}^*$
EXO-200 [47]	^{136}Xe	$T_{1/2}^{0\nu} > 1.1 \times 10^{25}\text{yrs}$
KamLAND-Zen [48]	^{136}Xe	$T_{1/2}^{0\nu} > 1.07 \times 10^{26}\text{yrs}$
SuperNEMO [44]	^{82}Se	$T_{1/2}^{0\nu} > 5.85 \times 10^{24}\text{yrs}^*$

Chapter 3

The SNO+ Detector

The SNO+ detector [45], illustrated in Figure 3.1, is a multi-purpose neutrino detector located in the Creighton mine, an active nickel mine in Sudbury, Ontario. Inheriting much of its electronics and infrastructure from the SNO experiment which shared the 2012 Nobel prize with Super-K for observations of neutrino oscillations, SNO+ consists of a transparent 6 m radius acrylic vessel (AV) suspended in a cavity filled with ultra-pure-water (UPW). The AV is surrounded by a PhotoMultiplier Tube (PMT) Support System (PSUP), an approximately spherical 8.4 m radius structure which houses ~ 9300 PMTs.

3.1 The SNO+ detector

3.1.1 The location

The SNO+ detector is located in the SNOLAB underground laboratory, on the 6800ft level of the Creighton mine, an active nickel mine in Sudbury, Ontario. This deep underground location creates an overburden of 2070 m of flat rock, or 6010 m.w.e., which results in a muon rate of $(0.286 \pm 0.009)\mu/\text{m}^2/\text{d}$ [50].

Mitigation of Radon in the mine air

Mine air contains a concentration of $123 \text{ Bq}/\text{m}^3$ of ^{222}Rn . ^{222}Rn ($T_{1/2} = 3.8 \text{ d}$) decays eventually into ^{210}Pb ($T_{1/2} = 22.2 \text{ y}$), many of which implant $\sim 100 \text{ nm}$ below the surface of the AV if the AV is exposed to the mine air. The AV was exposed to the air during commissioning, leaving ^{210}Pb atoms still in the AV surface, which decay into ^{210}Bi and ^{210}Po , both of which create backgrounds for several Physics analyses.

One reason for the water phase in SNO+ was to reduce the amount of mine air the AV was exposed to during commissioning, as UPW has less Radon than the laboratory air.




Figure 3.1: The SNO+ neutrino detector in Sudbury, Ontario [49]

The AV and cavity UPW was regularly recirculated and re-purified to keep the Radon levels as low as possible, recirculation of the cavity UPW is still ongoing. Similarly, to protect the scintillator from mine air whilst filling with scintillator the water was removed from the bottom of the AV as the scintillator was added into the top.

There are two cover gas systems in SNO+, the cavity cover gas system which is an area in the cavity above the UPW which is constantly flushed with nitrogen gas, and the AV cover gas system. The AV cover gas system is newly developed for SNO+ and is a layer of nitrogen above the scintillator in the neck of the AV, connected to three cover gas bags specially designed to protect the system from small pressure changes present during normal operations in the laboratory.

3.1.2 The Acrylic Vessel

The AV is a 6 m radius sphere of 5.5 cm thickness containing the target material. The neck of the AV is 7 m long, reaching the deck mounted to the cavity walls, offering the only access to the detector during operation.

The cavity in which the AV sits is filled with 7000 tonnes of UPW, acting as a barrier to shield the AV from backgrounds created by both the PSUP and the cavity walls. The AV is held securely in place by a system of hold-up and hold-down ropes attached to the top and bottom of the cavity, respectively. The hold-up ropes were installed by the SNO experiment, whilst the hold-down ropes are new for SNO+ to accommodate the scintillator




Figure 3.2: A depiction of the electronic set up of SNO+ [51]

cocktail, which is less dense than water. The hold-down rope net is made of 5 rope-pairs which criss-cross over the AV, and secure into the cavity floor at 20 points, each with alarms to ensure the tension is correct. Part of the UPW recirculation cools the water to maintain an approximate temperature of 12°C, reducing both the PMT dark rate and the temperature-related stretching of the ropes.

3.1.3 The PhotoMultiplier Tubes

The AV is surrounded by ~9300 8 inch Hamamatsu R1408 PhotoMultiplier Tubes (PMTs), which are arranged in the 9m radius PSUP (PMT Support System), which is supported by ropes anchored into the cavity walls. The vast majority of the PMTs in the PSUP face inwards to the AV, whilst 91 face outwards to the UPW in the cavity as a method of identifying muons. Each internally facing PMT has a ‘reflector’, a 27cm hyperbolic aluminium bucket, making the total PMT coverage 54% [50].

Two cables attach to each PMT, a readout cable and a high voltage (HV) cable. The

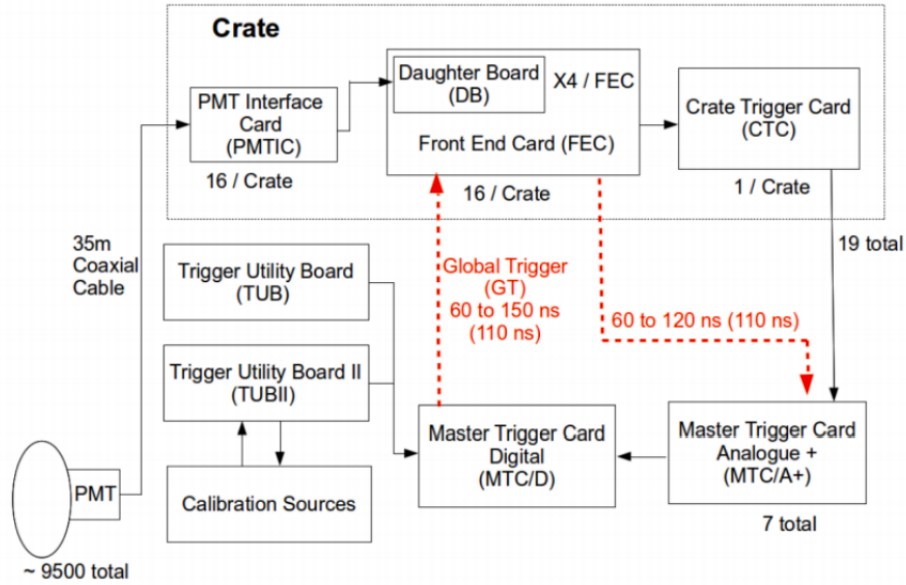


Figure 3.3: The flow of data through the SNO+ Data Acquisition System [52].

readout cables are fed into cards, as can be seen schematically in Figure 3.2. The cables are all identical length to remove any timing latency. There are 8 PMTs per paddle card, with 4 paddle cards per PMT Interface Card (PMTIC). The PMTIC controls the High Voltage supply for those 32 PMTs. Each paddle card also has an associated daughter card which stores the time and charge information for those 8 PMTs, and passes them forward to the Front End Card or Mother Board if they are part of an event. Sixteen of these Mother Boards make up one crate, making 512 PMTs per crate. There are 19 crates, meaning a possibility of 9782 channels are available. However 32 of those are reserved for calibration sources, and 91 channels are currently unused. This makes 9605 channels in the detector, mostly in normal (inward facing) PMTs, also with 91 Outward looking PMTs (OWLs), 4 Neck Tubes, and 49 Low Gain tubes.

3.1.4 The SNO+ Trigger System

Figure 3.3 shows the Data Acquisition (DAQ) system for SNO+. Everytime a PMT pulses, this pulse is sent to the PMTIC, which collects data for 8 PMTs to pass onto the Daughter Board (DB). The PMTIC has two relays, so if one PMT is having problems the corresponding relay can be opened remotely, resulting in only 4 PMTs being removed from data taking until underground repair work can take place, rather than all 8.

The DB then uses a combination of discriminators and integrators to determine if a channel is ‘hit’ or not. The threshold for a ‘hit’ is usually set to quarter of the average photoelec-

tron charge for that channel [50], as this has been determined to be the level where it is more likely to be a photon than noise. When a channel reaches this threshold, the Time-to-Amplitude-Converter (TAC) begins to ramp in the corresponding CMOS chip. The input charge signal is integrated over two time windows to at ‘high’ signal gain; 60 ns and 390 ns creating QHS and QHL charge measurements respectively, and a ‘low’ signal gain over 390 ns to create QLX charge measurement. The TAC ramps until a Global Trigger (GT) is sent to the DB, at which point it stores the data in the CMOS chip along with the trigger identification number (GTID). If there is no GT signal received within 410 ns, the memory on the CMOS chip is reset and the hit is discarded.

The Front End Cards (FECs) constantly poll the DB for signals to convert to digital using Analogue-to-Digital-Converters (ADCs) and stores them in first-in-first-out (FIFO) memory storage, to be read out if the GT is triggered.

Once a channel is considered hit, the DB sends the discriminator pulse to the relevant CTC, which sums up the pulses from the entire crate, and sends them to 7 MTC/A+ cards. Each of these MTC/A+ cards controls a trigger, the two main Physics triggers being N100 (the number of pulses within 100 ns), and N20 (the number of pulses within 40 ns, named historically). The MTC/A+ cards sum the pulses from every crate and then send them to the MTC/D card which compares this sum to the pre-defined trigger thresholds. If one of these trigger thresholds is met, the MTC/D issues a global trigger, prompting all hits in the last 410 ns to be read out. The MTC/D then places a lock for 420 ns on the GT so that no GTID can overlap.

The MTC/D can also be prompted to send a GT through one of two trigger utility boards. These trigger utility boards can be programmed, or can take input from calibration sources. An example of a programmed trigger is the Pulse Global Trigger (PGT) which is a forced GT at a rate of 10Hz, and can be used to look for PMT noise created by channels triggering too often, as will be explained in Chapter 4. The MTC/D can also make individual channels discriminators fire, which is used during some calibrations of the detector.

Every crate has an XL3 card, which is responsible for monitoring the health of the FEC cards (memory overload, temperature, voltage), as well as sending the PMT hit information from the FECs to ‘the builder’ if a GT is sent. The builder is a computer which collects the PMT information from the XL3s, and writes it out to a file, a ZEBRA databank (zdab) file structure. To avoid individual files becoming too large, everytime a file reaches ~1GB, the builder closes that file and starts a new one, creating multiple sub-runs per run. Once the run is complete, the builder sends all of the subruns to ‘the nearline’, a computer which

stores the raw data files for a short while, usually a few weeks, after the run and sends it to the main storage servers. Everytime a new run is transferred to the nearline, a series of automated jobs are run over the raw data files, to produce key run-level information including data quality, run information, the first pass of certain data cleaning checks, and the status of the PMTs in that run. The status of PMTs is taken from a number of PMT quality checks, including Channel Software Status, a run-by-run PMT calibration, which will be explained further in Chapter 4.

The nearline transfers the files to grid storage, where the main data processing occurs.

3.2 Data Processing in SNO+

Data processing has several steps, and produces a RAT data structure (ratds) file-type with several features not present in the raw ztab files including all calibrated PMT information, event reconstruction, and whether any data cleaning flags were applied to each event. Data processing also produces a scaled-down ntuple file-type for each run containing just key information most commonly used to analyse data such as reconstructed position, number of PMTs triggered (nhits), and GTID of the event, but removes PMT-specific data such as hit-times and PMT charge to reduce storage size.

Data processing uses the event reconstruction method, and combination of data cleaning cuts that are most broadly applicable to analyse the data with current detector conditions. Data can however be reprocessed on request using different fitting methods, or different data cleaning cuts if analysers need more specific processing to complete their analysis.

3.2.1 Data Cleaning in SNO+

Data cleaning aims to identify events that are either electronics effects, or muons which produce considerable light in the detector particularly in scintillator phase and often lead to multiple triggers and follower events. For the purpose of data cleaning, a cut is defined as a bitword flag on an event that fails certain data cleaning criteria. Data cleaning has a variety of masks for different analyses, which group together relevant cuts. Every event has two data cleaning bit words; dcApplied and dcFlagged which correspond to what data cleaning cuts were run over the event and which data cleaning cuts the event failed respectively. These bit words can then be used to remove events from processing or analysis. The majority of data cleaning cuts aim to identify events with unusual electronic effects, such that they can be excluded for reconstruction and analysis, for a complete list of these see Appendix A.

Muon based cuts in Data cleaning

The other cuts in data cleaning are primarily around identifying muons and the events immediately following them, there is also a check for atmospheric neutrinos and the events following them.

- **Atmospheric Cut** - This cut aims to flag atmospheric neutrino events by looking for high n_{hit} events, followed by a slightly less high n_{hit} event within $20\mu s$. It does this to try and find first the muon from the initial interaction and the decay electron following it. There must be less than 5 OutWard Looking (OWL) PMT hits on this initial event, to avoid overlap with the muon tagging. If this combination of events is found, the second pass of the cut runs, which triggers a 20 second cut on the data following the event.
- **Two Pass Muon Follower** - This cut first tags muon events, and then in a second pass, flags every following event in either a ‘long’ (60 s) or ‘short’ (20 s) time window as a follower so they can be removed from analysis. This cut needed significant tuning for the scintillator phase, which was carried out by the author, see Chapter 6.
- **Muon Tag** - This cut tags muons in the same way as Two Pass Muon Follower, but does not then complete a second pass to identify follower events. This cut was removed from processing in scintillator phase due to the overlap with the Two Pass Muon Follower.
- **NeutronTag** - This cut is new for the scintillator phase, developed by the author. This cut uses the events flagged as muons by the Two Pass Muon Follower Cut, and then searches a small time window ($20 - 2000\mu s$) following them for a coincident neutron capture.

3.2.2 Hit Cleaning in SNO+

After data cleaning, events are passed through a hit cleaning processor which aims to remove PMT cross talk. This processor looks for PMTs with high charge hits and then evaluates nearby PMTs to identify if the hits are cross talk or not. Cross talk events have low charge and appear slightly out of time with the PMT that experienced the high charge event. This updated value is stored in the `ratds` and `ntuple` files as `nhitsCleaned`.

3.3 Water Phase

SNO+ was filled with water partly to commission the detector, and partly to reduce the ingress of radon daughters into the AV, UPW is lower in radon than the mine air.

As well as providing the opportunity to calibrate the detector and measure intrinsic backgrounds, data from the water phase was used to make Physics measurements. The water phase refers to data taken from May 2017 to July 2019 with 905 tonnes of UPW in the AV.

3.3.1 Detector Specifics

During this phase, the heavy water (D_2O) of SNO was replaced with UPW. The detector acted as a low threshold water Cherenkov detector, observing charged particles.

Cherenkov light is produced when a particle travels faster than light travels through the same medium, producing a cone of light with the angle θ relating to the refractive index n and relativistic correction factor of the particle β as:

$$\cos \theta = \frac{1}{n\beta}$$

To obtain the number of photons emitted through this per unit track length (N_γ) per unit energy loss (E) in a material of permittivity (μ) and refractive index (n), the Frank-Tamm formula can be used:

$$\frac{d^2N_\gamma}{dx dE} = \frac{\alpha^2 Z^2}{\hbar c} \mu(E) \left(1 - \frac{1}{\beta^2 n(E)^2} \right).$$

3.3.2 Physics Results

The water phase of SNO+ allowed for measurements of solar neutrinos to be made that could be compared with those measured by SNO and Super-K. These measurements had the lowest backgrounds for any water Cherenkov detector taking measurements of this type, owing to the uniquely deep location of the detector [53].

Many Grand Unified Theories (GUTs) predict baryon number conservation violation as a possible explanation of the matter-antimatter asymmetry in the Universe. One possible example of this is invisible nucleon decay, where a nucleon decays into a lepton via undetectable decay particles. Whilst there are a number of models of this decay type, they all agree on the resulting lepton being in an excited state allowing for model-independent searches for this phenomenon. During the water phase, SNO+ set world leading limits on proton nucleon decay as well as proton-proton and proton-neutron dinucleon decay modes

[54].

The water phase also allowed for a measurement of the characteristic 2.2 MeV γ produced through neutron capture on hydrogen. These measurements were made using an AmBe calibration source which produces neutrons, usually accompanied by a simultaneous 4.4 MeV γ . Studying the delayed coincidence between the two showed a neutron detection efficiency of 46%, varying by only 1% throughout the detector, the highest known neutron efficiency for a water Cherenkov detector [55].

3.3.3 Backgrounds

Instrumental backgrounds from light emitted within a PMT, or electrical breakdowns was reduced by a variety of cuts based on cuts SNO used relying on charge and hit time. The total sacrifice of these cuts in the ROI was measured with a ^{16}N calibration source to be 1.7%.

One of the largest backgrounds for the nucleon decay search was the elastic scattering of ^8B solar neutrino interactions. This was largely reduced by a cut on reconstructed direction relative to the angle of the sun.

Due to the relatively low rate of muons, a cut on 20s of data after a muon reduced spallation products to less than one event in the ROI in the period of data taking [54].

3.3.4 Calibrations

On the deck of the detector there are two Umbilical Retrieval Mechanism (URM) devices which allow thoroughly cleaned sources to be safely lowered into the AV to take measurements.

^{16}N scans were regularly taken during the water phase using an internally deployed ^{16}N source at a series of points placed along the detector axes, as well as some measurements off-axis. This calibration allows any asymmetries in the detector and the reconstruction to be evaluated.

The water phase also used the ‘laserball’, a laser pulse diffuser able to run 337, 365, 385, 420, 450, and 500 nm wavelengths. This was also deployed internally, and takes runs across the z-axis to measure the PMT gain and electronic time delays.

The other internally deployed calibration source is the $^{241}\text{Am}^9\text{Be}$ neutron source, more commonly referred to as the AmBe source. This source mimics the electron antineutrino signal well, the ^{241}Am emits an α -particle which is then absorbed by the ^9Be particle, which decays through neutron emission into ^{12}C . The AmBe source provides a coincident

signal, with the de-excitation of the ^{12}C being the prompt signal, followed by the delayed neutron capture.

There is also an external calibration system called the ELLIE system [56], standing for Embedded LED Light Injection Entity, which allows light to be injected into the detector at known positions and wavelengths. There are two calibrations done using this, SMELLIE [57] and TELLIE [50] which both use the ELLIE system to evaluate the timing profiles of the PMTs.

