


**Cosmic muon-induced
neutrons in the SNO+ water
phase**



Billy Liggins

Department of Physics
Queen Mary University of London

This dissertation is submitted for the degree of
Doctor of Philosophy

April 2020

For my Grandad.

Declaration

I, Billy Liggins, 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.

Details of collaboration and publications:

This work was carried out by the author in collaboration with the SNO+ collaboration.

Signature:

Date:

Billy Liggins

April 2020

Acknowledgements

Firstly I would like to thank my supervisor Jeanne Wilson, without her guidance and direction this thesis would have never come to fruition. I have grown immeasurably during my PhD, thank you for taking a gamble on me, I can only hope that my contributions pay dividends. I would also like to thank Adrian Bevan for taking over supervisor duties in the final stages of my PhD.

Thank you to my examiners Peter Simon and Seth Zens for providing swift assessment and administration.

During my time at Queen Mary I have had the pleasure of being part of a lovely group of people. I thank everyone that I have come into contact with at Queen Mary, without discussions with Antony Fray, Shivesh Mandalia and Rodrodrigo Gamboa Goni deadlines would have been that little bit harder to meet and without Tom Charman, Jesal Mandalia, Nick Prouse, Tom Stevenson, Lewis Millward, Jake Oliver and Eddie Thorpe lunches would have been considerably shorter.

I would like to highlight the SNO+ group at Queen Mary who have all had some input on this thesis. Ashley Back, Stefanie Langrock, Lorna Nolan and Evelina Arushanova shared their knowledge with me through our time together, Evelina deserves a special thank you for passing over her work on CSS. Matt Mottram, Rebecca Lane and Mark Stringer were all superb post docs, helping me immensely at the drop of a hat. A special

thank you to Mark for his work in the final stages this work. Jeanne also deserves another mention here for her leadership of the group. I really enjoyed working with everyone.

My thanks to the whole of the SNO+ collaboration, living and working in Sudbury was a fantastic experience. Gaining knowledge from everyone and living life with some was a true pleasure. Some special thanks to Ed Leming, Jack Dunger, Esther Turner for being great colleagues and becoming even better friends. Thank you to Javi Caravaca for heading the PMT CAL group while I worked on CSS. Martti Nirkko is responsible for Fig. 1 showing how this thesis was compiled over time, on the hard days I found this plot very motivating, thank you.

Finally I would like to thank my family for their support throughout my education. To my Grandad for sparking my interest in physics and to my Mum for supporting me every step of the way.

Above all others I'll always be thankful to my Wife, Josie, who may not understand this thesis, but her understanding of me has been the most significant part of its compilation. During my PhD we have brought a bungalow, got married and had our first child Emilia. Maybe one day Emilia may read this, if she does she should know her Mum made this thesis possible.




Fig. 1: Word count, number of pages and file size of this thesis as a function of time, courtesy of a macro shared by Martti Nirkko.

Abstract

The SNO+ experiment is the successor to the Nobel prize winning SNO experiment. SNO+ will ultimately search for neutrinoless double beta decay in ^{130}Te . 1.3 tonnes of ^{130}Te will be dissolved into 780 tonnes of organic liquid scintillator (LAB). 9300 photo-multiplier tubes (PMTs) will observe the loaded scintillator held in an acrylic vessel of 6 m radius. This final stage of operation is preceded by a number of preparation periods, during which other physics may be probed. This thesis reports the status of a physics measurement in the first phase of operation with a detector filled with ultra-pure water. The work aims to measure the production of cosmic muon induced neutrons that spallate in water and capture on hydrogen releasing a 2.2 MeV gamma ray. This process presents as a large background in many weak signal searches, such as dark matter searches, knowledge of the induced neutron production rate will inform both theoretical and experimental considerations. All components of the analysis developed to date are detailed. A low level PMT anomaly detection system is motivated and implemented, it is shown to remove high noise and inactive PMTs from data, applicable to general SNO+ analyses but especially important for low energy events. Muon event reconstruction is implemented, resulting in a track length error of $\leq 5\%$ for tracks lengths > 10 m. A muon selection focused on a purity is detailed, producing a sample of muon candidates consistent with the rate observed

in SNO. A neutron candidate selection with a purity of $\approx 99\%$ is also defined. After applying both selections to a custom data processing, the neutron capture time is measured to be $134 \pm 16 \mu\text{s}$ which is inconsistent with the expected capture time of $206.03 \pm 0.44 \mu\text{s}$ as measured by an $^{241}\text{Am}^9\text{Be}$ calibration source in SNO+. This inconsistency is shown to be a result of electronic instability at short times after a cosmic muon event. This electronic instability is a previously unknown effect, leaving the analysis incomplete and requiring future work.

Table of contents

1	Introduction	1
2	Neutrino theory and cosmic-muon induced neutrons in SNO+	5
2.1	Neutrinos in the standard model	5
2.1.1	Brief history of the standard model neutrino	5
2.1.2	The Standard Model neutrino	7
2.1.3	New physics	9
2.2	Neutrinos beyond the Standard Model	9
2.2.1	The solar neutrino problem	9
2.2.2	The Atmospheric ν Anomaly	11
2.2.3	Flavour mixing	12
2.2.4	Summary of oscillation measurements	16
2.3	Neutrino mass	18
2.3.1	Mass generation	18
2.3.2	Mass measurements	21
2.4	Neutrinoless double-beta decay	22
2.4.1	$2\nu\beta\beta$ decay	23
2.4.2	$0\nu\beta\beta$ decay	24
2.5	Cosmic rays showers	28
2.5.1	Cosmic ray energy spectrum	30

2.5.2	Cosmic muon production	30
2.6	Muon Spallation	34
2.7	Current neutron yield status	36
3	The SNO+ Detector	38
3.1	Introduction	38
3.2	Detector overview	38
3.3	Physics goals	42
3.3.1	$0\nu\beta\beta$ search	42
3.3.2	Light water phase physics	44
3.3.3	Partial water phase	45
3.3.4	Pure scintillator phase	45
3.4	Detecting Principle	46
3.4.1	Cherenkov radiation	46
3.4.2	Scintillation	47
3.5	Data acquisition system	50
3.5.1	PMT configuration	51
3.5.2	Trigger system	52
3.5.3	Nearline processing	54
3.5.4	Grid processing	55
3.6	Event reconstruction	55
3.6.1	Position and time fit	56
3.6.2	Direction and Energy Fit	57
3.6.3	Drive correction	58
3.7	Detector setup across the water phase	59
3.8	Detector calibration in water	60
3.8.1	Electronic calibration	61
3.8.2	Optical calibration	62

3.9	Backgrounds	68
3.9.1	Radioactive backgrounds	69
3.9.2	Instrumental backgrounds	70
3.10	Software	71
3.10.1	RATDB	71
3.11	Summary	72
4	Channel Software Status	73
4.1	Introduction to CSS	74
4.2	Standard runs	75
4.3	Cut definitions	76
4.3.1	Population based	78
4.3.2	Channel based	88
4.4	Implementation flow	92
4.4.1	Reprocessing	95
4.5	Effect of CSS on data	97
4.5.1	High occupancy	97
4.5.2	Low occupancy	99
4.6	CSS over all water data	100
4.7	Conclusion	104
5	Data Processing	105
5.1	Raw data set	105
5.2	Data processing	106
5.3	Raw processing results	107
6	Simulation and Reconstruction of Cosmic Muons	110
6.1	Cosmic Muon Simulation	110
6.2	Through going muon reconstruction	114

6.3	Calibrating the fitter	118
6.4	Fitter performance	121
6.4.1	Path length error	122
6.5	Potential improvements	124
6.6	Conclusion	124
7	Cosmic Muons Selection	125
7.1	Introduction	125
7.2	127
7.3	Potential physical muon backgrounds	129
7.4	Muon Event Selection	129
7.4.1	Trigger Selection	130
7.4.2	Data Cleaning	130
7.4.3	Low Level Muon Cuts	132
7.4.4	Reconstruction Level Cuts	134
7.5	Selected muon data set	135
7.5.1	Muon selection breakdown	135
7.5.2	MC Comparison	136
7.5.3	Live Time Of Muon Data	138
7.6	Muon data set comparison to SNO results	139
7.6.1	Observed muon rate	139
7.7	Potential Improvements	144
7.8	Conclusion	144
8	Neutron Selection	147
8.1	Neutron selection	147
8.1.1	MC generation	148
8.1.2	Fake data distributions	153
8.2	Fiducial Volume selection	156

8.3	$\Delta_\mu t$ & $\Delta_\mu R$ selection optimisation	157
8.3.1	Neutron selection performance	159
8.4	Long lived muon spallation products	162
9	Neutron Detection Efficiency	165
9.1	Outline	165
9.2	Efficiency modelling	166
9.2.1	Efficiency performance	168
9.3	Neutron MC AmBe verification	169
10	Analysis Strategy and Instrumental Effects Following Cos-	
	mic Muons	170
10.1	Outline	170
10.1.1	Muon track length systematics	172
10.2	Time spectrum side-band check	172
10.3	Investigation with TELLIE	173
11	Conclusion	177
	References	180

Chapter 1

Introduction

A major component of a particle physics experiment is the detailed understanding of background events that could be confused with a desired signal. In the following thesis, an analysis looking to measure the production rate of neutrons from cosmic muon spallation in the SNO+ detector is detailed.

Cosmic muon spallation products present as a troublesome background for many weak signal searches as they are long lived and therefore are hard to eliminate with a simple veto after a muon without having a substantial detector deadtime. Neutron spallations is a special case of muon daughter production with a cross section two orders of magnitude higher than the other long lived spallation daughters. The expected number of neutrons produced for a given muon track length through a detector is given by the neutron production yield. The neutron production yield is also normalised by detector material density, which should provide a measurement independent of detecting material. Although for similar atomic masses this should be true, second order effects may be expected, therefore it is important to measure the yield for each material.

SNO+ is primarily concerned with observing neutrinoless double beta decay in ^{130}Te dissolved in organic liquid scintillator. However, in achieving this goal, SNO+ will operate in many phases. The work presented in the following thesis is completely based in the first phase of the experiment in which the detector is filled with ultra pure water. Being uniquely placed below 2 km of overburden, SNO+ has the access to a low rate high energy cosmic muon flux for which neutron production yield is under studied.

As through going cosmic muons are unique events in the SNO+ detector, the analysis follows the form of a coincidence analysis with the muon being the prompt event and searching for signal events in a following time window. Being a coincidence analysis, backgrounds can be measured from real data which are not associated with a prompt muon. This drastically reduces the input from Monte Carlo (MC) simulations in evaluating contributions from backgrounds.

This thesis motivates and develops a measurement of the cosmic muon induced neutron yield in SNO+ in its first phase of operation, that is spallation on oxygen nuclei.

In Chapter 2, the theory relevant to both the physics analysis and the physics case for the SNO+ detector is laid out. Neutrino physics is reviewed with attention paid to the outstanding questions such as the nature of neutrinos as Dirac or Majorana particles and the measurements of the PMNS parameter δ_{CP} and the sign of Δm_{31} . The chapter ends by reviewing the theory associated with cosmic muon generation and muon spallation before outlining the current state of the cosmic muon neutron yield measurements.

Chapter 3 introduces the SNO+ detector, its physics goals, outlines of operational modes, detector data acquisition system and calibration, event reconstruction and computing infrastructure.

As part of the membership of a collaboration such as SNO+, there is a requirement to complete a number of service tasks. These tasks include detector operating shifts for which there is a yearly quota and analysis tasks that contribute toward detector performance. Chapter 4 details the authors service task that built upon the work of Evelina Arushanova and Nasim Fatemighomi. The work focused on low-level channel anomaly detection over the PMTs used in SNO+. The ultimate effect of this work is to remove hits from unstable PMTs from the processing chain, which have a pronounced effect on physics events with a low number of registered hits such as neutron captures. This system is now operational and effectively removes unstable channels across the water phase. The system is applied in the subsequent analysis.

In Chapter 5, the custom data processing required for the neutron yield analysis is outlined, with special consideration given to reconstruction of events coincident with and without a prompt muon.

Following work carried out on the SNO experiment, through going cosmic muon reconstruction is motivated, implemented and calibrated in Chapter 6. This step is essential in the event selection as a perpendicular event distance from the muon track was assumed to be a powerful method of background rejection. The fitter is shown to have a 5% error on tracks above 10 m in length.

Chapter 6 develops a muon selection drawing on the experience of the SNO detector and the previously defined muon reconstruction. The selection results in a through going muon rate consistent with that measured in SNO.

Chapter 8 details a neutron selection with a signal-to-background optimised selection in time since muon and perpendicular distance to the muon track. The default water reconstruction fit validity is shown to lead to a heavy sacrifice on the MC neutron signal. The selection is optimised and evaluated on muon MC doped with real data random backgrounds. backgrounds.

In Chapter 9, a method to evaluate the neutron detection efficiency for each individual muon through MC studies is provided. Much of the work completed in this chapter was completed by Dr Mark Stringer, however it is detailed as it is of paramount importance to the analysis. The author was partially responsible for developing the method but shows results from simulations completed by Dr Mark Stringer.

Chapter 10 motivates a likelihood method for extracting a final yield measurement and tackles the leading order systematics before showing the timing structure of events following a muon. At this point it becomes clear that the understanding of the detector DAQ and electronics on $\mathcal{O}(\text{ms})$ is not adequate. Details of the utilisation of a calibration sub system, TELLIE, to study the effects of a large amounts of light injected into the detector on the electronics is given. This study tentatively suggests that the electronic enter an unstable period which has the effect of changing the neutron selection efficiency as a function of time. This study shows a possible path to completion via incorporating this effect into the simulated efficiency. However, this effect was found at the eleventh hour and is beyond the time scope of this PhD thesis.

Chapter 11 concludes the thesis with a summary of the work carried out and it's implications for future work.

Chapter 2

Neutrino theory and cosmic-muon induced neutrons in SNO+

2.1 Neutrinos in the standard model

2.1.1 Brief history of the standard model neutrino

The neutrino was first proposed by Pauli in 1930 to save energy conservation in nuclear β decays. Chadwick [1] had shown in 1914 that the assumed two-body β decays had a continuous energy distribution, although the relatively low mass β particle and the heavy recoiled decaying nucleus under a two body decay would have exact energies. Neutral γ rays were ruled out experimentally by Meitner [2], which led to ideas of the missing energy being accounted for by a new particle, with Bohr famously suggesting that the energy is only conserved in a “statistical sense” [3].


Pauli announced his solution to the β decay problem in a letter beginning ‘Dear Radioactive Ladies and Gentlemen’. In the letter he proposed the missing energy was carried away by a neutral weakly interacting particle created as a decay product [4]. He predicted that the mass would be ‘the same order as the electron mass’ and gave it a spin of $\frac{1}{2}$ to account for the angular momentum conservation. Pauli named the particle the neutron but the particle was renamed the neutrino at a later date.

In 1934, Fermi included the neutrino in his theory of the weak interaction. The subsequent success in describing the β spectrum gave the neutrino a grounding in theory.

Bethe and Peierls estimated the cross section for inverse beta decay to be $< 10^{-44} \text{ cm}^2$ at 1 MeV leading to major experimental issues given the accelerator technologies at the time [2]. However, in 1956, Cowan and Reines detected $\bar{\nu}_e$ through inverse beta decay using the intense $\bar{\nu}$ source of nuclear reactors [5]. This work led to Reines being awarded the 1995 Nobel prize in physics, as well as confirming the existence of the neutrinos. In 1962, Lederman, Schwartz and Steinberger detected a neutrino associated with muon interactions, the muon neutrino ν_μ [6]. The tau neutrino, ν_τ , was detected by the DONUT collaboration in 2000 [7].

In 1956, Garwin, Lederman and Weinrich [8] found parity violation in μ decay and Wu [9] showed that β decay also violated parity. Wu’s group observed β decay in cooled ^{60}Co atoms with their spins aligned via an applied magnetic field. The β particle was shown to be preferentially emitted anti-parallel to the parent’s spin. The spin is invariant under parity transform but the β momentum isn’t. Therefore, parity was violated and spin conservation also implied that the weak interaction preferentially produced left hand neutrinos.

Further experimental evidence came from the Goldhaber experiment [10]. It measured the helicity of the neutrino produced in electron captures on ^{152}Eu :



The group found that the polarization of γ s emitted by the relaxed Sn nucleus was always aligned such that the ν_e spin was always anti-parallel to its momentum, to within experimental error. This showed that the ν helicity was the same in all inertial frames and therefore concluded that the ν_e was massless.

This result informed the development of the Standard Model with all neutrinos taken as left handed particles and their antiparticle counterparts taken as right handed.

2.1.2 The Standard Model neutrino

The Standard Model is a theory of sub atomic particles and there interactions developed in the latter half of the twentieth century. The Standard Model describes 17 fundamental particles interacting through three fundamental forces¹. Fig. 2.1 is a diagram grouping the particles into a table with subgroups aligned along both columns and rows. Standard Model particles are either fermions or bosons depending on whether the particles spin is half-integer or integer respectively. The fermions are further grouped into two sectors; Quarks and Leptons. Quarks are the only grouping governed by Quantum Chromodynamics (QCD) which is non-abelian $SU(3)$ gauge group. Both Quarks and Leptons are arranged

¹The fourth fundamental force of gravitation is not yet included in the Standard Model.




Fig. 2.1: Diagram grouping the 17 particles of the Standard Model in various subgroups, adapted from [11].

into three generations grouping, within each sector the generations form doublets displaying the $SU(2)$ symmetry of the Standard Model.

Bosons are force carriers in the Standard Model, gluons propagate the QCD force, photons propagate the quantum electrodynamics (QED) with all particles with a charge interacting under QED and the Z and W bosons propagating the weak force. The Standard model is a chiral gauge theory which requires spontaneous symmetry breaking by the Higgs mechanism to generated mass terms for the Z and W bosons. The Higgs mechanism results in an extra degree of freedom in the model presenting as a particle known as the Higgs particle. The Higgs then couples via Yukawa couplings to fermions producing their mass, these terms are sometimes referred to as Dirac mass terms.

The neutrino

The neutrino is a fundamental particle in the Standard Model. There are three types of neutrino that form doublets with their associated leptons to form the leptonic sector of the Standard Model. All neutrinos have spin $1/2$ and interact via the weak force. In the Standard Model the neutrino violates parity maximally and as such does not form Dirac mass terms easily and therefore the neutrino states are massless, see section 2.3.1 for details.

2.1.3 New physics

The Standard Model description of the neutrino was unquestioned until experiments at the turn of the millennium showed that neutrinos must have mass. This was first evidence of physics beyond the Standard Model.

2.2 Neutrinos beyond the Standard Model

The discovery that neutrinos have mass was made by solving two experimental anomalies. Namely measured deficits in the expected neutrino production from the Sun and in the Earth's atmosphere. In 2015, the Nobel prize was awarded to Kajita and McDonald for their leadership of two experiments instrumental in the discovery, Super-Kamiokande and SNO respectively.

2.2.1 The solar neutrino problem

In 1964, Davis and Bahcall, building on work by Pontecorvo, proposed an experiment to measure ν_e produced in the Sun.




Fig. 2.2: Solar neutrino energy spectrum as calculated by Bahcall and Serenelli [12].

The Standard Solar Model (SSM) developed between the 1920's and 1960's predicted fluxes of neutrinos from hydrogen burning in the centre of the Sun. The hydrogen burning chain is made from many intermediate steps that produce neutrinos of differing energies and fluxes. These fluxes are referred to by the most massive parent particle in the reaction, e.g. ^8B neutrinos are produced via:



Fig 2.2 shows the spectrum of solar neutrinos and their production chains. The proposed experiment aimed to detect neutrino captures on chlorine:



This reaction has a threshold of 0.8 MeV making it sensitive to ^8B , pep and CNO neutrinos. The experiment took place in the Homestake

mine in South Dakota and measured the rate of neutrino reactions by chemically separating out the ^{37}Ar . Initial results confirmed the SSM prediction for the temperature at the centre of the Sun as well as ruling out the CNO cycle as a dominant fusion mode. However, the flux of neutrinos was found to be a factor of two smaller than the SSM prediction. This discrepancy is known as the Solar Neutrino Problem [13–15].

This result was independently confirmed by the Kamiokande experiment and again later by the Super Kamiokande experiment. Both Kamiokande and Super Kamiokande (SK) were water Cherenkov detectors and as such were sensitive to elastic scattering at a higher energy threshold than the chlorine experiment. SK was able to observe that the ^8B flux was $0.358_{-0.008}^{+0.009}(\text{stat})_{-0.010}^{+0.014}(\text{syst})$ of the SSM prediction [16].

2.2.2 The Atmospheric ν Anomaly

Cosmic ray interactions in the atmosphere produce charge pions and muons, which upon decaying produce neutrinos. The physics of cosmic rays and their interactions is expanded on in §2.5, but for the purposes of this section the charged pions decay to a muon with a branching ratio of 99%. If the muon energy is below 1 GeV, it will likely decay before reaching the Earth’s surface, the chain can be represented by:

$$\pi^\pm \rightarrow \mu^\pm + \overset{(-)}{\nu}_\mu \rightarrow e^\pm + \overset{(-)}{\nu}_\mu + \overset{(-)}{\nu}_e \quad (2.4)$$

This chain provides three neutrinos with two ν_μ produced to every ν_e . However when the IMB, Kamiokande and Super-Kamiokande experiments measured the ratio they found a deficit of ν_μ interactions for higher incident angles associated with longer path-lengths. The interactions were

observed via a lepton produced by elastic scattering, (ES) and charged current, (CC), interactions on water:

$$\nu_l + N \rightarrow l^- + N^{+*} \quad CC \quad (2.5)$$

$$\nu_e + e^- \rightarrow \nu_e + e^- \quad ES. \quad (2.6)$$

The deficit of ν_μ interaction, along with the solar neutrino problem, could be explained by a simple model of flavour mixing but at the cost of the neutrino being massive.

2.2.3 Flavour mixing

The solution to both the solar neutrino problem and the atmospheric ν anomaly is understood by the revelation that neutrinos change their flavour content as they propagate through space. The explanation of such an effect is provided by a simple quantum mechanical argument. The flavour eigenstate of the neutrinos are not aligned with the mass eigenstates, which exist if and only if neutrinos have mass.

The origins of the idea go back to Pontecorvo in the late 1950s, motivated by kaon oscillations. However, his idea was oscillations between ν and $\bar{\nu}$, but the principle was applied to flavour states by Maki, Nakagawa and Sakata in 1962. This led Pontecorvo to predict the solar neutrino problem before it was experimentally observed. Wolfenstein, Mikheev and Smirnov developed the theory of the effect of matter on the flavour oscillations in the 1970s and 1980s, known as the MSW effect.

The fundamental explanation of flavour oscillations comes from accumulation of phase differences between the mass states, which naturally occur as the mass states travel through space. After time, t , a flavour

state will have evolved by the action of the operator $\hat{T}(t)$:

$$|\nu(t)\rangle = \hat{T}(t) |\nu_\alpha\rangle \quad (2.7)$$

The projection into state $|\nu_\beta\rangle$ is given by:

$$\langle \nu_\beta | \hat{T}(t) | \nu_\alpha \rangle \quad (2.8)$$

Therefore, the probability that the state interacts as state $|\nu_\beta\rangle$ is given by:

$$P_{\alpha \rightarrow \beta} = \left| \langle \nu_\beta | \hat{T}(t) | \nu_\alpha \rangle \right|^2 \quad (2.9)$$

$\hat{T}(t)$ is non diagonal in the flavour basis and related to the diagonal mass basis by the unitary linear operator U, that is:

$$T_{\text{flavour}} = U T_{\text{mass}} U^\dagger$$

$$T_{\text{mass}} = \begin{pmatrix} e^{-iE_1 t} & 0 & 0 \\ 0 & e^{-iE_2 t} & 0 \\ 0 & 0 & e^{-iE_3 t} \end{pmatrix} \quad (2.10)$$

Where $E_{1,2,3}$ are the energies of the 3 mass components written in natural units. U is the PMNS matrix.

In practice, experiments are constrained by a choice of baseline, L, from a neutrino source to the detector. This often leads to a reduced framework of 2 flavour mixing as the third phase difference is not relevant on the chosen baseline. In the two-flavour case, the probability of Equation 2.9 is given by:

$$P_{\alpha \rightarrow \beta} = \sin^2(2\theta) \sin^2 \left(1.27 \frac{\Delta m^2 L}{E} \right), \quad (2.11)$$

where Δm^2 is the difference between the relevant mass states in eV^2 , L in km and E in GeV and θ is the mixing angle in the 2D mixing matrix[2]. It is important to note that the probability in equation 2.11 is independent of the sign of Δm .

Oscillations in matter complicate the simple model shown so far. Elastic forward scattering interactions between the propagating neutrinos and the matter's electron cloud creates an effective potential that affects the neutrino's transition probability. Specifically, the mass difference between the relevant mass eigenstates is affected. The effect is such that the sign of Δm is possible to detect and measured.

The MSW effect in the Sun is such that ν_e travelling from the dense core to the less dense surface with energy above 2 MeV can undergo adiabatic conversion. This phenomenon results in ν_e emerging in the pure ν_2 mass state. As mass states won't oscillate the probability of measuring the ν_e production is given by $\sin^2(2\theta)$.

Experimental confirmation of flavour mixing

In 1998 Super-Kamiokande measured the so called up-down asymmetry of the atmospheric ν_e flux:

$$A_{\mu}^{\text{up-down}} = \frac{N_{\mu}^{\text{up}} - N_{\mu}^{\text{down}}}{N_{\mu}^{\text{up}} + N_{\mu}^{\text{down}}}, \quad (2.12)$$

Where N_{μ}^{up} is the number of ν_{μ} interactions coming from below the Super-Kamiokande detector and N_{μ}^{down} is the number from above the detector. Super-Kamiokande measured a six sigma deviation from zero in the ν_{μ} channel and results consistent with zero in the ν_e [17]. One explanation for this effect is that the ν_{μ} state oscillates to ν_{τ} states which are not detected.




Fig. 2.3: Ratio of observed to predicted flux vs. L/E for atmospheric ν_μ events at Super-Kamiokande [18]. L is the inferred baseline and E is the neutrino energy. Red: the best fit to neutrino oscillations. Dotted line: the best fit to neutrino decoherence. Dashed line: the best fit to neutrino decay.

This result was confirmed in a second improved measurement which showed the number of detected ν_μ events as a function of L/E , Fig. 2.3. The agreement between the data and the oscillation hypothesis outweighed the agreement between two competing hypotheses of neutrino decay and neutrino decoherence. This result added considerable weight to the flavour oscillation explanation but more results were needed [19].

The SNO experiment, the predecessor to SNO+, set out to measure not only the ν_e flux but also the total solar ν flux through neutral current interactions on the D_2O target material:

$$\nu_x + D \rightarrow n + p + \nu_x . \quad (2.13)$$

The free neutrons were detected through captures on deuterium, loaded salt and by custom neutral current detectors brought online in the final stage of the experiment. In 2002, SNO published results which showed the total ν flux was consistent with the SSM, while the ν_e flux was roughly 1/3 of the SSM prediction [20]. The ν_e flux was measured independently

Table 2.1: Summary of experiments that have produced the current values of the PMNS matrix, also shown is the mode of detection and corresponding PMNS parameters.

Source	Experiments	Measurement	PMNS Parameters	Reference
Nuclear Reactors	RENO, Daya Bay, Double CHOOZ	ν_e appearance	$\sin^2\theta_{13}, \Delta m_{2e}^2 $	[23-25]
	KAMLAND		$\sin^2\theta_{12}, \Delta m_{12}^2 $	[26, 27]
Terrestrial Accelerators	T2K, NOvA	ν_μ disappearance, ν_e appearance	$\sin^2\theta_{13}, \Delta m_{23}^2 $	[28]
	MINOS	ν_μ disappearance	$\sin^2\theta_{12}, \Delta m_{12}^2 $	[29, 30]
Atmospheric Neutrinos	Super-Kamiokande, IceCube DeepCore, Antares	ν_μ disappearance, ν_e appearance	$\sin^2\theta_{23}, \Delta m_{23}^2 $	[31-35]
Solar Neutrinos	SNO	^8B ν_e flux, ^8B ν_μ flux		[36]
	Super-Kamiokande	^8B ν_e flux		[37, 38]
	Gallex, SAGE	pp+ pep + ^7Be , ^8B + CNO flux	$\sin^2\theta_{12}, \Delta m_{12}^2 $	[39, 40]
	Homestake	^8B + pep + CNO ν_e flux		[41, 42]
	Borexino	pp, pep, ^7Be , ^8B ν_e flux		[38, 43-45]

through charged current interactions on the deuterium. This result solved the solar neutrino problem: confirming ν flavour change, however, there still existed other possible explanations including non-standard interactions of massless neutrinos [19].

The KAMLAND experiment finally confirmed the flavour oscillations model by measuring $\bar{\nu}_e$ disappearance in the intense neutrino fluxes provided by fission reactors [21, 22].

The energy threshold in the SNO experiment meant that the main contribution to the ν_e flux was from ^8B neutrinos. This species have energy such that most of them leave the Sun in the ν_2 mass state. This led to tension between solar ν_e flux measurements made in gallium and chlorine detectors such as the Davis experiments previously mentioned. This tension was resolved by the confirmation of the MSW effect in the Sun and thus determined the sign of the relevant mass difference, Δm_{12} .

2.2.4 Summary of oscillation measurements

The PMNS matrix elements and two squared mass differences between the $\nu_{1,2,3}$ mass states are all accessible by measuring flavour appearance and disappearance from sources of definite flavour. Being a 3x3 unitary

matrix the PMNS has four degrees of freedom, for reasons laid out below the PMNS is sometimes parametrised:

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

where s_{ij} and c_{ij} are the sine and cosines of the mixing angle θ_{ij} . The four degrees of freedom are then written as $\theta_{12}, \theta_{23}, \theta_{13}$ and δ_{CP} . The three θ_{ij} s are the mixing angles between the mass states and δ_{CP} is a CP violating phase. This parametrisation has the advantage that variables in each sub-matrix are accessible by experiments looking at different neutrino sources in the simplified 2D mixing case. The far left matrix is probed with atmospheric neutrinos, the middle matrix parameters are found in the reactor and accelerator sector and the final matrix is characterised by solar neutrino measurements². If neutrinos are Majorana fields then the PMNS matrix has a fourth matrix of the form $\text{diag}(e^{i\alpha_1}, e^{i\alpha_2}, 1)$ where $\alpha_{1,2}$ are two CP violating Majorana phases. Table 2.1 provides a list of experiments that have probed and continue to probe the PMNS elements.

Of the two squared mass differences only Δm_{12} has had its sign measured, with both magnitudes measured to within 1–2%. This leaves two scenarios with $|\Delta m_{31}|$ taking either sign. The case where $m_1 < m_2 < m_3$ is referred to as the normal hierarchy³ (NH) and the $m_3 < m_1 < m_2$ case is known as the inverted hierarchy (IH). Fig 2.4 shows a diagram of the possible cases. The two mass hierarchies have small but detectable effects on the oscillation probabilities. At the time of writing the NH

²KamLAND is also sensitive to these parameters.