3.4 Scintillator Phase

The scintillator fill was completed in March 2021, and the scintillator phase will run from then until Tellurium loading begins.

3.4.1 Detector specifics

SNO+ is using a scintillator cocktail [58], the primary ingredient being LAB ($\text{C}_6\text{H}_5\text{C}_{12}\text{H}_{25}$). When charged particles deposit energy into the detector, LAB releases scintillation light. However this light would be immediately reabsorbed by the LAB. To mitigate this PPO (2,5-Diphenyloxazole) is added to function as a fluor at a concentration of 2g/L, emitting light with a longer wavelength, and hence preventing re-absorption. There is also a secondary wavelength shifter added to the cocktail, to put the wavelength in the best region for the PMTs.

Scintillation occurs when inelastic scattering creates fluorescence - light emitted after non-thermal excitation by ionising particles. This typically has a very short decay time, on the order of ns to μs . This light is isotropic, meaning that the direction of the particle creating it cannot be reconstructed as it could be in the water phase.

Whilst Cherenkov light is also expected as in water phase, it is not currently distinguishable compared to the more prominent scintillation light.

3.4.2 Physics Goals

Geoneutrinos

Geoneutrinos are neutrinos from the Earth and in recent years there has been increased collaboration between neutrino physicists and geophysicists in a bid to use these to explain the heat flow of the Earth [59]. The geoneutrinos from the thick continental crust come from the Uranium and Thorium decay chains, both of which are well understood. The

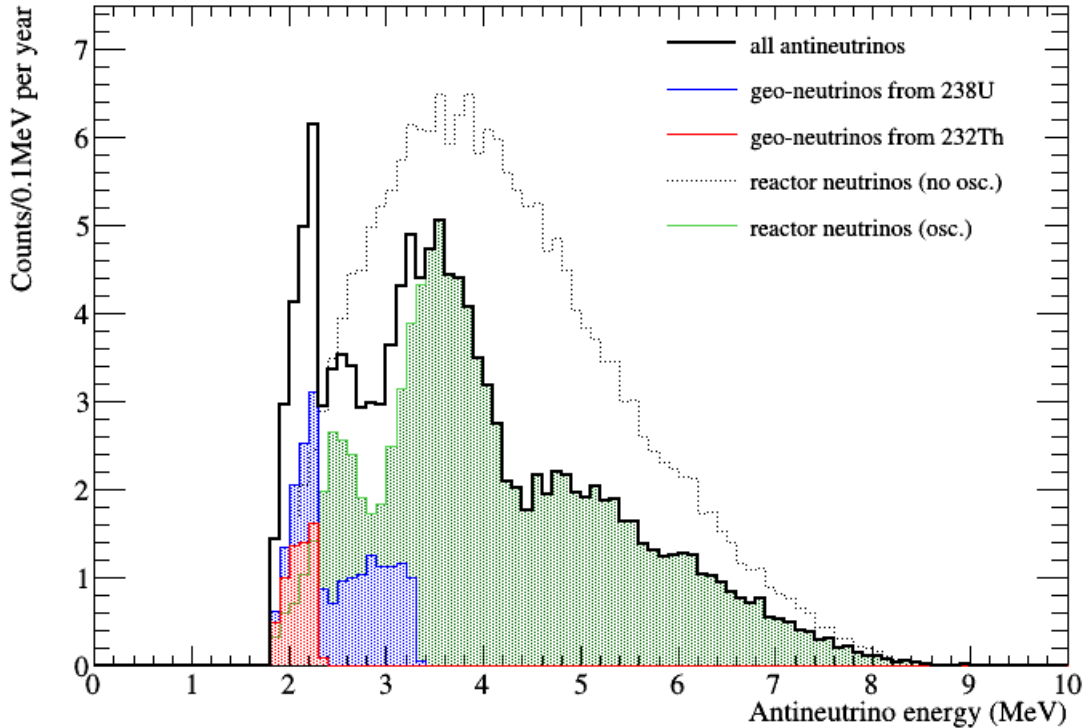


Figure 3.4: Predicted antineutrino spectrum in SNO+ scintillator phase [61]

data from this will be used for two separate studies, the first one where the ratio from the Uranium and Thorium chains will be fixed according to current geophysical models, and then a fit will be applied on the geoneutrinos to see how well they fit to the current models. The second study will be a global study combining the data from SNO+, KamLAND and Borexino [60]. This study will aim to separate the Uranium and Thorium contributions as well as the contributions from the mantle and the crust, with the aim to build a more detailed picture of the Earth's heat flow. Figure 3.4 shows the antineutrino spectrum, with the different types of geoneutrinos also highlighted.

Reactor neutrinos

The reactor anti-neutrinos are from nuclear reactors at a variety of baselines [45], three in Canada, the short baselines, and several in the United States of America and further afield which are at a longer baseline. The shorter baseline reactors in Canada are predicted to have a very well defined oscillation pattern and it is hoped that measurement of this will lead to constraining the Δm_{12} oscillation parameter [45]. The systematic uncertainties are expected to be very low in this measurement. The flux from the reactors should in general to be constant but the Canadian reactors are going to be undergoing maintenance over

the next few years, each maintenance period will reduce the flux by around 10%. This will allow the shorter baseline reactors to be separated from the longer baseline reactors as well as from other sources of anti neutrinos.

Supernovae neutrinos

Supernovae are sources of all types and flavours of neutrinos, meaning that detection of these would give an insight into the explosion mechanism. Neutrinos have only been detected from one supernovae so far, SN1987A [62]. SNO+ will be able to see neutrinos from supernovae at lower energies than some other detectors, and is going to be a part of the SuperNova Early Warning System (SNEWS) [63]. As a part of SNO+ joining SNEWS, a pre-SN monitoring system and a supernova burst trigger have been developed, and the livetime of the detector is kept as high as possible to maximise sensitivity.

Solar neutrinos

During the liquid scintillator phase, the study of high energy ^8B solar neutrinos carried out in the water phase can continue. Additionally, during this phase the detector will be sensitive to lower energy ^8B neutrinos as well as neutrinos from the CNO and *pep* chains. Measurement of the CNO chain could help to determine the metallicity of the inside of the sun. CNO neutrinos are predicted to depend linearly on the metal content of the sun, and combining this measurement with precision measurements of ^8B would help constrain it further [64].

Measurements of the *pep* chain and the lower energy ^8B chain could offer a chance to see new Physics, as these neutrinos are especially sensitive to matter effects. If the liquid scintillator is low enough in ^{14}C , this may allow for a measurement of *pp* neutrinos, which may result in a deeper understanding of reactions within the sun. A predicted spectrum of solar neutrinos in the pure scintillator phase can be seen in Figure 3.5.

3.4.3 Backgrounds

Figure 3.5 shows the predicted backgrounds for solar neutrino searches in SNO+, this assumes that Borexino's levels of backgrounds are achieved. The majority of the backgrounds shown in this figure are from the Thorium-232 and Uranium-238 decay chains, also present is Potassium-40. Not shown on this plot but a background for solar neutrino investigations is Carbon-11, caused by cosmogenic muon interactions with the liquid scintillator. This background and the identification of it will be discussed in more detail in

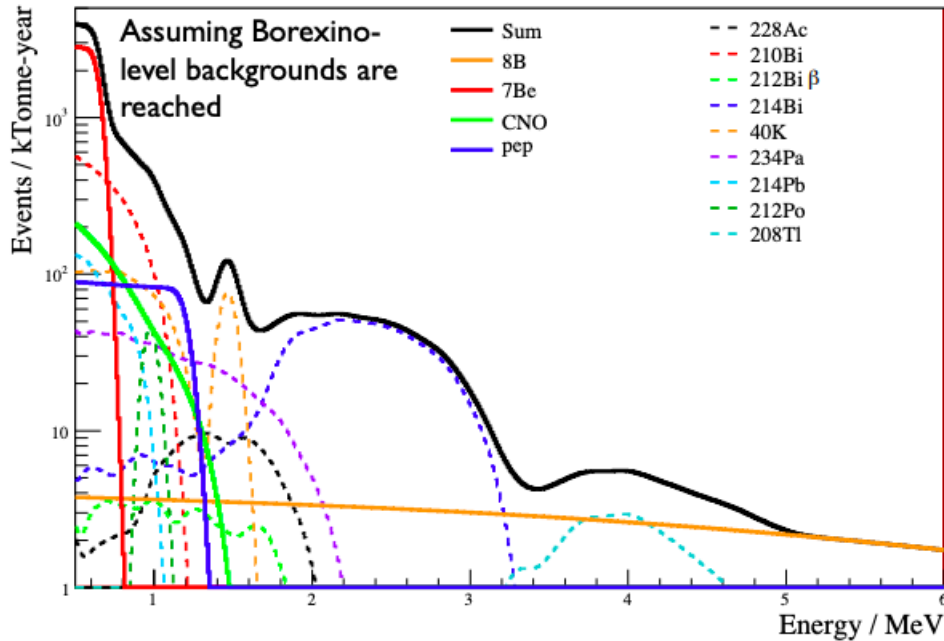


Figure 3.5: Predicted backgrounds for the solar neutrino searches at SNO+, assuming that Borexino levels of backgrounds [65] are achieved.

the rest of this thesis.

3.4.4 Calibrations

The external calibrations will be the same as the water phase, and the AmBe source has been used externally during the scintillator phase. However, there are more contamination risks associated with internal calibrations during the scintillator phase, and investigations are ongoing as to which sources can be cleaned to a high enough standard to use as an internal calibration source in the scintillator phase.

3.4.5 Scintillator Filling

Precautions must be taken at every step of the scintillator's journey to reduce the risk of any contamination, be that cosmogenic muon induced backgrounds, or exposure to mine dust.

Scintillator Production and Transportation to SNOLAB

Until the scintillator is underground, there is a risk of cosmogenic muon induced backgrounds building up through higher exposure to muons at surface level. The scintillator is produced from materials that have been stored underground, to minimise the backgrounds

present before the scintillator is even produced. After production, the scintillator must be transported to SNOLAB via a large tanker truck, which presents no opportunity to shield from muon interactions. The Cepsa plant in Becancour, Quebec is used for scintillator production, and it is just 900 km away from SNOLAB, reducing the time the scintillator is exposed to muons during above-ground transportation.

Once it arrives at SNOLAB, the truck is cleaned and a sample of the LAB tested (more details on the testing in Section 3.4.5) to ensure it is of sufficiently high quality before the transfer begins.

Transporting Scintillator Underground

After the scintillator passes the relevant checks, it is transferred from the truck into a large capacity holding tank in the surface building where it is stored before being transferred into railcars to be shipped underground. There are 6 railcars, each of which can store over a tonne. Once the railcars reach the entrance of the underground laboratory, they are cleaned before being brought into the outer area of the car wash, where they are cleaned again. The railcar fittings are covered with plastic to protect them from the mine dust, and after the entire railcar has been cleaned, these plastic bags are removed and the fittings carefully cleaned with UPW before the LAB is transferred to two 60 tonne capacity holding tanks in the scintillator purification plant.

PPO Loading

The fluor PPO is added in powder form to a small quantity of pure LAB to create a highly concentrated mixture termed ‘master solution’ which is then slowly mixed inline with pure LAB before it reaches the detector.

Upwards of 30 kg of PPO powder is mixed with approximately half a tonne of pure LAB per batch of master solution. Before it is sent to the scintillator purification plant, the master solution is mixed and sparged with nitrogen for at least 2 hours, before undergoing 3 water extractions to remove known contaminants such as heavy metals and many insoluble compounds which are present in the powder form of PPO.

The Scintillator Purification Plant

A large capacity scintillator purification plant has been constructed underground in order to fill the SNO+ detector, referred to here as the scintillator plant. The plant includes a multi-stage distillation column (primary distillation), a single stage distillation process

(secondary distillation) as well as a stripping tower, storage for LAB ('the 60 tonne tanks'), and a system for preparing the distilled LAB with PPO. The primary distillation and stripping systems focus on purifying the stock LAB, putting it through the process of multi-stage distillation and Nitrogen stripping whilst continuously removing waste LAB from the system. This aims to remove all contaminants, including the cosmogenic muon induced isotopes produced in the scintillator during surface-level transportation. The secondary distillation process purifies the master solution before mixing it inline with the pure LAB from primary distillation before it reaches the detector.

Quality Assurance (QA)

In order to ensure the continued quality of the product entering the detector, a scintillator quality assurance team regularly samples the scintillator at the step just before it goes into the detector and tests it for common contaminants. During the author's time on site, this was the primary responsibility undertaken, and the author took on a supervisory role in some QA shifts, and assisted in training newer QA members.

The first test performed is to measure the density of the scintillator using a densitometer, this is expected to be close to 0.85 g/cm^3 @ 21°C , and a reading significantly deviated from that can be indicative of water contamination in the scintillator. Additionally the density reading can also be used to calculate the PPO concentration of the sample during inline mixing¹.

The second test is the turbidity of the scintillator, which is expected to be below 0.06 NTU. As with the density, a significantly different reading can be an indication that water has mixed with the scintillator.

The third test is the UV-Vis spectroscopy of the scintillator. This can test the PPO concentration of the sample, as well as if water has got in the scintillator, and the optical clarity of the sample. To build this spectrum up, a sample of around 10ml is loaded into a UV-Vis spectrometer, which scans over a range of 200 – 600 nm, taking approximately 5 minutes². One of the key things that can be identified by this test is oxygen contamination. This shows up clearly as a peak around 380 nm as can be seen in Figure 3.6[b], as compared with a sample of excellent optical clarity, Figure 3.6[a]. The comparison samples

¹In February 2020 the densitometer broke, and the decision was made to suspend this test rather than replace it as the results offer little information unique from the other two tests. It will be reinstated if the laboratory decides to replace the densitometer.

²The range the UV-Vis spectrometer scans over was changed in February 2020 from 200-1000 nm to increase the speed of the test, as the higher wavelengths are no longer used in the tests.

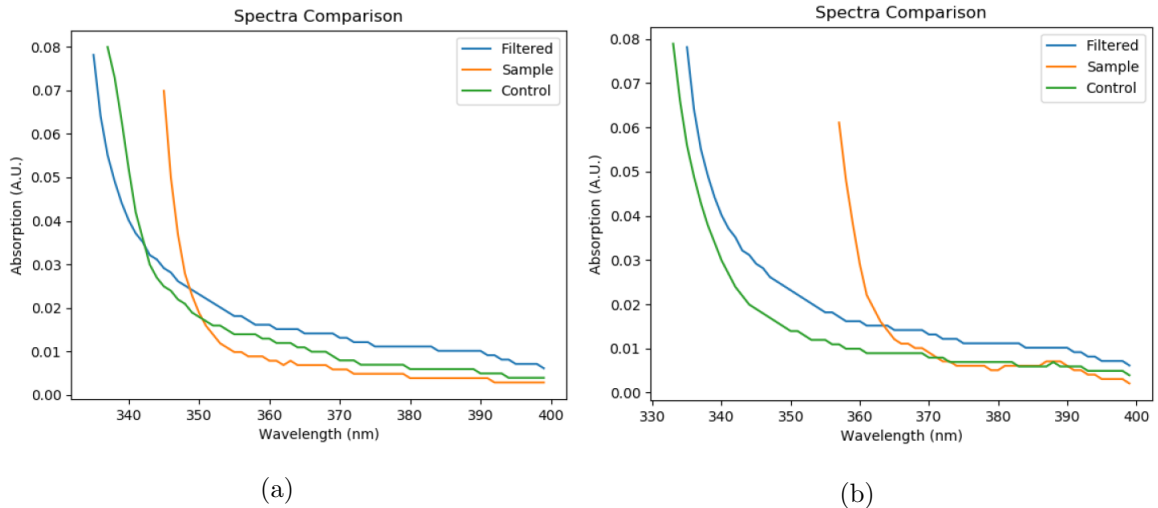


Figure 3.6: UV-Vis spectra for two samples of LAB-PPO. [a] shows a good sample with excellent optical clarity, whilst [b] demonstrates a slight build up of contaminants as indicated by the peak appearing at 380 nm. The baseline shift seen in [b] characterised by the curve beginning at 360 nm instead of 330 nm is due to the higher PPO concentration in that sample and is no indication of the optical quality.

that can be seen on the graph are control LAB, which is a sample of LAB known to be of excellent quality, and is our aim of LAB standard to put in the detector. The other is filtered LAB which is LAB of a slightly lower quality as a comparison. As well as testing the optical clarity and the build up of backgrounds, the UV-Vis spectrometer can also calculate the PPO concentration of the sample during inline mixing.

When the scintillator plant is preparing to fill the detector again after a pause in fill for any reason, at least one sample must be taken and pass all three tests before fill can be approved by a QA expert. During fill, these tests are completed at least every few hours and any indication of increasing contamination is reported back to the scintillator team who attempt to correct it by altering the amount of waste LAB being removed from the system. In the rare event this does not work and the contamination is building up, fill is stopped until QA returns to acceptable levels.

The master solution is also tested in the same manner after the water extractions, before it is sent to secondary distillation. Once the master solution has been distilled in the secondary distillation process, the post secondary PPO is tested once again before it recombines with pure LAB and goes to the detector. However, due to the fact that the secondary distillation system operates under vacuum, the distilled PPO is unable to be

continually sampled throughout the week. Therefore, to check that contaminants are not building up and to check the PPO concentration, a sample is taken and analysed first thing every Monday morning before the plant resumes normal vacuum conditions in secondary distillation.

3.4.6 Stable Partial Fill Phase

In March 2020 the Covid-19 pandemic caused a halt to the scintillator filling as access to the mine was restricted to limit transmission of the virus. This delay lasted several months until filling could resume in October 2020, resulting in several months of data taking with the detector in a stable configuration with 360 tonnes of scintillator at a PPO concentration of 0.53g/L, hereafter known as the Stable Partial Fill Phase. This phase allowed for analysis cut tuning to begin, and backgrounds to be studied in preparation for the scintillator phase without the unstable effects of filling operations on the detector.

Calibrations

There were no internal calibrations taken in the stable partial fill phase, but external AmBe calibrations were performed, as well as TELLIE and SMELLIE calibrations.

3.5 Tellurium-Loaded Phase

3.5.1 Detector Specifics

Once the scintillator phase is completed, Tellurium will be added to the AV. The Tellurium will first be dissolved into Diol before being mixed with the scintillator in a dedicated processing plant built in the underground laboratory. An additional wavelength shifter (bisMSB) will also be added to ensure the light emitted by the scintillator is still in the optimum range for the PMTs. Plans are to load the detector with 0.5% natural Tellurium-130 [58].

Among a number of factors in the decision to chose Tellurium are the high natural abundance (34%, meaning natural Tellurium can be used with no need for enrichment) and the compatibility of it with the scintillator. Tellurium has a predicted half-life of $(7.0 \pm 0.9(\text{stat}) \pm 1.1(\text{syst})) \times 10^{20} \text{yrs}$ for $2\nu\beta\beta$, one of the longest. This was another reason Tellurium was chosen as the energy resolution of SNO+ is not as high as other experiments.




Figure 3.7: A plot showing the normal and inverted hierarchy, as well as some of the current experimental limits [66].

3.5.2 Physics Goals

The experimental method to measure the Majorana mass from this decay is to measure the rate of the decay, which corresponds to the half-life.

SNO+ aims to be sensitive to a lower limit of $T_{1/2}^{0\nu} > 2 \times 10^{26}$ yrs, which corresponds to a Majorana mass of $m_{\beta\beta} \approx 40 - 90$ meV [58].

After this phase is completed, the aim is to load with more Tellurium, bringing the concentration up to 5%. This will reduce the light output, but during this phase the PMTs will be upgraded and the reflective concentrators which will increase the light yield by a third, bringing it to a similar level as the first Tellurium phase. This phase aims to set a lower limit of $T_{1/2}^{0\nu} > 7 \times 10^{26}$ yrs, corresponding to a Majorana mass of $m_{\beta\beta} \approx 19 - 46$ meV [45]. Figure 3.7 shows the inverted and normal hierarchy and the current experimental limits.

3.5.3 Backgrounds

Figure 3.8 shows the predicted backgrounds for the neutrinoless double beta decay search in the Tellurium phase of SNO+. The Boron-8 solar neutrinos are the largest background displayed here, and without the directionality of a water cherenkov detector, harder to remove. The internal Uranium-238 and Thorium-232 chains are able to be measured and constrained as much as possible in the scintillator phase, the two neutrino double beta




Figure 3.8: Predicted backgrounds for the SNO+ Tellurium phase in the neutrinoless double beta ROI.

decay background is constrained as much as possible by using an asymmetric ROI, and the cosmogenic component is very low as the telluric acid is already being stored in the laboratory to cool off before it is added to the detector.

At the time of writing this thesis, the SNO+ detector is in the scintillator phase, and as such the work presented within will be tailored to the scintillator physics goals, and includes data from the stable partial fill phase and the scintillator phase.

Chapter 4

Channel Software Status

The goal of the Channel Software Status (CSS) processor is to identify channels that are not performing as expected between electronic calibrations. When the author took over the task towards the end of the water phase of SNO+, CSS had been integrated onto nearline for automatic calibration, and ran two occupancy-based calibrations on a run-by-run basis. CSS is considered a service task for SNO+, this chapter will detail the author's contributions to the task, which were focused on developing a new charge based CSS check and monitoring CSS for the scintillator phase.

4.1 Overview of Channel Software Status

CSS runs over all runs and flags channels that are not performing optimally between calibrations. This can be due to either an incorrect calibration of the channel, or a hardware failure on the channel after the most recent calibration. Flagged channels can then be removed from analysis as required. Channels are flagged if they fail any of the CSS checks, and each channel can fail more than one check. Each check evaluates the channel's performance relating to occupancy and charge over an entire run.

The occupancy cut values are recalibrated after every Electronics Calibration (ECA) by the nearline client *CSSProc*. *CSSProc* calls the RAT processor *ChanSWStatusProc* to calculate an updated high and low occupancy cut value from the first good Physics run meeting basic run selection criteria on run length, online channels, and data quality following an ECA and outputs it to a RATDB table called *CSS_CUTS*. It then sets the *CSSCalib* bit in the database, allowing the other CSS nearline processor, *CSSCalib*, to run. *CSS-Calib* is the step which flags the channels, it calls the RAT processor *ChanSWStatusCalib* to compare every channel's performance to the occupancy cuts calculated by *CSSProc*.

CSSCalib also produces a RATDB table, `CSS_RESULT`, which contains an array called `channel_words` which contains a bitword for every channel in the detector. The bit word is updated for a channel each time it fails a check.

4.1.1 Standard Run

The first good Physics run after an ECA is known as a Standard Run (SR) and must meet certain criteria.

A SR in water phase and stable partial fill phase must be a minimum of 50 minutes long, pass all data quality checks, have all crates online, and at least 7000 channels.

The author made some adjustments to speed up the processor to adjust for the increased light yield in scintillator phase which mean only quarter of the events are used per run (details of this change can be found in Section 4.2). This change removes the justification for a SR to be 50 minutes, so the definition was changed to require a SR to be at least 30 minutes long, in line with current run selection criteria.

The water phase requirement of a SR having all crates online results in some ECAs not having an associated SR, and if a crate is offline for a prolonged period of time (for example such as was the case during the restricted laboratory access caused by the Covid-19 pandemic), a SR that no longer reflects the detector conditions can be used to process hundreds of runs. To avoid this, the author has dropped the requirement for all crates to be online for a run to be considered as a SR. This allows the calibration to better reflect current detector conditions, as an ECA is always run as soon as possible after a crate is repaired even if this is not on the usual schedule, which would trigger another SR to be found which would reflect the updated detector conditions.

4.2 Changes to CSS for Scintillator Phase

Evaluation of each CSS test for the scintillator phase is discussed below, but some changes to the processor itself were also required for the change to scintillator phase. Due to CSS operating on a run-by-run basis, speed is important to ensure it is not creating a backlog on the nearline machine. Owing to the significantly increased light yield in scintillator phase compared to water phase, the processor began taking too long to process a run. To counter this, the author introduced a pre-scaling into the processor so it only uses a quarter of the events per run to speed it up. This was tested on multiple runs, and showed consistency with the original high and low occupancy distributions, and showed no significant change to the numbers of channels flagged by any of the three CSS tests.

In the code, histograms are created of timing and three different charge measures for all channels, but only one of the charge measures is used at this time by CSS, so the unused ones were commented out of the code to further speed the processor up.

In addition to speed issues, one of the tables that the processor uses to assess whether a run is good enough to be a SR is no longer regularly produced in scintillator phase. Instead the author altered the processor using code from the run selection processor to conduct brief tests on the run to assess whether it can be a SR. To count as a SR the run duration must be at least half an hour, and data quality checks such as channel breakdowns, noise rate, data collection errors, and unusual laboratory activity in line with current run selection requirements.

4.3 Channel Software Status Tests

4.3.1 High Occupancy

To test for channels at high occupancy, Pulsed Global Trigger (PGT) events registered by the channel are used. PGT is a 50Hz forced trigger on the baseline detector state, making it sensitive to noise on the channel.

A channel is considered high occupancy if it is triggering significantly more than the other channels. This can affect the data quality; too many channels triggering that aren't connected to the event can cause poor reconstruction. For this reason, it is important these channels are flagged so they can be removed during processing before events are reconstructed.

CSSProc calculates the high occupancy cut value by plotting the high occupancy statistic for every channel. The high occupancy statistic is the number of PGT events registered by that channel over the run divided by the total number of PGT hits registered across the whole detector over the run. The high occupancy statistic (Γ_{occ}^{high}) for any channel i is defined as

$$\Gamma_{occ_i}^{high} = \frac{n_{pgt_i}}{N_{TotalPGT}}, i \epsilon pmts$$

where n_{pgt_i} is the PGT events registered on channel i over the run, and $N_{TotalPGT}$ is the total PGT events registered on all channels in the run. The high occupancy statistic is plotted for all channels and fitted with a gaussian on the first half of the peak, and 175 FWHM above the mean of this fit is taken as the high occupancy cut value. This cut value was already established before this author began work on CSS, for a more detailed account




Figure 4.1: High Occupancy Plot for run 207417, a Standard Run in the water phase. High Occupancy Statistic refers to the number of Pulsed Global Trigger hits registered on that channel divided by all Pulsed Global Trigger hits registered across the entire detector. Both of these quantities are over a whole run, and this statistic is plotted for every online channel in the run. The shaded area above the peak shows the region marked as high occupancy

of the tuning please see Chapter 4 of [67]. This cut was tuned by plotting the average number of channels flagged over many runs at a variety of FWHM above the mean. This showed a plateau starting at 150 FWHM above the mean, so the decision was taken to use 175 FWHM above the mean as the cut value. The results of this for SR 253333, a run in the water phase, can be seen in Figure 4.1, where the shaded blue area represents the values where channels would be cut. *CSSCalib* calculates the high occupancy statistic for every channel and if it is higher than the cut calculated by *CSSProc*, it sets the bitmask for high occupancy on that channel.

High Occupancy in scintillator phase

The distribution of high occupancy statistics has not changed with the addition of scintillator, as the PGT is unaffected by the detector medium. Figure 4.2 shows the high occupancy distributions for a SR in water phase, stable partial fill phase, during filling, and scintillator phase. Although a small variation can be seen in the scintillator run, it is not significant to motivate retuning the cut. The average number of channels flagged by this cut is still consistent with water and partial fill phases.

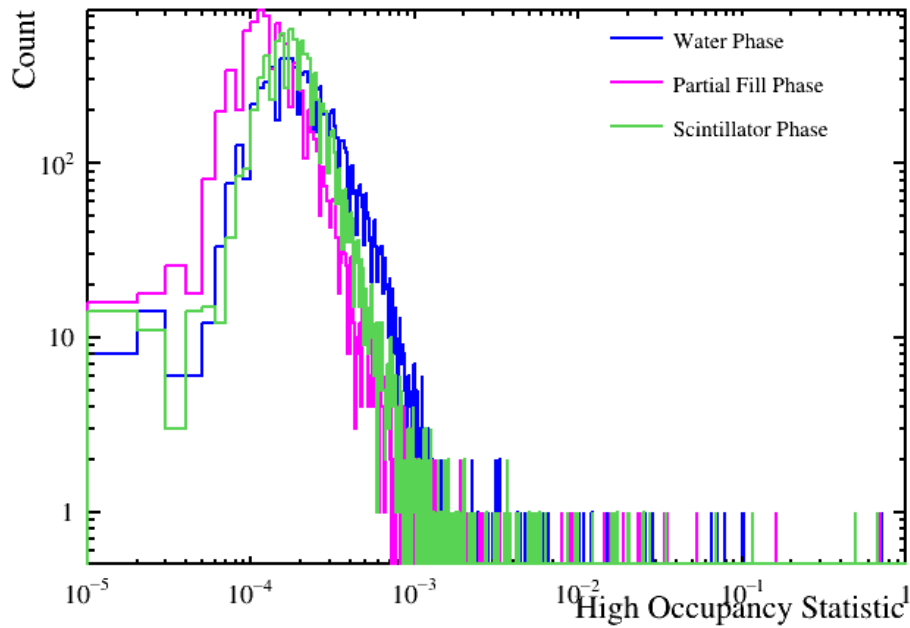


Figure 4.2: The High Occupancy statistic distribution for water phase, partial fill phase, and scintillator phase. It shows no significant variation between the phases that would require the cut to be returned.

4.3.2 Low Occupancy

A channel is considered low occupancy if it registers light at a significantly lower rate than the other channels. To determine this it is necessary to use all physical triggers unlike for high occupancy, making the low occupancy statistic for a channel the total hits registered on that channel across the run divided by the total hits registered in the whole detector across the run. To calculate this cut, *CSSProc* plots the low occupancy statistic for each channel, fits the beginning of this peak with a gaussian, and then records the cut value as 3.5 FWHM's (in water phase) below the mean of this fit. Figure 4.3 shows this distribution and corresponding cut value for a water phase SR.

Low Occupancy in Stable Partial Fill Phase

During the stable partial fill phase, the low occupancy distribution was no longer a single smooth peak, as PMTs towards the top of the detector were seeing far more light than PMTs at the bottom of the detector due to the split detector medium of water and scintillator. This created a thin peak of PMTs that were primarily seeing light only from events in the water, and a secondary, wider peak formed of the PMTs that were seeing more scintillation light. This altered distribution resulted in *CSSProc* not being able to




Figure 4.3: Low Occupancy Plot for run 207417, a Standard Run in water phase. Low Occupancy Statistic refers to the number of hits registered of any trigger type on that channel divided by all hits registered across the entire detector of any trigger type. Both of these quantities are over a whole run, and this statistic is plotted for every online channel in the run. The shaded area below the peak shows the region marked as low occupancy.

calculate an effective cut value, and instead almost all PMTs in the bottom half of the detector were flagged as low occupancy.

Figures 4.4, and 4.5 show the distribution of channels flagged as low occupancy for water phase and stable partial fill phase respectively. These plots demonstrate the issue with this cut whilst there were two mediums in the detector.

Low Occupancy in Scintillator Phase

Once scintillator filling was complete, the low occupancy distribution returned to a single peak, higher and thinner than the water phase. Figure 4.6 shows the low occupancy distribution for a standard run in water phase, and a standard run in scintillator phase. Unlike in partial fill phase, the scintillator phase distribution can be easily fitted, but using the water phase calculations of the cut value results in more channels than expected failing this cut.

The 3.5 FWHM's below the mean used to calculate the water phase cut value was chosen by plotting the average number of channels cut against the number of FWHM below the peak for a number of runs and identifying the earliest plateau, for verification of this cut value see Chapter 4 of [67]. To recalculate this cut value for scintillator phase, this

Low Occupancy for Run 253324 (quarter scaling)

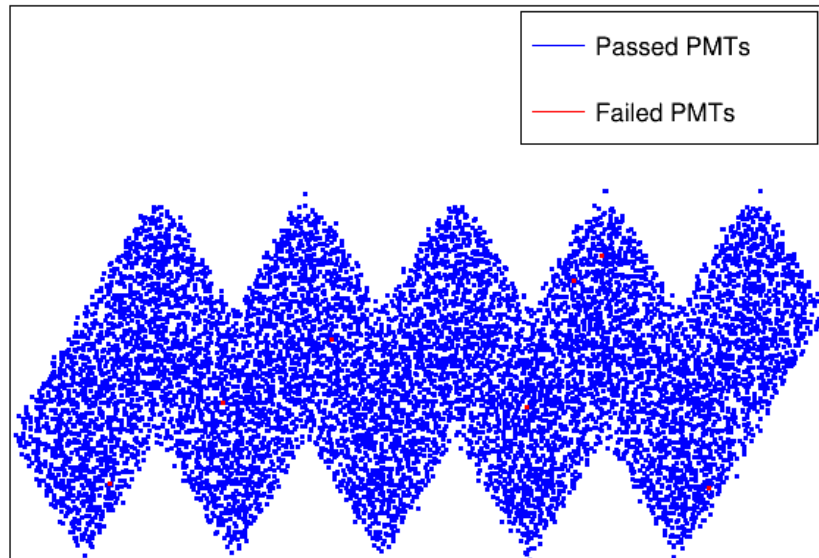


Figure 4.4: This figure shows a flat map of the channels failed (red) and passed (blue) by the CSS low occupancy cut for a run in water phase. The flat map matches the flat map of the electronics view in the SNO+ event display.

Low Occupancy for Run 260434

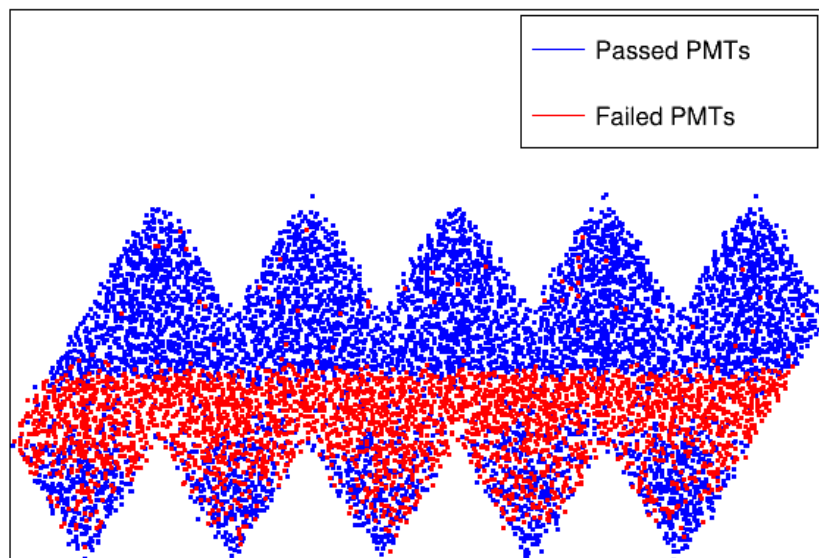


Figure 4.5: This figure shows a flat map of the channels failed (red) and passed (blue) by the CSS low occupancy cut for a run in the stable partial fill phase. The flat map matches the flat map of the electronics view in the SNO+ event display

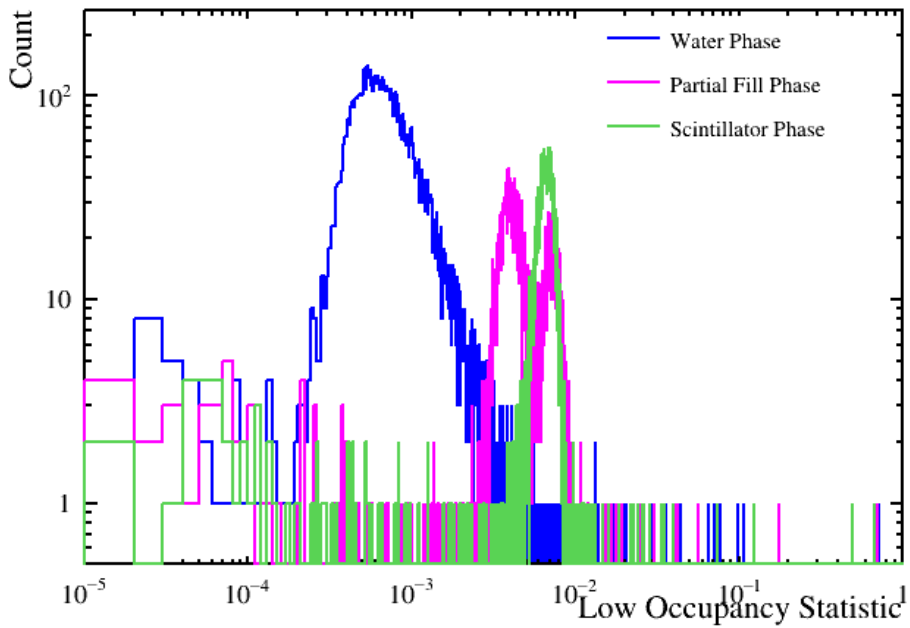


Figure 4.6: The Low Occupancy statistic distribution for water phase, partial fill phase, and scintillator phase. It shows the progression of the single peak into two, back into one higher, thinner peak as fill was finished.

plot was reproduced in Figure 4.7, motivating a cut of 5 FWHM as this is the lowest value of the plateau. Figure 4.9 shows the low occupancy statistic for a run with 3.5 FWHM and 5 FWHM below the mean marked on them.