³Hierarchy is sometimes referred to as ordering in the literature.

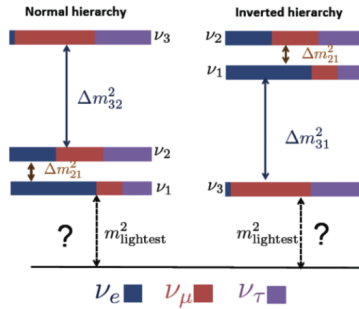


Fig. 2.4: The two possible neutrino mass hierarchies due to the currently unknown sign of Δm_{31}^2 . The sign has an effect on both oscillation results as well as affecting the value of a possible effective Majorana mass [46].

is preferred but only to approximately 3σ [47]. This result is primarily informed by the long baseline accelerator experiments T2K [48] and Nova [49, 50]. It is possible for a combined T2K/Nova measurement to reach the 5σ in the mass hierarchy measurement, however, the DUNE [51, 52] and Hyper-Kamiokande [53] experiments will be able to reach 5σ in both mass hierarchy and δ_{CP} .

Table 2.2 is a summary of the current best fit values and their 3σ ranges for the flavour oscillation parameters. Results for both mass hierarchies are shown either in a specific row or by brackets with the inverted hierarchy results shown in the brackets.

2.3 Neutrino mass

2.3.1 Mass generation

The experimental evidence for massive neutrinos is conclusive, however, a theoretical understanding of the form of massive neutrinos is an active field. Of the many theoretical ideas, two scenarios have gained popular interest.

Table 2.2: Summary of the current best fit values and their 3σ range (2σ range for δ) taken from the particle data group [54]. Both mass hierarchy scenarios are shown either with a dedicated entry or with the IH shown in brackets.

Parameter	best-fit	3σ
Δm_{21} [10^{-5}eV^2]	7.37	6.93 - 7.96
$\Delta m_{31(23)}$ [10^{-3}eV^2]	2.56 (2.54)	2.45 - 2.69 (2.42 - 2.66)
$\sin^2 \theta_{12}$	0.297	0.250 - 0.354
$\sin^2 \theta_{23}$, NH	0.425	0.381 - 0.615
$\sin^2 \theta_{23}$, IH	0.589	0.384 - 0.636
$\sin^2 \theta_{13}$ NH	0.0215	0.0190 - 0.0240
$\sin^2 \theta_{13}$ NH	0.0216	0.0190 - 0.0242
δ_{CP}/π	1.38 (1.31)	2σ : (1.0 - 1.9) (2σ : (0.92-1.88))

The first is the existence of heavy sterile neutrinos which provide, yet to be observed, right handed fields that interact minimally with the weak force⁴. The right handed fields result in neutrino mass generation through the Higgs mechanisms via a standard right, ψ_R , and left, ψ_L , handed mixing term in the Lagrangian:

$$\mathcal{L}^D = -m_D (\bar{\psi}_L \psi_R + \bar{\psi}_R \psi_L) , \quad (2.15)$$

where m_D is the Dirac mass parameter.

The second scenario comes with the observation that neutrinos are the only neutral fermions and therefore it is possible for them to be described by Majorana fields. In 1937, Ettore Majorana found that two degrees of freedom could be removed from a field, ψ , by insisting that ψ is its own charge conjugated field. This Majorana condition is:

$$\psi = -\gamma^0 \mathcal{C} \psi^* = \psi^{(c)} , \quad (2.16)$$

⁴Sterile neutrino fields will not be completely sterile as they will contain a small left handed component.

where γ^0 is one of the γ matrices from the chiral representation of the Dirac equation, \mathcal{C} is the charge conjugation operator and ψ^* is the complex conjugate of the field. Majorana fields obeying Equation 2.16 are interpreted as fields that are their own anti-particles. It has been shown that such fields can lead to mass terms in the Lagrangian, written with left handed components, of the form:

$$\mathcal{L}^M = -m_L \left(\overline{\psi_L^{(c)}} \psi_L + \text{H.c.} \right) . \quad (2.17)$$

This leads to the term producing non-renormalisable terms. This result means that the existence of Majorana particles shows that the Standard Model is an effective field theory of some beyond the standard model (BSM) theory resulting from the spontaneous symmetry breaking of a larger symmetry group. Another result of Majorana neutrinos is the violation of lepton number in the Standard Model in reactions such as neutrinoless double-beta decay, again insisting the existence of BSM theories.

One concern with the heavy sterile neutrino description of neutrino mass is the uniquely small Yukawa coupling, g , required for neutrino masses lower than the current experimental upper bounds. The addition of Majorana fields can resolve this via the see-saw mechanism.

The see-saw mechanism is boardly described as taking the total mass contribution as the sum of Dirac and Majorana terms in the Lagrangian. Diagonalising the resulting matrix then gives rise to both a heavy particle with mass of order of some new physics scale, \mathcal{M} , and a light mass state with mass:

$$m_\nu = \frac{gv^2}{\mathcal{M}} , \quad (2.18)$$

where v is the particles vacuum expectation value. The light mass is then not a consequence of a low Yukawa coupling but naturally occurs as a result of \mathcal{M} in the denominator of equation 2.18. This attractive feature makes Majorana fields a suitable candidate for the neutrino fields via the see-saw mechanism.

2.3.2 Mass measurements

Indirect mass measurements

Cosmological constraints are made through observations of the affect of a finite mass on the large scale structure of the universe. In the well verified Λ CDM[55] model, neutrinos are created in the big bang and freeze out just before recombination. The matter density of neutrinos then contributes to the seeding of the large scale universe. The neutrino matter density can also be observed in the power spectrum of the cosmic microwave background. A recent combining analysis across the field concluded an upper limit on the sum of all 3 masses [56]:

$$\sum_{i=1}^3 m_i < 0.17 \text{ eV} . \quad (2.19)$$

at 95% confidence.

Direct mass measurements

The absolute neutrino mass measurements are beyond the realm of oscillation measurements other than setting a lower limit for the combined masses of $\sum_{i=1}^3 m_i > 0.05 \pm 0.1 \text{ eV}$ implied from the squared mass measurements [2].

The total kinetic energy released in a β decay is given by:

$$Q_\beta = Q_N - m_e - m_{\nu_e}, \quad (2.20)$$

where Q_N is the total energy released in the decay, m_e is the mass of the electron and m_{ν_e} is the mass of the ν_e . A non-zero neutrino mass presents as a constant offset in the decay energy spectrum. This offset is measurable near the decay's end point with the offset equal to:

$$-m_{\nu_e}^2 = -\sum_{i=0}^3 |U_{ei}^2| m_i^2, \quad (2.21)$$

The KATRIN experiment [57] aims to measure such a shift in tritium decays using an electromagnetic spectrometer. The Project8 experiment [58] will do the same by measuring the Larmor frequencies of the electrons in a magnetic field and the ECHO experiment [59] will use a low temperature calorimeter to measure the energy spectrum of ${}^{63}\text{Ho}$. All experiments aim to be sensitive to a m_{ν_e} of $\mathcal{O}(100\text{meV})$.

If neutrinos are Majorana in nature then the so called Majorana mass, denoted here as $m_{\beta\beta}$, can be found through measurements of the yet to be observed neutrinoless double β decay, the subject of the next section.

2.4 Neutrinoless double-beta decay

The only viable experimental test of the Majorana nature of neutrinos is the observation of neutrinoless double- β decay ($0\nu\beta\beta$). The following section motivates the exotic $0\nu\beta\beta$ decay as well as reviewing the field's current status.




Fig. 2.5: Diagram showing how the nuclear energy of double beta decaying isotopes changes as the atomic mass changes as a result of β or $2\nu\beta\beta$ decays. The blue curves show isobars of nuclear energy, with the black dots showing certain isotope energies. It is possible that certain isotopes are energetically unable to decay by a single β decay but can $2\nu\beta\beta$ decay.

2.4.1 $2\nu\beta\beta$ decay

Double β decay ($2\nu\beta\beta$) is a rare form of β decay that occurs when the energy of the decaying nucleus increases with a single β decay but decreases with two simultaneous decays. Fig. 2.5 illustrates how nuclear energy changes along two contours of equal energy representing nuclei with even Z (atomic number) and N (mass number) and odd Z and N . The odd Z odd N curve has a higher energy as the paired nucleons are more strongly bonded.

β^\pm decays move isotopes between curves but about the curves minimum energy, a transition from even Z even N isotopes to odd Z odd N is energetically disfavoured however two simultaneous decays is possible, this is the double beta decay. Fig. 2.6(a) shows the process at quark level, the process has two e^- and two $\bar{\nu}_e$ in its final state which share the released energy.




Fig. 2.6: Quark level Feynman diagrams for both $2\nu\beta\beta$ decay, (a) and $0\nu\beta\beta$, (b).




Fig. 2.7: A pseudo Feynman diagram demonstrating the Schechter-Valle theorem [61], adapted from [62].

It is theoretically possible for 35 isotopes to undergo $2\nu\beta\beta$ however only 20 isotopes have been observed to decay this way. One of these isotopes is the ^{130}Te , which will be used in SNO+, that has a measured $2\nu\beta\beta$ half-life of $8.2 \pm 0.2(\text{stat}) \pm 0.6(\text{sys}) \times 10^{20}$ years [60].

2.4.2 $0\nu\beta\beta$ decay

If neutrinos are Majorana in nature then one can draw a diagram in which the $\bar{\nu}_e$ s in the final state can be joined into a single virtual massive neutrino. Fig 2.6(b) shows such a diagram but one may draw many BSM diagrams. Such a process changes the lepton number by two units and is therefore not allowed in the SM.

It has been shown that if $0\nu\beta\beta$ decays are observed then neutrinos are Majorana, this is proven by the black box or Schechter-Valle theorem [61],

depicted in Fig 2.7. The theorem states that one can always draw a diagram showing the equivalences of $\bar{\nu}$ and ν in $0\nu\beta\beta$ regardless of mechanism. For light neutrino exchange mechanisms such as the one shown in Fig 2.6(b), the rate depends on the nuclear physics and the effective Majorana mass:

$$m_{\beta\beta} = \sum_{i=1}^3 U_{ei}^2 m_i . \quad (2.22)$$

One may write the half life of the process, $T_{1/2}$, with phase space factor $G^{0\nu}$ and nuclear matrix element $M^{0\nu}$ as :

$$\frac{1}{T_{1/2}} = G^{0\nu} \|M^{0\nu}\|^2 \left(\frac{|m_{\beta\beta}|}{m_e} \right)^2 , \quad (2.23)$$

where m_e is the electron mass.

Nuclear matrix element

Nuclear decays of high Z elements are complicated systems and it is not possible to write down the corresponding analytic matrix element. Therefore approximations must be made; a range of methods such as lattice QCD and Ab Initio nuclear-structure theories are utilised to produce results [63].

This nuclear uncertainty presents as a factor of above three over results produced by a number of groups. The ultimate affect of this uncertainty is that when using Equation 2.23 to convert the experimental observations of $T_{1/2}$ to effective Majorana mass measurements, a range of values are provided.




Fig. 2.8: The so called “lobster plot” showing the effective Majorana mass as a function of the lowest neutrino mass. Both mass hierarchy scenarios are shown, with the bands provided by the uncertainty of both the Majorana phases and the PMNS elements. Courtesy of Martti Nirliko.

Mass hierarchy separation

$m_{\beta\beta}$ contains all of the neutrino physics in Equation 2.23 and depends on a number of variables; the PMNS elements, all three neutrino masses and the two Majorana phases. Fig 2.8 shows $m_{\beta\beta}$ as a function of the lightest neutrino mass⁵. Both mass hierarchies are shown with the bands of each scenario provided by the possible $[0, 2\pi]$ range of the Majorana phases and the PDG PMNS values with their 1–3 σ variations.

At $m_0 < 0.1$ eV there is a difference of 2–4 orders of magnitude in $0\nu\beta\beta$ rate. The dip tending to zero in the normal hierarchy is the result of the free parameters used to make the plot conspiring to sum the three complex numbers in Equation 2.22 to zero. It is not observable in the inverted hierarchy as the free parameters don’t allow such a cancelling triangle to exist.

⁵ This is the so called lobster plot.

Experimental searches

Both $2\nu\beta\beta$ and $0\nu\beta\beta$ release the same amount of energy but the latter has fewer particles in the final state. This means that this energy will be split between fewer particles, importantly between the detectable particles. The virtual neutrino in $0\nu\beta\beta$ means that the summed energy of the remaining electrons will be higher than in $2\nu\beta\beta$ decays. This means that a typical $0\nu\beta\beta$ search focuses on both the shape and counts in a region of interest in energy about the $2\nu\beta\beta$ Q-value.

With a 100 meV $m_{\beta\beta}$, experiments such as SNO+ only expect to observe ~ 2.5 $0\nu\beta\beta$ counts/100kg/year. In the event that SNO+ doesn't observe a $0\nu\beta\beta$ signal an exclusion limit will be produced. For a simple counting experiment with a non-zero background rate, r_B , the 90% confidence interval is given by:

$$T_{\frac{1}{2}}^{90\%} = \frac{\ln 2}{1.28\sqrt{r_B}} \cdot \epsilon \cdot \sqrt{t} \cdot N_{\beta\beta} , \quad (2.24)$$

where ϵ is the signal detection efficiency, t is the exposure time and $N_{\beta\beta}$ is the number of $2\nu\beta\beta$ nuclei present. It is clear that background rejection and large amounts of $2\nu\beta\beta$ isotope are important to be optimally sensitive.

Given the number of different isotopes and different detector technologies, there are a number of experiments that have searched and produced limits on $0\nu\beta\beta$ decay. Table 2.3 provides a list of experiments with their latest limit for the relevant isotopes. Also shown in Table 2.3, where relevant, is the particle identification (PID) technique utilised. These PID techniques add to an already high background rejection.

The KAMLAND-Zen experiment holds the strongest limit on $0\nu\beta\beta$ decay. KAMLAND-Zen use liquid scintillator to observe 750 kg of ^{136}Xe in a sus-

Table 2.3: Summary of some $0\nu\beta\beta$ limits set by experiments in recent years across a number of isotopes.

Experiment	Isotope	Isotope Mass/kg	Technology	Particle ID	Latest Limit (90% confidence)
CUORE	^{130}Te	206	Crystal Bolometer	-	1.5×10^{25}
GERDA	^{76}Ge	35	Segmented Germanium Detectors	Pulse shape discrimination	8.0×10^{25}
EXO-200	^{136}Xe	160	Time Projection Chamber	Single/multi-site discrimination	1.1×10^{25}
	^{100}Mo	6.9			4.6×10^{23}
	^{82}Se	0.9			1.0×10^{23}
	^{116}Cd	0.4			5.0×10^{21}
NEMO-3	^{100}Mo	-	Tracking Calorimeter	Full tracking	5.8×10^{23}
	^{150}Nd	-			1.8×10^{22}
	^{82}Se	-			2.1×10^{23}
	^{96}Zr	-			8.6×10^{22}
	^{48}Ca	-			1.3×10^{22}
KamLAND-Zen	^{136}Xe	750	Liquid scintillator	Pulse shape discrimination	1.07×10^{26}

pended balloon in the centre of a spherical photo multiplier arrangement. At the time of writing the most stringent KAMLAND-Zen limit is given by [64]:

$$T_{1/2} > 1.07 \times 10^{26} \text{ years} , \quad (2.25)$$

Corresponding to a effective Majorana mass of:

$$m_{\beta\beta} < (61 - 165) \text{ eV} , \quad (2.26)$$

where the range of values is a reflection of the uncertainty of the nuclear matrix element. SNO+, the subject of this thesis, has a projected $0\nu\beta\beta$ sensitivity of $T_{1/2}^{0\nu\beta\beta} \geq 1.78 \times 10^{26}$ years at 90% confidence.

2.5 Cosmic rays showers

The adage of the vacuum of space is somewhat misleading as there exists a non-zero density of charged particles pervading space. These particles are known as cosmic rays and consist of mainly protons. Cosmic rays range

in energy from non-relativistic to ultra-relativistic (10^{20} eV). The origin of cosmic rays is an area of active study with the high energy astronomical objects such as supernovae and quasars thought to be responsible. A recent review can be found in [65]. One notable result is the 2017 publication by the Pierre Auger Collaboration [66], which hinted that the highest energy cosmic rays are extragalactic in origin.

Cosmic rays with energies greater than 1 TeV interact hadronically with the Earth's atmosphere producing secondary particles that then go on to interact producing more secondary particles. Repetition of this process produces a shower. The initial strong interactions in these showers produce, among other things, pions and kaons, the resulting shower is known as a hadronic shower. The hadronic components may decay, leading to the creation of muons, or interact depending on the particle's energy and position in the atmosphere. The ratio of expected interaction length and decay lifetime of the particle determines the rate of change of the shower components. The expected interaction length is a function of air density, therefore, the shower components evolve as the shower propagates. An example of this is at lower altitudes the density is such that the hadronic components are attenuated, explaining why pions were discovered on mountain top experiments and not sea level experiments.

As well as charged pions, hadronic showers also contain a π^0 component. These π^0 s decay electromagnetically via:

$$\pi^0 \rightarrow \gamma + \gamma . \tag{2.27}$$

Table 2.4: Fitted values of γ over the energy range shown in fig 2.9, the fitted values are taken from [67].

E ($\log_{10}(eV)$)	γ
<15.5	1.7
15.5 - 17.7	2.0
17.7 - 18.5	2.3
18.5<	1.7

The resulting γ s can then seed showers of electrons via nuclei pair production:

$$\gamma + N \rightarrow e^- + e^+ + N . \quad (2.28)$$

The electrons and positrons can then produce photons via Bremsstrahlung. These two processes can then cycle to create an electromagnetic shower.

Both the electromagnetic shower and hadronic showers overlap geometrically but have different properties.

2.5.1 Cosmic ray energy spectrum

Fig 2.9 shows the cosmic ray energy distribution measured by the Auger observatory. The distribution is well described by an inverse power law:

$$\frac{dN}{dE} \propto E^{-(\gamma-1)} , \quad (2.29)$$

with the fit parameter γ given by table 2.4. The variation of γ over the 13 orders of magnitude in energy is expected to be a result of the different acceleration mechanisms.

2.5.2 Cosmic muon production

As hadronic showers propagate, a muon component grows as the result of hadron decays. At lower cosmic ray energies, pion decay is the dominant




Fig. 2.9: Cosmic ray energy distribution as measured by the Auger observatory [67]. The approximate integral flux (per steradian) is also shown (green dashed line).

muon production mode, however, when the cosmic ray energy is ~ 100 GeV kaon decay becomes relevant. The population of muons is referred to as cosmic muons, a misleading name as they are produced in the atmosphere and are not extraterrestrial.

Charged pions in the hadronic shower can decay whilst in flight via:

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

The pion decays with a lifetime of 26 ns [54]. Muons can also be produced through kaon decay via :

$$K^+ \rightarrow \mu^+ + \bar{\nu}_\mu \quad (\sim 63\% \text{ BR}) \quad (2.30)$$

The muons in turn decay via:

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

The muons have a lifetime of $\sim 2.2 \mu s$, in this time the muons can naturally reach the Earth's surface from its atmospheric origins, a consequence of time dilation.

Cosmic muons underground

The electromagnetic shower can be completely attenuated by a few meters of the Earth's surface, therefore this component never reaches the depths of SNO+, see Fig 2.10. The same is true for hadronic showers, however, muons can penetrate deeper than other components with the highest energy muons reaching the SNO+ detector.

As muons travel through the earth's crust they lose energy through interactions with the rock. The mode of energy loss is energy dependent;

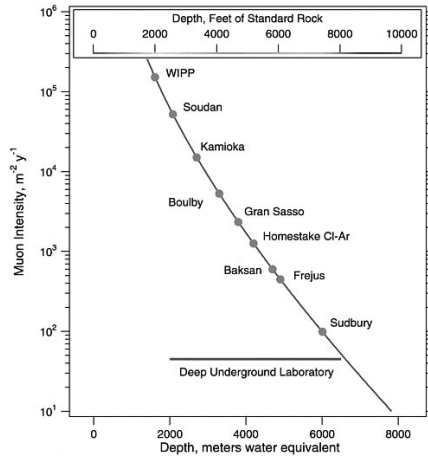


Fig. 2.10: Cosmic muon intensity as a function of experiment depth given in meters water equivalent. SNOLAB in Sudbury is the deepest operational laboratory [71].

at low energy ionisation of the rock is dominant whereas discrete radiative interactions dominate at high energy. The energy at which the modes are equal is called the critical energy, which for muons is ~ 1 TeV [68]. Above the critical energy, the muon loses energy through Bremsstrahlung, direct e^+e^- pair production and electromagnetic interactions with nuclei [69]. Losses through electromagnetic interactions are an order of magnitude smaller than the approximately equal contributions from Bremsstrahlung and e^+e^- pair production.

Muon propagation is simulated stochastically as the high energy losses are discrete. Uncertainties on each of the muon energy loss modes are in the range 1–2%, with the exception of the photonuclear interactions which have an uncertainty of 10–20% [70].

The energy loss mechanisms are proportional to the density of the material the muon is travelling through. In order to compare different experiments and their depths, the overburden density is quoted in “meters of water equivalent” (m.w.e). Figure 2.10 shows the m.w.e of a number of

experiments, SNOLAB is at the time of writing the deepest operational laboratory in the world. Not shown on Fig 2.10 is the Jinping laboratory in Sichuan, China, which has a depth of 6720 m.w.e ⁶ [72] which is yet to be operational.

SNOLAB depth not only reduces the rate of muons that reach the laboratory but also selects muons with higher energy. Therefore, experiments in SNOLAB are uniquely placed to measure cosmic muon related physics. One of these physics topics is the rate of neutron production via nuclear spallation. These neutrons present as backgrounds in rare signal experiments.

2.6 Muon Spallation

The average cosmic muon energy loss is given by [68]:

$$\frac{dE}{dx} = \alpha(E) + \beta(E) \cdot E , \quad (2.31)$$

where the $\alpha(E)$ term accounts for the continuous loss of energy through ionisation of atomic electrons. $\alpha(E)$ has a typical value of $2 \text{ MeV cm}^{-2} \text{ g}^{-1}$ and varies little with energy [68]. The $\beta(E)$ term is the contribution due to radiative losses which at hundreds of GeV is mainly due to pair production and Bremsstrahlung. Pair production is a continuous energy loss process whereas Bremsstrahlung fluctuates largely.

As muons of energy E travel through matter, they are accompanied by a flux of virtual photons whose energy spectrum is given by:

$$N_{\mu}(E, \omega) = \frac{\alpha}{\pi\omega} \left[(1 + \epsilon^2) \ln \frac{2\epsilon E^2}{m_{\mu}\omega} - 2\epsilon \right] , \quad (2.32)$$

⁶Meter water equivalent.

where ω is the energy of the virtual photon, $\epsilon = (E - \omega)/E$, α is the fine structure constant, and m_μ is the mass of the muon (106 MeV). These virtual photons can be absorbed by the target nucleus which disintegrates in a process known as spallation. Spallation produces a number of daughter products depending on the original nucleus and the muon energy.

In ultra pure water (UPW) cosmic muons spallate on ^{16}O and produce a number of long lived daughters, importantly for the work presented in this thesis, including free neutrons. After an initial period of non-linear thermalisation, these neutrons random walk until they capture on hydrogen, releasing a 2.2 MeV γ . Traditionally this low energy γ presents a detection challenge in water Cherenkov detectors, due to high detector thresholds to combat high background levels. However, the extremely low background seen in SNO+ presents the possibility of detecting neutron captures in standard physics running.

Li and Beacom produced two detailed Monte Carlo studies [68, 69] of the expected muon spallated induced background rate. In reference [68] the expected background yields for both the long lived daughters and neutrons are given for the SK detector. Although SK is considerably larger than SNO+ and receives a different mean muon energy, the reported neutron yield can be used to give a first order estimation of the number of neutrons to be expected in SNO+.

Taking a muon of maximum track length in SNO+ of 1700 cm and the reported yield, one finds that 0.35 free neutrons are expected per muon per kT of water. With the expected muon rate as measured in SNO one would expect 9000 free neutrons per year in the SNO+ detector.




Fig. 2.11: The neutron yield as measured by a number of experiments across a number of mean muon energies [73]. Also shown is the global fit to the data along with two theoretical predictions. There is a 30% disagreement at a muon energy of 350 GeV expected in SNO+.

2.7 Current neutron yield status

Much of the work focusing on neutron yield measurements has been carried out in scintillator base experiments [74–82]. This is primarily due to the low energy thresholds found in scintillator experiments.

Experiments typically report the neutron yield which is the number of neutrons expected to be spallated for a given muon track length in a given material density. As spallation events are inelastic electromagnetic scatters, the neutron yield should be expected to increase with muon energy. Fig 2.11 shows a number of neutron yields for a range of mean muon energies.

The plot also shows a global fit to the data and the prediction of two theoretical models. There is a clear disagreement between the global fit and theoretical prediction at high energy. At 350 GeV, the mean energy of muons passing through SNO+, there is a 30% disagreement between

theory and the global fit. The plot also shows the prediction of two MC models at three points corresponding to three detecting halls in the Daya Bay experiment. There is a clear disagreement between both the MC models and between the models and the global fit. To resolve this discrepancy, the community requires measurements such as the one motivated in this thesis.

Chapter 3

The SNO+ Detector

3.1 Introduction

The following chapter introduces the SNO+ detector, a general purpose liquid scintillator neutrino detector. At the time of writing, SNO+ is in its second phase of operation, half way to achieving its ultimate goal of searching for neutrinoless double beta decay in ^{130}Te loaded into the scintillator. The following section introduces the SNO+ detector, covering the main physics goals and phases of operation as well as reviewing the electronics setup and detector calibration methods. A general overview is given with details provided on topics important in the remaining chapters, that is the first phase with a detector filled with ultra pure water.

3.2 Detector overview

The SNO+ detector is situated approximately 2 km, 5890 ± 10 meter water equivalent (m.w.e) [83], underground in the active Creighton nickel mine. The large overburden reduces cosmogenic backgrounds considerably, the cosmic muon flux has been measured to be [84]:




Fig. 3.1: Cartoon of the SNO+ experimental setup, see section 3.2 for explanation.

$$(3.31 \pm 0.01(stat) \pm 0.09(sys)) \times 10^{-10} \mu\text{s}^{-1} \text{cm}^{-2} \quad (3.1)$$

SNO+ inherits a considerable amount of its infrastructure from the SNO experiment. Although primarily tasked with a different physics goal, much of the detector hardware, electronics and fill systems have been repurposed. Fig 3.1 shows a schematic of the SNO+ detector set up. Working inwards we have:

- **Cavity**: The whole detector is submerged in a 22 x 34m cylindrical void sealed with urylon, containing 7 kt of ultra pure water (UPW). This buffer of UPW acts as a shield against penetrating radioactive and cosmogenically induced events. This body is referred to as external water hereafter.

- **PMTs support structure (PSUP) & Photomultiplier tubes (PMTs)**:

Approximately 9300 active Hamamatsu 8 inch R1408 PMTs face the inner detector. The PMTs are mounted onto panels, which come together between a steel frame to form a geodesic sphere of inner radius 8.5 m, this is the PSUP. Each PMT is placed in a hyperbolic bucket made from aluminium reflective petals, which increases the PMT's collection efficiency. With these additions (adjusting for reflection effects) the total coverage of PMTs is 54%. The PSUP is held in place by supporting cables anchored in the surrounding cavity. The PSUP also supports 90 outward facing PMTs (OWLs) used to monitor activity in the cavity water e.g. Cherenkov light from cosmic muons.

- **Inner UPW volume**: The space between the PSUP and the inner detector is filled with UPW.
- **Holding ropes**: Depending on the experiment's target material, a range of buoyancies are catered for with a system of ropes which both suspends (artefact from the SNO experiment) and holds down the detecting volume. Rope tensions are adjusted to compensate for the positive or negative buoyancies of the detector.
- **Acrylic vessel (AV)**: The detecting material is housed in a sphere of acrylic (AV) with a opening at the top referred to as the AV neck or neck. The neck is the only passage from the outer laboratory to the centre of SNO+. The neck is used to fill the detector and to deploy calibration sources.
- **Universal interface (UI) and electronics**: Day to day maintenance is carried out on a platform, the deck, which seals the cavity. All of the detector's electronic read out is placed on the platform as well as an ultra clean room¹ housing the universal interface (UI) which seals the top of the neck. The UI houses many monitoring instruments as well as 4 PMTs pointing down the neck, these are the neck PMTs.




Fig. 3.2: $2\nu\beta\beta$ isotope natural abundance against the their $2\nu\beta\beta$ Q-values. Also shown are the end points two β decaying backgrounds, ^{214}Bi and ^{208}Tl . Taken from [85].

3.3 Physics goals

3.3.1 $0\nu\beta\beta$ search

SNO+ is primarily concerned with observing $0\nu\beta\beta$ in ^{130}Te . ^{130}Te undergoes $2\nu\beta\beta$ with a half life of:

$$T_{1/2}^{2\nu\beta\beta} = (8.2 \pm 0.2(\text{stat}) \pm 0.6(\text{syst})) \times 10^{20} \text{ years [60]} \quad (3.2)$$

and has a end point energy:

$$Q_{2\beta} = 2527.518 \pm 0.013 \text{ keV [86]}. \quad (3.3)$$

There are various advantages to using ^{130}Te as a possible $0\nu\beta\beta$ isotope. Fig 3.2 shows the natural abundance and $2\nu\beta\beta$ end points for the eleven viable $0\nu\beta\beta$ isotopes. ^{130}Te occupies a unique position with a high natural

¹The author played a part in the construction of this room.




Fig. 3.3: Projected energy spectrum after 5 years of ^{130}Te loaded SNO+ data with all current selections and background reductions included. The histogram is split into various contributing components with the red component being that of a 100 meV $m_{\beta\beta}$ $0\nu\beta\beta$ signal [88].

abundance and decay end point. ^{130}Te has the largest natural abundance, $\approx 34\%$ of natural tellurium [83]. This natural abundance means that no enrichment is needed to achieve a high target mass. The relatively high end point of the ^{130}Te means that contaminations from low energy backgrounds are reduced. Also shown on Fig 3.2 is the end point of ^{208}Tl which presents close to $Q_{2\beta}$, the rejection of this background has been studied to a large degree, using the decay products of ^{208}Tl to tag the intruding γ ray, and is continuously improved upon [87]. ^{130}Te also has a relatively high $T_{1/2}^{2\nu\beta\beta}$ compared to other $2\nu\beta\beta$ isotopes, this means that the rate of the irreducible $2\nu\beta\beta$ background is reduced.

^{130}Te will be loaded into a scintillator cocktail; the loading is initially scheduled to be at 0.5% ^{130}Te by weight in the 780 t scintillator mass. This scintillator, discussed further in section 3.4.2, provides a high light yield required for good energy resolution. Fig 3.3 shows a stacked histogram of the projected reconstructed energy spectrum with all current background rejection techniques applied. The red component of the histogram is

the expected signal from a 100 meV $0\nu\beta\beta$ candidate, the background components are discussed further in section 3.9.