4.3.3 Short Window Charge Cuts

The author developed CSS checks on the short-time-window charge (QHS). Unlike the occupancy cuts, these cut values are static instead of being recalculated after every ECA. This is because, unlike the occupancy cuts, the charge cuts are just looking at the charge distribution on the individual channel rather than comparing it to other channels. A first step to developing these cuts was to understand what a QHS distribution for a channel performing optimally is supposed to look like over a run. Figure 4.9 shows the expected distribution, a narrow peak just above zero showing that the pedestal has been correctly calculated for that channel, as well as the possible deviations from this distribution that can be seen, along with the percentage of channels with this issue averaged over several runs.

The negative peak can be caused either if the pedestal calculated by the ECA is too high, or if the channel's charge distribution has moved an unusually large amount since




Figure 4.7: Average number of channels flagged as low occupancy at a variety of FWHMs over multiple runs. Errors are statistical from averaging over multiple runs.

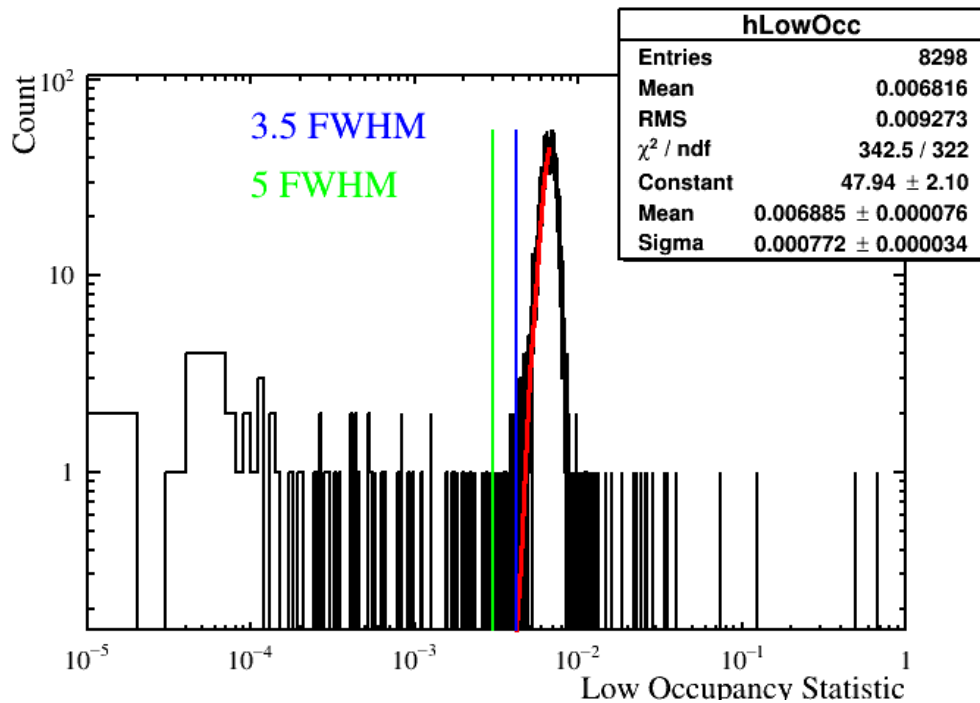


Figure 4.8: Low Occupancy statistic plotted for a run with 3.5 FWHMs (blue line) and 5 FWHMs (green line) below the fitted mean shown to compare the water phase and scintillator phase thresholds. The red line is a gaussian fit of the first half of the peak.

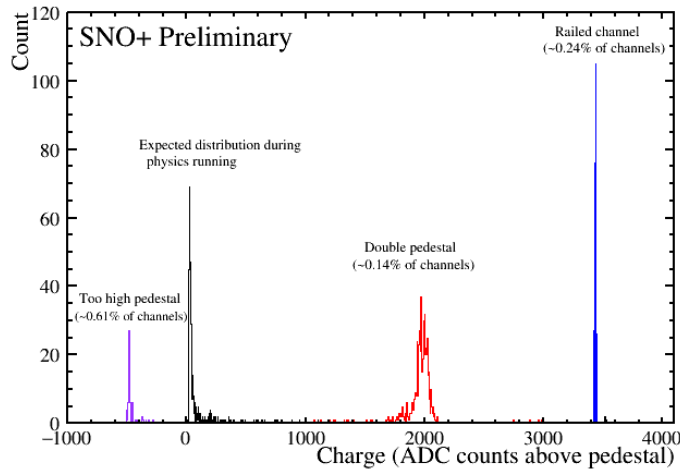


Figure 4.9: This shows the expected distribution of QHS charge per channel over an entire run (black histogram), and the possible deviations from this. The purple histogram represents a channel with a too high pedestal, registering negative charge. Red histogram in conjunction with the expected black peak means a double pedestal has been recorded for this channel. The blue peak represents a railed channel, which has become saturated.

Table 4.1: A table showing the CSS charge cut criteria. A channel is flagged by the calibration if the percentage of charge in any of these regions meets the criteria compared to the amount of charge across the entire channel.

Range of Charge (in ADC counts above threshold)	Flag Boundary
-1000 to -15	>15% of whole
-15 to 500	<50% of whole
500 to 1500	>15% of whole
1500 to 2500	>15% of whole
3300 to 4000	>10% of whole

the last ECA. The double peak occurs if the ECA mistakenly records two pedestals for a single channel. The highest charge channel is saturated or railed, and always appears as the majority of the charge at that specific value. It is usual to see a channel record charge at a saturated level for a particularly high energy event such as a muon, but if it records charge at that level across the entire run, the charge is not useable and the channel should be flagged.

Table 4.1 shows the cut thresholds for the CSS QHS checks. This check will flag the channel if any of the flag boundaries are exceeded. For instance if greater than 15% of the

charge is below -15 ADC counts above threshold, this indicates the pedestal is too large for the channel, and the channel will be flagged by CSS.

Although water phase and partial fill phase are discussed in this chapter the focus of this thesis is mostly data from the pure scintillator phase. CSS is one of many components in place to ensure that the data used for analyses in all phases is high quality.

CSS will continue to be used by SNO+ to assist in ensuring the data is of high quality before being processed and help to identify channels that may need investigated. CSS can be used in every phase of the experiment, and there is scope to expand the tests to monitor more aspects of the channels including time, and long window charge measures.

Chapter 5

What did Carbon-11 ever do to us?

5.1 Introduction

Muons that pass through the detector can interact with the detector medium causing spallation neutrons and isotopes (for more detail see Chapter 2), both of which can be backgrounds for physics searches, meaning identification of these are important. There is particular focus on ^{11}C in this chapter as it is a background for solar searches in SNO+, and forms the bulk of the author's analysis, but all expected cosmogenic muon induced backgrounds with a half life above one second that are expected to be seen in SNO+ scintillator phase [68] are explained. For more information on expected backgrounds in the Tellurium phase please see [69].

Understanding the spallation background from high energy muons from measurements taken by SNO+ is important as the spallation neutrons pose a background for dark matter searches being carried out by other experiments at SNOLAB.

A carefully tagged background can also be used as a calibration source, for example ^{11}C has a well understood positron spectrum which, if identified correctly, could be used to test the detector response, or tags of other isotopes.

5.2 Cosmogenic Muons

5.2.1 Primary Cosmic Rays

Primary cosmic rays come directly from astrophysical sources such as stars and supernovae, and consist of protons and helium, as well as other particles produced in stars. With the

exception of solar flares, these cosmic rays are produced outside of our solar system. As can be seen in Figure 5.1, showing the spectra for different nuclei of the primary cosmic rays, primary cosmic rays can be very energetic as they are accelerated by magnetic fields as they travel through the universe to the Earth.

The distribution of the cosmic ray energy as measured at the Auger observatory can be seen in Figure 5.2, which shows the flux peaks around 1 GeV, but cosmic rays are observed at energies up to 10^{21} eV.




Figure 5.1: Primary cosmic ray nuclei fluxes in particles per energy-per-nucleus [70].

5.2.2 Secondary Cosmic rays and Cosmic muon production

Higher energy cosmic rays can interact with the Earth’s atmosphere, creating secondary cosmic rays which in turn go on to interact and create further particles. The initial interactions are strong interactions and produce particles such as pions and kaons. If the secondary particles keep interacting and producing more secondary particles, this creates a hadronic shower, as can be seen in Figure 5.3. Depending on their energy and their location in the atmosphere, the hadrons can either decay into muons or interact further. At lower altitudes, the hadronic components are attenuated, meaning that pions are more likely to be present at higher altitudes.




Figure 5.2: Cosmic ray energy distribution. Measurements taken by the Auger observatory [71]. The green line represents the approximate integral flux.




Figure 5.3: A depiction of a hadronic shower, showing both the hadronic component and the electromagnetic component. Figure from [72].

Muon production can occur in two ways, either through pion decay:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu$$

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu$$

or through kaon decay, which occurs when the cosmic rays have energies higher than around 100 GeV:

$$K^+ \rightarrow \mu^+ + \nu_\mu$$

This muon production occurs at altitudes of around 15 km meaning that, with a decay time of $2.2 \mu s$, only relativistic muons can reach the Earth's surface. The mean energy of the muon by the time it reaches the surface of the Earth is 4 GeV. This means that many muons do decay in the atmosphere by

$$\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu$$

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

which produces another source of the atmospheric neutrino flux in the form of ν_e and ν_μ .

5.2.3 Muons in the SNO+ detector

Both the electromagnetic and hadronic showers do not reach the levels of SNO+, due to the depth of the detector. However, muons are able to travel further through the Earth, meaning some still reach the depths of SNOLAB. A comparison of the muon flux at different depths is shown in Figure 5.4.

As cosmic muons travel through matter, they lose energy with the average energy loss given by Equation 5.1, where the $\alpha(E)$ term gives the continuous energy loss by the ionisation of the atomic electrons. This term has a usual value of $2 \text{ MeV cm}^{-2} \text{ g}^{-1}$ and does not vary significantly with energy. The $\beta(E)$ term is the radiative energy loss, and is primarily from bremsstrahlung which varies considerably and pair production which is continuous energy loss.

$$\frac{dE}{dx} = \alpha(E) + \beta(E) \cdot E \tag{5.1}$$

This energy loss as muons travel through the Earth means the majority of muons that reach the depths of SNOLAB are high energy vertical muons as they have to travel through less of the Earth than sideways ones. Figure 5.5 [74] shows the power-law fit of




Figure 5.4: The variation of muon flux with depth of several underground laboratories [73].




Figure 5.5: A power-law fit of the cosmogenic neutron yield in liquid scintillator (by the Daya Bay Collaboration). There are from experiments at different depths resulting in different average muon energies. Results from SNO Phase 1 are overlaid but not included in the fit, as the measurement is in water not liquid scintillator [74].

cosmogenic neutron yield (detailed in Section 5.3.1) for liquid scintillator experiments with the results from SNO overlaid, showing the average muon energy at SNOLAB is around 400 GeV. The scale of this in the SNO+ detector, assuming 10m muon path length, gives $2\text{MeV cm}^{-2}\text{g}^{-1} \times 0.86\text{g cm}^{-3} \times 1000\text{cm} \approx 1.7\text{GeV}$ energy deposit. 11900 photons per MeV are expected in the SNO+ scintillator cocktail [50], meaning hundreds of thousands of hits are expected per muon. This would show as multiple hits in the PMTs and many PMTs being saturated.

Muon Reconstruction in SNO+

Due to the high amount of light they emit in the detector, muon reconstruction is a very computationally intensive task. Light in scintillator is isotropic, so muon reconstruction relies on the charge distributions of the hit PMTs to try to determine entry and possibly exit points.

Whilst new methods are being tested to reduce the time it takes to reconstruct muons, the framework for tagging muon induced backgrounds is currently being developed without muon reconstruction. Muon reconstruction would have to be significantly faster to be incorporated, just a few minutes per muon.

5.3 Muon Induced Backgrounds in Scintillator

The rest of this chapter will explain the mechanism of cosmogenic muon induced background production, before detailing in brief the expected backgrounds from this mechanism in SNO+. There will then be a focus on ^{11}C , to explain why this isotope was a particular focus for the author's analysis.

5.3.1 Production method

As muons pass through matter they create a flux of virtual photons, with an energy spectrum given in Equation 5.2. In this equation, E is the energy of the muon, α is the fine structure constant, m_μ is the muon mass and ω is the energy of the virtual photon.

$$N_\mu(E, \omega) = \frac{\alpha}{\pi\omega} \left[(1 + \epsilon^2) \ln \frac{2\epsilon E^2}{m_\mu \omega} - 2\epsilon \right] \quad (5.2)$$

As the muon travels through the matter, in this case liquid scintillator, nuclei within the medium can disintegrate through absorbing these virtual photons. This process is known as muon spallation and creates a variety of isotopes and free neutrons. More information on these isotopes, which can present as backgrounds for physics searches, is presented in Sections 5.3.3 and 5.3.2.

Cosmogenic backgrounds in liquid scintillator can be created through muon interactions as they pass through the detector as well as at sea level before the scintillator is brought underground.

5.3.2 Long Lived Isotopes

Some isotopes created through muon interactions with the liquid scintillator are long lived, on the order of months, and can be a background whether produced before or after the scintillator is in the detector, hence one of the goals of the scintillator purification process (discussed in more detail in Chapter 3) is to reduce these backgrounds before the scintillator is loaded into the detector [75]. There is significantly higher production of cosmogenic backgrounds on the surface due to the higher muon flux.

^{14}C

^{14}C is produced by cosmogenic muon interactions with ^{14}N and has a half life of 5700 years. Production of this isotope would occur before the scintillator is produced. 4.08×10^9 ^{14}C decays are expected per year once the detector is full of scintillator. Rather than being a

specific background to any particular physics analyses planned by SNO+, ^{14}C is expected to contribute to pile up backgrounds.

^7Be

^7Be is created through spallation neutrons interacting with the ^{12}C molecules in the scintillator. It has a half life of 53.2 days and decays through electron capture into ^7Li with a Q-Value of 0.48 MeV. The majority of this isotope SNO+ expects to see will be created during transport and storing of the scintillator, i.e. before it is transported underground. From this part of the production, we expect 1.84×10^9 ^7Be particles in the detector at the beginning of fill [68]. Combining this starting value with the rate of production once underground, SNO+ expects 9360 decays of ^7Be per year once fully filled with scintillator.

^{39}Ar

^{39}Ar is produced by cosmogenic muon interactions with ^{40}Ar . The ^{40}Ar can get into the scintillator from the air, so a first step to limit this background is to limit the scintillator's exposure to air and degass the scintillator once underground. It decays through Beta decay to the ground state of ^{39}K with a Q-Value of 0.585 MeV. It has a longer half life than most isotopes produced in this way, 269 years. Whilst this is not a background for the planned solar searches in scintillator phase, it carries the risk of pile up backgrounds if too much of it is produced. During scintillator phase of SNO+, 85483 events per year are expected [68].

5.3.3 Short Lived Isotopes

Table 5.1 shows the short lived muon induced backgrounds found in Borexino and Super-K, their lifetime, Q-Value and decay mode. Borexino is a liquid scintillator detector whilst Super-K uses water as its detection medium. Both are shallower depth than SNO+, meaning they experience a higher flux of muons, and with different energy spectra, meaning the background composition of SNO+ will differ from these, but is expected to be made of some combination of these short-lived modes.

Due to the short lived nature of the isotopes presented in this section, only those created once the scintillator is in the detector will present as a background, as any created during production and transportation of the scintillator will decay before the scintillator is loaded into the detector. In an attempt to reduce these backgrounds in data taking, two vetos are in place following a muon passing through the detector, one to allow the

‘short’ followers to decay (20 s) and one to allow the ‘long’ followers to decay (60 s).

Table 5.1: Cosmogenically induced backgrounds found in Super-K [76] and Borexino [65].

Isotope	Lifetime	Q-Value (MeV)	Decay Type
^{11}Li	12.0 ms	16.4	β^-
^{13}O	13.0 ms	20.6	β^+
^{12}N	15.9 ms	17.3	β^-
^{14}B	20.0 ms	14.6	β^-
^{13}B	25.0 ms	13.4	β^-
^{12}B	29.1 ms	13.4	β^+
^{12}Be	34.0 ms	11.7	β^-
^8He	171.7 ms	10.7	β^-
^9C	182.5 ms	16.5	β^+
^9Li	257.2 ms	13.6	β^-
^{16}C	1.08 s	4.0	β^-
^8B	1.11 s	18.0	β^+
^6He	1.16 s	3.51	β^-
^8Li	1.21 s	16.0	β^-
^{15}C	3.53 s	9.8	β^-
^{16}N	10.3 s	10.4	β^-
^{11}Be	19.9 s	11.5	β^-
^{10}C	27.8 s	3.65	β^+

5.3.4 Muon Induced Isotopes in SNO+

^{11}C

Carbon-11 is most commonly produced by a cosmic muon interacting with a ^{12}C atom, producing a ^{11}C and a neutron that is later captured. There are 6 possible ways that this isotope can be produced through interactions with cosmic muons, and in 95% of the time they are produced alongside a neutron, with the other 5% of decays considered invisible channels, as they produce no characteristic neutron [77]. The neutron is captured on Hydrogen 99% of the time, resulting in a 2.2 MeV γ , with the other 1% captured on Carbon, producing a higher energy γ , 4.9 MeV [78]. ^{11}C decays via




Figure 5.6: ^{11}C Spectrum shown alongside the expected low energy solar signals.[61]



The expected decay rate of ^{11}C in SNO+ is $(1.14 \pm 0.21) \times 10^3$ decays /kt/yr, from [78], found by extrapolating data obtained by KamLand.

The half life of ^{11}C is 20.3 minutes, too large to be removed from SNO+ by instilling a deadtime after a muon passes through the detector. This makes it a prime candidate for the three-fold tagging identification which will be expanded upon in section 5.4.

Due to the decay energy of ^{11}C it is expected to be a problematic background for solar neutrino studies, specifically for *pep* and CNO neutrinos, during the scintillator phase of SNO+, as can be seen in Figure 5.6.

^{16}N

^{16}N is produced by cosmogenic muon interactions with ^{16}O . It has a half life of 7.13 seconds and decays through β^- decay with a Q-Value of 10.4 MeV. Due to this relatively short half life, this background is expected to be removed via the previously discussed muon follower vetos.

^8Li and ^8B

^8Li and ^8B are often combined in searches [65] due to their similar half lives (1.21 s and 1.11 s respectively) meaning they occur in the same time window following a muon. They can be differentiated by their different Q-Values (16 MeV and 18 MeV respectively) and decay modes (β^- and β^+). They are both produced as a result of cosmogenic interactions with ^{16}O . They are distinguished from other isotopes with similar Q-Values but shorter half-lives by a lower bound on the time window after the muon, in KamLAND taken to be 0.6 s. [79].

 ^6He

Although ^6He has a similar half life (1.16 s) to ^8Li and ^8B , presenting some overlap in searches, it is easier to distinguish from the other two due to its much lower Q-Value, 3.51 MeV. It is produced as a result of cosmogenic interactions with ^{16}O .

 ^{10}C

^{10}C , much like ^{11}C has a free neutron produced with it, allowing it to be identified with a three fold coincidence. It has a half-life of 27.8 s and decays through β^+ decay with a Q-value of 3.65 MeV.

5.4 Three-fold coincidence Tagging

As detailed earlier in this chapter, a short veto is applied following a muon to tag short lived isotopes as a result of that interaction. However, for longer lived isotopes such as Carbon-11, this approach is not appropriate, as the veto would result in too significant of a loss of livetime. For these isotopes with a long half life and a coincident neutron capture involved in their decay, a three-fold coincidence approach is required, firstly identifying the muon passing through the detector, next tagging the coincident neutron capture, and finally identifying isotope decays from near the neutron capture. This approach is used in other experiments such as Borexino [80], who reconstruct the muon track and then only consider neutron captures within a cylinder around the muon track to be coincident. The approach the author has developed for SNO+ will be outlined in the following three chapters.

Chapter 6

Partial Fill

The partial fill phase gives the first opportunity to tune the cuts required for this analysis on a period of stable data. This is only the first stage of tuning required due to the detector changes during filling and PPO-loading periods.

6.1 Water Phase Muon Tag in Partial Fill Data

In the SNO+ water phase, for an event to be classed as a muon, it must have at least 150 normal PMT hits (nhits) and 5 Outward Looking PMT hits (OWL PMT hits). The water phase muon tagging also vetos high-charge events that have more than 2 of the 4 PMTs that face the neck of the detector (neck PMTs) triggered to try to prevent breakdowns or high noise events being tagged as muons. A natural result of adding scintillator into the top of the detector is that energetic events such as muons are more likely to trigger neck PMTs than they were in water phase, meaning that the water phase tagging vetos several muons that pass through the scintillator when run on partial fill data. Running the water phase criteria without the neck veto over partial fill data produces a much higher average muon yield than would be expected in SNO+. For the purposes of this chapter, running the water phase criteria over partial fill data refers to the water phase criteria without the neck cut. For a detailed outline of the water phase criteria and a comparison to the retuned partial fill criteria see Table 6.1.

Figure 6.1 shows the OWL PMT hit distribution in water phase (blue) compared to the water criteria run over partial fill data (black). As the OWL PMTs detect Cherenkov light in the external water, this distribution should not have been affected by changing the detector medium, implying the criteria is falsely tagging other, lower energy events as muons. This is supported by Figure 6.2, showing the nhits distribution of the water phase




Figure 6.1: OWL PMT hits distribution of events identified as muons by running water phase criteria on water phase data (blue), and partial fill data (black). Area normalised for comparison.




Figure 6.2: Nhits distribution of events identified as muons by running water phase criteria on water phase data (blue), and partial fill data (black). Area normalised for comparison.




Figure 6.3: Nhits against OWL PMT hits of events identified as muons by running water phase criteria on partial fill data. The red line is at 8 OWL PMT hits, to demonstrate the potential impact that increasing the OWL PMT threshold could have.

Table 6.1: A comparison of the criteria for a muon between partial fill and water phase.

	Water Phase	Partial Fill Phase
OWL hits	5	5
Neck hits (<5000 nhits only)	Maximum 2 neck PMT hits AND occuring earlier than nhits at bottom of detector	None
Definition of OWL event	Any with 5+ OWL hits	Any with 5+ OWL hits
Time since last OWL event	5000 ns	5000 ns
Nhits	150	750
Spread of Z position of OWL hits	None	<25 mm
Spread of TAC of OWL hits	None	<25

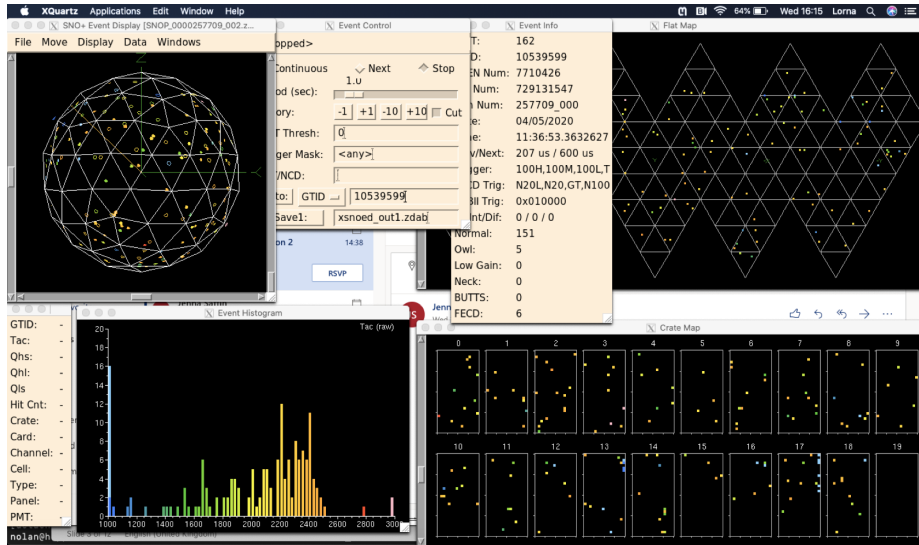


Figure 6.4: An example of an event falsely tagged as a muon as viewed in the SNO+ event display.

(blue) and the water phase criteria run over partial fill data (black). Figure 6.2 shows that there is a higher proportion of lower energy events being tagged as muons in partial fill phase.

Figure 6.3 shows the OWL PMT hits against the nhits for the water phase cuts applied to partial fill data with the red line showing an example increased OWL PMT cut at 8 OWL PMT hits to show the effect this would have. Although this shows that the majority of the lower OWL PMT hit events are low energy, a higher cut on the OWL PMT hits would still leave a lot of lower energy events, and risks missing true muons. Upon further investigation of the events in this low nhit/OWL PMT hit region, most of the OWL PMTs in these events are located near the region where small holes have been drilled in the PSUP as part of the instillation of the hold-down ropes. This means there is a possibility a small amount of scintillation light is leaking into the external cavity water, which would partially account for the marked increase of low OWL PMT hit events being tagged as muons.

6.2 Altering the Muon Tag for Partial Fill Phase

In order to alter the muon tagging to work for the partial fill phase, a large sample of events tagged as muons in partial fill data by the water phase criteria was created. A portion of this dataset was then handscanned to identify a sample of events that were falsely tagged as muons, and a sample of events that appear to be true muons. This was done by looking at the events in XSNOED, the SNO+ event display. This display shows

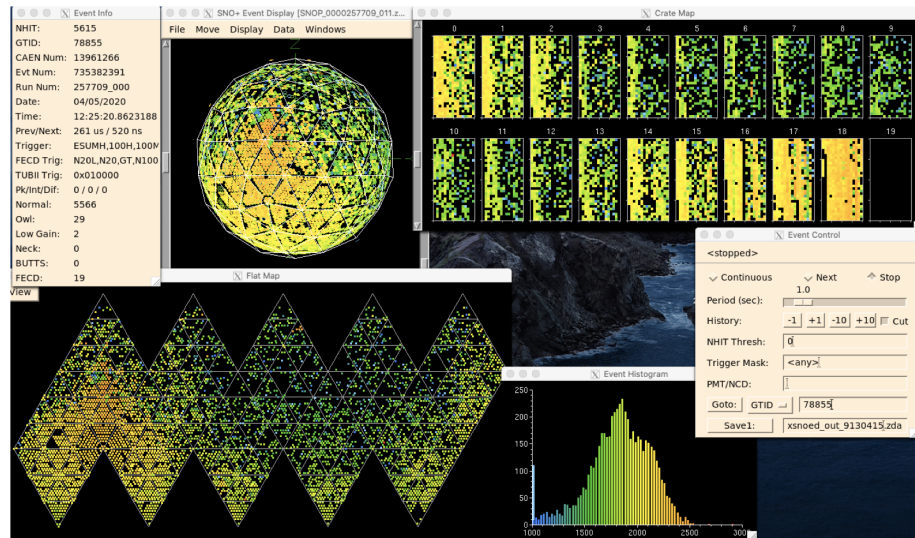


Figure 6.5: An example of an event correctly tagged as a muon as viewed in the SNO+ event display.

the PMT hits in a 3D spherical view, a crate map view, and a detector flat map. It allows these views to be in terms of PMT charge, or first time the PMT was hit in the event, also showing a histogram of these values. Whilst SNO+ has no physical equipment on site to formally classify a muon as it passes through the detector, the event display tool offers valuable information allowing muons to be identified with a reasonable degree of certainty. Figure 6.5 shows an event from this dataset that appears to be a true muon. It is fairly high energy, the majority of the OWL PMT hits are nearby the earliest events in the AV. Additionally, a time correlation can be seen, consistent with a particle travelling through the detector. For contrast, Figure 6.4 shows an event from this dataset that appears to be an event falsely tagged as a muon. The most stark difference is how many fewer PMTs this event has triggered, but there are also other features inconsistent with it being a muon; the five OWL PMT hits are very separated from each other, and there is no obvious time or spatial correlation between the hits as would be expected with a muon.

Using this partially tagged dataset, several features of the OWL PMT hits were investigated, the most relevant are the RMS of the position of the OWL PMTs triggered in Z (the vertical direction), the mean Z position of the OWL PMTs triggered, and the RMS of the timing of the OWL PMTs triggered (TAC, a raw event timing, for more detail see Chapter 3). Figure 6.6 shows the mean position in Z of the OWL PMTs triggered in the event. As can be seen from this plot, the events manually identified as muons have a higher mean Z position, consistent with most of SNO+'s muons being downgoing. Despite this correlation, the decision was made to not use this as a muon identification, due the

risk of excluding true upgoing muons which may not have been seen in this sample but would have a lower mean Z position.

Figure 6.7 shows the spread of the timing of the OWL hits for this dataset, showing a group of falsely tagged events and unclassified events with very limited spread in timing. This is in contrast with the muons, and several false muons. TAC starts ramping when the first PMT is triggered, and continues throughout the event, meaning events with such a small distribution of TAC indicate that all the PMTs in the event triggered almost simultaneously, and hence are more consistent with electronic noise events than they are with a true physics event. This cut was designed to be conservative, focussing on accepting all muons over rejecting all electronic noise events. For this reason, a minimal cut was placed on the spread of the timing such that if the event has less than 25 RMS of TAC then it will not be tagged as a muon.

Figure 6.8[a] shows the spread of the position in Z of the OWL PMTs triggered in that event. Similar to the spread of the TAC, it seems like the true muons have a larger spread. This is consistent with expectations, as when a muon passes through the external cavity, there is likely to be a larger spatial spread in OWL PMTs that are triggered by the resultant Cherenkov light compared to electronic noise which could trigger just a small amount of PMTs directly next to each other. Figure 6.8[b] shows this same plot zoomed in to determine an appropriate cut placement. As this zoomed in plot shows, occasional muons do have a small spread of Z position. To minimise the chance of rejecting true muons, a cut was placed on the spread of the Z position such that if the event has less than 25 mm RMS of the OWL PMTs Z position, it will not be tagged as a muon.

Nhits threshold adjustment

As discussed above, changing the internal detector medium gives no motivation to increase the external PMT threshold. It does motivate increasing the internal hit threshold. Figure 6.9 shows the effects on the average hourly muon candidate rate with varying nhit thresholds. This is the value averaged over multiple runs. As can be seen, there is a significant drop at 750 nhits to much closer to the 3 muons per hour expected. There is a far less significant change at 850 nhits, so the nhits threshold was set to be at 750 nhits, to ensure a reasonable rate without going higher than required to minimise the risk of missing muons that merely clip the scintillator before passing through the water.




Figure 6.6: Mean of the position in Z of OWL PMT hits. Red line shows events believed to be muons, blue line shows events believed to have been falsely tagged whilst the black line has not been identified either way.




Figure 6.7: RMS of the TAC of OWL PMT hits. TAC is a measure of the timing of PMT hits within an event, for more detail see Chapter 3. Red line shows events believed to be muons, blue line shows events believed to have been falsely tagged whilst the black line has not been identified either way.

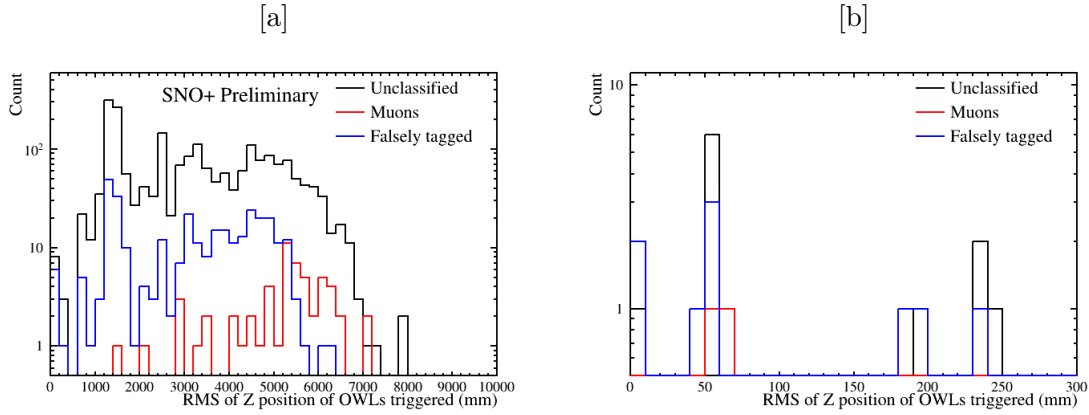


Figure 6.8: RMS of the position in Z of OWL PMT hits, with [b] zoomed in to show only the relevant section of [a]. Red line shows events believed to be muons, blue line shows events believed to have been falsely tagged whilst the black line has not been identified either way

6.3 Results of these modifications

6.3.1 In stable partial Fill

The muonTag results have been evaluated by comparing to water phase results in 3 different metrics to measure the rate, which naturally should not change between phases. The water plots for comparison are from Billy Liggins' thesis [67] looking at muons in the water phase. The plots in [67] include some cuts on the path length through the detector which the comparison plots in partial fill do not have, so a slightly increased rate is expected. Figure 6.10 shows the muons per hour in hour long runs only, [a] shows water phase whilst [b] shows the partial fill phase. This is a slightly biased estimator of the rate, as it excludes shorter runs, but it shows a consistent rate with water phase.

Figure 6.11 shows the time between muons (s), with [a] showing the water phase, and [b] the partial fill. This only includes muons within the same run, so it is also a slightly biased estimator, as this cuts the tail of the exponential. The effect of this can be seen by comparing Figure 6.11[a] with Figure 6.12[a], where Figure 6.12 shows the time between muons but including different runs. When the time between runs is included, it shows a better match with the SNO result [81]. As can be seen through all three of these measures, although the rate in partial fill is slightly higher, it is not significantly higher than water phase, and is a sensible rate.

Figure 6.13 shows the nhits [a] and OWL PMT hits [b] distributions for the retuned partial fill cuts. The nhits distribution shows a large amount of the muons in this sample




Figure 6.9: Average hourly muon tagging rate over a sample of partial fill data with varying minimum nhits. Error bars too small to be seen, error is statistical from averaging over multiple runs.