Current sensitivity studies project SNO+ can set a limit on the $0\nu\beta\beta$ half-life:

$$T_{1/2}^{0\nu\beta\beta} \geq 1.7807 \times 10^{26} \text{ years} \quad (90\% \text{ C.L.}) \quad (3.4)$$

Although $0\nu\beta\beta$ is the ultimate goal of SNO+, the experiment will transition through various phases of operation in which relevant $0\nu\beta\beta$ backgrounds will be measured and systems tested. There are three distinct phases before the $0\nu\beta\beta$ phase, in chronological order these are; light water fill, partial scintillator fill and full scintillator fill.

3.3.2 Light water phase physics

In this phase the AV was filled with UPW and throughout the phase the external background level were measured. A number of physics analyses are also possible with the water filled detector²:

- Invisible nuclear decay[89]: Some beyond-the-standard-model (BSM) theories predict processes where by a nucleon can decay leaving an excited state which then emits a delayed signal. SNO+ can search for this signal, producing world leading limits on three double and single nucleon decay modes.
- Low background ^8B solar neutrino flux measurement [90]: With the low levels of low energy background present in SNO+ ^8B neutrinos interactions can be observed exploring the low energy region of the spectrum. The result is statistically limited but nevertheless consistent with measurements made by Super Kamiokande [91] and

²At the time of writing some of these analysis are published as indicated by the references.

SNO [92]. More importantly it shows that the detector is well understood with very low backgrounds.

- Solar axion search: axion-like particles potentially produced in the Sun can be searched for via various interactions in SNO+ e.g. Compton conversion [93].
- Anti-neutrino search: Neutron detection in UPW is technically challenging, however the low background and high PMT coverage of SNO+ allows for proof of principle reactor anti-neutrino studies to be carried out.
- Cosmic muon induced neutron production in H_2O , the subject of this thesis.

3.3.3 Partial water phase

In this phase of the experiment pure scintillator will be filled from the neck of the detector while the same volume of water is removed from the bottom of the AV. This fill procedure insures that the newly filled scintillator will not be contaminated by unwanted backgrounds such as radon in the mine air. In this phase the physics program is reduced but there exists ideas to measure cosmic muon flux with the increased energy resolution of scintillator with the added directionality of the water Cherenkov radiation.

3.3.4 Pure scintillator phase

This phase of the experiment heralds the start of a new state of detector operation with the scintillator introducing both a greater amount of light to detect as well as a projected order of magnitude increase in event rate

due to the lower energy threshold. All systems will be tested in this phase including the full detector readout and operation, calibration systems and grid processing chain. The internal background levels and (low energy) external backgrounds will be measured in this phase.

With the lower energy threshold there is a selection of physics that can be probed, these include:

- Low energy solar neutrino measurements³.
- Reactor $\bar{\nu}$ measurements of Δm_{21} and θ_{21} .
- Geo-neutrino searches.
- Supernova detection : If a supernova is observed, SNO+ will be able to detect any low energy ν_e component of the immense flux associated with the shock wave [94].
- Cosmic muon induced neutron production in LABPPO scintillator.

3.4 Detecting Principle

Over the four phases of operations, SNO+ will be filled with different detecting media, these being UPW and liquid scintillator. Both of these media have different detection mechanisms.

3.4.1 Cherenkov radiation

When charged particles travel faster than light in a material, the particle can be thought of as travelling faster than its own electric field. This process results in a shock wave like wave front, presenting as a cone

³Depending on internal background levels, solar neutrinos other than ^8B interaction may be detectable.

of radiation with the angle of the cone's apex, θ , being related to the particle's energy by:

$$\cos \theta = \frac{1}{n\beta}, \quad (3.5)$$

where n is the index of refraction of the material and β is the particle's relativistic correction factor.

The Frank-Tamm formula gives the number of emitted photons per unit track length, N_γ , per unit energy loss, E , produced by a particle travelling through a material of refractive index n and permeability μ :

$$\frac{d^2 N_\gamma}{dx dE} = \frac{\alpha^2 Z^2}{\hbar c} \mu(E) \left(1 - \frac{1}{\beta^2 n(E)^2} \right) \quad (3.6)$$

This leads to a typical energy resolution of tens of PMT hits per MeV deposited in the water.

Cherenkov radiation is the only known detecting mechanism in water. Such behaviour is also expected in scintillator but current technologies can't separate the relatively small Cherenkov light from the more prolific scintillation light.

3.4.2 Scintillation

Scintillation is an example of fluorescence caused by inelastic scattering or indirectly by ionisation and recombination into an excited state. Fluorescence is defined as light emitted after non-thermal excitation produced by ionizing particles. Fluorescence is characterised by the short decay time of ns to μs , whereas phosphorescence decays on the order of ms to s. The emitted light is isotropic and therefore can not be correlated with the detected particle's direction.

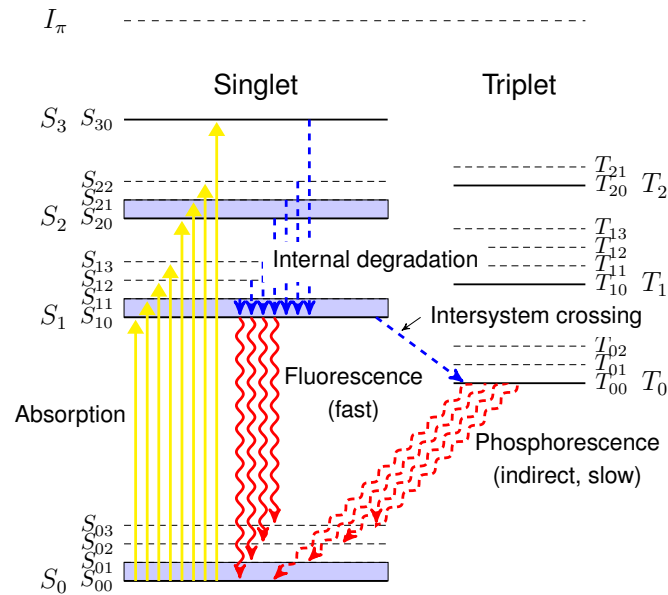


Fig. 3.4: Jablonski diagram of electron energy levels in a molecular π orbital, adapted from [95]

Organic liquids scintillate due to the formation of π bonds as a consequence of the hybridisation of aromatic rings. The scintillation process is caused by the de-excitation of highly de-localised electrons in the aromatic rings. Fig. 3.4 shows the typical energy levels in a π bond.

Electrons are promoted to the energetic states via either inelastic scatters or recombination of ionised electrons. Once in an excited state, electrons relax into lower energy levels, emitting photons in the process. The energy and time constant associated with the photon depends on the various pathways through Fig. 3.4.

Inelastic scatters promote electrons to singlet states, which then typically decay via vibrational relaxation or internal degradation to the primary scintillation level S_{10} . The transition from S_{10} to S_{0x} nominally emits a photon of the primary scintillating wavelength, this occurs within picoseconds. It is possible for the transition from S_{10} to S_{0x} to be

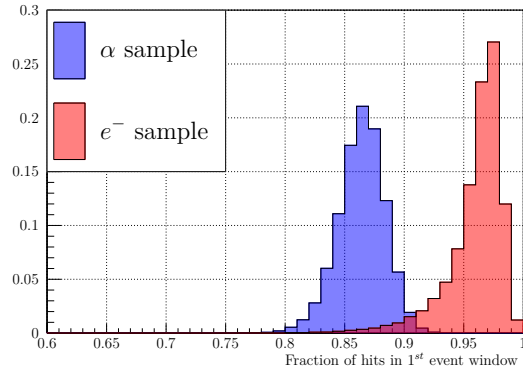


Fig. 3.5: Fraction of hits contained in the primary event window in SNO+ for both a sample of α and e^- . α particles have approximately 10% less hits in the primary event window than electrons. This translates to an underestimate of the reconstructed alpha energy.

achieved via non-radiative transitions, the fraction of excitations that lead to radiative transitions is known as the scintillator’s fluorescence quantum yield.

Triplet state can be indirectly created from singlet states via inter-system crossing. Triplet states typically decay on the phosphorescence timescales, this is due to T_{00} being stable due to spin and parity selection rules. Although the inverse process leads to S_{10} decays but on an extended time scale, referred to as delayed fluorescence or late light.

Triplet states are produced 75% of the time as a results of ionisation-recombination interactions. It is therefore expected that heavy ionising particles, e.g. α particles, will have a large late light component in their scintillator response time. This difference in the response time can be used to classify the particle associated with each event, particle identification.

The late light tail can “leak” beyond the time window used for SNO+ events, leading to a loss of light in the recorded event. Fig. 3.5 shows the fraction of hits registered in the primary event window for α and e^- samples. The α fraction is approximately 10% less than the electron

sample, leading to an underestimation of the alpha energy as hits are an important input to the reconstructions, as explained in section 3.6.2.

Heavy ions such as α particles typically lead to regions of dense ionization, these regions are dense enough for inter-molecule interactions leading to non-radiative decays, this is known as ionization quenching. A complete description of ionization quenching is not available but the resulting light yield is described by Birk's law [96].

Materials that scintillate exhibit a Stokes shift, which stops the primary scintillation being reabsorbed through energy conservation. This shift can be enhanced by doping the scintillator with a fluor. The purpose of a fluor is to achieve an excited state at the expense of the excited scintillator state, on time scales shorter than the typical primary fluorescence. The fluor then decays emitting a different wavelength, the efficiency of the transfer is such that the quantum efficiency is enhanced.

SNO+ will use the organic liquid scintillator compound LAB and will use the primary fluor 2,5-Diphenyloxazole (PPO) doped at 2 g/L. During the tellurium phase a second fluor, 1,4-Bis(2-methylstyryl)benzene (bisMSB) will be added, as well as surfactant to aid dissolving the ^{130}Te via telluric acid $\text{Te}(\text{OH})_6$ [83].

3.5 Data acquisition system

SNO+ detects radiation from events inside the PSUP via PMT triggers. PMTs detect photons by amplifying an initial photoelectron emitted by the photoelectric effect through a series of alternating anode and cathode plates. The analogue output of a PMT is typically a pulse in voltage; if the magnitude of the pulse exceeds a predetermined threshold then the PMT is said to have triggered. Collecting and processing inputs from all

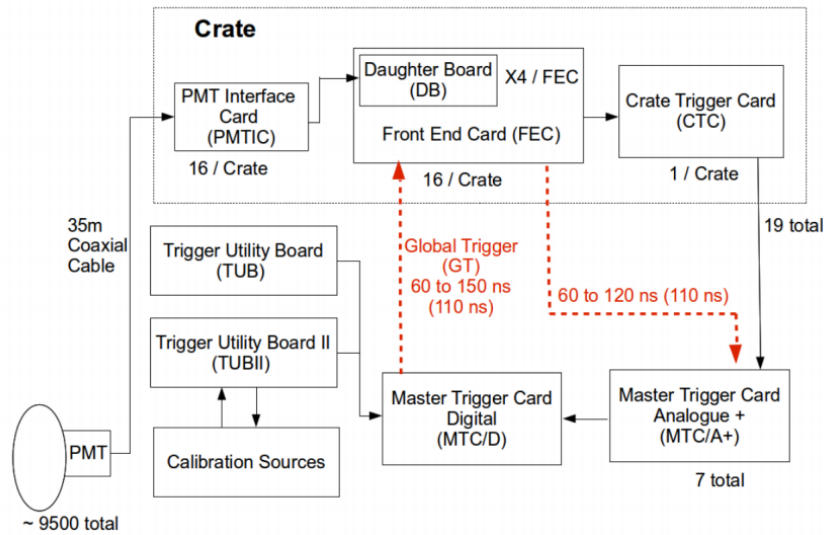


Fig. 3.6: A diagram of the data acquisition system used in SNO+ [97].

PMTs is handled by the data acquisition system shown systematically in Fig. 3.6 and explained in the rest of this section.

3.5.1 PMT configuration

The PMT base is connected to the rest of the electronics by a single cable running from the PMT's base to the deck. SNO+ arranges its front end electronics into 18 crates of 512 channels. Here a channel is defined as the complete electric pathway from the PMT to the final readout. Inside each crate the 512 channels are arranged into 16 cards of 32 channels, these cards have three components. PMTs are collected into groups of eight and fed into a PMT interface card (PMTIC), which as well as supplying the PMT high voltage also feeds into a daughter board taking the same eight channels. Four daughter boards (DB) are collected into on Front end card (FEC). There are 16 FECs per crate which feed into a signal Crate Trigger Card (CTC).

3.5.2 Trigger system

Triggering in SNO+ is a multi stage process, with analogue signals being digitised and collated before a decision is made on whether to read out the raw event data stored in buffers. Readout is triggered by a Global Trigger (GT), which is the result of the current state of the detector meeting a predefined trigger threshold. Working from the raw PMT signals to the GT signal the trigger system works as follows:

- A PMT pulse is passed through its associated PMTIC. This stage seems trivial but the PMTIC is responsible for providing voltage to groups of four PMTs via four relays. It is on the PMTIC levels that faulty PMTs are removed from the detector, typically via the relay and later by physical disconnection thus in the best case saving the three remaining PMTs on the relay. Each PMTIC collects signals from eight PMTs and forwards the signals on to the relevant DB.
- Each DB is responsible for discrimination and integration of the PMT signals. Each DB is fitted with two 4 channel discriminators, two eight channel integrators and eight one channel CMOS chips. When a channel reaches threshold the hit is registered and the Time-to-Analogue-Converter (TAC) inside the CMOS chip starts ramping. The input signal is integrated in parallel over two time windows, 60 and 390 ns respectively, producing two charge measurements QHS and QHL respectively with a 'high' signal gain, and one measure of signal integrated over 390 ns with 'low' signal gain, QLX.

The purpose of the TAC is to measure the time between the PMT being hit and the issuing of a GT further up the chain. Upon

receiving a GT, the TAC ramp is stopped. Its level, and the signal integrations, are then stored in analogue memory cells in the CMOS chip with other pieces of information e.g. global trigger identification (GTID) held in digital memory. If no GT is issued within 410 ns the memory is reset with the TAC and integrations discharged.

- When a channel crosses threshold the DB injects two analogue signals of constant amplitudes and varying time widths into trigger sums, N20 and N100. The N100 signal is ≈ 100 ns long whereas the N20 pulse is ≈ 40 ns long⁴. These sums are collected on the crate level by the CTC. All CTCs feed into seven MTC/A+ boards that collate the sums into a detector wide sum. The MTC/A+ amplifies the signals by applying various gains and digitised them, the result is five trigger signals from the original N20 and N100 signals: N20, N20LB, N100L, N100M, N100H. These signals are then passed on to the MTC/D which compares the trigger sums to predefined thresholds. If the trigger sum's threshold is met and in a requested mask then a GT is issued. The GT time is taken from the next count of the detector's 50 MHz clock. The GT leads to all data collected in the previous 410 ns being read out. Once a GT is issued, a lockout of 420 ns is placed on the MTC/D so that no GTs can overlap.
- The MTC/D can be forced to issue GTs via inputs from two trigger utility boards. These can be fed by calibration sources or controlled programmatically. PULGT events are forced GTs issued at a rate of 10 Hz. The MTC/D can also force individual channel discriminators

⁴N20 is named historically.

to fire. The resulting output is known as the pedestal of the channel, used to calibrate the channel.

- The FEC constantly polls the DBs for signals to convert to digital with analogue to digital converters (ADCs). Once a signal and previously calculated quantities are digitised they are sorted in FIFO⁵ memory as part of PMT bundles that contain, among other things, the addresses of the PMTs hit.
- Each crate contains an XL3 card, which gathers PMT bundles from all the FECs. As well as monitoring the FEC conditions e.g. FEC temperatures, the XL3 bus relays the PMT bundles over Ethernet to dedicated servers.

3.5.3 Nearline processing

Raw detector data is processed by a machine known as the ‘builder’, which writes the events to data files. Typically, a number of 1 Gb large files are required to contain all the events in an hour-long run. These files are then stored on a nearline machine for an undefined amount of time. Whilst on the nearline machine the data is processed at the run level. Each time a new run is transferred to the machine a scheduler spawns a number of processes that process the data in some way, often producing tables of information which are then pushed to a databases for reference later. An example of a nearline process is the automated CSS recalibration explained in Chapter 4. Many of the nearline processes interact with a postgres database known as the nearline database and a high speed NoSQL caching redis instance.

⁵First In First Out

3.5.4 Grid processing

The raw data is also pushed to various Grid processing sites, where they are put through the SNO+ processing chain. This chain consists of many different processes. Combining the raw data, referencing databases (nominally `ratdb`, see section 3.10.1) and applying reconstruction methods, the data is processed into files containing fully reconstructed events ready for analysis. Given a change to the processing chain, data can be reprocessed and passed through custom macros which are analysis specific.

3.6 Event reconstruction

The physics results produced by SNO+ are completely dependent on converting raw data into reconstructed information about the events that produce it. Event reconstruction has many layers ranging from simply counting the number of PMT hits (`Nhits`) in an event to simultaneous event energy, position and time estimations.

Reconstruction methods vary slightly from operational phase to operational phase, however the methods are primarily driven by an event's PMT hit times. The rest of this thesis is primarily concerned with reconstruction in the water phase, which is described here, but much of the method is applicable across other phases.

The `WaterFitter` is a collection of sequential methods which produces the final fit result. Taking the distributions of PMT hits, the event's timing reconstruction is governed by the rising edge of the distribution, whereas the position reconstruction takes the distribution's shape into account. Photons emitted from the centre of the detector reach the PMTs at the same time, whereas an event near the PSUP will have a range

of hit times depending on the particle's direction of travel. The event's energy estimations take the fitted position as an input and are governed by the number of PMT hits in the event.

3.6.1 Position and time fit

A theoretical detector with perfect timing information would be able to find the position and time of an event with the hit times of four PMTs. Each PMT hit would have a hit time, t_h , and position, \vec{x}_h , that are related to the photon's originating time, t_e and position, \vec{x}_e , by the equation:

$$|\vec{x}_h - \vec{x}_e| = c(t_h - t_e), \quad (3.7)$$

where c is the speed of light in the detecting material. The four PMT hits would provide four equations with four unknowns, t_e , \vec{x}_e , which can be solved without uncertainty.

However, noise hits and PMT jitter means that the fit result would depend on the four PMT picked. The `QuadFitter` takes a large number of PMT combinations, solves the Equations 3.7 and takes the medians of the resulting cloud of results. The fit is tuned by setting c to minimise the radial bias. This `QuadFitter` result is used as a seed to the next fitting step.

`PositionTimeLikelihood` extends Equation 3.7 to account for the difference in propagation time introduced by the AV and maximises the likelihood of the event's time residuals being from a water Cherenkov event.

A PMT's time residual, t_{res} , is effectively the difference between the measured and expected hit time due to effects such as PMT timing jitter and photon scattering. t_{res} is defined as:

$$t_{\text{res}}(\vec{x}_{ev}, t_{ev})_i = t_{h_i} - t_{t.o.f}(\vec{x}_{ev}) - t_{ev} \quad i \in \text{PMTs}, \quad (3.8)$$

where t_h is the PMT's trigger time, $t_{t.o.f}$ is the calculated time of flight for the emitted photon travelling direct to the PMT and t_{ev} is the event's time. For a postulated event time and position, a PMT's t_{res} is calculated, with the $t_{t.o.f}$ evaluated using the group velocities of 400 nm photons in water and acrylic [98]. With a set of time residuals one can build a likelihood by evaluating the probability, $P(t_{\text{res}_i})$, of each time residual with the PDF shown in Fig. 3.7. The log likelihood is defined as:

$$\log(\mathcal{L}) = \sum_i^N \log P(t_{\text{res}_i}) \quad (3.9)$$

$\log(\mathcal{L})$ is constructed with PMTs that are hit within 50 ns of the median PMT hit time and then maximised with the Powell method [99] while floating the event's position and time. Fig 3.8 shows the x position resolution for an ^{16}N calibration run for both tuned MC and source data. The fit is seen to perform well with a negligible bias and a non-zero standard deviation.

3.6.2 Direction and Energy Fit

A point-like event's direction is fitted through a likelihood fit⁶. The likelihood is evaluated with the probability of the hit PMT's angle to the postulated direction compared to a PDF of PMT angles for a 10 MeV electron.

⁶Cosmic muons are not point-like events but their followers are.




Fig. 3.7: Time residuals PDF for 6 MeV electrons uniformly filling the detector within a 5.5 m radius. This PDF is used to evaluate the likelihood in the `PositionTimeLikelihood` method. Taken from the final version of RAT, section 3.10, used for water.

The energy fit is governed by the nhits registered in an event. Using the fitted position and direction, the energy fit can be thought of as a conversion from nhits to deposited energy. Details can be found in [100].

One fundamental limit to such an energy estimation taken from the number of PMT hits is introduced by the Poisson nature of PMT hits. The number of hit PMTs in an event, N_{hits} , is a Poisson distributed random variable and as such the number is only known to within $\sqrt{N_{hits}}$. Taking a perfect reconstruction with an energy estimation, E , the fractional energy error goes like $1/\sqrt{N_{hits}}$. This limit of knowledge is referred to as Poisson limit beyond which a hit based energy reconstruction can't improve. The SNO+ water phase was found to have a rough nhits to energy conversion of 8 nhits/MeV [100].

3.6.3 Drive correction

Both the event's position and time are floated in the optimisation of `PositionTimeLikelihood` method. This leads to the reconstructed vertex being unconstrained along the axis of the particle's direction. This is due

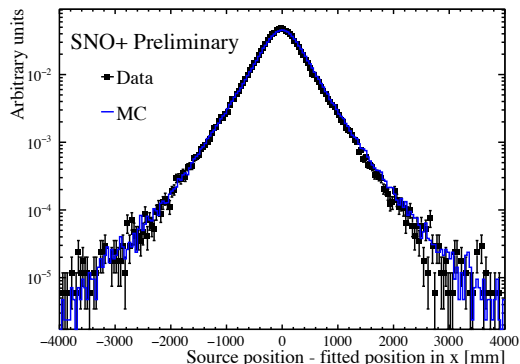


Fig. 3.8: Difference between the expected fit x position and the source x position for both ^{16}N MC and real data for an ^{16}N , see section 3.8.2, calibration run in the center on the detector. The MC spectrum has been tuned as described in section 3.8.

to a shift in position being correlated with a shift in the event's time. As observed in SNO, photons in the Cherenkov cone can be scattered outside of the cone, effectively increasing the cone's angle and increasing the hit PMT's time residual. This introduces a bias in the fit. This effect is corrected for with a transformation of the form:

$$\vec{x}_{Corrected} = \alpha \cdot \vec{x}_{fit} + \beta \cdot \vec{u}_{fit}, \quad (3.10)$$

where α and β are parameters selected by minimising the position bias. Details of this correction can be found in [101].

3.7 Detector setup across the water phase

Water physics data taking began on the fourth of May 2017 [100]. Initially the channel thresholds were set low, leading to a detector effect known as dropout. This effectively increases the number of channels expected in an event. These hits are removed at a later stage in processing but the

detector trigger settings have to reflect this in their triggering thresholds. The lowest trigger threshold was placed at ≈ 15 hits on the N100L signal. This phase of running was in operation from run 100,000 to 104,613 [102].

Later, channel thresholds were increased by 1 ADC count, which reduced the dropout in the detector such that the detector trigger thresholds could be lowered. For runs post 104,613, the lowest trigger was set to about 7 hits. At such a low thresholds the possibility of detecting low energy signals was realised such as neutron capture on hydrogen. This subject is explored through this work.

Data was typically taken in chunks referred to as runs, lasting a maximum of 1 hour. The detector was monitored at all times by an operator whose job it was to monitor the detector for electronic faults and record detailed information about the detector running.

Each run on physics data is assessed by the run selection group. The group compiles information about the each run citing numerical checks, hardware checks and the detector operator log. This information is then used to decide if the run quality is high enough to be used for physics analyses. Runs that are deemed good enough are added to the golden run list, this list of runs is then used by physics analyses including the one contained in this thesis.

3.8 Detector calibration in water

One of the corner stones of experimental science is understanding the apparatus being used to carry out the experiment. In particle physics the detectors and sub detectors are characterised in a number of calibration campaigns. Calibration is the process of measuring a detector's response against a well characterised source of events. When these sources are in-

roduced, the detector response is measured, with the detector electronics being fine tuned or the detector's resolving power estimated.

Calibration is carried out with both sources of light and charged particles in the detector as well as artificially injected events in the front end electronics.

3.8.1 Electronic calibration

ECAL

Custom PMT thresholds are determined in ECAL campaigns. Among other things the ECAL process pulses each channel at a fixed rate while varying the PMT trigger threshold. The threshold is adjusted until the channel hits are at the expected number given the pulser rate. The threshold is placed one ADC count higher than the step where hits start to register. Details of the ECAL procedure can be found in [103].

ECA

Electronic calibration (ECA) campaigns calibrate the hit timing and charge information at channel level. ECA runs follow the general pattern of firing channel discriminators and issuing artificial GT (PULGT) at some variable time later. This crate channel level readout is used to characterise the channel's calibration.

Charge calibration: Dark noise causes a channel to have a non-zero charge integration without the channel being triggered by photons. To assess this baseline, each channel is triggered with a PED trigger, 150 ns later a PULGT is issued. Averaging over this 150 ns window provides a measure of the channel's baseline in the QHL, QHS and QLX charges.

With this repeated 10 times on each channel, the channel's expected baseline pedestal is found as the median baseline charge.

Time calibration: Like the charge calibrations, the time calibration begins with a PED trigger that starts the TAC voltage ramp. PULGT are then issued at varying times. These times are chosen to sample the full TAC slope over a 500 ns time window in 50 10 ns bins. A linear model is fitted between neighbouring points leading to a conversion from TAC voltage to nanoseconds. The TAC is also shifted by -500 ns to chronologically correct times e.g. PMTs with high TACs were hit early in the event.

These two steps are nominally carried out once a week, with two dedicated 10–20 minute long runs. Once the runs are complete, analysis and production of relevant constants is performed. After the constants are inspected by an expert, the constants are uploaded to a database. During this analysis stage a lock is placed on the rest of the analysis chain, known as the processing lock. This lock pauses the standard processing while the new constants are uploaded.

3.8.2 Optical calibration

Light sources come in a variation of guises e.g. in-situ background events, but for the most part SNO+ chooses to introduce into the detector sources of light or particles. There exists two categories of calibrations sources in SNO+; deployed sources and in-situ sources.

Deployed sources are typically radioactive sources or light fibres contained in a housing being lowered into the detector through the neck and positioned at various points in the detector. In-situ sources come in the form of LED and laser based systems permanently mounted to the PSUP

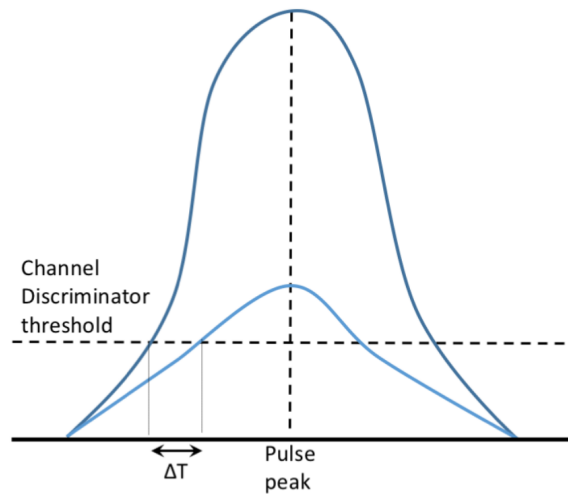


Fig. 3.9: Schematic showing the ‘time walk’ effect measured over all PMTs during the PCA process.

which can then emit light across the detector at a variety of different angles and trajectories depending on the emitting node chosen.

PCA

PMT calibration (PCA) aims to calibrate parameters which are PMT specific. One effect that is measured is a PMT’s cable delay, which is the time a PMT signal takes to propagate to the MTC/D.

Another effect that is measured in a PCA analysis is the so called ‘time walk’ of the PMTs. This effect is shown schematically in Fig. 3.9. The time taken for a PMT to reach its discriminator threshold depends on the magnitude of the charge deposited. This effect can be measured by correlating the calibrated QHS with the trigger time in the low charge region [104].

A PCA analysis requires a tagged light source to introduce well timed light into the detector. This light is typically provided by the laserball[105] but can also come from the TELLIE system, as described in section 3.8.2.

The laserball is a quartz sphere, 10.9 cm in diameter, filled with 50 μ m diameter air filled glass beads suspended in soft silicone gel. A nitrogen laser on the deck above the detector is passed through one of several dye resonators to produce light at several optical wavelengths. Fibre optic cables feed this laser light into the laserball to produce near isotropic light. The laser intensity is tuned to produce only single photo-electron (pe) hits on any PMT and, by forcing the detector to trigger asynchronously, the laserball produces light with a fixed, known relationship to the GT. The laserball can be positioned in various locations in the detector which also allows for PMT angular response measurement.

PCA campaigns are carried out when they are required and have no set routine as the constants are relatively stable over time, across the water phase PCAs were carried out approximately every three months.

Optical characterisation

Laserball PCAs are carried out with a single wavelength, once complete the optical properties of the detector media e.g. the AV can be studied. This is done by scanning the range of laserball wavelengths in various positions in the detector. It is then possible to measure the effect of the media on the light, thus providing a detailed detector response, which can then be folded into event reconstruction.

^{16}N source

Energy calibrations are generally taken with radioactive sources. The ^{16}N source is a radioactive calibration system with a variable rate depending on the rate of nitrogen gas passing through the source⁷. ^{16}N decays to

⁷In practise the nitrogen flow rate is set as high as possible to achieve the maximum event rate.

Table 3.1: Radioactive specifications for deployed sources in the water phase, at the time of writing.

Isotope	Particle	Dominant decay mode energy (MeV)	Calibrates
^{16}N	γ	6.1	Energy scale and resolution
$^{241}\text{Am}^9\text{Be}$	γ & n	4.4, 2.2	Neutron detector response

^{16}O emitting a 6.1 MeV γ ray with 66.2% branching ratio. As the source can be studied on the bench top, knowledge of the amount of energy deposited in the detector by the source is well understood. Once this source is deployed and data is taken, the MC model of the source can be tuned to best fit the detector response. The angular response found with the laserball is also tested with the ^{16}N source. The MC model includes systematic effects such as PMT angular response. The result of the analysis is a fitted detector response at the energy deposited by the source. Through deploying this source into the detector at different positions a map of the detector's energy resolution is found.

Following in the footsteps of SNO, SNO+ deployed a ^{16}N source into the detector [106]. Table 3.1 shows some details of the source. The source was deployed at various points in the detector and the detector response was parametrised by position in detector. Fig 3.10 shows the agreement between the MC model and a central ^{16}N run.