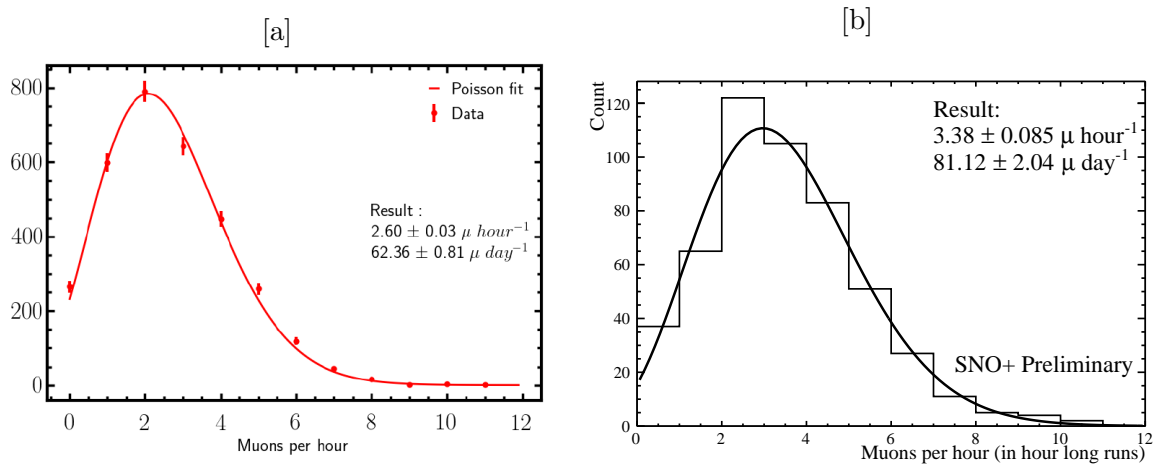


Figure 6.10: The muon rate per hour in water phase [a] compared to the muon rate per hour in partial fill with the adapted cuts [b]. Data from the stable partial fill period of summer 2020.

saturate the detector, which was seen more rarely in water phase, but is expected due to the higher light yield of scintillator. The OWL PMT hits distribution on the other hand is now very consistent with what was seen in water phase.

OWL hit distribution in partial fill

Figure 6.13 shows the OWL hits distribution [b] with the adapted partial fill cuts. [b] is overlaid with the OWL distribution from the water phase. A longer tail can be seen on

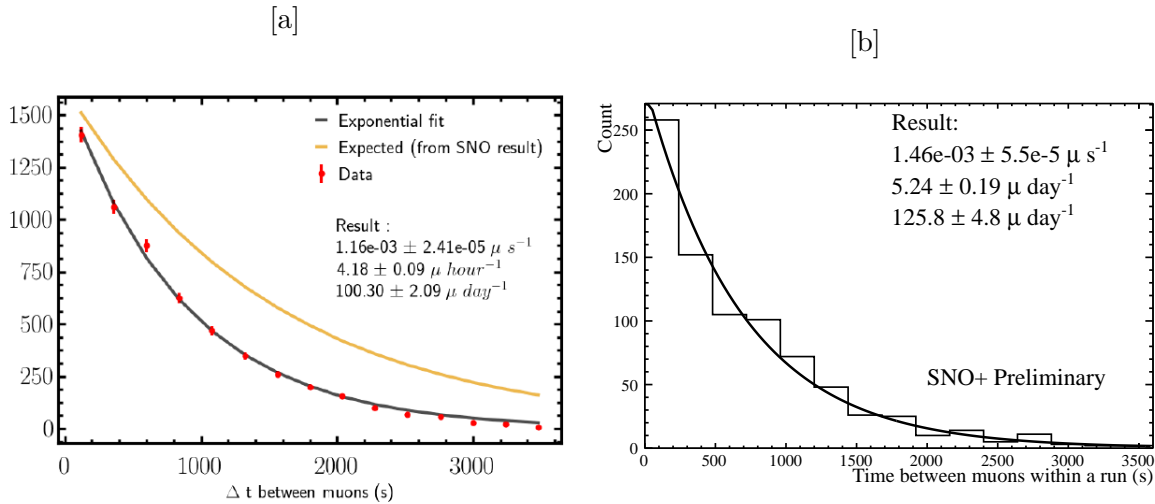


Figure 6.11: The time between muons in water phase [a] compared to the time between muons in partial fill with the adapted cuts [b]. This is only muons within a run, the time between muons from other runs is excluded. Data from the stable partial fill period of summer 2020.

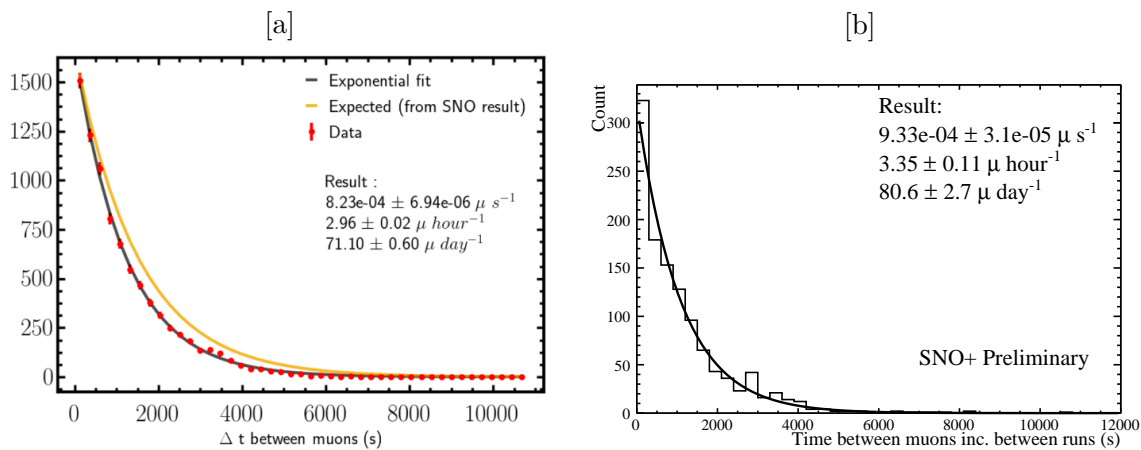


Figure 6.12: The time between muons in water phase [a] compared to the time between muons in partial fill with the adapted cuts [b]. This includes the time between muons even if they're from a different run. Data from the stable partial fill period of summer 2020.

the partial fill distribution, as well as a few events that trigger every OWL PMT. The explanation for the increased tail appears to be that there is a small amount of scintillation light leaking into the cavity on particularly energetic events. The events that saturate the OWLs are breakdowns, an example can be seen in Figure 6.14. Breakdowns such as these trigger every PMT which would include the neck PMTs, explaining why these events are tagged as muons now the neck veto has been removed. These events are also tagged as breakdowns, and the DC group agrees there is no harm in tagging occasional breakdowns

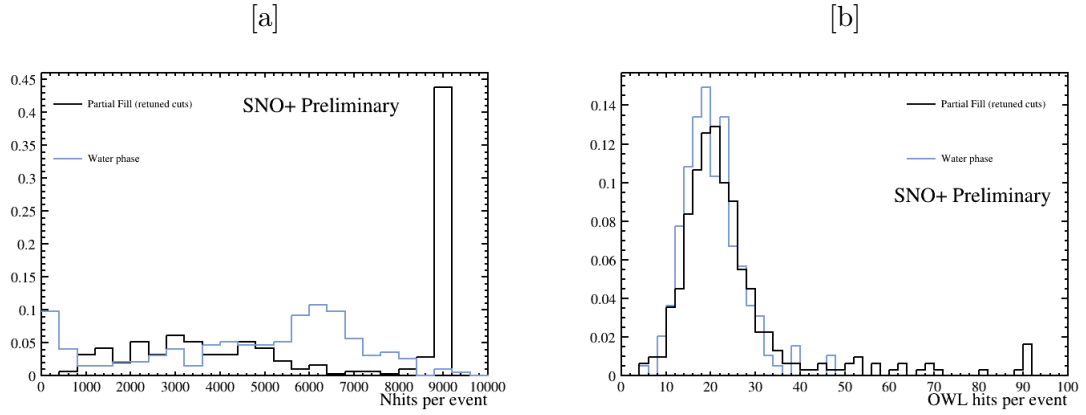


Figure 6.13: Comparison of water phase muons and partial fill muons with retuned cuts for nhits [a] and OWL PMT hits [b].

as muons and breakdowns.

6.4 Neutron Tag

The data cleaning processor NeutronTag is a new processor developed by the author for partial fill/scintillator phase, and tags coincident neutron captures following muons. This processor follows the framework of the second pass of the existing data cleaning step called TPMuonFollower.

This processor flags an event as a coincident neutron capture if the event is between $20 - 2000 \mu\text{s}$ from a muon, and has between 510 and 750 nhits. The upper bound on the timing between the muon and coincident neutron capture aims to maximise the amount of coincident neutron captures detected whilst minimising the introduction of random backgrounds into the sample. This value is based on the expected neutron lifetime ($\tilde{200} \mu\text{s}$), as well as the values used for similar studies by Borexino [65] and the water phase of SNO+ [67]. The lower bound on the timing is to counter the electronic effects seen following a muon passing through SNO+, known as muon ringing. Chapters 8 and 10 of [67] detail this effect as well as showing a brief investigation into the length of this effect using the TELLIE calibration system. The nhits window for the coincident neutron captures was determined using Monte Carlo simulations of neutron captures, the nhits window was chosen to include 99% of the simulated events.

The results of this processor ran over stable partial fill data can be seen in Figure 6.15. Figure 6.15[a] shows the time between the muon and the neutron capture candidate in nanoseconds, with a τ of $186 \pm 19 \mu\text{s}$, consistent with expectations. Figure 6.15[b] shows the nhits distribution for events selected with this criteria (black histogram) compared

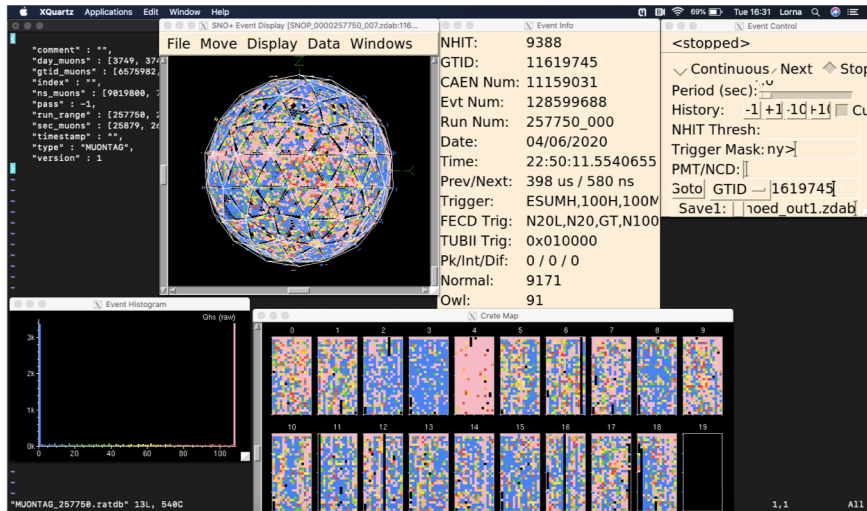


Figure 6.14: An example of an event that saturated the OWLs viewed in xsnoed. The charge in crate 4 indicates this is a breakdown rather than a Physics event, as almost the whole crate has a very high charge, much higher than the majority of the rest of the PMTs. In a true Physics event we see channels which are close spatially have similar charge distributions, rather than just channels that are in the same crate. The number of OWL PMT hits are shown in the textbox in the event viewer as well as the small dots on the very edge of the sphere.

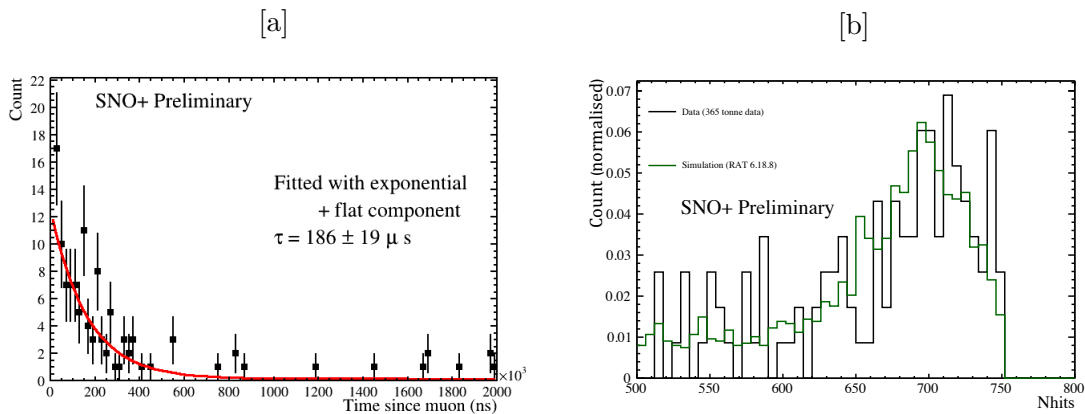


Figure 6.15: The results of the neutronTag processor over stable partial fill data. [a] shows the time in nanoseconds between the muon and the neutron capture candidate, whilst [b] shows the nhits of the neutron capture candidates compared to the simulation.

with simulated events (green histogram). The Monte Carlo simulation used in this section was made with the updated partial fill geometry in RAT-6.18.8.

6.5 Integrating these changes into the DC framework

When the author began work on these changes, there was no way to implement them into the data cleaning framework, as the framework was only built to hold one version of each data cleaning cut. This meant that in order to integrate the changes to the muon tagging algorithm, we would have to remove the ability to use the water phase version of the cut. This structure also meant that if the cut was updated to the partial fill version, the same issue would be encountered if any changes were required for the scintillator phase. It is not just the muon tagging cut that would benefit from the functionality of having multiple versions of the same cut able to be run, several of the data cleaning cuts rely on an nhit threshold which renders the water phase tuning of the cuts unusable in partial fill and full scintillator phase.

To solve this problem, the author has added the provisions needed in the framework to have multiple different definitions of cuts depending on which data is being analysed. This is done through additional indexed tables of cut values, the relevant index is called in the processing step that starts the data cleaning on the run. This allows the current index to be used in the immediate run-by-run data processing, whilst still allowing analysers to run the appropriate data cleaning over their data if they need to reprocess any older runs.

In this chapter we have described the steps needed to develop and tune a framework to tag muons and muon-induced neutrons, and we have demonstrated that the tagged neutrons show the expected time distribution after the muon. However, the statistics in the partial fill data set is somewhat limited and as we were unable to fully calibrate this period, the focus of this thesis is on the full scintillator phase which will be discussed in the next chapter.

The work discussed in this chapter sets the foundation for the author's analysis with the full scintillator data, as well as allowing all data cleaning cuts to be retuned as the detector medium changes without removing the ability to use that cut on any older data. The muon tagging, neutron tagging, and the changes made to the data cleaning framework are currently being used for every run in the SNO+ experiment, and will continue to be utilised by the experiment in the future.

Chapter 7

Tagging Muons and Coincident Neutron Captures in Scintillator Phase

The tools developed for tagging muons and coincident neutron captures in the stable partial fill phase needed to be tested on the scintillator phase to see if the increased scintillator level and altered PPO concentration resulted in a need to change the cuts.

Livetime and run list

For this thesis analysis, the most stable period of data available at the time of analysis was chosen, from May 2021 - July 2022. Two run lists were used for this analysis, all runs from the Preliminary Partial Gold run list from run 270918 to run 282007 were used, with a total livetime of 5005 hours 50 minutes. All runs from the Preliminary Scintillator Gold run list from run 300000 to run 302068 were used, with a total livetime of 1107 hours 37 minutes. There is no difference between these run lists regarding the criteria used to select the runs, but a new run list was created to reflect the formal beginning of the SNO+ scintillator phase at run 300000. The two run lists combined gave a total livetime of 6113 hours 27 minutes. Both Preliminary Scintillator Gold and Preliminary Partial Gold are internal SNO+ run list names to refer to our gold standard runs that pass the highest level of run selection tests. To make these lists a run must have all electronic crates online, be at least half an hour in length, and pass a variety of data quality checks including that there is no unusual activity underground, no electronics alarms during the run, and minimal 'orphaned' events which are events that are not assigned an event ID.




Figure 7.1: Time between muons (in seconds) for muons in scintillator phase of SNO+.

7.1 Muon Tagging

The muon tag was outlined in the previous chapter whilst detailing tuning it for partial fill phase. The tag did not require further tuning for the full-fill scintillator phase, as can be seen in Figure 7.1, the result is consistent with the partial fill phase.

Although no further changes were required at this time, the muon rate should be monitored to ensure it can be retuned as detector conditions change.

7.2 Neutron Tagging

After the filling was complete there was a period of PPO top up, resulting in an inconsistent PPO level throughout the dataset, as PPO acts as a fluor in the scintillation cocktail, the PPO concentration affects the light yield in the detector (for more details see Section 3.4 of Chapter 3). The required nhits thresholds can be calibrated using the results of regular BiPo214 tagging in the detector. ^{214}Bi is a decay product in the ^{238}U decay chain, which can be used to monitor levels of ^{222}Rn in the detector. ^{214}Bi decays into ^{214}Po through β decay, which in turn decays through α decay into ^{210}Pb . The ^{214}Po has a decay lifetime of $164.3\ \mu\text{s}$, giving a distinct coincident signal to the BiPo214 decay. As well as allowing for a straightforward removal of the background, this distinct coincidence signal also allows a




Figure 7.2: Time between muons and coincident neutron captures in scintillator phase of SNO+. The red line shows a fit using an exponential plus a constant. Tau represents the neutron capture lifetime.

very pure sample of BiPo214 events to be tagged to monitor radon levels in the detector. These tagged events are plotted in terms of nhits and shown regularly in calibration calls to allow analysts to tune their cuts if required, as well as to allow the whole collaboration to understand how the light yield was changing.

7.2.1 Data Cleaning Level Cuts

The time since the muon condition of the cut remained the same as the partial fill phase (20-2000 μs), but the nhits requirement for the coincident neutron capture was changed to 400 - 800 nhits. With these cuts alone run over 3 weeks of data to test, an unusually high number of candidates in early time were seen, as can be seen in Figure 7.2.