$^{241}\text{Am}^9\text{Be}$ neutron source

After adjusted the channel triggers detailed in §3.7 and given the low background environment in SNO+, the possibility of detecting neutrons was tabled. A $^{241}\text{Am}^9\text{Be}$ source was used to first assess the possibility of detecting neutrons and then measuring the detection efficiency. As detailed in table 3.1 the $^{241}\text{Am}^9\text{Be}$ source decayed producing a 4.4 MeV




Fig. 3.10: Agreement between ^{16}N MC model and data for a central calibration run, taken from [106].

prompt γ , along with a neutron. After the neutron thermalises, it then captures on hydrogen in the light water producing a 2.2 MeV γ .

The prompt γ followed by the potential delayed neutron signal was assessed in a coincidence analysis. The analysis was based on detector trigger timing and PMT hits in a timing window after a detected prompt signal. Details of the analysis can be found in [107].

The neutron signal is observed in the event timing spectrum following prompt candidate events as shown in Fig. 3.11(a), a large achievement in a water Cherenkov detector. Fitting the timing spectrum with a model accounting for all three combinations of possible prompt/delayed contributions e.g. true prompt event followed by a false neutron event, allows extraction of the number of signal candidates observed. The contribution of these fitted populations are shown relative to the total nhit spectrum in fig 3.11(b). The analysis produced a fitted neutron capture time, τ_n :

$$\tau_n = 206.03 \pm 0.44 \mu\text{s}, \quad (3.11)$$

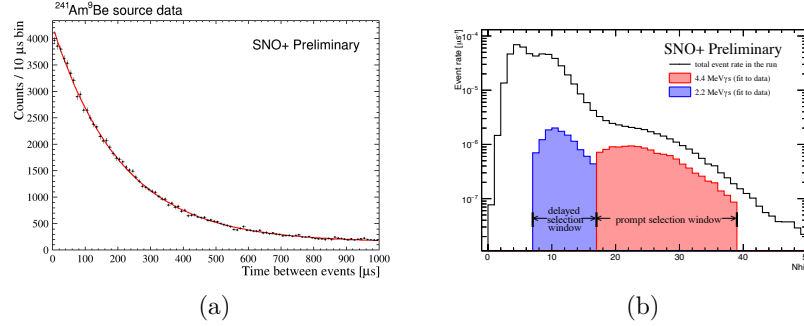


Fig. 3.11: Results from the $^{241}\text{Am}^9\text{Be}$ source analysis. Fig (a) shows the trigger times of events following a prompt candidate for a central AmBe run. The analysis fitted the neutron capture time of $206.03 \pm 0.44 \mu\text{s}$, which is consistent with the expected time. Fig (b) shows the total nhit spectrum of events during an $^{241}\text{Am}^9\text{Be}$ run, also shown are the fitted nhit contributions from both the prompt and delayed candidates.

with a detection efficiency, ϵ_n , at the centre of the detector found to be:

$$\epsilon_n \geq 48.3 \pm 0.2 \%, \quad (3.12)$$

where ϵ_n is given as a lower bound due to unoptimised candidate selections. A detection efficiency of 48% is a technical achievement in a water Cherenkov detector running with nominal physics triggers. This result motivates the analysis presented in this thesis after chapter 4.

ELLIE sub system

The **E**mbded **L**ED **L**ight **I**njection **E**ntity (ELLIE) is a set of sub systems mounted in-situ on the PSUP. The system consists of three sub systems each injecting a unique light source. Each system injects light from various nodes on the PSUP, with a variety of injection angles depending on the node. A schematic diagram is shown in fig 3.12.

TELLIE : The **T**iming module, consisting of $\mathcal{O}(\text{ns})$ LED drivers and 92 fibres entering the detector aligned so that the produced light passes




Fig. 3.12: Schematic of the ELLIE system.

through the centre of the AV. The alignments are such that all PMTs can be triggered by light tuned to 1 pe intensities. As a tagged source, one of the aims is to carry out PCA analyses without the need for laserball deployment with its associated risk of contamination.

SMELLIE : The **Scattering Module** is designed to measure Rayleigh scattering in the detecting material. SMELLIE has 12 fibres which can take light from 4 fixed wavelength lasers and a super continuum laser able to produce lasing across a range of wavelengths, 400–700 nm.

AMELLIE : The **Attenuation Module** consists of 8 angled LED fed fibres across 4 nodes which aims to measure attenuation caused by contamination.

At the time of writing all sub systems have been commissioned in water with analysis still on going.

3.9 Backgrounds

One of the most important parts of an experiment's capability is the understanding of the backgrounds inherently present in the detector.

Table 3.2: Target background levels in g/g and corresponding decay rate for the internal ^{238}U and ^{232}Th chains in a water filled SNO+ [83].

Source	Target [g/g]	Decays/yr
^{238}U chain	3.5×10^{-14}	1.2×10^7
^{232}Th chain	3.5×10^{-15}	4.1×10^5

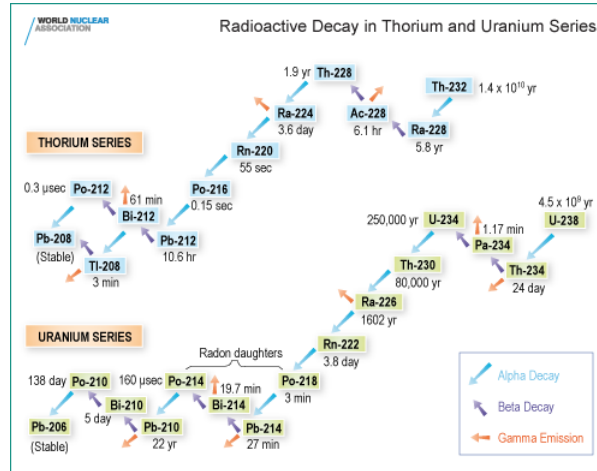


Fig. 3.13: The thorium and uranium radioactive decay chains [108].

Backgrounds can be defined as any event in the data that can be confused with a signal event. In SNO+, the general backgrounds originate from either radioactive events inherent in the detector and surroundings or from electronic disturbances in the electronic hardware.

3.9.1 Radioactive backgrounds

As with many weak signal experiments, radioactive decays in the vicinity of the detector constitute the majority of background events. Natural radioactive backgrounds are the result of two decay chains shown in Fig. 3.13: the thorium chain and the uranium chain. Table 3.2 provides the target background levels which need to be achieved in the water phase.

Both chains produce a wide range of event energies affecting the majority of studies carried out with the detector. However, the general trend is that at low energy the radioactive background level is larger than at higher energies. The analysis presented in this thesis aims to measure a low energy signal in this high background region. The reason this is possible is due to the coincidence of the potential signal with a prompt muon. This greatly reduces the number of expected radioactive background events.

3.9.2 Instrumental backgrounds

The hardware used in SNO+ is susceptible periods of instability. An example of such a period is a high voltage breakdown on a PMT base. Such a breakdown can create a chain of events which leads to a false trigger in the detector. The resulting event is completely unphysical and has the potential to contaminate physics analyses. Such events are referred to as instrumental backgrounds⁸ as their origin is the hardware itself. Similar to the radioactive backgrounds the reduction of rate of such instrumental backgrounds is important in the detectors operation.

Unlike the radioactive backgrounds, instrumental backgrounds can potentially be tagged via tests on the events topology in PMT space⁹ and timing. The process of tagging potential instrumental events is known as data cleaning; each event is passed through a number of binary checks, the result of these checks are then compared against a data cleaning mask. The data cleaning mask is set on an analysis by analysis basis. The data cleaning masks used in this thesis is detailed in sections 5.2 and 7.4.2.

⁸Sometimes referred to as instrumentals.

⁹Also known as create space.

3.10 Software

The primary software tool of SNO+ is based on the **RAT** (Reactor Analysis Tools) package originally developed by Stan Siebert for the Braidwood collaboration. The software is multifaceted, being responsible for data processing, event reconstruction, MC generation and calibration analysis.

MC generation is based on a **Geant4**[109] full detector geometry model with a bespoke physics list. Particle propagation is carried out with the **GLG4sim** package [110], with all particles being tracked. PMTs are full 3D objects with the front end and trigger systems simulated down to discriminator level on each channel. Effects such as noise hits and trigger signal jitter are included in the front end simulation.

MC is reconstructed and processed with the same codebase used for data, thereby excluding processing differences. General **RAT** output has two types: **ratds** and **ntuple**, both in the **ROOT** file format. The **ratds** data structure contains all entries for all information that can be produced, the **ntuple** is a flat **TTree** containing a subset of information contained in the **ratds**.

3.10.1 RATDB

RATDB is a central Postgres database used to store information such as calibration constants and reconstruction PDFs relevant to the **RAT** job at hand. Loading from this database allows simulations to run with the latest information, which is the default behaviour. However, it is also possible to freeze the database in tagged releases. Tags can then be specified at runtime to consistently load the same detector setup. Tags

are typically used with analysis iterations to make sure that comparisons are valid.

Adding information to the database is subject to very strict entry conditions, which exclude the possibility of having conflicting information for a given run. An example of this is the `run_range` field that governs which table is valid for each run.

3.11 Summary

The SNO+ detector is a world leading neutrino observatory with the ability to measure a world leading $0\nu\beta\beta$ limit in ^{130}Te . This chapter reviewed all of the major steps that will be taken to get to this result. The versatility of SNO+ was shown by reviewing a selection of physics topics that can also be analysed throughout the phases of operation.

The extensive calibration campaign required for successful operation was reviewed as well as the workings of the DAQ system. Finally, the event reconstruction and software resources were explained. Special attention has been given to topics which are relevant for the remaining chapters.

Chapter 4

Channel Software Status

Channel Software Status (CSS) is a procedure inherited from SNO, which looks to identify channels that can not be relied upon to provide accurate data. These channels are then removed from the analysis chain providing confidence in the PMT level information contained and used in physics event reconstruction. This ultimately leads to smaller systematic uncertainty on physics results. CSS can also be used to identify channels that may require hardware repairs which can then be carried out.

CSS is considered a service task on SNO+ and has had two previous developers. When the author inherited the task, the first steps towards integrating CSS into RAT had been completed [104]. This chapter details the work carried out by the author in completing the integration of CSS into the SNO+ processing chain. This includes defining a cut tuning methodology, adapting the efficiency of the inherited code and implementing an automatic calibration system on the nearline machine.

4.1 Introduction to CSS

CSS aims to monitor and flag channels that have degraded between calibrations. Any flagged channels can then be removed from analysis (via the `PMTCalSelector`) or brought to the attention of an electronics expert. In assessing whether a channel needs to be flagged, CSS passes a channel through a number of ‘checks’. Checks are evaluated on a run-by-run basis, with the channel being assessed over a whole run.

CSS is comprised of two RAT processors. The first is referred to as `ChanSWStatusProc`, the purpose of this processor is to calibrate all of the CSS checks. As some of the checks depend on the detector state i.e. channel occupancy depends on the current channel trigger thresholds, CSS must be calibrated on a regular basis. To ensure the best charge calibration, CSS is tuned after every ECA. CSS is run over a ‘standard run’, which is defined as the first suitable physics run after an ECA. The criteria for a run to be a suitable standard run is defined in section 4.2. `ChanSWStatusProc` outputs a RATDB table, `CSS_CUTS`. This table contains a list of cut values which are subsequently used to judge whether a channel should be flagged¹.

With the suitable CSS constants uploaded to RATDB, all runs need to be assessed for unstable channels. This assessment is carried out with the second processor referred to as `ChanSWStatusCalib`. Like `ChanSWStatusProc`, the `ChanSWStatusCalib` processor calculates the same check statistics and compares a channel’s statistic against the relevant cut value. The results of all checks are stored in a channel word. These channel words are then outputted to RATDB table, `CSS_RESULT`.

¹The parametrisation of these criteria are stored in another RATDB table, `CSS_CUT_VALUES`.

Table 4.1: Standard run definitions for both runs before and after CSS table generation came online (run number = 114631).

Standard run criteria	run number < 114631	run number \geq 114631
Physics run	✓	✓
Online channels > 7000	✓	✓
Run direction > 50 minutes	✓	✓
Data quality checks		✓

This table is stored in `RATDB` and read by the `PMTCalSelector` when processing data. Any given check can be used to remove a channel from the analysis chain, but as detailed below only selected checks are used to remove channels. These checks are specified by a bitmask. This bitmask is built from a human readable `RATDB` table, `CSS_SELECTION`.

The checks implemented by CSS fall into two categories, individual channel checks and population based checks. Channel checks assess the channel's behaviour to that of the same channel in the standard run. The population based checks calculate a test statistic for all channels. The distribution of these test statistics are then used to find outliers.

4.2 Standard runs

Standard runs are reference runs that are used to find cut values that go on to be used by `ChanSWStatusCalib`. As CSS is concerned with flagging bad channels in physics runs a standard run should itself be a physics run. CSS relies on data from channels, therefore a standard run should last long enough to maximize the chance for a channel to trigger and have enough channels online to ensure good coverage. The standard runs are selected on the nearline machine, at this point in the processing chain a suite of low level data quality checks are made by a completely independent process. It was decided that a standard run would also need

Table 4.2: List of standard runs used to generate `CSS_CUTS` tables in the water data. This list is correct at the time of writing. An up to data list can be found at <https://snopl.us/monitoring/css-proc>.

Standard Runs							
100002	100206	102210	104619	105235	105841	108309	105482
106248	111254	111467	105813	103886	107516	103417	109349
103577	103144	100244	106830	108643	106652	112134	100585
100173	104734	101183	108783	102460	106459	102958	106893
111607	107278	104366	107670	111883	108275	105708	101482
102738	104126	110090	101687	106026	109053	100525	104995
112364	114306	115824	112565	112746	112980	113116	113343
113573	113848	114010	116307	116546	116031	117572	117306
118508	118900	117833	117895	118288	116707	118652	118087

to pass all of these checks. CSS was brought online at run 114,631, at the time of writing it was not possible to include the data quality checks into the standard run definition for run numbers before 114,631. Future work may be required to merge the data quality checks into the standard run definition for all runs. However, each standard run selected before 114,631 was checked by hand and found to be suitable. All standard runs after 114,631 were subject to the data quality checks. Table 4.1 summarises the definition of a standard run for both runs before and after CSS table generation was brought online.

Table 4.2 lists the current standard run list for all pure water data with the criteria defined in section 4.1².

4.3 Cut definitions

To have confidence that the checks carried out in CSS are serving their purpose, well motivated methods to find reliable and repeatable cut values must be found. These cut values must be parametrised with respect to

²Latest list can be found at <https://snopl.us/monitoring/css-proc>.

the data gathered in the standard run e.g. the RMS of a statistical distribution. This is due to a lack of MC as the RAT simulation uses a simplified detector model³, it is therefore important to parametrise the distribution with parameters which are not sensitive to the desired outliers. For example, if the distribution is gaussian in nature then the full width half max (FWHM) is a better suited parameter than the corresponding fitted standard deviation as outliers do not effect the FWHM.

Population based checks, discussed in section 4.3.1, have cuts which generally find a fixed point in the distribution and step away from this point in steps until the check is flagging channels of interest.

The second category is channel based, which looks to compare a characteristic of a channel to the same characteristic in the standard run. For example if the distribution of a channel's charge changes from the standard run distribution then something may have happened to the channel that may require attention from an electronics expert. In general, a channel flagged by a channel based cut doesn't need to be removed from an analysis as individual PMT charges are used minimally in event reconstruction.

The exact form of these cut values can only be motivated after observing a large amount of data. As such the data used to tune the cuts consists of 434 physics runs spanning the run range [100000 - 107436].

³The RAT detector model fails in that it doesn't simulate 100% of the detector behaviour. This is justified by tools such as CSS which look to reduce simulation short comings.

4.3.1 Population based

High Occupancy

The high occupancy check aims to flag channels with hit activities far higher than expected from a well calibrated channel. This increase in activity is often caused by high channel noise rates⁴. Noise rates are dependent on many variables including PMT temperature and operating voltage. However, the high noise channels targeted by the high occupancy check will typically arise due to some electronics issue on the channel e.g. a blown resistor on the PMT base.

High noise rates can lead, in the worst cases, to real physics events being contaminated by instrumental backgrounds, ultimately leading to an increase in the reconstruction energy scale and potential loss of real physics. To estimate the magnitude of the effects of high noise channels a simple MC study was performed. Fig. 4.1(a) and 4.1(b) show how the mean and spread of the nhits distribution in a simulated ¹⁶N calibration changes with a number of artificially injected high noise rate channels. Both the mean and standard deviation of the distribution increases linearly with the number of high channels⁵. Fig. 4.1(c) and 4.1(d) show the corresponding reconstruction energy mean and standard deviation for the same simulation. The spread is seen to increase, whereas the mean energy estimation decreases. This is initially confusing, as the energy estimation is directly proportional to nhits. However, the energy estimation method takes noise rates into account through an independent

⁴SNO+ PMTs operate at ≈ 500 Hz, caused by dark noise hits [111].

⁵That is the mean and standard deviation of a gaussian fitted to the distribution.

measure, this correction introduces the observed slope in Fig. 4.1(c)⁶. Still it can be seen that 100 high channels can lead to $\approx 0.2\%$ energy shift and $\approx 2\%$ increase in energy resolution at ≈ 5 MeV.

For a given run a channel will have an occupancy, defined as the fraction of times a channel is triggered in an event compared to the total number of events. The occupancy of a given channel is dependent on the threshold for that channel, with the thresholds being set when the detector is undergoing channel threshold tuning. As such, there is a large range of occupancies which can be considered to be acceptable, but channels do exist with occupancies far and above the “good” population; such channels are referred to as screamers. To assess the noise rate of a channel, we select the PGT occupancy. A PGT is a global trigger which is artificially triggered at a nominal rate of 10 Hz. Contributions from channels in these events should be solely based on noise hits⁷. The test statistic Γ_{occ}^{high} is defined as:

$$\Gamma_{occ_i}^{high} = \frac{n_{pgt_i}}{N_{Total_PGT}}, \quad i \in \text{pmts}, \quad (4.1)$$

where n_{pgt_i} is the number of PGT hits registered in the i^{th} channel and N_{Total_PGT} is the total number of pulsed global triggers in the run. Fig (4.2) shows the Γ_{occ}^{high} distribution from run 102147 with contributions from all active channels.

⁶This correction takes channel by channel noise estimates into account. However high noise channels are unstable and therefore it is desirable to remove these channels via CSS.

⁷The probability of events at 1 kHz to pileup with the PGT rate of 10 Hz is small, approximately 0.06%

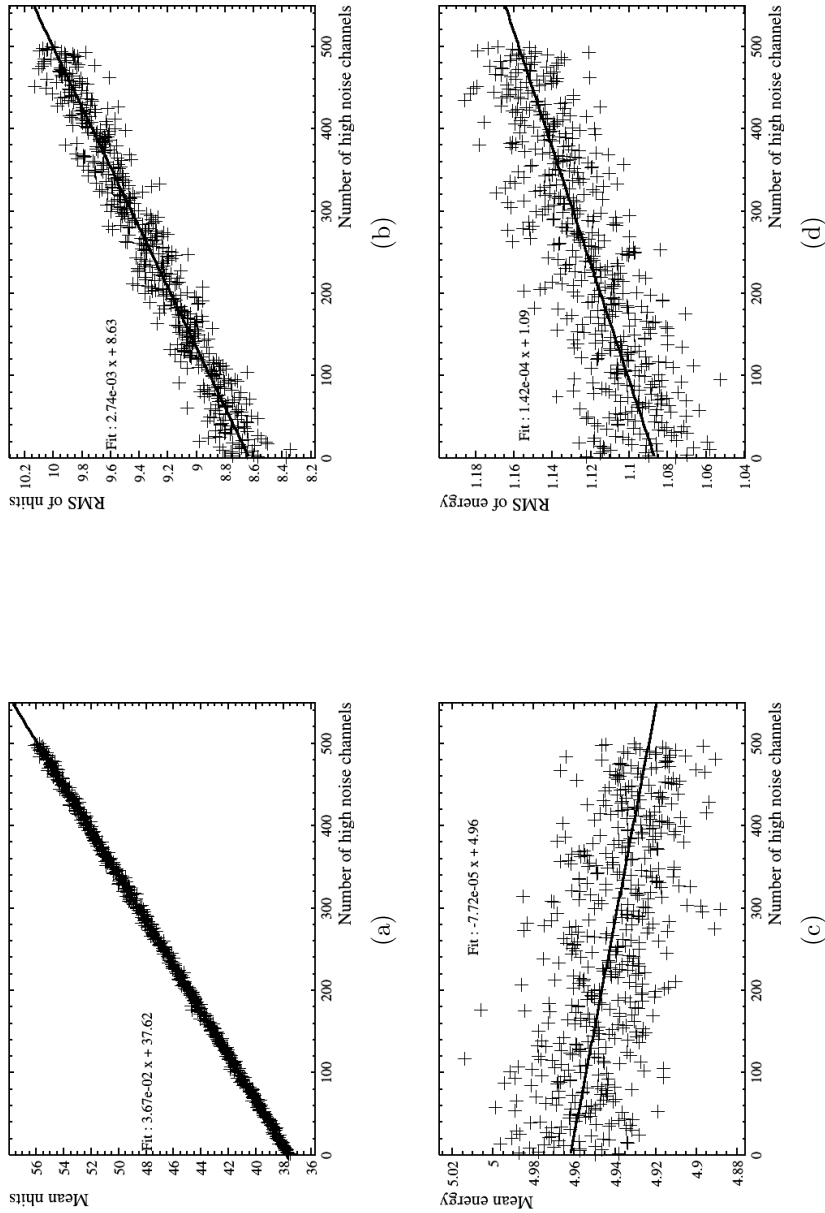


Fig. 4.1: These plots show results of a simple MC study into the effect of high noise PMTs on the performance of the detector. The MC study uses a central ^{16}N calibration run with 10^5 events as input. Fig. (a) and (b) show how the variation of the mean and RMS of the nhits distribution changes as the number of high noise tubes changes, respectively. Fig. (c) and (d) shows the mean and RMS of the corresponding energy estimation, respectively. The parameters of linear fits are also shown on each plot.




Fig. 4.2: Example of a high occupancy distribution taken from run 102147.

Cut Tuning

When parametrising the high occupancy distribution it is important to select parameters that are stable when describing similar distributions. It was found that the mean, μ , and full width half maximum, Δ_{FWHM} ⁸, were sufficiently stable to define a cut on the distribution. Fig (4.3) shows both the collected values and variation in run number of the μ and Δ_{FWHM} of high occupancy distributions, as well as the RMS of the same distributions. The μ and Δ_{FWHM} show stability on short time scales as well as values which are approximately gaussian distributed, whereas the standard deviation shows large variation.

The high occupancy cut takes the form $\mu + x \cdot \Delta_{\text{FWHM}}$, where x is the tunable variable defining the number of Δ_{FWHM} above the μ to cut. Attention now turns to what value x should take to remove only high occupancy channels. To approach this question a scan of x is carried out,

⁸The FMWH is found on distributions with the same binning as used in Fig. (4.2)

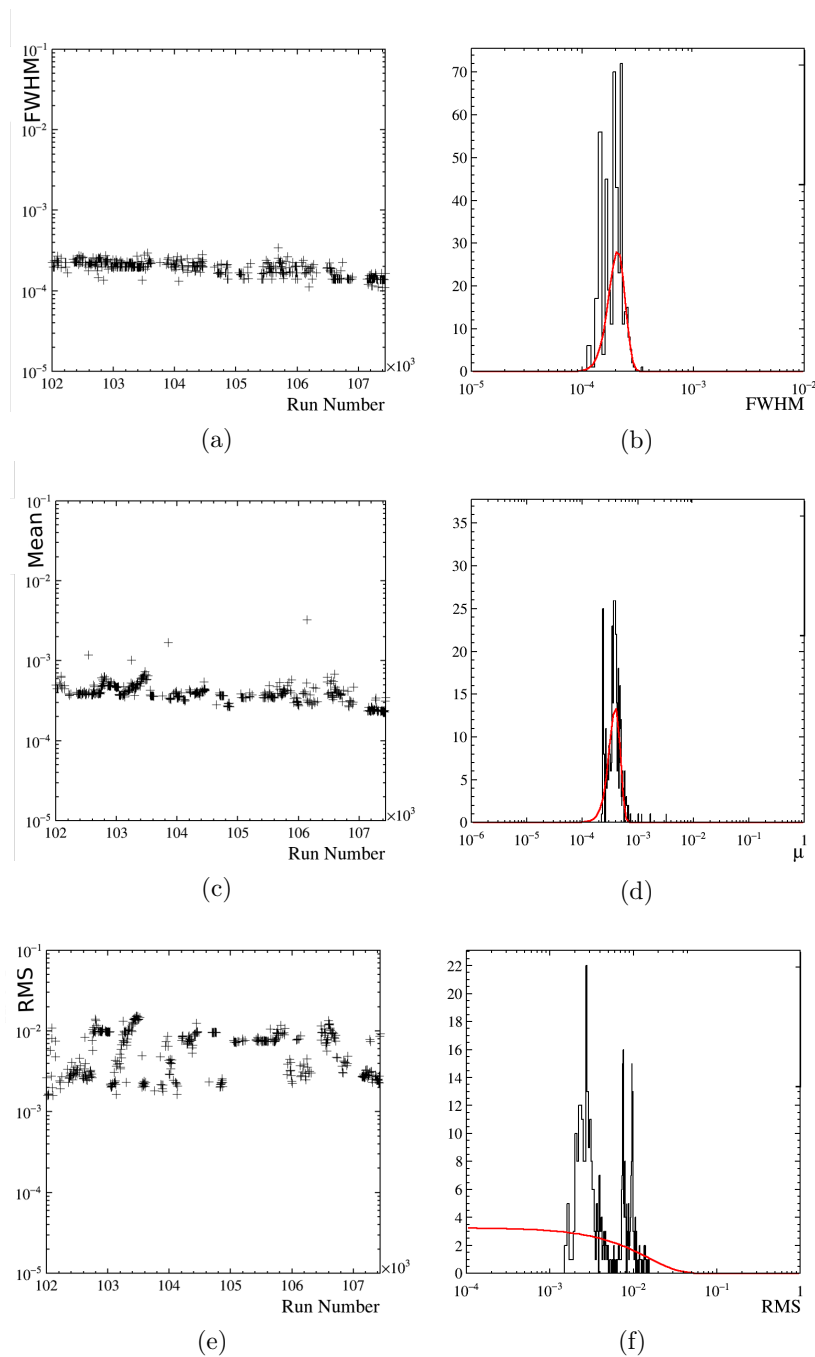


Fig. 4.3: These plots show how certain parametrisations of the high occupancy distribution vary from run to run. Figs. (a) shows the variation in FWHM of the high occupancy distribution, with (b) showing the collected FWHM values. Figs. (c) and (d) shows the same but for the mean, μ , of the high occupancy distribution and Fig. (e) and (f) shows the same for the RMS of the distribution. Both the mean and FWHM show stability, whereas the RMS shows large variations. Fig. (c) shows four outlying points, these are examples of bad physics runs which are not included in any final physics analyses. Figs. (b), (d) and (f) all show an attempted gaussian fit (red), this fit doesn't converge for Fig. (f).




Fig. 4.4: A scan of the mean number of channels flagged with a high occupancy cut value tuned with the number of FWHMs indicated on the x axis. The number of channels flagged is found by running the high occupancy cut over all runs, binning the resulting number of flagged channels and fitting a gaussian to the distribution. The fit values are then used as the mean and error on the graph.

assessing the number of channels flagged, at each x across our range. If the cut is removing channels which are constantly high then one should expect a decrease in the number of channels flagged which then plateaus before decreasing again. The cut should then be defined by the value of x for which the number of flagged channels is approximately constant. Fig (4.4) shows such a scan of x , where the x is used to tune the cut on all runs and then showing the mean and standard deviation of the resulting distribution of the number of flagged channels. The first plateau can be seen to be in the region 160–180 FWHM from μ , therefore it was agreed that the cut will be placed with $x = 175$.

Low occupancy

The low occupancy check is designed to flag channels that have substantially lower activity than expected. The removal of these channels

is well motivated as these channels are not simulated in our MC, and therefore can lead to shifts in the scale and resolution of the energy reconstruction⁹. Channels may have lower activity than expected for various reasons including bad calibration or hardware failure. As the check is designed to flag under active channels, the test statistic takes all, non artificial, triggers into account. Therefore, we define the low occupancy test statistic:

$$\Gamma_{occ_i}^{low} = \frac{n_i}{N_{Total}}, \quad i \in \text{pmts}, \quad (4.2)$$

where n_i is the number of triggers registered on i^{th} channel and N_{Total} is the total number of triggered events over the same run.

As with the high noise channels in section 4.3.1, a simple MC study can show the effect that a number of low occupancy channels will have on energy estimation. Simulating a high statistics ^{16}N source run and disabling a number of channels (not accounted for in the reconstruction), it is possible to gauge how much of an effect low occupancy channels have. Fig 4.5(c) and 4.5(d) show the mean and spread of the energy estimation. The mean energy estimation decreases with the number of low channels. With 100 low channels, the simulation predicts a 1% energy scale and an increase of 1% on the energy resolution. By flagging these channels the reconstruction will take the number of low channels into account, therefore correcting for this effect.

Unlike the high occupancy cut, the low occupancy cut value is bounded by zero and therefore great care must be taken in not generating negative cut values. Following an investigation of various methods, it was decided to generate the cut values with parameters that are produced by fitting a

⁹Position resolution will also be affected but less so due to the shorter PMT timing window.

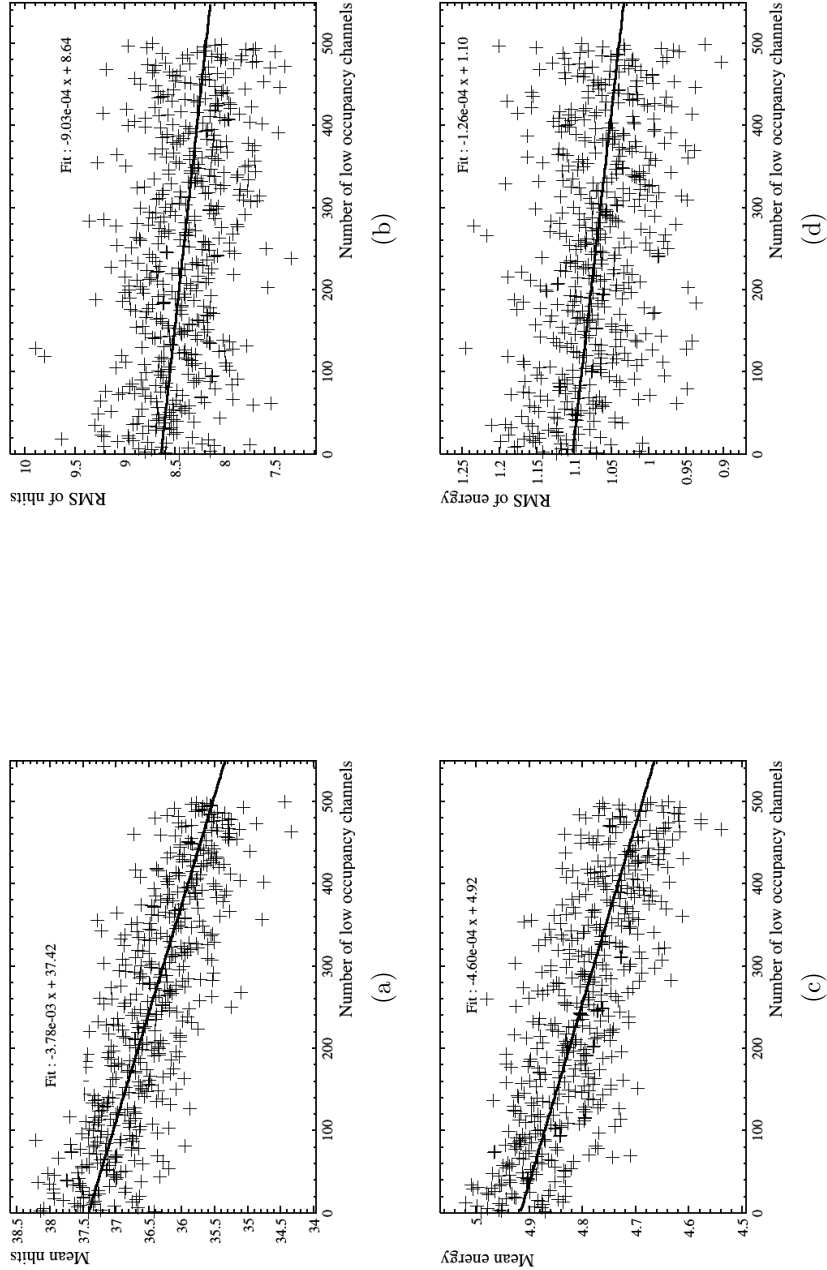


Fig. 4.5: Mean and spread of nhits, Figs. (a) & (b) respectively, and energy, Figs. (c) & (d) respectively, distributions of a ^{16}N simulation with a varying number of channels removed from the simulation. The channels are removed in such a way that the reconstruction algorithm has no knowledge of which channels are removed.

gaussian to the lower edge of the main population of the low occupancy distribution. An example of a low occupancy distribution is shown in Fig. 4.6. The raw distribution is shown in Fig. 4.6(a), it is easily observed that there is a cluster of channels below the general population and it is those channels that the low occupancy cut aims to flag. Fig. 4.6(b) shows the same distribution with the gaussian fit applied.




Fig. 4.6: An example of a low occupancy distribution taken from run 102147. Fig. 4.6(a) shows the raw distribution and Fig. 4.6(b) shows the same distribution but with the gaussian fit used to define the cut value.

Following a similar path as for the high occupancy cut presented in section 4.3.1, the low occupancy cut takes the form $\mu - x \cdot \sigma$. Where μ is the mean of the gaussian fitted, σ is the standard deviation of the same gaussian and x is the tunable parameter.

Fig. 4.7 shows the collected values of the fit parameters and their dependence on run number across the whole run range. Two distinct regions can be observed, these regions correspond to periods of running with different detector trigger thresholds. As in the case of the high occupancy cut, the fitted parameters show stability over short time spans with values approximately following a gaussian distribution.

Following the same procedure as presented in section 4.3.1, the parameter x is scanned across while accessing the number of channels cut.

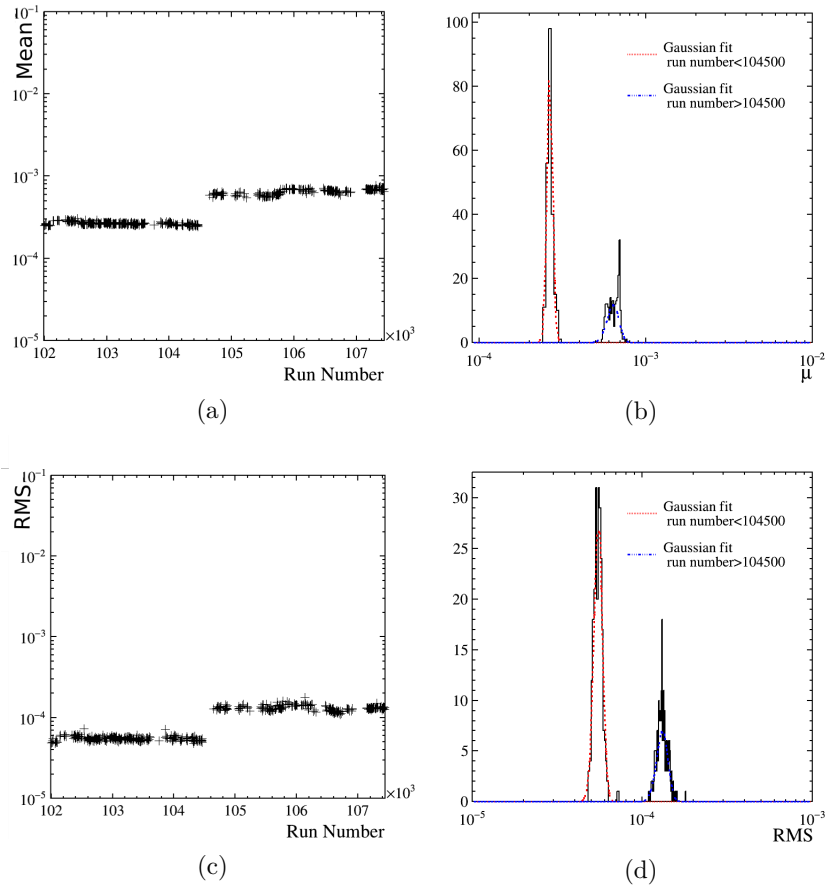


Fig. 4.7: These plots show how certain parametrisations of the low occupancy distribution vary from run to run. Fig. 4.7(a) shows the variation in mean of a gaussian fitted to the lower edge of the low occupancy distribution, with Fig. 4.7(b) showing the collected mean values. Figs. 4.7(c) & 4.7(d) shows the same but for the standard deviation of the same fitted gaussian. Two distinct regions are defined by a change in channel trigger thresholds.

The value of x for which the number of cut channels plateaus is then the cut value. Fig. 4.8 shows such a scan and motivates the decision to set $x = 3.5$ for the low occupancy cut.




Fig. 4.8: The mean number of channels flagged as low occupancy against cut value. The errors are the standard deviation of the collected number of flagged channels in the run sample.

It is possible to check that the low occupancy cut isn't just selecting channels with high thresholds. Fig. 4.9 shows all low occupancy statistics against that channel's threshold above its base noise level, that is the channel threshold in ADC counts above the dark noise level. No channels are seen to have high thresholds and low occupancies. This provides confidence in the initial electronic calibration and demonstrates that the low occupancy cut only selects low occupancy channels which are expected to have nominal activity.

4.3.2 Channel based

Channel based checks aim to compare a channel's behaviour to its output in the standard run. Typically channel based checks will take the PMT charges and TAC as inputs to the check. These low level inputs will




Fig. 4.9: The collected low occupancy values for each channel over the complete run range plotted against the channels threshold above their base noise level. Also displayed are the collected cut values (horizontal lines). It is observed that there are no channels with high thresholds and low occupancies.

favour diagnostic checks as compared to channel removing checks. At the time of writing no channel based checks have been implemented. However, outlines of the checks have been developed.

The general pattern of distribution based checks is envisaged to be comparing the bin entries in one region of the distribution to another region or to the total number of events in the distribution. These kind of checks will be applicable to all PMT charges and TAC records. In the following, the QHL distribution will be the focus in order to motivate some possible checks.

Fig. 4.10 shows the calibrated QHL distribution of the first 400 channels in a randomly selected run. The distributions are binned along the y-axis with the channel number on the x-axis. The distribution includes entries from all real events¹⁰ in which the channel was triggered. There are various features that can be observed. The great majority of

¹⁰excluding pedestals, artificial triggers.

the channels are highly peaked around 0 ADC. This is due to the large number of instrumental backgrounds present in the distributions, which typically produce hits that contain no light and therefore no charge is recorded. A simple fractional population check in the region about zero could indicate whether the channel's calibration is still accurate.

In the calibration process (ECA) each channel is artificially triggered to find the typical background charge that would be recorded in the event of a pedestal trigger. The pedestal value, typically 600 ADC, is then subtracted from each recorded charge to give the charge caused by light in the PMT. The charge values are digitised on a 12 bit digitiser, meaning the maximum values the calibrated charge record can take is ≈ 3500 ADC. These maximum charge hits are referred to as railed hits and can be consistently observed over many of the channels in Fig. 4.10. However, some channels between channel numbers 155-170 have an order of magnitude increase in probability to have a railed hit. Whereas some channels between channel numbers 20-30 have no railed hits. Both could be symptomatic of a faulty PMT or invalid calibration. A check on the portion of entries in a, to be defined, 'railed' region could highlight these channels.

There are a number of channels between channel numbers 165 and 190 which have a large number of negative charge hits. These hits are situated around ≈ -600 ADC counts¹¹. These channels were identified by hand in the water phase, however, a simple check counting the number of negative hits could also flag such channels.

Future work will focus on implementing and tuning such checks on the charge and TAC distributions.

¹¹It is presumed that these channels had a faulty ECA calibration.




Fig. 4.10: The collected calibrated QHL distributions for 400 channels for the run 105755. Many structures can be observed, with the most obvious feature being the negative charge hits between channels 165 and 190. One could envisage a check looking for these channels.

4.4 Implementation flow

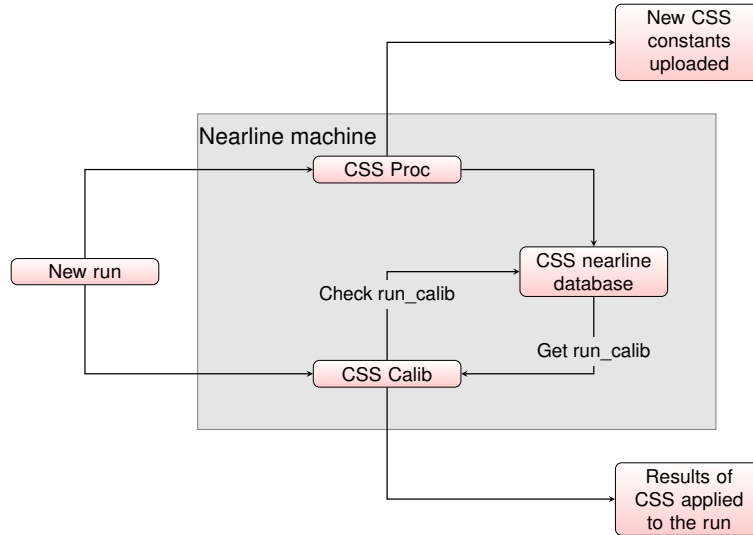


Fig. 4.11: High level overview of the CSSPROC and CSSCALIB processors and their interactions with the nearline database.

Both stages of the CSS process have been automated on the nearline machine. The implementation uses two processes with one being dependent on the output of the other. Fig. 4.11 shows a high level overview of the flow of the CSS processes on the nearline machine. For each run, the nearline scheduler will spawn both the CSSPROC and CSSCALIB processes. The CSSPROC is responsible for monitoring when new CSS constants should be produced and the CSSCALIB is responsible for applying valid cuts to each run to produce the channel cut list for each run.

Fig. (4.12) shows the logic used in the CSSPROC process. To maximise the confidence in any charge base checks, new CSS constants should be generated on the first ‘good’ physics run after an ECA, where ‘good’ is defined by the standard run requirements, see section 4.2. Nominally,

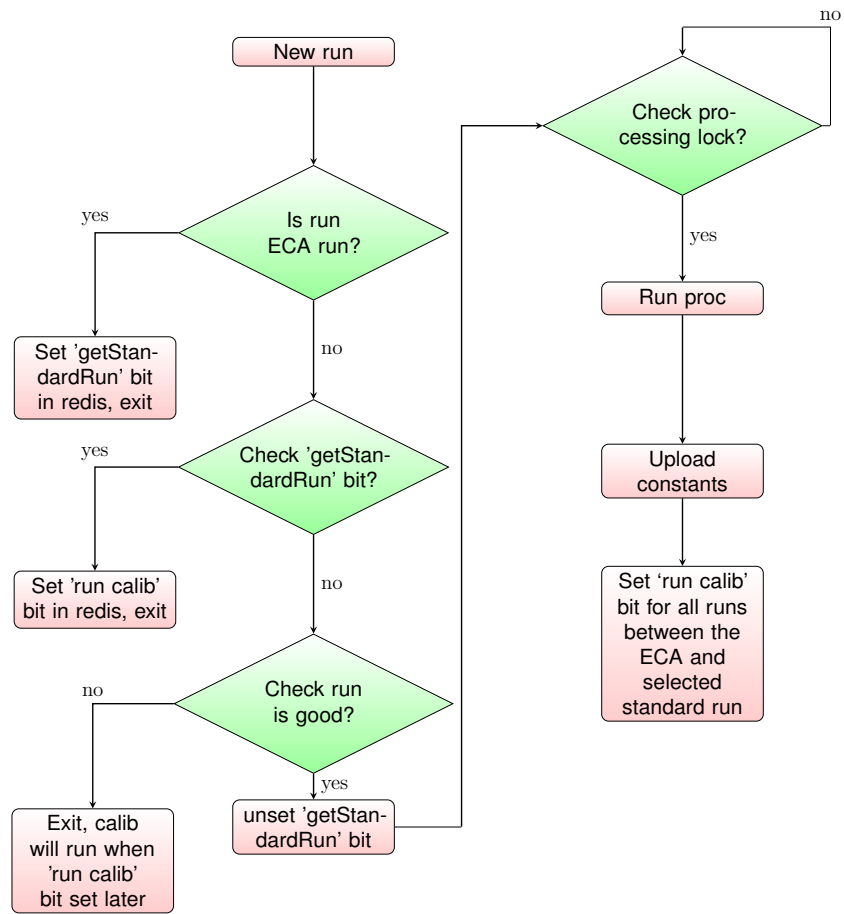


Fig. 4.12: Overview of the logic used in the CSSPROC processor.

ECA calibrations are taken once a week. Therefore, for the most part the CSSPROC process will have nothing to do and will complete without taking any action. However when a ECA run is taken the temporary bit, `'css_proc_standardRunNeeded'`, is set in the redis nearline database. This bit then allows any subsequent CSSPROC processes to search for a standard run fitting the criteria laid out in section 4.2. When a suitable run is identified and the new ECA constants are uploaded, the `ChanSWStatusProc` RAT processor processes the run. After some checks are made on the output, the subsequent `CSS_CUTS` table is uploaded to RATDB.

Upon completion of the CSSPROC process, an entry is pushed to the `cssproc` table in the nearline database. This entry records various pieces of information including the flag `run_calib`.

It is important that if new constants are produced and uploaded, they are used when applying the cut values to the run. For this reason the `CSSCALIB` process waits for the `run_calib` bit to be set in the relevant database entry. When `CSSCALIB` finds the `run_calib` bit has been set, it is released to process the run with the `ChanSWStatusCalib` processor, pushing the resulting `CSS_RESULT` table to RATDB.

There are various situations in which CSSPROC will complete without setting the `run_calib` bit for the run in question. One such case is when a physics run takes place after an ECA run but is not deemed good enough to be a standard run i.e. less than 50 minutes long. This run still needs CSS to be applied to it. However the latest CSS constants will not have been uploaded yet. Therefore the `run_calib` bit isn't set. Only after a standard run is found and processed does CSSPROC set the `run_calib` bit for all runs between the latest ECA and the standard run.

The waiting CSSCALIB processes can then be processed with the new CSS constants.

One edge case that should be highlighted is the possibility of a dummy ECA run, e.g. a testing ECA run. This would cause an unnecessary CSS recalibration. It is impossible to guard against an unplanned false ECA run, however, to accommodate testing, CSS recalibration can be switched off via a flag in the redis database.

4.4.1 Reprocessing

Historically, the incarnation of CSS used in SNO was a tool developed in reaction to observations made about the detector electronics. Checks were developed at different points in time and the same should be expected for SNO+. Any new checks will then need to be applied to all of the runs retrospectively. Any change to the standard run list will also require a full CSS reprocessing. However, reprocessing is not trivial as the initial generation of CSS constants depends on the processing lock which isn't available when reprocessing. Details concerning RATDB such as run validity and pass number also complicate matters. Interaction with the CSS table on the nearline database makes reprocessing possible and scripts have been written to aid reprocessing¹².

The general procedure for carrying out reprocessing is for the expert to locally download all of the raw standard run files, produce and upload the new `CSS_CUTS` tables. Then reprocess all of the data on the grid with a custom macro, which will update the run-by-run `CSS_RESULT` tables. Reprocessing will then be complete with the results placed in RATDB ready for reading into `PMTCalSelector` when needed.

¹²These scripts are placed in the CSSPROC nearline source folder.

Interaction with RATDB

The CSS relies on tables stored in RATDB. The `CSS_CUTS` table stores various cut values extracted from a standard run. The `CSS_RESULT` table is a run by run table storing the results of applying the checks to the run in question.

Each table has a `run_range` field which specifies the range of runs for which the table can be used. The importance of this `run_range` is that any updated table must be entered into the database with the exact `run_range` present in the current table.

The `CSS_CUTS` tables are generated and applied on a run by run basis, therefore these tables have a simple `run_range` including only the run number in question.

The `CSS_RESULT` table is different. It stores cut values to be applied over a `run_range`. The `run_range` used to store `CSS_RESULT` tables is defined as the last ECA run of the calibration in question until the run before the last ECA of the next calibration¹³. An important edge case is the latest `CSS_RESULT` table, which, by definition does not have a next ECA campaign. In this case the table is pushed into the database valid from the last ECA run to `RATDBINFINITE`¹⁴. The pass number for these tables must be set to -2. When a new ECA run is taken, this table's `run_range` will be clipped down. This has the effect that if this table needs to be updated, the replacement table has to have the exact `run_range` with a pass number set to -1. See section 4.4.1 for details of the reprocessing framework implemented with current scripts.

¹³ECA calibrations are composed of multiple runs, for the rest of this document an ECA run number will refer to the last run in the ECA campaign.

¹⁴`RATDBINFINITE` is a large number used as a reference to the end of database. In the special case that the `run_range` includes `RATDBINFINITE` then the `run_range` for this table can be shortened by uploading another table with a `run_range` up to `RATDBINFINITE`.

If at any point a `run_range` is incorrectly defined e.g. a fake ECA run is taken while automatic recalibration is active, then the only way to fix this is to increase the version number of the table before reprocessing and push the new tables to `RATDB`. This increase in version number is only required for the incorrect tables in question, therefore all tables do not need to be uploaded again. However, once the version number is increased, as a rule of thumb, any reprocessing should be carried out with the new version number.

4.5 Effect of CSS on data

It is important to observe the numbers of channels flagged and ultimately removed from the data. This provides reassurance that CSS is removing only the channels it is designed to remove. For an adequately calibrated detector, and well tuned cuts, the number of channels flagged is expected to be small.

4.5.1 High occupancy

Fig. 4.13(a) shows all high occupancy distributions projected vertically against run number. The black lines show the cut value that would be applied and the bottom plot shows the number of channels flagged for each run. As expected, the number of channels flagged is of order ~ 10 , reaffirming the detector calibration and stability. Fig 4.13(b) shows the same data but with channels flagged exclusively by the high occupancy check, that is channels which are not already flagged by `PMTCalSelector`¹⁵. It can be observed that \sim two extra channels are flagged by the high occupancy check.

¹⁵The `PMTCalSelector` takes input from many tests (ECA, PCA etc).

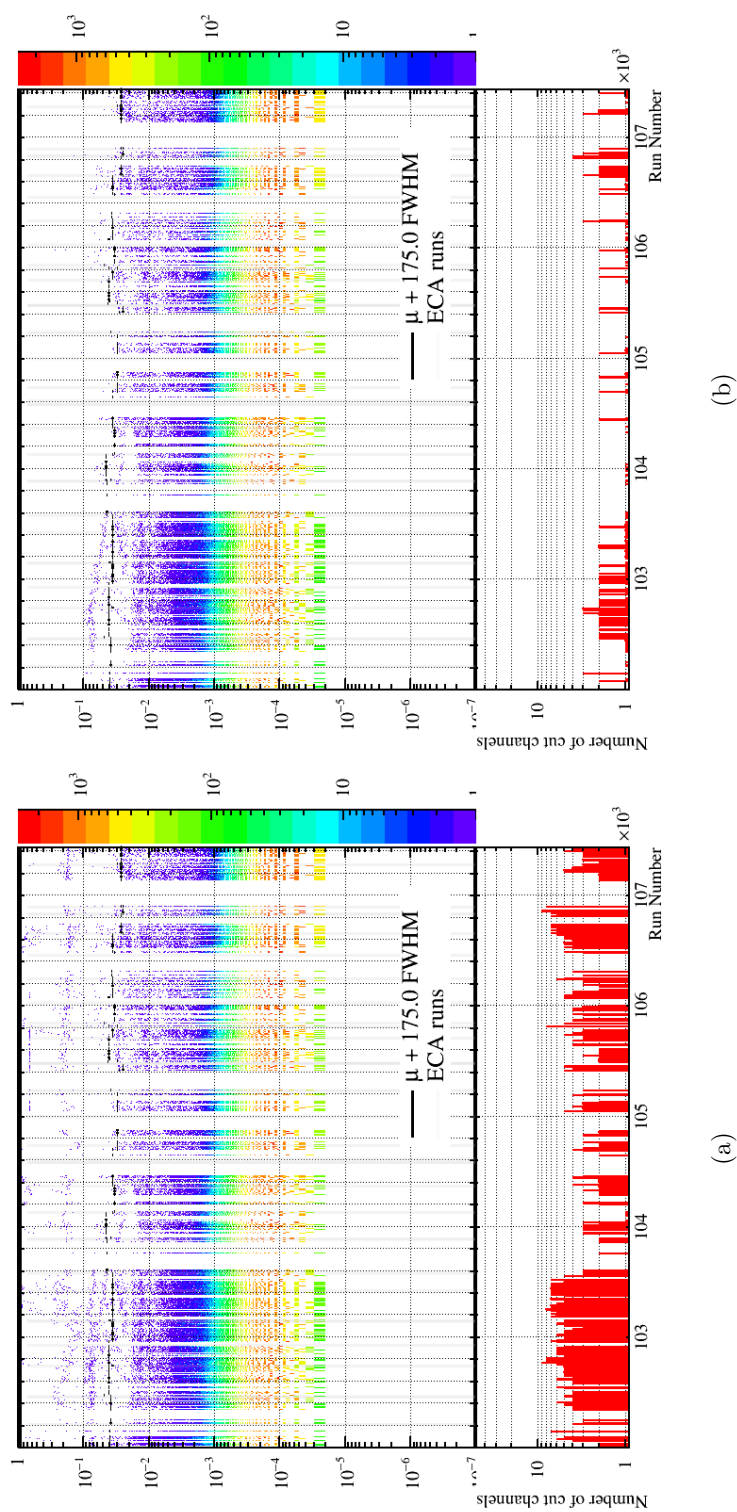


Fig. 4.13: High occupancy distribution for a number of runs with the CSS cut value shown by the black lines. The bottom plot shows the number of channels flagged in that run. Fig. 4.13(a) includes all channels whereas Fig. 4.13(b) shows only channels which pass all of the other checks included in the `PMTCalSelector`. The gaps in run number are due to periods of non-physics running.

Channels to be excluded due to the high occupancy check are expected to have high occupancies for extended periods of time, therefore these channels should be flagged repeatedly. Fig 4.14 shows all channels which are flagged over the run range with the run in which they are flagged shown vertically. Vertical lines on the plot represent the channel being continuously flagged. It can be observed that the majority of channels are flagged multiple times, providing confidence in the effectiveness of the cut. It can also be observed that over the run range only 44 channels are flagged.




Fig. 4.14: All channels flagged by the high occupancy check over the run range and the runs in which each channel is flagged. The white space shows that the channel didn't register as high occupancy, there are many reasons for this including channel being removed from the detector or being repaired.

4.5.2 Low occupancy

Fig. 4.15 shows a plot similar to Fig. 4.13(b) but for low occupancy distributions; that is channels flagged exclusively by the low occupancy cut. Again the horizontal black line lines show the cut value applied and the bottom plot shows the number of channels flagged by the low

occupancy cut for each run. The vertical shift that occurs around the run number about 104,600 is due to a change in the channel trigger threshold.




Fig. 4.15: Low occupancy distribution for a number of runs with the cut value imposed. The bottom plot shows the number of channels flagged exclusively by the low occupancy check in that run.

Fig. 4.16 shows the low occupancy analogous to fig 4.14. The two blocks between the channels 2352–2366 and 5768–5791 can be attributed to two badly calibrated FECs, which were repaired. A total of 140 channels were flagged over the run range considered.

4.6 CSS over all water data

CSS was introduced into the processing chain through the `PMTCalSelector`, a utility which removes channels based on a number of inputs of which CSS is one. As CSS was finalised independently from the other inputs, there will be channels that are flagged by multiple checks. For this reason




Fig. 4.16: All channels flagged by the low occupancy check over the run range and the runs in which each channel is flagged.

the effect of introducing CSS is quantified by the channels which are flagged solely by CSS. These channels are identified by comparing flagged channels before and after CSS is enabled.

With CSS fully integrated it is possible to collate the number of channels flagged across a much larger run range than that used to tune the cut value optimisation in section 4.5.

Fig. 4.17 shows the number of channels flagged by the high occupancy check (top row) and low occupancy (bottom row) for both the absolute number of channels flagged (left) and the number of channels flagged solely by CSS (right). All of the runs shown are golden physics runs, see section 3.7.

The high occupancy cut flags an extra 3 channels per run on average. However, tens of channels are flagged in some runs. 86 channels were flagged as high occupancy over the run range. The low number of high occupancy channels in any given run will have a negligible effect on event reconstruction, as predicted in section 4.3.1, however cutting these channels removes their unpredictable output and increases the confidence in the data recorded.

The low occupancy cut has a higher variance compared to the high occupancy cut. Some structures can be attributed to whole FEC boards developing a fault and the relevant channels becoming underactive. It is generally the case that the number of low channels in any given run is below 50, but for a select number of runs that number rises to 100–200. Removing these channels is essential in these runs, and will allow the reconstruction methods to account for missing hits, thus improving the fitter’s performance.

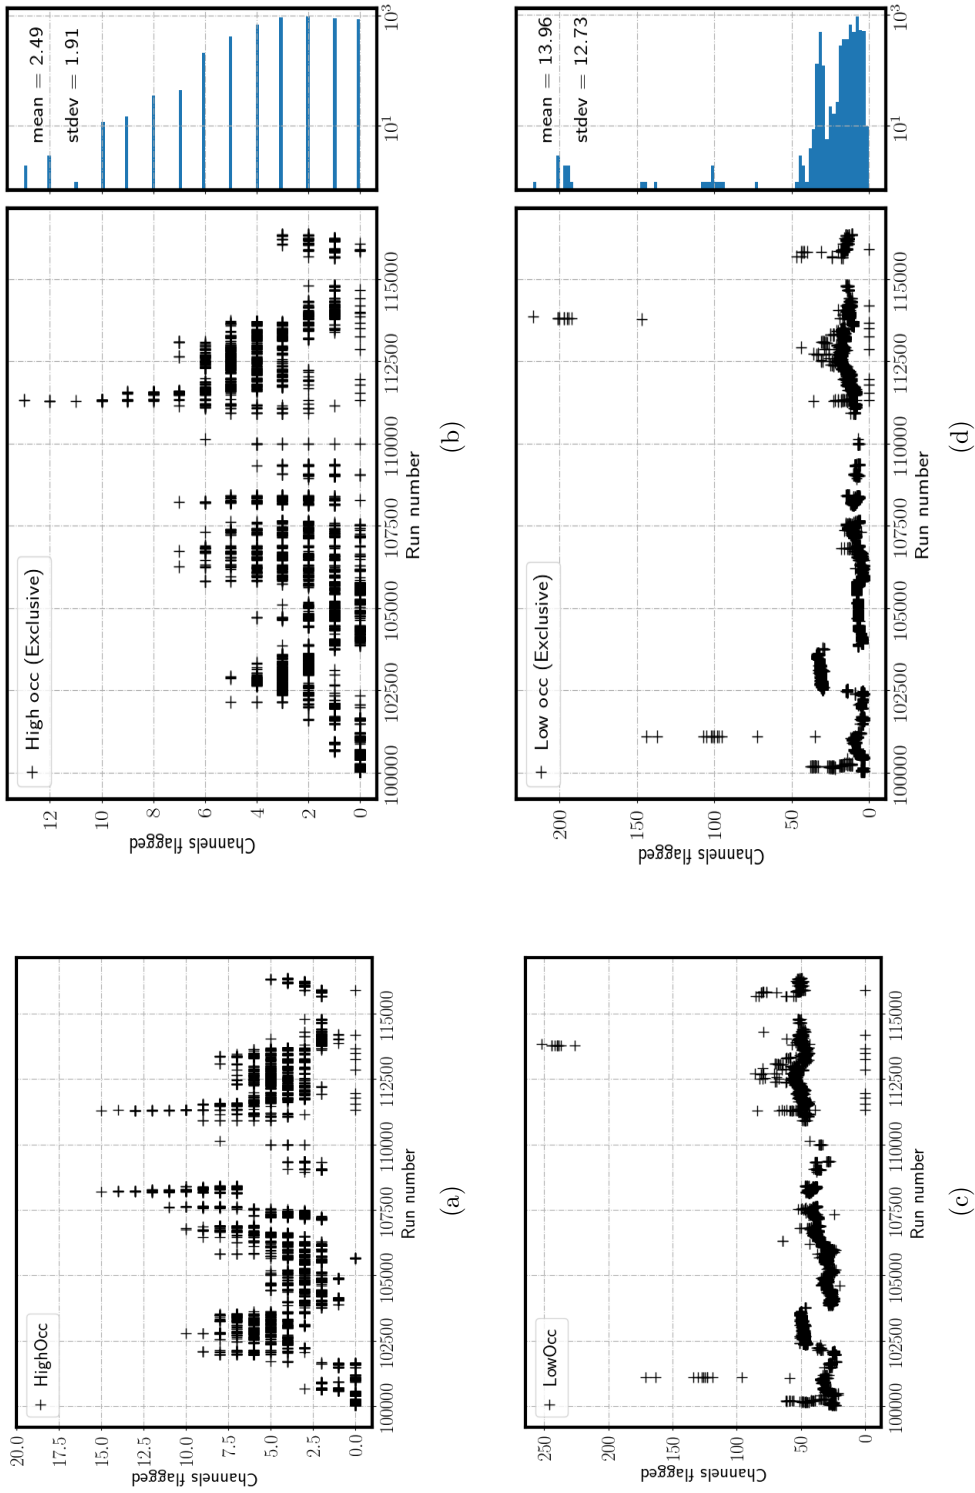


Fig. 4.17: High and low occupancy cut performance over golden water physics run list. (a) and (c) show the total number of channels flagged by the high and low occupancy cuts respectively. (b) and (d) show the number of channels flagged exclusively by the high and low occupancy cuts respectively. That is the channels that are cut due solely to the cut in question.