Secondary Cuts to Address High Volume of Early Time Neutron Capture Candidates

When a muon passes through the SNO+ detector it creates a high amount of light, usually lighting up every PMT in the detector, Figure 7.3 shows the distance travelled through the detector in metres against the number of PMTs triggered in the event for simulated

muons. In addition to this, many muons entirely saturate the charge on many PMTs. As a result of this there is a settling down process over the next few milliseconds known as muon ringing where the events the detector is registering are just the electronics settling back down to normal. These events can be easily recognised in the event display, as they are ‘confetti’ events or flat TAC events. Because they aren’t physics events, purely electronic noise, the PMTs trigger randomly throughout the event, so the resultant event looks flat in time. Examining a sample of the early time events showed several events of this nature. Because these events are mostly just electronic noise, the `nhitsCleaned` (the `nhits` of an event after data processing removes any electronic effects, for more information see Chapter 3) of the event will be very small. To investigate this, `nhitsCleaned/nhits` was plotted for all the coincident neutron capture events, with the understanding that true Physics events would be close to one, whilst confetti events would fall lower on the scale. This plot can be seen in Figure 7.4, which shows three bands of events. The events which are in the bottom band are electronic effects, and the middle band is likely to be events which are true Physics but still have some residual electronic effects in. After examining this plot, a secondary post data cleaning cut was imposed, specifying that the event must have at least 70 `nhitsCleaned`. The purpose of this cut was to remove the lower band of events, the electronic effects. With the `nhits` range being used for this tag, an `nhitsCleaned` threshold of 70 ensures that all events have over 15% of their `nhits` remain after cleaning, removing events in the lower band from the sample. The effects of this cut over the same 3 week dataset can be seen in Figure 7.5, which shows a much improved distribution, and a neutron capture time (τ) consistent with expectations.

At the time of writing these cuts are being run as a part of the data cleaning section of data processing. The resulting tables of muon and neutron events are uploaded to the internal database, and used for the Carbon-11 likelihood classifier, which will be detailed in the next chapter.




Figure 7.3: Distance a muon travels through the AV (m) plotted against the number of PMTs triggered in the event. Made from MC simulations in RAT-6.18.12.




Figure 7.4: NhitsCleaned divided by nhits for all coincident neutron capture candidates in a 3 week dataset. Events which score low are ‘confetti’ events, events which are purely electronic noise. These events have no correlation in time unlike Physics events which makes them appear multi-coloured in the SNO+ event display, resulting in their nickname of confetti events.




Figure 7.5: Time between muon and coincident neutron capture for a 3 week dataset in scintillator phase of SNO+. This plot contains only events with a minimum `nHitsCleaned` of 70. The red line shows a fit using an exponential plus a constant. Tau represents the neutron capture livetime.

Chapter 8

Carbon-11 Likelihood Classifier

8.1 The Framework

The first two steps of the three-fold coincidence have been outlined in the previous two chapters. The tables from these steps are read in by a likelihood algorithm to determine the likelihood of an event being Carbon-11.

A likelihood algorithm was picked in order to allow for both a high efficiency background rejection whilst also allowing a pure sample of ^{11}C to be obtained. Due to this being a background for the solar analysis, a Fiducial Volume (FV) was set at 5.5m, mirroring the FV used in the water phase solar analysis.

The likelihood algorithm uses a look-back approach, by first identifying events in the energy range of a ^{11}C (0.95 - 1.95 MeV) and then looking at previous runs to identify a coincident neutron following a muon in the 142.1 minutes (7 half-lives of ^{11}C) beforehand.

$$L_{^{11}\text{C}} = P(^{11}\text{C}/nhits) \times P(^{11}\text{C}/\delta t) \times P(^{11}\text{C}/\delta d) \quad (8.1)$$

Equation 8.1 shows the likelihood equation for the ^{11}C candidates, where $L_{^{11}\text{C}}$ is the likelihood the event is a ^{11}C decay. The algorithm uses a PDF to calculate the likelihood of each of the components, $^{11}\text{C}/nhits$ for the nhits, $^{11}\text{C}/\delta t$ for the time between the muon and the candidate decay, and $^{11}\text{C}/\delta d$ for the spatial separation of the candidate and the coincident neutron capture.

8.1.1 The PDFs

All of the PDFs are normalised such that the maximum likelihood is 1.




Figure 8.1: The nhits PDF of ^{11}C candidates. This PDF contains simulated events, made using RAT-6.18.12. A fiducial volume of 5.5 m was set on these simulations to match the fiducial volume used by the likelihood algorithm.

Nhits

The nhits PDF, seen in Figure 8.1, was made from MC simulations. Due to the binning, 10 nhits per bin, this PDF will create a small amount of blockiness in the likelihood spectrum.

Time

The time PDF, Figure 8.2, was made by using a standard exponential function with a τ of 20.3 minutes, consistent with the half life of ^{11}C .

Spatial

The PDF for the spatial separation of the ^{11}C candidate and the coincident neutron capture was created from MC simulations of muons to get the MC value for how far from the muon the coincident neutron capture wanders. The initial output from this was then smeared using the reconstruction uncertainties of both the ^{11}C and coincident neutron captures. Using the MC simulations, the difference between the MC truth position and the reconstructed position in x, y, and z were separately plotted for both the ^{11}C and the neutron capture. To smear the PDF, a position was randomly generated within a 6 m sphere, representing the ^{11}C decay, then a second point was generated at a random distance away sampled from the MC plot of distance the neutron wanders before being captured. Once those two points were generated, they were adjusted in x, y, and z by a random value drawn from the respective reconstruction uncertainty plot. The distance




Figure 8.2: The time PDF of ^{11}C candidates. Made from exponential function with τ of 20.3 minutes.




Figure 8.3: The spatial separation PDF of ^{11}C candidates and coincident neutron captures. This was made from MC simulations of coincident neutron captures, and was smeared with the uncertainties of the ^{11}C and coincident neutron capture reconstructions.




Figure 8.4: A sample of 10% of the data sample for ^{11}C candidates, showing the likelihood distribution for this sample.

between the two points after this adjustment was then plotted, this process was repeated to give Figure 8.3.

8.2 Tuning the cuts

8.2.1 10 percent of sample

In order to begin determining the appropriate cut threshold (alpha value) for the likelihood, a small sample of 10% of the data was examined, this sample was sufficient to see the likelihood distribution forming, as shown in Figure 8.4.

Figure 8.5 shows the time between ^{11}C candidate and preceding coincident neutron capture for all events with a likelihood of 0.2 or greater in this 10% sample in black, and all events in the sample with a likelihood of 0.4 or greater in red. Whilst the statistics are too low on such a small sample to draw any meaningful conclusions, the distribution with the cut at 0.4 has more candidates in the first time bin compared to the sample with the cut at 0.2, indicating the sample with the cut at 0.4 is likely to be a better exponential, as expected with a purer sample.

Tuning the alpha cut value

In order to determine a sensible α cut value the likelihood of random coincidences was compared to the likelihood of true ^{11}C . The theoretical ^{11}C likelihood spectrum was created




Figure 8.5: A sample of 10% of the data sample for ^{11}C candidates, showing the timing distribution for events with a likelihood of 0.2 or above in black, and events with a likelihood of 0.4 or above in red.




Figure 8.6: The black line shows the theoretical likelihood distribution of ^{11}C , made from randomly sampling values for nhits, spatial separation, and time since coincident neutron capture from simulations and then running those values through the likelihood algorithm. The red line shows the random coincidence distribution, found by selecting events below the ^{11}C nhits region, and running these events through the likelihood algorithm with an nhits value randomly sampled from the ^{11}C simulations.

by randomly generating values from the MC simulations of spatial separation, nhits, and time separation. These values were then run through the likelihood algorithm.

The random coincidence spectrum was created by selecting events below the ^{11}C nhits region. An nhits value was randomly sampled from the ^{11}C MC simulations for each of these events, which were then run through the likelihood algorithm to see the likelihood of a non- ^{11}C event being close enough in time and space to a coincident neutron capture.

Figure 8.6 shows the theoretical ^{11}C likelihood distribution (black) and the random coincidence distribution (red).

These two distributions were used to create the purity and efficiency at a variety of potential cut values, as can be seen in Figure 8.7, where the red line is the efficiency and the black line is the purity. As this is to remove a background for the solar neutrino studies, it was also important to ensure that any cuts do not also remove a significant fraction of the solar neutrinos. The fraction of the solar neutrinos removed by each cut value can also be seen in Figure 8.7, the green line. As this shows, the fraction of solar neutrinos removed by the cuts drops off sharply at a cut value of just 0.01.

For the background removal cut, the goal is to remove as many ^{11}C candidates as possible, whilst preserving signal. This will be used in analyses such as solar analyses. For this background rejection, an α value of 0.01 was selected. This is a sacrifice of just 1%




Figure 8.7: The red line is the efficiency and the black line is the purity. As this is to remove a background for the solar neutrino studies, it was also important to ensure that any cuts do not also remove a significant fraction of the solar neutrinos. The fraction of the solar neutrinos removed by each cut value can also be seen in the green line.

of signal, whilst rejecting 94% of the ^{11}C .

In order to develop a pure sample of ^{11}C , ideally the cut value would be at a very high purity. As can be seen from these figures, a higher purity results in a much lower efficiency resulting in a much smaller sample. Due to the limited statistics available for this thesis, to gain as high purity as would be ideal, the sample size would be too small to draw any conclusions from. For this reason 2 cut values (0.2, 0.4) are presented in this thesis, to demonstrate the merits of this analysis with the data available, whilst also demonstrating what we can see from the higher purity samples, showing what would be possible if more data was available.

8.3 Systematic considerations

As outlined in the previous section, this analysis ended up being statistically limited (7.9% for the sample of 0.4 or above, 5% for the sample of 0.2 or above). However, when there is more data, the following systematics should be considered.

1) Uncertainty on the separation between ^{11}C and neutron capture. This uncertainty comes from the potential effect of recirculation on the position of the ^{11}C . Owing to the small time difference between the muon capture and the coincident neutron capture it is unlikely to have any significant impact on the location of the neutron, but as the time difference of the ^{11}C is much greater, this could have a potentially significant impact. It

could be possible to only use data where none of the runs had recirculation in the analysis to remove this uncertainty, but given that the time window for ^{11}C stretches over three consecutive runs, this could limit the useable livetime. This systematic is expected to have the largest impact.

2) Uncertainty on the neutron reconstruction in x , y , and z . The AmBe source is deployed using guide tubes so the neutrons are emitted at known locations. This can be used to test how the SNO+ reconstruction algorithms perform on neutron captures, but this is not a straight forward comparison because the energy of the AmBe neutrons is different to the energy of the coincident neutrons in this analysis. This is expected to be a small systematic uncertainty.

3) Uncertainty on the tagging efficiency of the muons. This is a difficult systematic to calibrate, as SNO+ has no other muon identification available to compare, such as an external muon monitor on the deck of the experiment. There are also additional considerations regarding the MC for muons which affect its ability to be used to quantify this uncertainty. The cavity SNO+ sits in is carved from rock and painted in white waterproofing paint, as such it has bumps and rough patches. However, the MC models the cavity as a perfect cylinder and does not account for any reflections that may occur from the paint. Given the importance of the OWL PMT hits in the muon tagging, this means the MC is limited in its ability to assist with constraining this uncertainty. During the last phase of the SNO experiment, an External Muon System (EMuS) was installed above the detector which took data for 95 live days in order to calibrate their muon reconstruction algorithms [82]. Due to space constraints in the laboratory, the system was installed 12m above the detector, 3m west of centre, and at a 55 angle. After event level cuts, only 62 muons passing through EMuS were selected for further analysis of which, after correlating with the SNO event display, it was found only 32 passed within the PSUP. Due to the goal being to calibrate the reconstruction, there is not sufficient discussion of whether these events were identified as muons to use this to estimate this uncertainty, but it does show that temporary installation of an external muon system is a potential method by which SNO+ could evaluate this uncertainty. This systematic source is expected to have the smallest contribution to the error of the three outlined.

8.4 The results

For the sample at an ideal purity, any candidates with a likelihood of 0.4 or greater, 161 events were tagged across the dataset. Figure 8.8 shows the nhits of all ^{11}C candidates

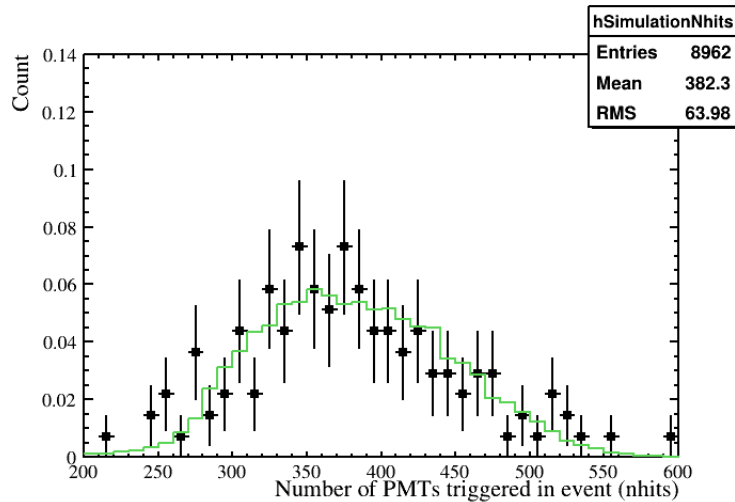


Figure 8.8: The nhits for all ^{11}C candidates with a probability of 0.4 or greater in black, in green is the MC simulation nhits distribution for ^{11}C . Both are normalised to area under the histogram to be one for comparison.

in the sample with a likelihood of 0.4 or greater in the black, in the green is the MC simulation distribution, both normalised for comparison. As discussed earlier this is a low statistics sample, but it has a better agreement with the MC than the higher statistics, less pure sample.

For the less pure sample, any candidates with a likelihood of 0.2 or greater, 397 events were tagged across the dataset. Figure 8.9 shows the nhits of all ^{11}C candidates in the sample with a likelihood of 0.2 or greater in the black, in the green is the MC simulation distribution, both normalised for comparison. Due to the less pure sample being selected for the purposes of statistics, it doesn't have as good agreement with the MC as the purer sample seen in Figure 8.8. This is seen mostly in both the high and low tails of the distribution, with the data showing a higher tail in both than the MC.

Figure 8.10 shows the distribution of the time between muon and ^{11}C candidate for all events with a likelihood of 0.2 or higher, whilst Figure 8.11 shows the same but for events with a likelihood of 0.4 or more. As can be seen, the 0.4 sample appears to have a better agreement with the fit, despite it being low statistics. This is to be expected as it is the purer sample, the 0.2 sample has more random backgrounds in, which is evidenced by the messier exponential.

As seen in this chapter, whilst a good background rejection was possible, the low statistics available at the time of writing this thesis made it difficult to obtain a pure sample of ^{11}C with high enough statistics. For this reason two cut values are presented in this thesis, a cut value of likelihood of 0.2 or higher which isn't as pure as desired for this




Figure 8.9: The nhits for all ^{11}C candidates with a probability of 0.2 or greater in black, in green is the MC simulation nhits distribution for ^{11}C . Both are normalised to area under the histogram to be one for comparison.




Figure 8.10: The timing distribution for the time between muon and ^{11}C candidate for events with a likelihood of 0.2 or higher.




Figure 8.11: The timing distribution for the time between muon and ^{11}C candidate for events with a likelihood of 0.4 or higher.

sample, but did yield reasonable statistics. The second cut was 0.4 likelihood or higher, creating a very pure sample, but of very low statistics, showing what would be possible with this analysis once more data is available. The purer sample has a stronger agreement with the MC distribution of nhits, with the higher statistics sample having more data in the edges of the distribution.

Chapter 9

Conclusion

Despite the deep underground location of SNO+, muons and induced backgrounds still pose a significant background to physics analyses. One such example is long-lived isotope ^{11}C , with a half-life of 20.3 minutes it poses a background to solar neutrino studies, and cannot be removed with the existing muon-follower veto.

SNO+ is now in its scintillator phase, but this thesis presents results from both the scintillator phase and the partial fill phase, a period of several months when the detector was filled with 365 tonnes of liquid scintillator and the rest water.

This thesis presents the current efforts in the three-fold coincidence method to identify and remove ^{11}C from analyses. This has involved changes to the muon identification to adjust the detection to scintillator instead of water. A newly developed neutron tagging process is also outlined in this thesis, which was developed on partial fill data and then re-tested on scintillator data.

This thesis then presents the development and current results of a ^{11}C likelihood algorithm, this is currently only presented for this isotope, but the framework could be adapted to other long lived muon induced isotopes if required. Using this likelihood algorithm, a background rejection cut was determined at a likelihood of 0.01, this rejects 94% of the ^{11}C whilst only sacrificing 1% of the signal.

Then, two higher purity cuts are presented, a likelihood of 0.4 or greater, which has a purity of 0.9, and a likelihood of 0.2 or greater, which has a lower purity of 0.6. The higher purity sample shows the potential of this analysis for a MC calibration, but is unfortunately statistically limited at the time of this thesis. The lower purity has enough statistics to make a more meaningful comparison to the expected results, but has a higher rate of random coincidences in the sample.

The systematic sources that would need to be considered to fully complete this study

if the statistics were available are also outlined and discussed.

Although the samples presented in this thesis are statistically limited, the framework can be used to complete a high statistics sample of ^{11}C that can be used as a data-MC calibration source. Additionally, the changes made by the author to the data cleaning framework to allow changes to the muon identification will allow the collaboration to retune all data cleaning cuts for the scintillator and Tellurium phases of the experiment.

This thesis also presents the work the author carried out on adding charge based checks to the CSS processor, a run-by-run software based PMT calibration. CSS runs on all data to flag channels that are behaving anomalously so they can be removed from the data as appropriate. This will continue to be used by SNO+ on a run-by-run basis throughout the scintillator and tellurium phases, and can be built on further to include tests on other PMT metrics such as timing.