4.7 Conclusion

It is the purpose of CSS to flag channels that have degraded between calibrations. This degradation can lead to the data from the channel being unreliable. In the cases where the channel is not reliable the channel is removed from the list of PMTs that feed the reconstruction algorithms used in SNO+. As shown in section 4.3 tens of reliable channels can lead to a percentage level systematic biases in reconstruction results. Removing these channels increases confidence in the reconstruction results and ultimately smaller systematic error estimates on physics results. In this chapter, details of the work done to fully integrate CSS into RAT and the SNO+ processing chain are given. Two PMT removing CSS checks were motivated, implemented and tuned. The high occupancy check has been shown to remove ~ 3 extra channels from the SNO+ processing chain. The low occupancy cut has been shown to remove ~ 13 extra channels, although there is a standard deviation of ~ 12 on this number. In both cases there are runs in which a larger number of channels will be removed. Both these number of channels are on the order expected and these two checks will be integrated into any reprocessing of the SNO+ data. The automatic generation of the operating cut values used in SNO+ has been explained in detail, including both the first pass updating and any reprocessing that is needed in the event of a new check coming online. Future work on CSS will including implementing charge based checks and implementing any checks that will be needed in the scintillator phase. Both the high and low occupancy checks are included in the analysis presented in the remainder of this thesis.

Chapter 5

Data Processing

5.1 Raw data set

The SNO+ detector started taking data on the fourth of May 2017. However the low energy search that is presented in the remaining chapters requires data after the trigger threshold change that took place in August 2017.

The raw data analysed in this thesis is taken from the golden run list. The dataset spans the run range 104,617 to 116,519 and has a raw livetime of 260.9 days.

It was also required that all runs had less than 20 data cleaning muon tags. This removed five runs from the data set. The nominal muon rate of 3 per hour provides a Poisson probability of $\approx 1.18 \times 10^{-11}$ for 20 muons to occur. This removed five runs from the data set, these runs were probably contaminated with bursts of electronic breakdowns that could be mistaken as cosmic muons. The highest number of muons found in one run was twelve¹.

¹It is almost certain that these twelve tags are not true muons: it is the goal of the rest of this chapter to remove the potential false tags.

5.2 Data processing

The low energy muon coincidence analysis presents a number of challenges with respect to data processing, these challenges are not satisfied by the default data processing therefore a custom processing is required.

The custom processing was set up to reconstruct all events above six PMT hits within 20 seconds of a muon, where the prompt muon is flagged by the data cleaning “muon tag”². As introduced in chapter 3, data cleaning is a suite of tests, which attempt to flag events that are potentially undesirable e.g. instrumentals or muon events. Each muon following event is passed through the water fitter regardless of data cleaning.

One of the benefits of carrying out a coincidence analysis is that the background rate can be assessed in data before and after the prompt signal. The background samples are taken on a run by run basis and will be used to find the expected background rate in each run for each muon³. As the neutron signal is low energy, the background sample should also focus on low energy events. This presents a problem as the sheer number of background events at low energy is too large to process. Therefore 10% of the background events that pass a data cleaning check, with $6 < \text{nhits} < 15$ were processed. The data cleaning mask was taken from the standard processing chain with the muon specific checks removed, the resulting mask is `0x73ffe`⁴. 10% of the data still provides more than enough data to have sufficiently low errors after scaling the background

²The purity of this flag is low as the threshold to be tagged as a muon is low. However the muons tagging efficiency is high, expected to be close to 100%, although an exact efficiency is currently unavailable.

³The signal selection is muon dependent.

⁴The data cleaning mask is a set of bit which dictate which data cleaning checks are applied to events in an analysis, more detail is given in section 3.9.2.




Fig. 5.1: Raw detector event rate following all events tagged as muons by data cleaning. (a) and (b) show the events over a period of 1 sec and 100 ms respectively. A large number of triggers are seen at $\Delta_{\mu}t < 30$ ms.

estimate by 10. A background event is selected at processing level by any random event not following an event tagged as a muon by data cleaning.

5.3 Raw processing results

Fig 5.1(a) shows the raw event rate in two time windows following all of the muons in the processed data set, where $\Delta_{\mu}t$ is the time difference between the muon's and event's trigger times. Fig 5.2(b) shows the 100 ms time spectrum. There is clearly a large number of triggers which stretch out to ~ 30 ms after the muons. Fig 5.1(a) shows the same data in an extended 1 s window. The event rate is observed to stabilise after ~ 40 ms representing the nominal detector background event rate.


Random background comparisons

As the random background processing has a preselection on nhits and data cleaning, it is only possible to compare the post muon data after the same selection. The background spectrum is constructed from ten random event windows for each muon. With the background windows sampled from each run in proportion to the number of muons contained


in that run. Fig 5.2 shows this comparison over three time ranges, 1s, 100ms and $1\mu\text{s}$. Fig 5.2(a) shows the comparison over 1 s, the background is constant as expected with the muon trace agreeing at $\Delta_\mu t > 100\text{ms}$. Fig 5.2(b) shows the 100 ms window in which there is a clear disagreement at $\Delta_\mu t < 30\text{ms}$. The reason for this is hard to explain but is expected to have a negligible effect on a neutron yield analysis as an upper time selection of $\mathcal{O}(1\text{ms})$ should remove this. It should be noted that the non-constant event rate observed in Fig 5.1(b) appears to be considerably reduced by the inclusion of data cleaning and the nhit selection.

Fig 5.2(c) shows the spectra on the order of the expected timing selection. As thermal neutrons have a capture time of $209\mu\text{s}$, a falling excess of events should be expected. At first glance this seems to be present in the muon data which once again falls to the expected background at large times. A more detailed analysis of Fig 5.2(c) shows an excess of events which disappears below about $30\mu\text{s}$, this is not the expected behaviour for a neutron capture excess⁵. Nevertheless the excess is motivation enough to develop the analysis to a point to obtain a pure neutron candidates timing spectrum.


⁵Neutron thermalisation takes place on $\mathcal{O}(10-100)\text{ns}$ and therefore doesn't explain the dip.



(a)



(b)



(c)

Fig. 5.2: Time comparisons of data both with and without a prompt muon over three different time windows. The random background data is normalised to account for the 10% prescale and is sampled from each run in proportion to the number of muons in that run. Fig 5.2(a) shows the event rate in a 1 s time window after the prompt event. Fig 5.2(b) shows the event rate in a 100 ms time window after the prompt event. Fig 5.2(c) shows the event rate in a 1 ms time window after the prompt event.

Chapter 6

Simulation and Reconstruction of Cosmic Muons

This chapter details a how cosmic muons are simulated and reconstructed in the water phase of SNO+. Much of the work follows in the footsteps of the SNO experiment; the expected cosmic muon energy spectrum is recycled as an input and the reconstruction method described below follows the SNO procedure closely.

6.1 Cosmic Muon Simulation

The SNO experiment conducted a surface to detector MC study producing a two dimensional PDF in $\cos(\theta_z)$, where θ_z is the zenith angle of the muon, and muon energy (E_μ). The study was carried out with the MUSIC software [84] and the resulting PDF is shown in Fig. 6.1. SNO+ starts generating cosmic muons by sampling from Fig.6.1. The PDF is sampled




Fig. 6.1: Cosmic muon energy direction correlation, as found for the SNO experiment, predetermined with Gaisser's parametrization at surface, passed through matter with the MUSIC simulation software [112].

using the acceptance-rejection method and results in a muon energy and direction.

Clearly all muons with the same $\cos(\theta_z)$ have an equal probability to enter the detector at any point on the PSUP. To generate the entry positions, the muon is postulated to pass directly through the centre on the PSUP. The direction vector is used to construct a circular plane perpendicular to the muons path. A random point, \vec{x}_{new} , on this plane is then sampled. Using Pythagoras theorem, the point at which the muon first crosses the PSUP is found. This point is then randomly rotated azimuthally to generate the postulated entry position. Fig.6.2 shows a schematic representation of the generation process. As a result each generated muon has an entry position on the PSUP and momentum, these parameters are passed to GSim and the muon is propagated through the full RAT simulation.

The simulation is started from the inner surface of the PSUP, at a radius of 8.5 m. This reduces the simulation time per muon as light




Fig. 6.2: Schematic overview of the generation of muons in RAT. The first step is to sample a muon momentum from the distribution shown in Fig. 6.1, this muon is then postulated to travel through the centre of the detector, subfigure 1. This muon track is then randomly shifted along the perpendicular circular plane, subfigure 2. A simple geometric construction can then be used to find the point on the inner surface of the PSUP this muon will have crossed, subfigure 3. This point is then rotated by a random angle in ϕ . This point is then used as the starting position and momenta are then passed to GSim in RAT to complete the muon simulation.




Fig. 6.3: Three dimensional distribution of 10000 muons generated along the direction $\phi = \frac{\pi}{3}$ and $\cos(\theta_z) = 0.6$ generated by the RAT cosmic muon generator. Shown is the same surface seen from four different angles. The generator can be seen to distribute muons uniformly over the inner surface of the PSUP.

generated in the cavity before the muon enters the PSUP is not propagated through the simulation. One by-product of this decision is that secondary particles produced by the muon in the cavity are not accounted for in the MC model. The effects of such secondary particles on an analysis will have to be assessed independently and accounted for as a systematic error.

Fig.6.3 shows four views of the three dimensional distribution of 10^5 muons generated with the generator. All of the muons are generated along the direction $\phi = \frac{\pi}{3}$ and $\cos(\theta_z) = 0.6$. The generated positions are observed to be uniformly distributed over the inner surface of the PSUP.

6.2 Through going muon reconstruction

Cosmic muons are distinct physical events in SNO+. SNO+ typically looks to detect point like interactions within the detector, due to the relatively short path length of the charged particles detected. Cosmic muons are an exception as they typically travel through the detector, with path lengths of several meters, emitting Cherenkov radiation at every point. This leads to large amounts of light being deposited in the detector, leading to a large number of hit PMTs. These high nhit events are quite distinct and therefore require their own reconstruction method.

Cosmic muons which reach SNO+ have energies of $\mathcal{O}(100)$ GeV as seen in Fig.6.1. At these energies, path altering muon interactions are expected to be minimal. Therefore, all muon tracks are postulated to be straight lines, entering and exiting the PSUP at two distinct points. There are two classes of cosmic muon in SNO+: through going and stopping muons. Through going muons are the norm in SNO+ but stopping muons are observed. As the names suggest the difference between the two classes

are that stopping muons lose all their energy and come to rest inside the detector before decaying producing a relatively high energy electron. The method developed in the following will not reconstruct stopping muons, as PMT hits produced about the exit point of the muon are prerequisite for the method to succeed. The stopping muons are removed from the data using the selection defined in Chapter 7.

The approximate spherical symmetry of SNO+ can be exploited for reconstruction. Taking a muon path as a straight line intersecting a sphere, one can observe that two muons which have the same distance from the centre of the detector to the point of closest approach can be related by a rotation. This distance from the centre of the detector to the point of closest approach is named the impact parameter, b_μ , of the muon.

The generally idea of the reconstruction laid out below is to assess all muons with the same b_μ on the same footing. This is achieved by defining a ‘special frame’ in which the correlations will be evaluated and each muon will be ‘fitted’. This frame is muon dependent and therefore will have a unique transformation from SNO+ coordinates. After estimates are made in the special frame, the unique transformation can be used to transform to SNO+ coordinates thereby estimating the real world coordinates of the reconstructed parameters: entry position on the PSUP and the muon direction. Such a procedure will reduce systematic errors introduced by the real world SNO+ geometry by taking many event hit patterns into account and taking the average responses as the estimate.

Fig.6.4 shows the PMT hit patterns in spherical coordinates for two muons with the same impact parameter. Fig.6.4(a) shows the hit pattern in SNO+ coordinates, the two events are seen to have different hit patterns. Fig.6.4(b) shows the same two muons after being transformed into the

special frame. The hit patterns are seen to overlap to a great extent. This demonstrates that in the special frame these two muons would be treated on the same footing.

Fig.6.5(a) shows a muon contained in the x-z plane of the SNO+ PSUP entering approximately at \vec{t}_{vec} and travelling in direction $\vec{U} \cdot \vec{t}_{vec}$ is the estimated vector to the muons entrance position. To find the entrance position, prompt PMT hits are used. The earliest 80 prompt PMT hits are selected¹, these hits are used to construct a mean of the PMT positions, \vec{x}_{80} . The angle between each of the 80 tubes and \vec{x}_{80} is then calculated, with tubes displaced by more than 26° being removed from the list. This step removes spurious early noise hits not correlated with the muon entering the PSUP. Of the remaining prompt channels, the earliest 20 are taken and the mean PMT position of these channels projected onto the inner PSUP is taken as the initial muon entrance position estimate.


\vec{q}_{vec} is the estimated exit position of the muon. This is found by taking a charge weighted average position of all of the hit PMTs in the event.

$$\vec{q}_{vec} = \frac{\sum_{PMT} Q_{PMT} \cdot \vec{x}_{PMT}}{|\sum_{PMT} Q_{PMT} \cdot \vec{x}_{PMT}|}, \quad (6.1)$$


where Q_{PMT} is the QHL and \vec{x}_{PMT} is the position of the PMT is question.

Once \vec{t}_{vec} and \vec{q}_{vec} are found, the angle they subtend, α , can be found. It is always possible to calculate the rotation matrix taking \vec{t}_{vec} and \vec{q}_{vec} from SNO+ coordinates into the special frame. The special frame is a two dimensional plane with coordinates x_s and z_s . Fig. 6.5(b) shows the muons from 6.5(a) in the $x_s - z_s$ frame. The frame is defined as the

¹Chapter 7 shows that selected muons will have $\mathcal{O}(10^3)$ nhits therefore we can always find 80 PMT hits



(a)



(b)

Fig. 6.4: The PMT hit patterns for two muons, $\mu 1$ and $\mu 2$, with the same impact parameter in SNO+ coordinates, Fig. (a) and in the special frame, Fig. (b). When in the SNO+ coordinates the patterns are separable, however the pattern can be seen to overlay each other in the special frame.

frame in which \vec{q}_{vec} points in the $-z_s$ direction and \vec{t}_{vec} being completely contained in the x_s axis.

Once the system has been rotated into the special frame, \vec{t}_{vec} and \vec{U} only require a single parameter to be characterised. In the case of \vec{t}_{vec} , Z_r is chosen, with $X_r = \sqrt{r_{\text{PSUP}}^2 - Z_r^2}$. \vec{U} is defined by the zenith angle, θ_z , it subtends. Fig. 6.5(c) shows the definitions of parameters Z_r and θ_z relating to \vec{t}_{vec} and \vec{U} respectively for the same muon as in Fig.6.5(b).

The ultimate aim of this reconstruction step is to relate the estimated angle α to the expected Z_r and θ_z in the special frame. A study can be made to find the expected Z_r and θ_z from rotating the MC truth muon entrance and direction vectors into the special frame. Fig. 6.6 shows the correlations between estimated $\cos(\alpha)$ and the true Z_r and θ_z variables.

6.3 Calibrating the fitter

The reconstruction method described above is dependent on the correlation between the estimated $\cos(\alpha_{QT})$ and the expected observables in the special frame. As is typically the case with calibration, MC truth information can be used to characterise the calibration curves. In the case of the muon fitter, the MC truth information is used to find the expected observables in the special frame.

To find the calibration curves, 10^5 randomly distributed muons were simulated with RAT². Once the Cherenkov light is registered in the simulation, the $\cos(\alpha_{QT})$ (the dot product between the \vec{t}_{vec} and \widehat{qVec} vectors) of the event is found alongside the true muon position and direction

²As will be expanded on in chapter 7, these simulation were set to effectively only simulate the Cherenkov radiation from the muon itself. This was achieved by tracking particles for 500 ns after their creation.




Fig. 6.5: Transformation from SNO+ coordinates to the special fitter frame used to estimate the muon path. Fig. (a) shows an example muon contained completely in the $x-z$ plane of SNO+. The vectors \vec{t}_{vec} and \vec{q}_{vec} point to the entrance and exit points of the muon. The angle α between \vec{t}_{vec} and \vec{q}_{vec} is of course preserved when rotating from SNO+ coordinates to the special frame. Fig. (b) shows the same muon transformed into the special frame. In this frame the \vec{q}_{vec} points along the $-z_s$ axis. Fig. (c) shows the definitions of variables Z_r and $\cos(\theta_z)$ which are used to set up the final reconstructed entrance and direction vectors in the special frame.




Fig. 6.6: Muon fitter calibration curves relating the observable $\cos(\alpha_{QT})$ to the special frame coordinates; $\cos(\theta_z)$ and Z_r . The red lines show the polynomial fits with their coefficients shown in the legend.

being rotated into the special frame. Fig.6.6 shows the correlation found between $\cos(\alpha_{QT})$ and the two desired variables Z_r and $\cos(\theta_z)$.

For a quick and robust relation between $\cos(\alpha_{QT})$ and the desired observables, simple polynomial fits are made. $\cos(\theta_z)$ is fitted with an 11th order polynomial whereas Z_r is fitted with a fifth order polynomial. The order of these polynomials were taken from the SNO implementation[113]. The fit values and the best fit lines are shown on their respective curves in Fig.6.6. Fig.6.6(a) has an interesting structure, at $\cos(\alpha_{QT}) > 0.7$ the calibration points become double valued. The cause of this is hard to explain, however a $\cos(\alpha_{QT}) > 0.7$ corresponds to an impact parameter of 780 cm. At high impact parameter it should be expected that the PMT hit patterns of the entry and exit wounds of the muons may overlap to some extent. It is postulated that the structure in $\cos(\theta_z)$ is the result of the mixing of hit patterns. In any case, muons with such impact parameters will have short path lengths and therefore were expected to have a smaller contribution to the final analysis. Therefore, these calibration curves were

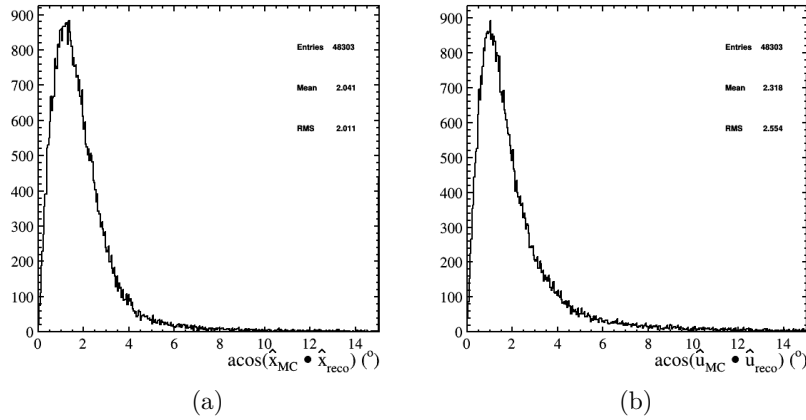


Fig. 6.7: The fit errors of the reconstruction given as the number of degrees between the true vector and the reconstructed vector for both the muon's entry position, \vec{x} , and direction, \vec{u} .

adequate as inputs for the reconstruction. In the case of $\cos(\alpha_{QT}) > 0.7$, the reconstruction should be expected to systematically overestimate the muon path length. The hypothesis that neck reflections are causing the double valued structure in Fig.6.6(a) was tested by making the same plot with events that enter the lower half of the detector. The same structure was observed.

Further investigations should focus on isolating the events in the $\cos(\alpha_{QT}) > 0.6$, $\cos(\theta_z) < 0.5$ region and comparing the reconstructed values to the truth information.

6.4 Fitter performance

As the fitter produces, in principle, independent measurements of a muon's position, \vec{x}_{reco} , and direction, \vec{u}_{reco} , there are two errors to measure. As the observables are both vectors of fixed length, the best way to quantify the error is to take the angle between the true and reconstructed vectors. Fig.6.7 shows the angle subtended by the two vectors describing the

sample of fake muons compared to their true values. The curves are roughly, but not exactly, described by a single sided crystal ball function, which is just a piecewise gaussian with an exponential tail.

To get a figure of merit on the expected error in each variable, the distributions were integrated until 68% of the curve was covered, with the upper limit taken as the figure. Values of 2.1° and 2.2° were found for the position and directions respectively.

6.4.1 Path length error

The reconstructed muon track length is used in both the selection of cosmic muons and as input to the final analysis, therefore its error must be assessed. This can be achieved by comparing truth to reconstructed information. Fig. 6.8(a) and 6.8(c) show how the absolute and fractional errors of the reconstructed track length compared to the truth information varies with reconstructed track length. It can be seen that at shorter track lengths, the error of the reconstruction is both large and biased. This bias and variance decreases as the reconstructed track gets longer.

It is typical to take the ‘profile’ of a two dimensional distribution to assess the mean and variance of the distribution projected along the y axis in each x bin. The relevant profile error bars are also shown on both Fig. 6.8(a) and 6.8(c). These errors are found by fitting gaussians to the one dimensional projection along the y axis in each path length bin. Figs. 6.8(b) and 6.8(d) show two examples of these one dimensional projections, both distributions also show their gaussian fits. Both distributions can be observed to depart from the gaussian shape, however, it was decided that as the fits underestimate the peak, the standard deviation of the gaussian is in this case a conservative estimate of the expected error.

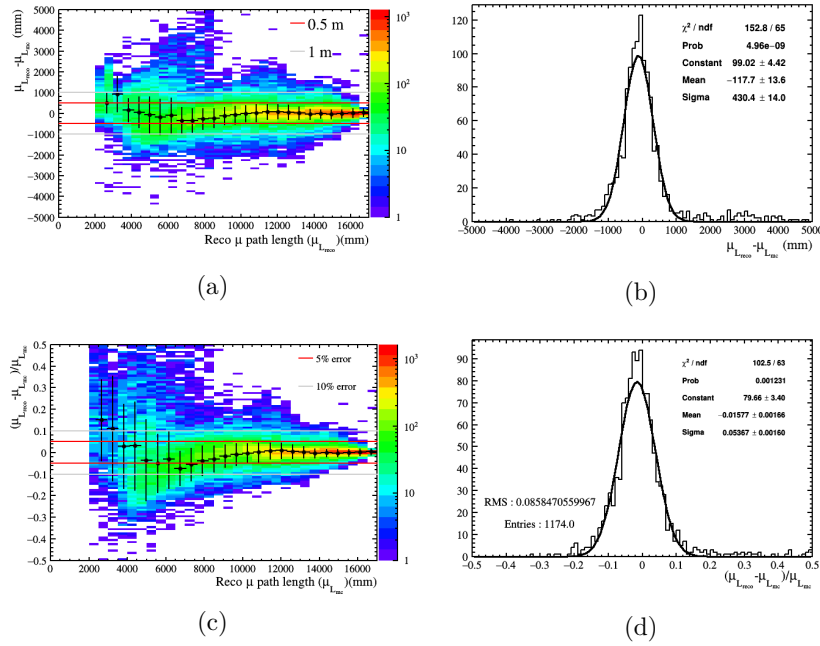


Fig. 6.8: Detailing the absolute and fractional error on reconstructed muon path length. Fig. (a) shows the distribution of the absolute path length error against the reconstructed path length. Fig. (c) shows the distribution of the fractional path length error against the reconstructed path length. Both Figs. (a) and (c) show error bars which are found from gaussian fits along the y projections in each x bin. The error bars show the mean and standard deviation of each of the fits. Figs. (b) and (d) show example marginalised distributions in arbitrary x bins each with their gaussian fit. The distribution can be seen to depart from being a gaussian but, as explained in the text, is taken to be a conservative estimate of the error for reconstructed path lengths greater than 10 m.

The errors in Fig.6.8(a) show that for the most part the reconstruction is accurate to within 1 m. At approximately 10 m the absolute error falls to < 0.5 m. The fractional errors shown in Fig. 6.8(c) show that the fractional error becomes $< 5\%$ for tracks > 10 m. Track lengths lower than 10 m have an observed bias and range in fractional error from 10% to 25%.

The method described above produces no figure of merit on the goodness of fit. This means that the only goodness of fit parameter is

the reconstructed parameters themselves. Ultimately, this means that muons selected for any physics analysis will have to be reconstructed above some threshold. The analysis presented in the rest of this thesis sets its threshold at 10 m and takes the 5% error as a systematic.

6.5 Potential improvements

As stated above, the reconstruction method implemented here is only the seed step of the SNO muon fitter. The SNO fitter took the seed and then performed a likelihood maximisation with the likelihood constructed from the PMT time residuals and a PDF lookup table based on impact parameter. Such a fitter requires an extensive MC study to produce the lookup table but would result in a reduction in both good track threshold and absolute error.

The SNO fitter also estimated the energy lost over the track length, $\frac{dE}{dx}$, such an observable would help in selecting muons that induce a shower in the detector.

6.6 Conclusion

The chapter gives an overview of the muon fitter implemented for this thesis. The fitter utilises the spherical symmetry of the SNO+ PSUP to develop the method based on the prompt time hits and collected deposited charge. The result is a fitter with a 5% uncertainty on muon tracks over 10 m. This uncertainty is accounted for in the final analysis by folding it into the measurements systematic error.

Chapter 7

Cosmic Muons Selection

In the following chapter the methodology behind the selection of through going cosmic muon candidates is motivated and detailed. Through going muons are those that enter and exit the detector opposed to muons that come to rest and subsequently decay inside the detector, leading to no muon in the final state, these are referred to as stopping muons.

Work from many different collaborators is included, mainly from the run selection, reconstruction and data cleaning groups. Input from these groups are combined with the fit results given by the muon reconstruction algorithm detailed in Chapter 6.

7.1 Introduction

When carrying out any particle physics measurement it is of great importance to understand which subset of data to analyse. The SNO+ collaboration has three sub groups which feed into data selection. These are the run selection, data quality and data cleaning groups.

Typically data is taken and passed through a number of ‘data quality’ checks. These checks are carried out on the nearline machine and are the

same as those used in the standard run selection in CSS, see Chapter 4 for details. Low levels checks such as the electronic coverage of the detector and run length duration are made. It was decided that for physics data to be considered for general processing it must be taken from a run which lasts at least 30 minutes¹. The work presented in this thesis follows the same rule, that is that all runs analysed last for greater than 30 minutes. High level checks on observables such as time between events and trigger rate are also carried out. These data quality checks are then collated and passed on to the run selection group. As a starting point for the analysis presented in this thesis, only runs on the golden run list as specified by the run selection group are analysed.

In general all physics runs that last longer than 30 minutes are processed with a general physics analysis macro. However, due to special considerations in the following analysis, a special macro was developed with only specified runs being processed. One of the steps of this processing is to evaluate a number of event level data cleaning flags, developed by the data cleaning group. These flags tag instrumental effects in the data as well as events which can cause backgrounds to a typical SNO+ analysis. An example of this is the TPMuonFollower cut, which tags all events in a 20 second window after a candidate cosmic muon has passed through the detector. The pass/fail result of these data cleaning flags are stored in the event's data cleaning word, this word can then be passed through a mask when analysing data offline to select events of interest. In the case of a muon induced neutron measurement, one would reverse the intended use of the TPMuonFollower flags to select data to analyse.

¹One reason for this is to have some confidence in low level checks such as the low occupancy which make some kind of statistical test.

As the analysis presented in this thesis is a coincidence analysis between muons and following events, the effect of reducing the number of primary events, muons, in the data only reduces detector livetime of the analysis, potentially potentially causing issues due to low statistics of the sample. As the reductions have no bearing on the physics that is being observed, a selection on the primary event shouldn't bias the measurement in any way.

7.2 Potential Muon Instrumental backgrounds

SNO+ is not unique in having the detector itself producing effects which present in the raw data as events of interest. These are collectively known as instrumentals and are undesirable to all physics measurements. This is especially the case in a coincidence analysis, where contamination can lead to false prompt events and therefore windows of data that are certain not to have the desired signal in e.g. searching for neutron captures after a random non neutron inducing instrumental event.

Cosmic muons introduce a large amount of light into the detector. There are four instrumental effects that produce a large amount of light in the detector: Wet end breakdown of PMTs (WEBs), PMT spontaneous light emission (flashers), light generated at the top of the detector (neck events) and light generated by the water level sensors (bubblers).

WEBs are high voltage breakdowns on the PMT base, these typically present as bursts of high nhit events. It is typical that a WEB will persistently contaminate the run. In this case the detector shifter will fix the problem and note the occurrence in the shift report, this run will

then typically fail run selection. However, it has been seen that WEBs can stop after a short period of time and possibly evade detection by the shift operator.

Flashers are similar to WEBs but produce less light. Both WEBs and flashers produce light from a single PMT. These differ from muon like events where the light is produced on the muon track. One possible way to separate these events would be to make some geometric test on the PMT hit positions present in the event. The event PMT timing can also be utilised.

Light coming from the top of the detector may come from light leaks on the deck of the detector. SNO+ has a number of PMTs placed at the top of the neck housed in the 'glove box'. A simple check on the number of neck PMTs hit in the event can tag these events.

The water system has a number of downward facing tubes filled with nitrogen gas which periodically release bubbles of nitrogen. The pressure needed to produce the bubbles can be correlated to the levels of the filling material. Unfortunately, it was found in SNO that high nhit events increase when the bubblers were operational. Such events could be mistaken for a muon signal.

There also exist electronic pick up instrumentals that cause large nhit events without producing light in the detector. There are various causes of these events; one common occurrence is the PMT discrimination can be triggered by high frequency noise induced on the PMT cables or front end electronics. As the resulting event is not caused by light in the detector the integrated charge in all PMTs involved in the event is expected to be low. This, along with the clustering of the PMT channels in crate space, can be used to tag these events.

All of these instrumentals have data cleaning tags which aim to remove these events from the data. The data cleaning tags which aim to remove the four instrumental categories described above are described in section 7.4.2. Typically a sacrifice study is carried out and the results are folded in to the final selection efficiency.

7.3 Potential physical muon backgrounds

Due to the point like nature of a lot of the physical events inside the PSUP in SNO+, most physics events are easily separable from expected muon topologies. However, there are some events which should be considered. Atmospheric neutrino events can lead to the emission of high energy muons causing large amounts of Cherenkov light. As these muons are not ‘cosmic’ in nature they are a background.

Stopping muons are by definition a background to through going muons. As explaining in Chapter 6, stopping muons can not be expected to be reconstructed with the current muon reconstruction.

7.4 Muon Event Selection

The SNO experiment developed a cut based approach to select muons from their general data. SNO+ can benefit from this information by developing a similar cut based selection. The primary goal of the muon follower analysis is not concerned with the rate of muons themselves but only the number of muons detected. Therefore, while maximising the number of muons in our data set is of great importance, sacrificing some muons and increasing muon purity only results in a reduction of effective livetime. All selection variables have this effect, only after the selection

is finalised can the observations on the rate of muons be assessed and compared to the expected rate.

7.4.1 Trigger Selection

Given the large amount of light emitted by cosmic muons, all physics triggers are expected to fire when a muon enters the PSUP. However, there is no explicit selection on trigger type, only one physics trigger is required for the event to be at least initially processed as a muon.

7.4.2 Data Cleaning

Data cleaning looks to identify, flag and ultimately remove typically undesirable events. These events are typically instrumental effects, such as a high voltage breakdown but DC is also used to flag certain outliers such as high nhit events caused by cosmic muons. As muons are not of primary interest to SNO+, the muon data cleaning tag has a low threshold. This leads to many false positives, however, the rest of the data cleaning suite reduces this contamination.

Muon tag

The muon tag flags events that have more than 150 inward facing PMTs hits and more than five OWL PMT hits. As the muon travels through the detector cavity it will produce Cherenkov light, which will register on the OWL PMTS. OWL hits in combination with a relatively high nhit event in the detector is a sign that a cosmic muon has entered the detector.

To flag events caused by cosmogenic isotopes, all triggered events that take place in a 60 second period are tagged as muon followers. This muon follower tag is split into two sub flags, the first flags events within 20

seconds of the initial event and the second flag includes all events within 60 seconds.

Neck event tag

Neck events are events that originate in the neck, typically causing some light to enter the detector and hit PMTs at the bottom of the detector. Over the SNO+ water physics phase, the number of PMTs placed at the top of the detector neck changed. Initially there were three PMTs which then reduced to two². The neck cut flags any event that has two or more neck PMT hits. If one neck PMT is hit then if the average PMT hit times for PMTs below the detector equator is 85 ns or less then the event is also flagged as a neck event.

Burst tag

The front-end has a limited buffer that can fill up during high rate bursts of events such as small electronics break downs. Once these buffer are filled, no new hits can be recorded, this is a concern for other analyses, however, in the muon selection such a tag can be used to find small breakdowns that can be initially selected as muons.

Atmospheric tag

The goal of the atmospheric cut is to tag atmospheric neutrino events by looking for a high nhit event, ≥ 200 nhits, followed by a slightly less high nhit event, ≥ 100 nhits, within $20 \mu s$. Such a pattern is a signature of an atmospheric neutrino interaction producing a muon that decays to

²After the data period discussed in this thesis, repairs were made to the neck tubes so that the number increased to four.




Fig. 7.1: The correlation between muon cleaned nhits and muon track length from MC.

a high energy electron (Michel electron). This flag will look to remove these events from our dataset.

Retrigger tag

The retrigger tag looks for two events within 3 microseconds of each other, flagging both events when they are found. This cut will remove so called invisible friendly WEBs: short bursts of electronic breakdown which typically lasts less than a second.

7.4.3 Low Level Muon Cuts

Cleaned Nhit cut

As muons take a range of track lengths through the detector it should be expected that the nhits of the muon should also take a range of values. Fig. 7.1 shows the MC correlation between track length and cleaned nhits³ of each event. As expected long tracks produce more light in the

³Cleaned nhits is the number of nhits an event contains corrected for the inclusion of bad PMT e.g. those flagged by CSS.




Fig. 7.2: (a) shows the RMS of the QHS distributions for MC muons, (b) show the RMS of the PMT hit timing distribution for the same muon MC. Cut lines are shown on both plots indicating the cut values used when selecting muons in the data.

detector and therefore more PMTs will be hit. It was observed that a cut of 500 cleaned nhits should be adequate to select muons.

Late after this decision was made it was observed that the final nhit spectrum observed in the SNO muon analysis has no entries below 1238 nhits. Fig. 7.3 shows the final SNO muon nhit spectrum. Although the SNO muon analysis is closely followed some of the variables used in the selection are not available to SNO+. It was decided that the nhit cut should be increased to 1238 cleaned nhits. This increase in the cleaned nhits threshold was a practical decision that attempted to improve the muon selection to mirror that made on SNO. As we are raising the threshold we only risk removing low nhit muons which will have small track lengths and therefore contribution little to the final analysis. In the best case scenario the increase will provide some of the benefits of the SNO selection that is not included in the SNO+ muon selection, resulting in a purer muon selection.

QHS spectrum

When the muon exits the detector it deposits a large amount of light in the tubes nearest the exit point. This leads to a peaked QHS distribution. Fig. 7.2(a) shows the RMS of the MC muon simulations. Following the SNO muon analysis, it was decided to remove events with a RMS of the QHS distributions less than 4.5 Gain Scaled Units (GSU). Where GSU units are just the calibrated QHS charge registered on each tube divided by the high half point on that channel. The high half point of a channel is a measure of the gain applied to each PMT. It is defined as the upper point defining the full width half maximum of the QHS spectrum.

PMT timing spectrum

Some instrumental backgrounds have been observed to have PMT hit times that are uniform in time. This is opposed to Cherenkov events which are peaked in time. One measure of the flatness of a spectrum is the RMS. Fig. 7.2(b) shows the RMS of the raw PMT hit times for MC muon simulations. Following the SNO muon selection an RMS of less than 38 ns should be used to identify muons[113].

7.4.4 Reconstruction Level Cuts

Reconstructed Muon track length

As shown in Chapter 6, the muon reconstruction can only be trusted to within a 5% error if the reconstructed track length is greater than 10 meters. Therefore, only track lengths of greater than 10 meters are selected.

Table 7.1: Summary of muon cutting scheme. Each cut is applied sequentially, the order of the cuts was not optimised.

	Livetime (days)	Candidates remaining	% remaining
Pure Muon tag	140.51	14495	N/A
Secondary DC tag	140.39	13577	93.67
nhitCleaned ≥ 1238	132.60	9292	64.11
Charge cut	132.48	9202	63.48
PMT timing cut	132.48	9194	63.43
Burst flag	132.48	9184	63.36
Atmospheric flag	132.48	9183	63.35
Retrigger flag	132.33	9147	63.10
Neck PMT flag	132.29	9140	63.06
$\cos(\theta_z) \geq 0.4$	130.29	8901	61.41
Track length cut ≥ 10 m	117.08	6396	44.13

Reconstructed Muon Direction

The current muon selection is geared towards selecting downward going muons. Defining a muon's zenith angle, θ , as the angle from the z axis one can select downward going muons. Fig. 6.1 shows that downward going muons are expected to enter the detector from a direction $\cos(\theta) \geq 0.4$, therefore, this cut value is acceptable.

7.5 Selected muon data set

7.5.1 Muon selection breakdown

Table 7.1 shows a breakdown of the effect of each cut when applied to the gross muon tag data set. An event is only selected if it satisfies all on the selection requirements. Fig. 7.3 shows a breakdown of the muon raw nhit spectrum. Again, each cut is applied sequentially. Also shown on the figure is the reported SNO muon spectrum after applying their more mature selection⁴. The SNO spectrum is scaled to the normalisation of

⁴The SNO spectrum was digitized by hand using an online tool.




Fig. 7.3: Breakdown of muon nhits spectrum by each cut grouping. The cuts are applied in the order shown in the legend. Two nhitCleaned cuts are shown to demonstrate the effect of imposing a higher nhitCleaned cut after observing the lowest SNO spectrum bin, see section 7.4.3 for details.

the final SNO+ dataset. Great care should be taken when comparing the SNO and SNO+ nhit distributions as although much of the detector is similar, it is crucially not the same. The SNO selection also includes more cuts on variables not currently available in this analysis. However, in general the agreement between the final SNO and SNO+ distributions is good and gives confidence in the SNO+ muon selection.

7.5.2 MC Comparison

It is, of course, desirable to have a perfect MC model which reproduces the data exactly, but such models don't exist. It is common place to reweight MC distributions to agree well with data taken from a calibration source. High confidence in a MC model allows for decisions, such as a signal selection, to be made and evaluated on MC with confidence the same result will hold on data. Any systematic error between the MC and the data is taken as a systematic and tracked through to the final result.




Fig. 7.4: These figures show the normalised comparisons of the muon data set and the expected distributions from MC after applying the selection provided in Table 7.1. Figs (a) and (b) show the low level cuts QHS_{RMS} and T_{RMS} respectively. The disagreement between the data and MC is attributed to the MC model simulating the electronics perfectly. Figs (c) and (d) show the cleaned nhit and reconstructed track length comparison respectively.

In the case of the selection presented in this chapter, all of the selection variables and cut values were either inherited from SNO or motivated from reconstructed parameters. As this is the case the MC/data agreement is not as relevant as in typical analyses. However, the MC/data comparisons are shown for completeness and as a reference for future work⁵. RAT is tuned for the main physics focus, hence the ^{16}N calibration and the lack of a cosmic muon source leaves the muon MC model suboptimal.

Fig. 7.4 shows the normalised MC comparisons for four observables. Fig. 7.4(a) and 7.4(b) show the RMS of the QHS and raw PMT timing

⁵Although at the time of writing there is no plan to analyse muons in the water phase.

distributions respectively. The data has narrower PMT timings and the peak of the QHS RMS distribution is shifted towards lower distributions widths. This is explained by multi photon hits on the PMTs. When PMTs are hit with many thousands of photons, the detector's analogue to digital converters for both time and charge may return a null readings. In this case the PMT hit will not register as having a good time and charge record and will not be included in the time or charge distributions. The analogue to digital converter effect is not simulated in RAT and therefore the MC distributions typically have wider distributions, hence the long tail out to high QHS RMS.

Fig. 7.4(c) shows the cleaned n_{hit} distribution for both MC and data. There is a clear disagreement between the two distributions, however the agreement is sufficient for the muon selection as the cleaned n_{hit} cut ultimately comes from an observation of the SNO muon selection.

Fig. 7.4(d) shows the comparison between the reconstructed track length for both MC and data shown on a log scale. The agreement between the two distributions is adequate and was deemed sufficient given the 5% error on the track length reconstruction.

7.5.3 Live Time Of Muon Data

As shown in Table 7.1 the selected muons were detected over a combined period of 117 days ranging over eight months. Although spanning three seasons, no seasonal variations were found in muon rate.

7.6 Muon data set comparison to SNO results

Making a comparison with the SNO muon selection is difficult as the improved muon reconstruction in SNO allows for a shorter muon track length threshold. However, the reconstruction presented in Chapter 6 still measures track lengths down to 4 m with at most a 15% resolution and a 10% bias. Therefore, although not acceptable for the final analysis a lower track length cut can be used to compare against SNO.

7.6.1 Observed muon rate

When making any event selection it is important to make independent cross checks to test the purity of the selection. In the case of the muon selection, cross checks can be made against results produced from the SNO detector. One of the interesting parameters measured by SNO was the rate of downward going cosmic muons. Of course a similar measurement on SNO+ should be consistent. SNO quoted a measurement of:

$$R_{SNO}^{\mu} = 62.9 \pm 0.2 \mu/\text{day}, \quad (7.1)$$

where these muons have an impact parameter of ≤ 830 cm [84]. Although SNO+ doesn't explicitly select on impact parameter, the muon track length cut of 4 m corresponding to an impact parameter of 826 cm, therefore both selections are comparable.

There are two independent methods to measure the muon rate. The first is to count the number of muons selected per unit time. The second is to fit an exponential to the time difference between muons.




Fig. 7.5: Number of muons in all hour long runs with a Poisson fitted with a mean of 2.6 muons per hour.

Raw muons per hour

As SNO+ has a nominal run length of 1 hour it is relatively easy to count the number of events that pass the muon selection in runs which last exactly an hour. Of the 3496 runs in the data set, 2927 runs have a duration of 1 hour. Only these runs are used for the rest of this subsection. Fitting a Poisson function to the distribution of counted muons will give the expected muons per hour.

Fig. 7.5 shows the number of muons counted per run in the reduced data set. The data points are shown with statistical errors only. The fitted Poisson fit finds:

$$R_{SNO+day}^{\mu} = 62.4 \pm 0.8 \mu/\text{day}, \quad (7.2)$$

where the quoted error is the fit error, the fit has a $\chi^2/\text{dof} = 24.2/10$. It can be said that the fit result is compatible with the SNO measurement.

Time difference between consecutive muons

As cosmic muons are poissonian in time, the time different between two consecutive muons will obey a decaying exponential with a time constant equal to the event rate. To evaluate the time difference between two muons, runs with two or more muons are selected and then the time differences are evaluated in run, that is time differences are only taken between muons in the same run. However, this method is inherently biased due to the hour long samples of data. This arises because limiting the runs to one hour leads to an under sampling of long time differences, therefore more time differences register at short times leading to higher event rate measurements.

To demonstrate this bias, a simple toy MC model can show the effect on the resulting fit; this can then be used to correct the fitted result. Fig. 7.6 shows the fitted event rate for the toy MC with varying run lengths. Also shown is the true rate used to generate the Poisson times. With the expected muon rate from SNO, R_{SNO}^{μ} , a raw unbiased result would require a run duration approximately 20 hours.

It is simple to numerically relate the fitted rate to the truth value. Fig. 7.7 shows the correlation between true event rate and the rate found by fitting the time difference between consecutive muons. The toy model takes a true event rate and simulates 10,000 hour long runs. Events are sampled at times governed by the true event rate and only events that fall into that run are taken. The inter-event times within each run are then fitted with a decaying exponential and the fit results produced. This is repeated 10,000 times building a distribution of fit results, Fig. 7.7(a) shows an example distribution of fit results. The mean and standard deviation of this distribution are then shown as the expected results and




Fig. 7.6: Fitted event rate of Poisson events which are sampled in varying lengths of run. SNO+ has a nominal run length of 1 hour, therefore a fit result of around 100 should be expected.

error. This is done for each truth value thereby providing means and standard deviations of the resulting distribution, which can then be used to construct a 1σ confidence region. Fig. 7.7(b) shows the collected values of the toy model output. To construct a confidence region, linear fits were made to both the upper and lower error as well as the expected values, taking the region contained by the upper and lower fits as the confidence region. This confidence region is shown in Fig. 7.7.

Fig. 7.8(a) shows the distributions of time differences between muon candidates that pass the muon selection. The data is fitted with an exponential with a fitted muon rate of $100.3 \pm 2.1 \mu \text{ day}^{-1}$. Also shown for reference is the expected result from the SNO result scaled to have the same number of events as the data at the first bin centre.

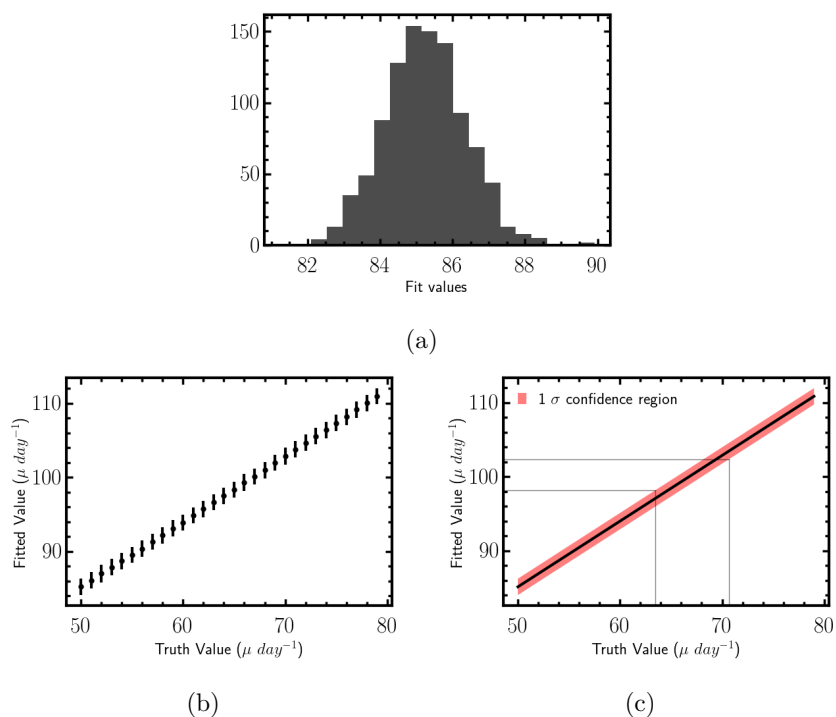


Fig. 7.7: Fig. (a) shows an example of the distribution of fit values used to estimate the expected fit value and error for a given true muon rate. Fig. (b) shows how the true event rate of a simple toy MC relates to the fitted event rate. The mean and standard deviation of 100 fit results are shown as the points and errors of the plot. Fig. (c) shows a 1σ confidence region constructed from linear fits to the extremities of the points on (b). Also shown is the mapping from the fit to the data shown in (c) where the lines represent the 1σ error on the fit value of $100.3 \mu \text{ day}^{-1}$.

As set out above, this result needs to be related to the underlying unbiased muon rate through the 1σ confidence region shown in fig 7.7(c). Indeed, the lines on the plot indicate the fit result on the y-axis and the underlying measurement on the x-axis. The final measurement is found to be :

$$R_{SNO+\Delta t}^\mu = 67.06_{-3.58}^{+3.60} \mu/\text{day} \quad (7.3)$$

$R_{SNO+\Delta t}^\mu$ is in line with the expected SNO muon rate.

In an attempt to minimise the long time difference under sampling it was decided that consecutive runs should be grouped to allow for the inter run time differences. Fig. 7.8(b) shows the exponential fit to such time difference. The fit rate is found to be $70 \mu \text{ day}^{-1}$, which is high and expected as the groupings have a variety of durations and therefore suffers from the same problem as the standard method.

7.7 Potential Improvements

Placing the muon track length cut at 10m introduces a large muon sacrifice. Although the longer tracks will contribute more to the final analysis, with improvements to the reconstruction, shorter tracks could also be included, ultimately increasing the sensitivity. The only way to incorporate these shorter tracks would be by increasing the accuracy of the reconstruction as highlighted in section 6.5.

Super Kamiokande highlighted in a 2015 paper[114], that the Cherenkov light yield could be measured as a function of time and showed that showering muons in which neutrons will be produced can be selected. Once again for a similar approach to be applied in SNO+ improvements in the reconstruction are required.

7.8 Conclusion

The result of this chapter is a pure sample of cosmic muons from the initial nine months of data, post channel threshold change. The muons are selected by drawing on time and charge based statistics, data cleaning and reconstructed quantities. The rate of selected muons is consistent with

that measured in SNO, therefore, providing confidence in the selection's purity. The rest of this thesis uses only the muons selected in this chapter.

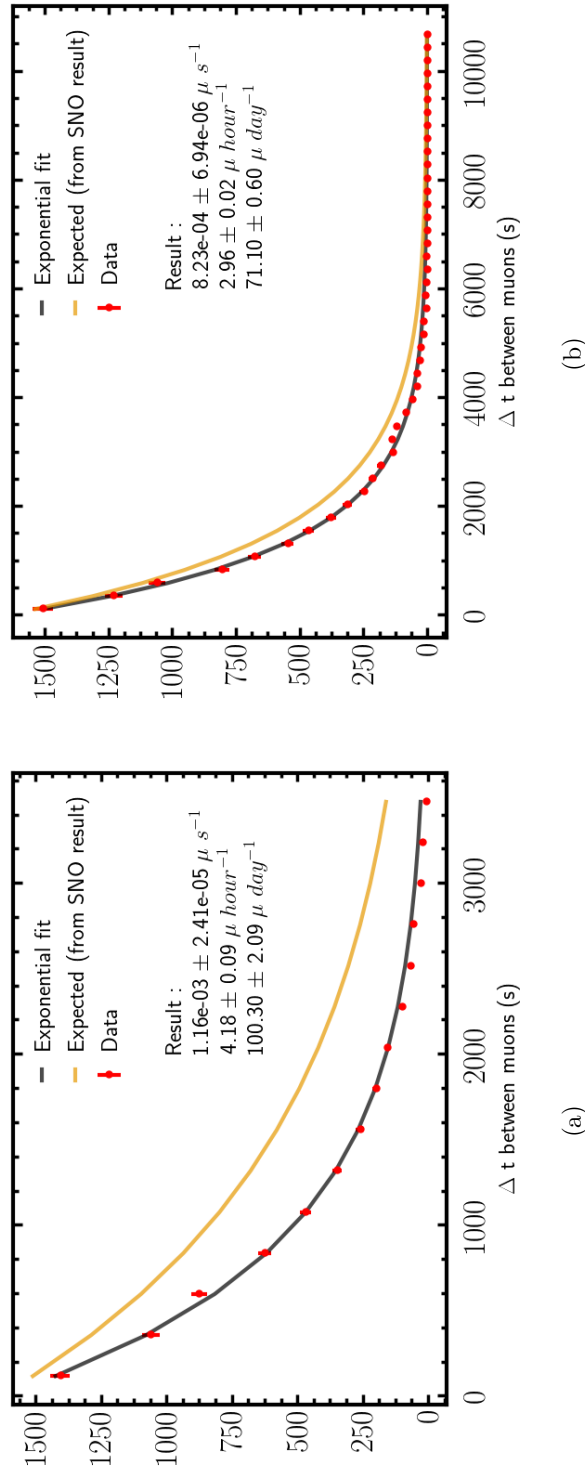


Fig. 7.8: The difference in time between consecutive muons for both in run, (a), and inter-run, (b). The red point show the histogram bin values with the fitted exponential shown by the red line. The expected exponential scaled to to fitted normalisation is shown by the yellow line. The fitted result is used to calculate the rate per second, hour and day which the results inset on both plots.

Chapter 8

Neutron Selection

The following chapter lays out the methodology behind the selection of neutron candidates following a muon. First the MC used to optimise the selection is motivated before showing the optimisation of the selection and its effects.

8.1 Neutron selection

Neutron capture on hydrogen produces a 2.2 MeV gamma ray, which then Compton scatters to produce the electrons emitting Cherenkov light. ^{16}N calibrations showed that SNO+ expects ≈ 8 nhit/MeV. The primary Compton electron carries an average of 1.8 MeV, which, after correcting for the electron Cherenkov threshold leads to 1.0 MeV of energy deposited in the detector. The corresponding event in the detector is an event with ≈ 8 nhits as shown in fig 8.1.

This is on the limit of the lowest trigger threshold (\tilde{N}_{100}) and leads to a trigger efficiency that is not optimal. In this nhits regime the general background event rate is higher than typically acceptable.




Fig. 8.1: The cleaned nhits of 2.2MeV gammas distributed randomly throughout the PSUP. Events are only included if the default water fitter reports a valid fit.

8.1.1 MC generation

In a typical SNO+ physics analysis expected backgrounds are studied on a component level, that is, each major background contributor is studied independently. This requires a large amount of MC. As alluded to in chapter 6, the RAT default cosmic muon simulations take a very long time to successfully complete and often fail resulting in loss of MC events. The log files indicate the errors often originate in Geant4’s light propagations model or it’s handling of long lived particles. The result of these problems was the full cosmic muon RAT simulation was not a viable option for MC generation. To this end an adapted MC generation method was developed. Fig 8.2 is a flow diagram showing the individual steps in this method. The rest of this section will explain the various steps.

Tracking of cosmic muons and daughter particles in the RAT simulation was found to be fast and robust. As thermal neutrons capture on hydrogen on $\mathcal{O}(10^5 \text{ ns})$ neutron daughters are expected to be distinct in

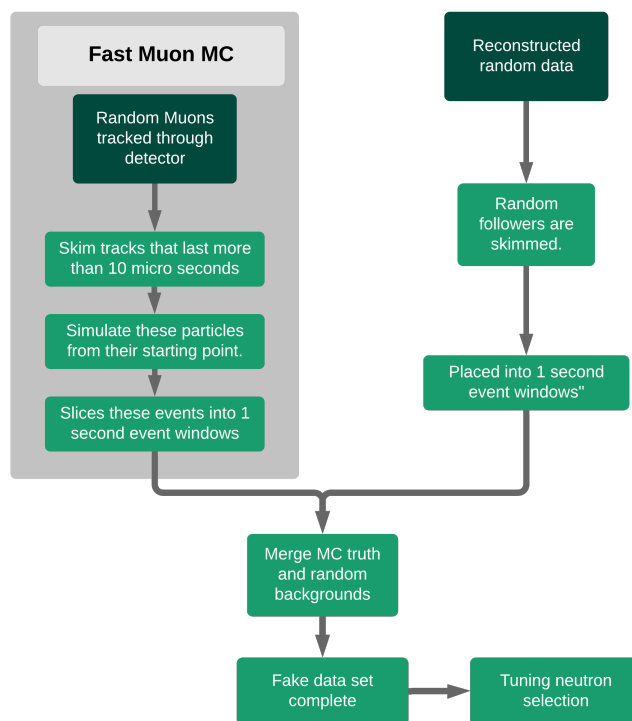


Fig. 8.2: Flow diagram showing the steps taken in the generation of the fake MC.

time with the majority of other daughters particles interacting at shorter time scales (primary and secondary scattering interactions) or on longer times scales (radioactive daughter decays).

Fig 8.3 shows a sample of cosmic muon induced neutron MC taken from RAT simulations. The plot shows the raw nhits and event time of all events in the MC. Interactions at short times have a large variance in nhit and are made up of two populations; at low nhit the majority of events are neutron captures with the high nhit events being explained as Michel electron events resulting from the decay of muons introduced by primary π from the neutron scattering on protons [115].




Fig. 8.3: Raw event nhits in time since the neutron generation time for a sample of muon induced neutrons generated in RAT. Courtesy of Mark Stringer.

There is a population of low nhit events between 10^{-6} – 10^{-3} s: these are neutron captures. This population appears to have a high (~ 50) nhit component not compatible with neutron captures. These are neutron captures on other atoms such as iron and carbon in the PMTs. Above 100 ms there is a large population of events these are the long lived spallation products detailed of which can be found in §8.4.

The neutrons capture time, $206.03 \pm 0.44 \mu\text{s}$ as measured by the $^{241}\text{Am}^9\text{Be}$ analysis, means that a preprocessing step could drastically reduce the number of tracks needed to be put through the full simulation. This preprocessing step was to track all daughter particles and select particles which have tracks which last for longer than $10 \mu\text{s}$. Once selected, the initial conditions of each track can then be passed through the full RAT simulation thereby achieving a MC of muon followers on the time scale of neutron signal neutron captures. Great care was taken in accepting all daughter particles which could last more than 10^4 ns. The

decision to place the timing cut off at 10^4 ns was informed by both the practical and theoretical considerations.

Practical Considerations

SNO experienced that the detector electronics suffer from various sources of pick up in the $1\ \mu\text{s}$ time window after a cosmic muon event, thus rendering all events in this region untrusted [113]. Such pick up is referred to as ringing and discussed further in chapter 10. Reducing these ringing effects would be a perfectly adequate avenue of research, however, setting a hard lower time limit removes the need for this with a small loss in sensitivity.

Theoretical Considerations

The decision to simulate selected followers from the start of the followers track was made to account short lived tracks created in the $10\ \mu\text{s}$ ringing period that could make it past the lower time cut. For example a neutron could capture on carbon which releases a relatively high energy γ . This γ will scatter an electron that could then register as an event passing a neutron selection. Fig 8.4 shows two cartoons of a standard neutron capture on hydrogen MC event, Fig 8.4(a), and a neutron capture on ^{15}C , Fig 8.4(b). Such a track wouldn't be included in the MC model set out above as the track will not have lasted $10\ \mu\text{s}$. However this inclusion comes at a cost; neutrons which are pruned and fully simulated, may capture before 10 ms and therefore not be in the final MC data set. This short coming is acceptable as the fake data set is to be used to optimise a selection and to measure the selection's performance. The lack of neutrons in the final data set will only lead to an under estimation of signal in the optimisation leading to conservative neutron selection.

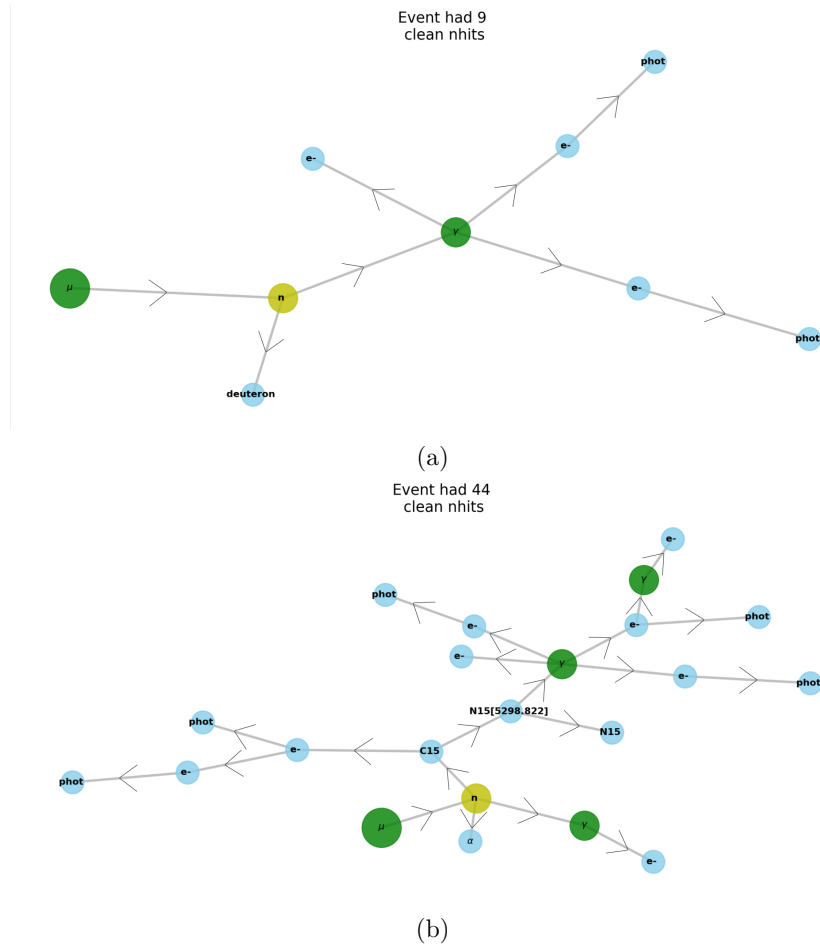


Fig. 8.4: Cartoon of two neutron capture events where parent particles are connected by arrowed lines to their daughter particles. The full event chain is reduced to show only particles of interest. Connections are limited to only particles of interest. Optical photons are represented by the label “phot” and intermediate resonance are defined by the mass point shown in square brackets. (a) shows a neutron capture on hydrogen. (b) shows a neutron capture on carbon.

Doping random backgrounds

As the cosmic muon rate in SNO+ is low, it is possible to assess the expected background levels for any muon by counting the backgrounds when no muon is present. To make a single fake MC data set, MC generated events were merged with windows of random backgrounds, from real data, lasting for one second. These random background windows were found by selecting a random event and collecting all events in the following second. Care was taken that these windows were not contaminated by muons and their followers by only selecting windows with no events tagged as muons or muon followers in data cleaning. The result is a fake data set with simulated tagged neutrons, muon induced backgrounds and real random background data present. This was arranged into a custom data structure based on the default ntuple output of a SNO+ RAT simulation. An example data set is shown in the next section.