Bibliography

- [1] David J Griffiths. *Introduction to elementary particles*. Wiley-VCH, 2008. xii, 3, 4, 5
- [2] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 2007. 3
- [3] Laurie M. Brown. The idea of the neutrino. *Physics Today*, 31(9):23–28, 1978. 4
- [4] J. Chadwick. Possible existence of a neutron. *Nature*, 129, 1932. 4
- [5] J. Chadwick. The existence of a neutron. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 136(830):692–708, 1932. 4
- [6] R. Brown, U. Camerini, P. H. Fowler, H. Muirhead, C. F. Powell, and D. M. Ritson. Observations with electron-sensitive plates exposed to cosmic radiation. *Nature*, 163(4132):47–51, 1949. 4
- [7] F. Reines and C. L. Cowan. Detection of the free neutrino. *Physical Review*, 92(3):830–831, 1953. 5
- [8] R. Davis and D. S. Harmer. Attempt to observe the $\text{cl}37(\text{nu e minus})\text{ar}37$ reaction induced by reactor antineutrinos. *Bull. Am. Phys. Soc.*, 4(217), 1959. 6
- [9] E. J. Konopinski and H. M. Mahmoud. The universal fermi interaction. *Physical Review*, 92(4):1045–1049, 1953. 6
- [10] G. Danby, J-M. Gaillard, K. Goulianos, L. M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9:36–44, Jul 1962. 7
- [11] K. Kodama, N. Ushida, C. Andreopoulos, N. Saoulidou, G. Tzanakos, P. Yager, B. Baller, D. Boehnlein, W. Freeman, B. Lundberg, and et al. Observation of tau neutrino interactions. *Physics Letters B*, 504(3):218–224, Apr 2001. 7

- [12] Giunti C. and Kim C. W. *Fundamentals of neutrino Physics and Astrophysics*. Oxford University Press, 2007. 8
- [13] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Physical Review*, 105:1413–1415, 1978. 10
- [14] Andrzej Wroblewski. The downfall of parity-the revolution that happened fifty years ago*. *Acta Physica Polonica B*, 39, 10 2008. xii, 10
- [15] John N. Bahcall. Neutrinos from the sun. *Scientific American*, 221(1):28–37, 1969. 10, 11
- [16] J. N. Bahcall. Solar neutrino cross sections and nuclear beta decay. *Physical Review*, 135(1B):B137–B146, 1964. 11
- [17] Aldo Ianni. Solar neutrinos and the solar model. *Physics of the Dark Universe*, 4:44–49, 2014. DARK TAUP2013. xii, 11
- [18] J. N. Bahcall and C. A. Barnes. Importance of a study of the decay of ca-37 to the problem of detecting solar neutrinos. *Physical Review Letters*, 12(1):48–49, 1964. 11
- [19] J. P. Cravens, K. Abe, T. Iida, K. Ishihara, J. Kameda, Y. Koshio, A. Minamino, C. Mitsuda, M. Miura, S. Moriyama, and et al. Solar neutrino measurements in super-kamiokande-ii. *Physical Review D*, 78(3), Aug 2008. 12, 16, 18
- [20] D. Casper and et. al. Measurement of atmospheric neutrino composition with the imb-3 detector. *Physical Review Letters*, 66(20):2561–2564, 1991. 12
- [21] Y. Fukuda, T. Hayakawa, E. Ichihara, K. Inoue, K. Ishihara, H. Ishino, Y. Itow, T. Kajita, J. Kameda, S. Kasuga, and et al. Measurement of a small atmospheric ν_μ/ν_e ratio. *Physics Letters B*, 433(1-2):9–18, Aug 1998. 12
- [22] Nicola Cabibbo. Unitary symmetry and leptonic decays. *Physical Review Letters*, 10(12):531–533, 1963. 12
- [23] R. R. Volkas R. Foot, H. Lew. Quark-lepton symmetry, 1992. 12
- [24] Makoto Kobayashi and Toshihide Maskawa. Cp-violation in the renormalizable theory of weak interaction. *Progress of Theoretical Physics*, 49(2):652–657, 1973. 12
- [25] B. Pontecorvo. Mesonium and antimesonium (translated). *Journal of Experimental and Theoretical Physics*, 6(2):429, 1957. 13

- [26] B. Pontecorvo. Neutrino experiments and the problem of conservation of leptonic charge. *Journal of Experimental and Theoretical Physics*, 26(5):984, 1967. 13
- [27] S. Bilenky. Neutrino oscillations: From a historical perspective to the present status. *Nuclear Physics B*, 908:2–13, 2016. 15
- [28] K. Abe et al. The t2k experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 659(1):106–135, dec 2011. 15
- [29] Francesca Di Lodovico and on behalf of the Hyper-Kamiokande Collaboration. The hyper-kamiokande experiment. *Journal of Physics: Conference Series*, 888, 2017. 15
- [30] B. Abi et al. Deep underground neutrino experiment (dune), far detector technical design report, volume i: Introduction to dune, 2020. 15
- [31] A Yu Smirnov. The msw effect and matter effects in neutrino oscillations. *Physica Scripta*, T121:57–64, 2005. 15
- [32] B. Aharmim, S. N. Ahmed, A. E. Anthony, N. Barros, E. W. Beier, A. Bellerive, B. Beltran, M. Bergevin, S. D. Biller, K. Boudjemline, and et al. Combined analysis of all three phases of solar neutrino data from the sudbury neutrino observatory. *Physical Review C*, 88(2), Aug 2013. 16
- [33] Y. Fukuda, T. Hayakawa, E. Ichihara, K. Inoue, K. Ishihara, H. Ishino, Y. Itow, T. Kajita, J. Kameda, S. Kasuga, and et al. Evidence for oscillation of atmospheric neutrinos. *Physical Review Letters*, 81(8):1562–1567, Aug 1998. 16
- [34] et al. (Particle Data Group) M. Tanabashi. Neutrino masses, mixing, and oscillations. *Physical Review D*, 98, 2018. xi, 17
- [35] Julien Lesgourgues and Sergio Pastor. Neutrino mass from cosmology. *Advances in High Energy Physics*, 2012:1–34, 2012. 18
- [36] Susanne Mertens. Status of the katrin experiment and prospects to search for keV-mass sterile neutrinos in tritium beta-decay. *Physics Procedia*, 61:267–273, 2015. 18
- [37] R. N. Cahn, D. A. Dwyer, S. J. Freedman, and et al. White paper: Measuring the neutrino mass hierarchy. *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*, 2013. xii, 18

- [38] Sanjib Kumar Agarwalla, Tracey Li, Olga Mena, and Sergio Palomares-Ruiz. Exploring the earth matter effect with atmospheric neutrinos in ice, 2012. 18
- [39] E.Kh. Akhmedov, G.C. Branco, and M.N. Rebelo. Seesaw mechanism and structure of neutrino mass matrix. *Physics Letters B*, 478(1-3):215–223, 2000. 20
- [40] M. Goeppert-Mayer. Double beta-disintegration. *Physical Review*, 48(6):512–516, 1935. 21
- [41] C. Patrignani and et al. Review of particle physics. *Chinese Physics C*, 40(10), 2016. 21
- [42] Frank T. Avignone, Steven R. Elliott, and Jonathan Engel. Double beta decay, majorana neutrinos, and neutrino mass. *Reviews of Modern Physics*, 80(2):481–516, Apr 2008. xii, 22
- [43] L. Cardani. Neutrinoless Double Beta Decay Overview. *SciPost Phys. Proc.*, page 24, 2019. xi, 21, 23
- [44] R. et al. Arnold. Probing new physics models of neutrinoless double beta decay with supernemo. *The European Physical Journal C*, 70(4):927–943, 2010. 21, 23
- [45] S. Andringa, E. Arushanova, S. Asahi, M. Askins, D. J. Auty, A. R. Back, Z. Barnard, N. Barros, E. W. Beier, A. Bialek, and et al. Current status and future prospects of the sno+ experiment. *Advances in High Energy Physics*, 2016:1–21, 2016. 23, 24, 34, 41
- [46] J. Martín-Albo, J. Muñoz Vidal, P. Ferrario, M. Nebot-Guinot, J. J. Gómez-Cadenas, V. Álvarez, C. D. R. Azevedo, F. I. G. Borges, S. Cárcel, and et al. Sensitivity of next-100 to neutrinoless double beta decay. *Journal of High Energy Physics*, 2016(5), May 2016. 23
- [47] J.B. Albert, G. Anton, I. Badhrees, P.S. Barbeau, R. Bayerlein, D. Beck, V. Belov, M. Breidenbach, T. Brunner, G.F. Cao, and et al. Search for neutrinoless double-beta decay with the upgraded exo-200 detector. *Physical Review Letters*, 120(7), Feb 2018. 23
- [48] A. Gando, Y. Gando, H. Hanakago, H. Ikeda, K. Inoue, K. Ishidoshiro, R. Kato, M. Koga, S. Matsuda, T. Mitsui, and et al. Limit on neutrinoless $\beta\beta$ decay of xe-136 from the first phase of kamland-zen and comparison with the positive claim in ge-76. *Physical Review Letters*, 110(6), Feb 2013. 23

- [49] K. Singh. Detector picture. *Internal SNO+ document: SNO+-doc-4331-v1*, 2017. 25
- [50] V. Albanese, R. Alves, M.R. Anderson, S. Andringa, L. Anselmo, E. Arushanova, S. Asahi, M. Askins, D.J. Auty, A.R. Back, and et al. The sno+ experiment. *Journal of Instrumentation*, 16(08):P08059, Aug 2021. 24, 26, 28, 33, 59
- [51] A. et. al. Mastbaum. Sno+ detector manual. *Internal SNO+ document*, 2017. xii, 26
- [52] James Waterfield. *Optical Calibration System for SNO+ and Sensitivity to Neutrinoless Double-Beta Decay*. PhD thesis, University of Sussex, September 2016. xii, 27
- [53] M. Anderson, S. Andringa, S. Asahi, M. Askins, D.J. Auty, N. Barros, D. Bartlett, F. Barão, R. Bayes, E.W. Beier, and et al. Measurement of the b8 solar neutrino flux in sno+ with very low backgrounds. *Physical Review D*, 99(1), Jan 2019. 31
- [54] M. Anderson, S. Andringa, E. Arushanova, S. Asahi, M. Askins, D.J. Auty, A.R. Back, Z. Barnard, N. Barros, D. Bartlett, and et al. Search for invisible modes of nucleon decay in water with the sno+ detector. *Physical Review D*, 99(3), Feb 2019. 32
- [55] Yan LIU. Ambe source calibration in the sno+ water phase. 2018. 32
- [56] E. Falk, J. Lidgard, M. I. Stringer, and E. Turner. Commissioning of ellie for sno+, 2017. 33
- [57] Esther TURNER. Smellie: A laser calibration system for sno+, June 2018. 33
- [58] Erica Caden. Status of the sno+ experiment, 2017. 33, 40, 41
- [59] G. Bellini, K. Inoue, F. Mantovani, A. Serafini, V. Strati, and H. Watanabe. Geoneutrinos and geoscience: an intriguing joint-venture. *La Rivista del Nuovo Cimento*, 45(1):1–105, 2022. 33
- [60] G. Fiorentini, G. L. Fogli, E. Lisi, F. Mantovani, and A. M. Rotunno. Mantle geoneutrinos in kamland and borexino. *Physical Review D*, 86(3), Aug 2012. 34
- [61] I. Semeneć. Antineutrino energy spectrum plot (geo-neutrino focus). *Internal SNO+ document: SNO+-doc-5527-v1*, 2019. xii, xiv, 34, 63
- [62] R. M. Bionta and et. al. Observation of a neutrino burst in coincidence with supernova 1987a in the large magellanic cloud. *Physical Review Letters*, 58(14):1494–1496, 1987. 35

- [63] S Al Kharusi et al. SNEWS 2.0: a next-generation supernova early warning system for multi-messenger astronomy. *New Journal of Physics*, 23(3):031201, mar 2021. 35
- [64] G. Settanta et al. First detection of cno neutrinos with borexino, 2021. 35
- [65] G. Bellini et al. (Borexino Collaboration). Cosmogenic backgrounds in borexino at 3800 m water-equivalent depth. *Journal of Cosmology and Astroparticle Physics*, 2013. xi, xii, 36, 62, 64, 75
- [66] A. Gando, Y. Gando, T. Hachiya, A. Hayashi, S. Hayashida, H. Ikeda, K. Inoue, K. Ishidoshiro, Y. Karino, M. Koga, and et al. Search for majorana neutrinos near the inverted mass hierarchy region with kamland-zen. *Physical Review Letters*, 117(8), Aug 2016. xiii, 41
- [67] Billy Liggins. *Cosmogenic muon induced neutrons in SNO+ water phase*. PhD thesis, Queen Mary University of London, September 2019. 46, 48, 72, 75
- [68] H.M. O’Keeffe and M.C. Chen. Expected radioactive backgrounds in the sno+ experiment. *Internal SNO+ document: SNO+-doc-0507*, 2010. 54, 61
- [69] V. Lozza and J. Petzoldt. Cosmogenic activation of a natural tellurium target. *Astroparticle Physics*, 61:62–71, Feb 2015. 54
- [70] J.J. Beatty and J. Matthews. Cosmic rays. *Particle Data Group Reviews*, 2011. xiv, 55
- [71] LUIS ANCHORDOQUI, THOMAS PAUL, STEPHEN REUCROFT, and JOHN SWAIN. Ultrahigh energy cosmic rays: The state of the art before the auger observatory. *International Journal of Modern Physics A*, 18(13):2229–2366, May 2003. xiv, 56
- [72] Petr Trávníček. *Detection of high-energy muons in cosmic ray showers*. PhD thesis, Charles University, July 2004. xiv, 56
- [73] The National Academies Press. *Neutrinos and Beyond: New Windows on Nature*. National Research Council., 2003. xiv, 58
- [74] B. Aharmim, S.N. Ahmed, A.E. Anthony, N. Barros, E.W. Beier, A. Bellerive, B. Beltran, M. Bergevin, S.D. Biller, R. Bonventre, and et al. Cosmogenic neutron production at the sudbury neutrino observatory. *Physical Review D*, 100(11), Dec 2019. xiv, 57, 59

- [75] A. Bialek, M. Chen, and R. Ford. Sno+ scintillator qa plans and specifications. *Internal SNO+ document: SNO+-doc-4899-v1*, 2018. 60
- [76] Y. Zhang et al. (Super-Kamiokande Collaboration). First measurement of radioactive isotope production through cosmic-ray muon spallation in super-kamiokande iv. *Phys. Rev. D*, 2016. xi, 62
- [77] Cristiano Galbiati, Andrea Pocar, Davide Franco, Aldo Ianni, Laura Cadonati, and Stefan Schönert. Cosmogenic c11 production and sensitivity of organic scintillator detectors to pep and cno neutrinos. *Physical Review C*, 71(5), 2005. 62
- [78] G Bellini, J Benziger, D Bick, G Bonfini, D Bravo, M. Buizza Avanzini, B Caccianiga, L Cadonati, F Calaprice, P Cavalcante, and et al. Cosmogenic backgrounds in borexino at 3800m water-equivalent depth. *Journal of Cosmology and Astroparticle Physics*, 2013(08):049–049, Aug 2013. 62, 63
- [79] S. Abe, S. Enomoto, K. Furuno, Y. Gando, H. Ikeda, K. Inoue, Y. Kibe, Y. Kishimoto, M. Koga, Y. Minekawa, and et al. Production of radioactive isotopes through cosmic muon spallation in kamland. *Physical Review C*, 81(2), Feb 2010. 64
- [80] (Borexino Collaboration). Cno and pep neutrino spectroscopy in borexino: Measurement of the deep-underground production of cosmogenic c11 in an organic liquid scintillator. *Physical Review C*, 74(4), 2006. 64
- [81] [SNO Collaboration]. Measurement of the cosmic ray and neutrino-induced muon flux at the sudbury neutrino observatory. *Physical Review D*, 80(1), jul 2009. 72
- [82] et. al T.J. Sonley. Calibration of muon reconstruction algorithms using an external muon tracking system at the sudbury neutrino observatory. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 648(1):92–99, August 2011. 91

Appendix A

Appendix A

- **Ring of Fire Cut** - Noise on the FECs can manifest as a "ring" of bad channels on the outside of a crate. This cut checks if most of the hits in an event are from the outside ring of channels on one crate, if they are it flags the event.
- **Crate Isotropy Cut** - Noise on the FECs can also manifest as an uneven distribution of hits, by causing noise on adjacent channels within the crate. This cut flags an event if most hits come from one crate, and 2 adjacent daughter boards on that crate.
- **Empty Crate Cut** - This cut flags high nhit events that have a crate with no hits at all, as this can indicate a breakdown.
- **Flasher Geometry Cut** - This cut aims to identify flasher events by looking for clusters of hits either on channels near each other or channels on the same FEC. If at least one event in this cluster has a bad charge reading or the average time of hits in the cluster varies enough from the rest of the detector, the event is considered a flasher.
- **In Time Channel Time Spread Cut** - This cut checks whether an event has most of its hits in one fixed time window, as Cherenkov events usually do. This cut is unlikely to be used in scintillator phase.
- **OWL Cut** - This cut tags events with a high rate of activity in the outer tubes.
- **Fitterless Time Spread Cut** - This cut looks for flashers using just the time spread of the events in the rest of the detector compared to the cluster of events, in cases where the charge readout is not anomalously high. This cut is unlikely to work in scintillator phase.

- **QvNhit Cut** - Events that are triggering on noise will typically have a much lower charge than events triggering on real Physics. This cut compares the charge of the event to the nhits of the event to try and identify noise events.
- **Junk Cut** - This cut flags any event where a channel has registered more than one hit, as this should not be possible.
- **QCluster Cut** - This cut looks for clusters of hits surrounding a high charge event which can be indicative of a flasher or a breakdown.
- **QvT Cut** - This cut aims to identify flashers with a lower charge than typically expected for a flasher, by looking for events with an above average charge that occurs much earlier than the mean hit time.
- **Neck Cut** - This cut looks for 'neck events' where multiple neck tubes are hit, or one neck tube with bad charge is hit significantly before the events in the bottom half of the detector.
- **Threshold Flasher Cut** - This cut looks for flashers where the nearby channels are at maximum threshold and hence would not see cross-talk events. This cut looks for one high charge event surrounded by channels at maximum threshold.
- **Polling Cut** - The PMTs are regularly polled throughout normal running to identify channels with poor readouts, this cut flags events that happen during the polling.
- **CAEN Missing Cut** - Sometimes the CAEN buffer is filled up before the event can be read out, this cut flags those events.
- **Pedestal Cut** - This cut flags events that occur during automatic checking of the pedestals that occurs regularly throughout usual running.
- **Missed Count Burst Hit** - The FEC has a limited buffer, which can be filled up during a burst of events. If a hit arrives and the FEC is full, it sets a 'missed count bit' on the channel. This cut searches for bursts of events with counts missed for many channels, and flags events for 1 second after the burst begins or until the burst ends.