8.1.2 Fake data distributions

Fig 8.5 shows the distributions of 4 reconstructed variables from a fake dataset of 10^6 uniformly distributed muons. Events are shown if and only if a muon follower was found following the above method. Fig 8.5(a) shows event trigger times after the muon event trigger, $\Delta_{\mu}t$. The plot shows bins of constant width over a logarithmic time axis, such that the flat random background presents as an exponential wall in time. The neutron signal, shown in fig 8.6(a), is exponential in time a fitted lifetime of $\tau_n = 200 \pm 3 \mu\text{s}$ with a $\chi^2/\text{dof} = 1.03$. The muon induced background e.g. neutrons captures on ^{15}C , is shown also and although relatively low in frequency appears approximately uniform in time. By eye, a timing selection between 10^4 and 10^6 ns should select the majority of neutrons.

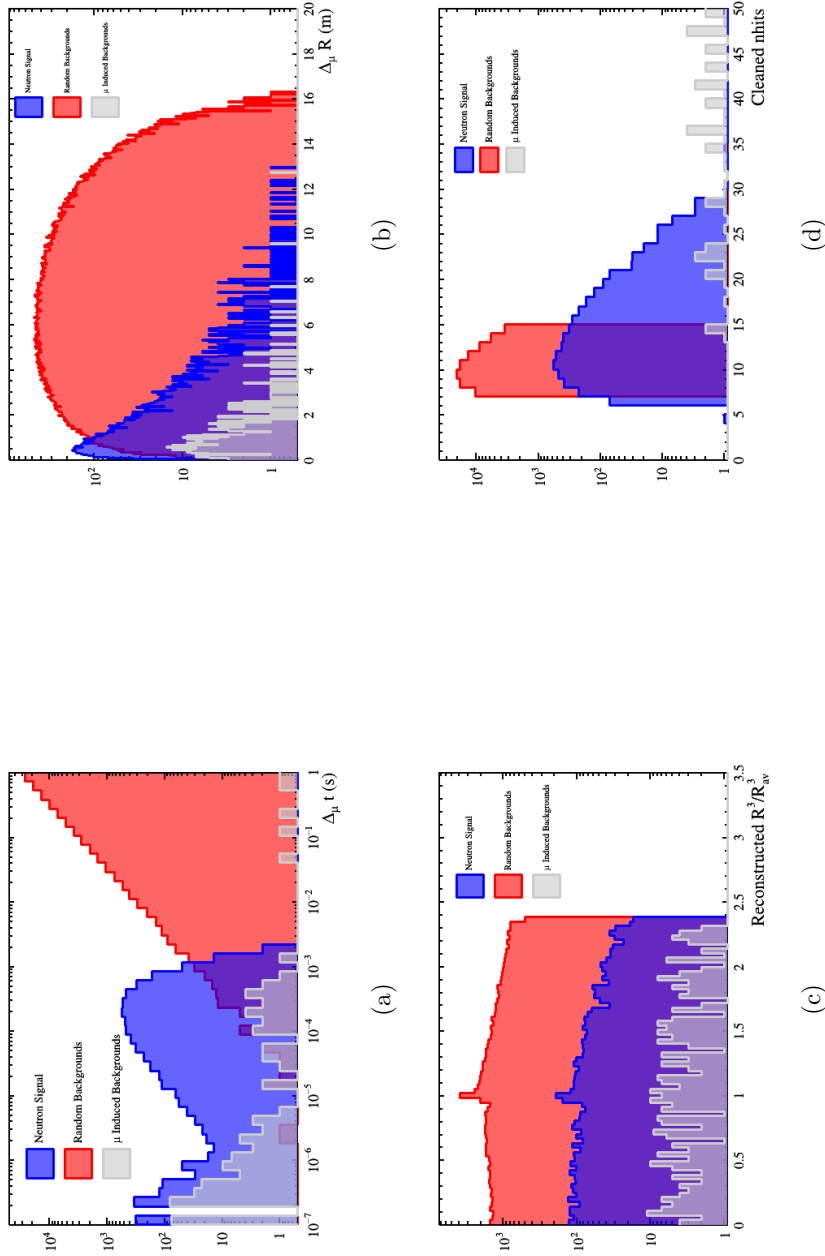


Fig. 8.5: Distribution of a sample of the custom muon follower MC, each figure contains contributions from the simulated neutron signal, simulated muon induced followers and the doped random backgrounds. (a) shows the populations in perpendicular distance from the muon trigger time and the muon trigger time. (b) shows the populations in perpendicular distance from the muon. (c) shows the distributions in normalised radius from the centre of SNO+. (d) shows the cleaned nhits of the three populations, the random background is only shown for $6 < \text{nhits} < 15$ for processing reasons explained in the text.




Fig. 8.6: (a) shows the time distribution of neutron captures taken from MC. The neutron lifetime is consistent with the expected value. (b) shows the perpendicular distance of random backgrounds from the events muon track.

Fig 8.5(b) shows the perpendicular distance from the muon track to each event's reconstructed position, $\Delta_{\mu}R$. It is expected that neutrons should random walk away from the muon track before capturing, 99% of the neutrons capture within 7.2 m. As the muon induced backgrounds are mostly neutrons capturing, on targets other than hydrogen a similar underpopulated distribution is observed. The random backgrounds are not distributed uniformly as might be initially expected this is explained by the spherical nature of the detecting volume in SNO+. The same $\Delta_{\mu}R$ distribution is shown on a linear scale in fig 8.6(b).

The two bottom plots in fig 8.5 show the cubed reconstructed radius in units of cubed AV radius and the cleaned nhits of the event respectively. Fig 8.5(c) shows a bump in both random backgrounds and neutron reconstruction positions at $(R/R_{AV})^3 = 1$, due to the water fitter mis-reconstructing events almost certainly due to poorly modelled optics in the AV region. Given that the muon induced backgrounds are almost entirely composed of neutron captures on elements other than hydrogen, it might

be expected that the muon induced background should be reconstructed on the PSUP in the PMTs. This is not found to be the case, instead the muon induced background are found to be reconstructed uniformly in radius. Although it is not fully understood these events will become negligible after the cleaned nhits cut is applied.

Fig 8.5(d) shows the cleaned nhits for all three groups of particle. The background events only have contributions from events with cleaned nhits larger than 6 and less than 15, this is due to practical processing decisions discussed in chapter 5. This ultimately governs the cleaned nhits selection on the following events. It can be seen that such a cut on the cleaned nhits greatly reduces the muon induced backgrounds and selects the largest region of neutron signal although a considerable neutron sacrifice is expected.

The variables $\Delta_{\mu}t$, $\Delta_{\mu}R$ and cleaned nhits are seen to have a large statistical power in attempting to separate the neutron signal. As the nhits cleaned selection window is determined by the processing strategy, only the $\Delta_{\mu}t$ and $\Delta_{\mu}R$ selections need to be determined.

8.2 Fiducial Volume selection

Many analyses in SNO+ incorporate a selection in radius, R , from the centre of the detector. This combats reconstruction effects and reduces known backgrounds from the AV and the hold up/down ropes. A typical cut might remove events 1 meter from the PSUP and 0.5 meters away from the AV in both directions.

Clearly this type of cut leads to a substantial reduction in detecting material. The decision not to include a cut in R in the neutron selection was primarily informed by the potential reduction in sensitivity due to

Table 8.1: Parameter limits used in the selection optimisation.

Parameter	Lower Limit	Upper Limit
T_{low}	10^4	10^6
T_{high}	10^5	10^7
$\Delta_\mu R$	0	18000

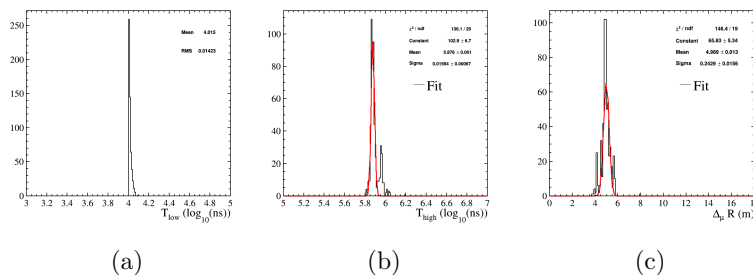


Fig. 8.7: Distributions of the optimised cut values from 100 optimisations each with random initial conditions. Figs. (a), (b) and (c) show the distribution of optimised values for the lower time, high time and perpendicular distance from the muon track cuts respectively. The distributions are shown to consistently converge showing the optimisation is finding a well defined maximum in $s/\sqrt{s+b}$.

the loss in detecting volume. The increase in background contributions was judged to be acceptable as the analysis is in coincidence with a prompt event therefore, the background will be acceptable after applying the powerful time and cleaned nhit cut. Including a fiducial volume selection also increases the complexity of the analysis, creating lots of work for little sensitivity gain.

8.3 $\Delta_\mu t$ & $\Delta_\mu R$ selection optimisation

To find the optimal selection window, typically the cut thresholds are floated while maximising the signal, s , to background, b , ratio. In the case that background levels are similar to that of the signal then the denominator contains contributions from the signal as well: $\frac{s}{\sqrt{s+b}}$. This quantity is referred to as the significance with the maximum value giving




Fig. 8.8: Sample of MC showing the neutron capture signal in blue with other muon induced background shown in grey. The red distribution shows the random doped background. The selection box from the optimisation is also shown. The perceived lack of grey entries is due to the low number of entries.

a measure of how well the signal can be separated from the background. To maximise the significance of the neutron selection, an optimisation was carried out in 3 variables: T_{low} , T_{high} and $\Delta_\mu R$. The parameter limits used in the optimisation are given in table 8.1. The optimisation was carried out with the simulated annealing algorithm. Fig 8.7 shows the collected fit results for 100 optimisations with differing initial conditions. All three parameters converge to similar values. The lower time cut fits to the lower bound of the parameter, since neutron captures exist in the MC before 10^4 ns, but the presence of muon ringing limits the lower time bound. Limiting the lower time threshold will have the effect of reducing the neutron capture detection efficiency¹. Table 8.2 shows the optimised selection parameter values and Fig. 8.8 shows the selection box bounded by the parameters.

¹If instrumental effects could be removed below $10\mu\text{s}$ then the lower limit could be pushed lower.

Table 8.2: Optimised neutron selection parameters used to select neutron candidates.

Parameter	Value
T_{low}	10 μ s
T_{high}	776 μ s
$\Delta_\mu R$	4600 mm

8.3.1 Neutron selection performance

Applying the selection to a general MC sample provides a quantitative measure of the expected signal sacrifice and background rejection. Table 8.3 shows a breakdown of the cumulative effect of each cut used in the selection. The figure shows that the selection is expected to accept 15.37% of all true neutrons simulated, this is the selection's signal efficiency, ϵ . The valid fit cut introduces the largest neutron sacrifice, which should be expected in this cleaned nhits regime. The same selection accepts only 0.01% of the possible backgrounds with the T_{high} cut being the cut with the largest discriminating power, with the suboptimal upper cleaned nhit cut also introducing a large sacrifice². Defining the rejection power as $R = \frac{B_0}{B_f}$, where B_0 and B_f are the number of background events before and after the selection respectively, the selection has a rejection power of 7609.

When defining a selection it is important to address confidence that a selected candidate is a signal event. This is a statistical question and requires prior knowledge of the expected signal and background levels. The purity, $P(S|selected)$, of a selection is the probability that a given selected candidate is the desired signal, S. Using the prior knowledge and Bayes theorem the purity is found to be :

²This short coming will be rectified on the analysis's next iterations.

Table 8.3: Neutron selection performance on neutrons and backgrounds from MC.

	Neutrons			Backgrounds		
	Candidates remaining	% remaining	% reduction	Candidates remaining	% remaining	% reduction
All Candidates	20139	N/A	N/A	152191	N/A	N/A
Valid fit	5706	28.33	71.67	101285	66.55	33.45
nhitsCleaned > 6	5631	27.96	1.31	101284	66.55	0.00
nhitsCleaned < 15	3556	17.66	36.85	101041	66.39	0.24
Distance from muon cut	3366	16.71	5.34	28861	18.96	71.44
Low time cut	3166	15.72	5.94	28803	18.93	0.20
High time cut	3095	15.37	2.24	20	0.01	99.93

$$P(S|selected) = \frac{1}{1 + \frac{\pi_B}{R \cdot \epsilon \cdot \pi_S}} \quad (8.1)$$

Where π_B and π_S are the prior probabilities that the event is a background or signal event respectively. Fig 8.9 shows the selection purity as a function of the ratio of the prior probabilities. This is shown to give various scenarios of background rates. Also shown are the evaluations of scenarios with background events 10 and 100 times more common than neutron capture events. With 100 times the background, the purity is 87% raising to 99% with a background level 10 times higher than the signal.

To estimate the theoretical purity of the selection one can use the neutron yield from the global fit on spallations in scintillator shown in 2.11. At the mean muon energy of SNO+ the neutron yield is $\sim 3.5 \times 10^{-4} \mu^{-1} \text{g}^{-2} \text{cm}^2$. Taking a maximum muon track length of 180 cm and applying the signal sacrifice from table 8.3, one finds an expected neutron rate of 0.063 neutrons per muon. To find the corresponding background rate a random background data set constructed to mirror the selected muon data set was found and the selection was applied. The resulting background rate per a muon was found to be 0.067.

The prior probability ratio in equation 8.1 can be approximated by the ratio of rates i.e. ~ 1 . Therefore the theoretical yield of the selection presented above is $\sim 99\%$.

Table 8.4 provides a summary of the selection found in this chapter.




Fig. 8.9: Selection purity as a function of $\frac{\pi_B}{\pi_S}$, the ratio between the prior probabilities that the event is a background or a signal event. Also shown are various scenarios of differing background and signal rates with their evaluated purity shown in the plot.

8.4 Long lived muon spallation products

Spallated neutrons are just one of the species of long lived particle produced by muon spallations. A wide range of isotopes both stable and unstable are produced. The unstable isotopes have a wide range of decay modes, end points and half-lives. In 2014, Li and Beacom published the expected yields of spallated isotopes in Super-Kamiokande [68], listing yields for all unstable isotopes, and specifying those relevant for analyses in Super-Kamiokande i.e. with an energy $\geq 3.5\text{MeV}$ as well as lower energy contributions. Although Super-Kamiokande has a lower mean muon energy than SNO+, these yields can be used to predict the expected number of isotopes to pass the neutron selection³.

³Ultimately this is justified as the predicted contributions to the final selection is still negligible after a 3 orders of magnitude increase in yield. Therefore it is a safe assumption that the uncertainty introduced by the difference in energy is compensated by the small yields.

Table 8.4: Summary of the neutron candidate selection motivated in this chapter.

Parameter	Cut value
fitValid	true
T_{low}	10 μ s
T_{high}	776 μ s
nhitsCleaned	>6
nhitsCleaned	<15
Δ_{μ} R	4.6 m

In fact, a subset of the full selection, table 8.4, is shown to produce negligible isotope contributions to the signal selection. Table 8.5 provides a list of all of the unstable isotopes listed in [68]. The probability of an event passing the time selection is found by integrating the decay curve between the time, limits in table 8.4, and normalising by the total integral, the probability is shown for each isotope in the 4th column of table 8.5. The 3rd column of table 8.5 shows the probability that an event passes the nhit cut, shown in table 8.4. To find this, 10^4 events were simulated in RAT for each unstable isotope, from which the expected nhit spectrum was found. The probability is then found by similar ratios of integrals as the timing probability. The last column shows the final probability of selection of each isotope per muon. The most likely contributions are from the isotopes ^{12}B and ^{15}O , however with the 6000 muons selected in chapter 7 the contributions for long lived isotopes are negligible, $\mathcal{O}(10^{-6})$.

Table 8.5: Prediction of numbers of long lived isotope expected in a reduced neutron selection.

Isotope	Yield ($10^7 \mu^{-1} \text{m}^2 \text{kg}^{-1}$)	Half-life (s)	Nhits window probability (10^{-3})	Timing window probability (10^{-4})	Events per muon (10^{-10})
^{18}N	0.02	0.62	2.00	15.85	0.11
^{17}N	0.59	4.17	4.17	2.37	0.99
^{16}N	18.00	7.13	1.71	1.39	7.26
^{16}C	0.02	0.75	10.59	13.24	0.48
^{15}C	0.82	2.45	2.84	4.04	1.60
^{14}B	0.02	0.01	0.00	691.76	0.00
^{13}O	0.26	0.01	1.59	1086.11	76.13
^{13}B	1.90	0.02	3.19	552.76	568.92
^{12}N	1.30	0.01	1.29	859.91	244.27
^{12}B	12.00	0.02	3.23	478.05	3150.06
^{12}Be	0.10	0.02	4.20	410.64	29.30
^{11}Be	0.81	13.80	2.94	0.72	0.29
^{11}Li	0.01	0.01	0.50	1098.14	0.94
^9C	0.89	0.13	1.60	77.64	18.81
^9Li	1.90	0.18	2.05	55.46	36.80
^8B	5.80	0.77	0.72	12.85	9.10
^8Li	13.00	0.84	0.83	11.81	21.54
^8He	0.23	0.12	4.97	82.84	16.11
^{15}O	351.00	122.24	750.00	0.08	3624.41
^{14}O	13.00	70.20	160.77	0.14	50.11
^{11}C	105.00	1220.04	1000.00	0.01	144.84
^{10}C	7.60	19.29	628.57	0.51	416.78
^{13}N	19.00	597.90	1000.00	0.02	53.48
Sum					8472.32

Chapter 9

Neutron Detection Efficiency

9.1 Outline

The number of neutron captures following a muon, N_{Obs} , can be modelled as:

$$N_{\text{Obs}_\mu} = \epsilon_\mu \cdot N_T + N_{\text{Bkg}} , \quad (9.1)$$

where ϵ_μ is neutron detection efficiency, N_T is the true number of neutron captures on hydrogen and N_{Bkg} is the number of background events passing the selection. The neutron detection efficiency is clearly an important part of such modelling, with the term responsible for accounting how many neutrons escaped detection through both trigger threshold effects and selection inefficiencies.

The lack of a cosmic muon calibration source in SNO+ leaves a MC study as the only way to assess the efficiency.




Fig. 9.1: Neutron direction and kinetic energy correlation found from tracking randomly distributed muons in RAT.

9.2 Efficiency modelling

The neutron detection efficiency is wholly dependent on the path of the spallating muon in question. That is, if the muons, skims the PSUP then a spallated neutron has a higher probability to randomly walk outside the PSUP and therefore not be detected. To get a quantitative measure of the neutron detection efficiency detailed MC studies were carried out for each selected muon.

The general goal of the MC study is to find a neutron detection efficiency estimate for each muon in the data. The reconstruction of each muon has a measured error in entry position and direction as set out in chapter 6. To find the neutron detection efficiency, the best fit muon track is perturbed in position and direction by an uncorrelated sampling of the distributions shown in Fig 6.7. While it is known that the entry position and direction errors are correlated, assuming that they are uncorrelated will lead to sampling a larger area of track errors, which will ultimately lead to more conservative systematic errors. As this will only make the systematic error larger, folding in the correlation wasn't pursued. The

new muon track is then propagated through the detector and the resulting line is then used as the starting point of a number neutrons generated along the line randomly. Each neutron is given a momentum sampled from the 2D distribution shown in Fig 9.1. The distribution was found by tracking neutrons spallated by muons in RAT, binning their initial kinetic energy and direction.

1,000 neutrons are generated for each perturbed muon track, these are then tracked through the full RAT detector simulation. Each data muon is perturbed 100 times to generate a cloud of neutron captures including the effect of the reconstruction errors.

By finding the number of neutron captures that should have been detected, N_0 and by counting the number of captures that pass the neutron selection, N_s , one obtains the neutron detection efficiency, ϵ_μ , as :

$$\epsilon_\mu = \frac{N_s}{N_0}, \quad (9.2)$$

with the statistical error of ϵ_μ , $\delta\epsilon_\mu$, given by the Poisson fluctuation of N_s :

$$\delta\epsilon_\mu = \frac{\sqrt{N_s}}{N_0}. \quad (9.3)$$

Completing the study on the muon data set required a large amount of computing time. One iteration of the study required 1 month of running on the 120 available cores in the Queen Mary physics department batch queue.




Fig. 9.2: The collected neutron efficiency calculations as a function of the muon path length, Fig (a), and the analogous impact parameter, Fig (b). The detection efficiency is seen to drop off at an impact parameter of 6, this is explained by the optical effects of the AV on the event reconstruction.

9.2.1 Efficiency performance

Although it is likely that ϵ_μ is unique for each muon, there should be some dependence on track length. Fig 9.2 shows the 2D distribution of ϵ_μ against both track length, Fig 9.2(a) and impact parameter, Fig 9.2(b). The detection efficiency drops off at shorter track lengths, Fig. 9.2(b) shows the drop off corresponds to the muons travelling through the detector in the vicinity of the AV. Reconstruction around the AV is suboptimal, resulting in a reduction in detection efficiency. By taking the collected efficiency calculations as a function of muon path length one can get a sense of the systematic spread of efficiency for a given track length.

As this study relies heavily on the GEANT4 neutron modelling it is important to verify the model.

9.3 Neutron MC AmBe verification

With no cosmic muon calibration source, a direct verification is not possible however the AmBe source data allows for a comparison at low energy which can then be extrapolated to the muon induced neutron energies. This work extends beyond the time limit of the authors PhD and is currently on going.

The general overview of the study is to evaluate the difference between the MC and AmBe source data in variables relevant to the neutron selection, namely the difference in capture time and nhits. The resulting disagreement will be folded into the muon by muon MC data set used to estimate the efficiency. This in turn will impact on the distribution shown in Fig 9.2 leading to a larger systematic error as found by the method laid out in §10.1.1.

Chapter 10

Analysis Strategy and Instrumental Effects Following Cosmic Muons

In the previous chapters the author has motivated and defined all operations on data needed to measure properties of cosmic muon induced neutrons. This final chapter brings all components together to produce results, however, an unaccounted for detector efficiency effect is discovered. The detectors trigger threshold is shown to shift following a high nhit event. This effect is found in a neutron lifetime side band fit and confirmed with a custom TELLIE analysis.

10.1 Outline

It is standard in the literature to report a measured neutron yield, Y_n , which can be expressed as:

$$Y_n = \frac{N_n}{\rho \sum_{\mu} L_{\mu}}, \quad (10.1)$$

where N_n is the number of spallated neutrons¹, N_μ is the number of muons observed, $\sum_\mu L_\mu$ is the summed muon track lengths and ρ is the density of the detecting medium. There are various ways to report such a result.

A 2017 publication by Daya bay[73] reported Y_n taking N_n by summing all neutron events, corrected for efficiency and estimated background, and approximating $\sum_\mu L_\mu \approx N_\mu \bar{L}_\mu$, where \bar{L}_μ is the mean muon track length in the Daya bay detectors. This is akin to placing all the muons end to end and taking all neutron detection as being produced by one long muon track. This approach is taken as Daya bay's muon reconstruction doesn't allow for track length estimation. Taking such an approach also inherently assumes knowledge of the underlying neutron multiplicity distribution. The estimator Y_n is effectively the mean of this distribution. However, the neutron multiplicity distribution is known to have long tails, therefore the mean is not the most useful observable when estimating the background contributions. Nevertheless, such a measurement is useful to gain knowledge about an understudied distribution.

SNO+ does have a muon track length estimation and therefore can calculate Y_n on a muon by muon basis. Binning each Y_n then provides a measure of not only the mean Y_n by also the spread of the sampling distribution.

As SNO+ could carry out both techniques it would be interesting to report both results.

¹Corrected for neutron captures on materials other than hydrogen.




Fig. 10.1: Time spectrum of muon followers and the appropriately scaled background. An exponential with a constant offset is fitted between 70 - 1000 μs , exponential's lifetime is fitted to $134 \pm 16 \mu s$ which is less than the expected neutron lifetime of 209 μs .

10.1.1 Muon track length systematics

The 5% error on the muon path length will correlate with the neutron detection efficiency as shown in Fig 9.2. The systematic error on the final Y_n will be found by carrying out the calculation many times while fluctuating the path length by 5% and then sampling the efficiency from the distribution shown in Fig 9.2(a). The width of the resulting distribution fit results will provide the systematic error introduced by the track length detection efficiency errors.

10.2 Time spectrum side-band check

Before carrying out the final analysis, a sanity check on the neutron lifetime was made, prior to applying any efficiency correction was carried out. This may be viewed as a side band fit as the neutron decay time shape is not relevant to the final analysis.

Fig 10.1 shows the time spectrum of events that pass the selection for both the muon data and the appropriately scaled background data set. As the background is constant and a neutron signal should decay exponentially in time, the muon data spectrum is fitted with the function:

$$f(t, x, y, \lambda) = x + y \cdot \exp \frac{-t}{\lambda} \quad (10.2)$$

The fitted mean lifetime is $134 \pm 16 \mu\text{s}$ which is 35% smaller than the expected neutron lifetime. It was also noticed that the second and third bins in the muon data spectrum are considerably lower than expected. It was hypothesised that these effects may be due to a change in detector response following very high energy events like muons.

10.3 Investigation with TELLIE

One of the challenges of a muon related study in SNO+ is the lack of input from a calibration source. However, the dynamic range inbuilt to the ELLIE system allowed the opportunity to use the relatively wide angled TELLIE beam to study high nhit effects.

Data were collected with the normal physics triggers and the TELLIE system injecting 2,000 events at a rate of 10Hz for 6 different intensities. The TELLIE light level was tuned for each nhit point to values, 1500, 2500, 3500, 4500, 5500 and 6000 nhits. With such high prompt nhit events, one can tag it and view events in a following time window, mirroring the muon analysis.

Fig 10.2 shows the time spectra of all TELLIE nhit points. As TELLIE events do not induce following physics events, a background flat in time is expected at times beyond after pulsing effects, $\mathcal{O}(\mu\text{s})$. However, the




Fig. 10.2: Collection of time spectra for each of the different TELLIE intensity runs (labelled by the nhits they were tuned to).

same dip at $\sim 50 \mu\text{s}$ is observed, but more importantly the excess of events is also observed across all TELLIE nhits. This strongly suggests some unexpected event contamination. There is also an observable dip at 300 ns, the explanation of all three effect's is of paramount importance.

One of the benefits of using the TELLIE data is that the event statistics are higher therefore effects such as the 300 ns dip is clearly observable relative to the muon data. It also allows the study of the effects cause.

It was noted that for events around the $\mathcal{O}(100) \mu\text{s}$ bump the number of in time² hits on N100 trigger was lower than expected. Fig 10.3 shows the in time N100 hits across the relevant time range for an example TELLIE run. The histogram's profile is also shown.

The mean in time hits clearly shifts in time, the mean in time hits appears to be completely anti correlated with the event rate shown in fig 10.2. This suggests that the increase and decrease in events at $\mathcal{O}(100 \mu\text{s})$ and $\mathcal{O}(300 \mu\text{s})$ respectively is due to a shift in the detector's effective

²In time hits are the number of PMT hits in the recorded trigger signal, e.g. The N100 in time hits are the number of hits in the 100 ns trigger time window.




Fig. 10.3: The in time N100 nhits distribution against time for an example TELLIE run with the histogram's profile super imposed. Clearly the profile shows structure incompatible with the expected flat line.

trigger threshold. A shift in detector threshold would present as either an increase or decrease in event rate. If the trigger threshold was to decrease, it would take less in time hits to trigger the detector and the converse is also true. This is exactly the observed behaviour. The effect is less pronounced in the $\mathcal{O}(300\mu s)$ since the event rate decreases as events in time hits increases however the rate exponential increases as in time hits decrease due to dark noise hits.

The cause of the threshold shifts is currently unknown and in any case is not wholly relevant as the effect is present in the low level data.

This shift in threshold effectively changes the detector's neutron detection efficiency over the selection window. This isn't included in the efficiency calculation presented in Chapter 9, thus rendering it inadequate in it's current state.

Possible solutions are being investigated with two avenues of research currently being assessed; Mark Stringer has made the observation that including a lower in time hit cut in the neutron selection could potential

remove the effects of any shift in efficiency by only accepting events which are not removed by the increased threshold shift at $300\ \mu\text{s}$. The in time cut would be assessed on the TELLIE data, however, the potential impact on the final result is yet to be assessed via the tools presented in Chapter 8. The other possible solution would be to further characterise the shape of the threshold shift in time and fold it into the efficiency calculation presented in chapter 9.

Although the discovery of the detector threshold shifts introduces the requirement for more investigation into a final analysis strategy. Its discovery is a major step forward in not only the analysis but also in the understanding of the SNOPLUS detector. In future phases of the experiment this effect may become relevant to a number of analyses.

Chapter 11

Conclusion

Weak signal searches such as dark matter experiments suffer from cosmic muon induced neutrons backgrounds. Even with overburdens such as those found at SNOLAB, the neutron background still provides a large background. Both theoretical and data driven models of neutron spallation at SNOLAB depths require experimental results to improve their predictions.

In this work, the current status of the measurement of cosmic muon induced neutrons has been presented. The analysis remains incomplete due to an unexpected detector effect which has implications for the detector's neutron capture detection efficiency. Unfortunately, the work required to overcome this effect is beyond the time limits of the authors PhD, however, active avenues of research have been laid out where relevant.

If this effect had not been present in the data, the majority of the tools, selections and studies had been developed to produce a result with partial systematics. Unless specified this is the authors work and their novel research contribution.

The CSS framework removes unstable channels from the analysis chain, which as implemented will have the largest effect for low n_{hit} events such as the 2.2 MeV γ from neutron captures.

The understanding of cosmic muons is clearly highly important to the induced neutron study. As cosmic muons are not the main focus of much of the research carried out on SNO+, much of the muon relevant infrastructure needed had not been developed at the start of the authors study. These included the muon MC, reconstruction and selection which, as detailed Chapters 6 and 7, were developed by the author. This resulted in a reconstruction method with a 5% track length error for muon tracks larger than 10 m and a pure muon selection consistent in rate with that measured in SNO.

A neutron candidate selection was developed with an innovative adapted fast muon follower MC. The resulting selection provided a theoretical selection purity of greater than 99%.

The author also developed sections of the final analysis strategy. These include the neutron detection efficiency MC calculation and the final likelihood analysis as well as obtaining TELLIE data.

Before the final analysis was completed, a side band check on the neutron capture lifetime was carrying out. During this check a previously unknown detector effect was discovered. The detectors trigger threshold was found to shift after high n_{hit} events i.e. cosmic muons. This effect was observed to last for $\mathcal{O}(\text{ms})$ therefore biasing the lifetime measurement. Further investigation with TELLIE, found that this threshold shifting was indeed an artefact of the detectors electronics. The cause of this artefact is still an open question, but in any case a solution would probably involve extensive electronics work. This new effect has an impact on final analysis which is not fully understood at this time, however work is

continuing investigating the possibility of introducing a neutron selection which although reduces the neutron detection efficiency will mitigate the detector threshold shifts.

Although the goal of this thesis is not yet complete, the work presented has set the foundations and will be built upon to produce a result. The field of high energy cosmic muon neutron spallation yields and the possibility of a neutron multiplicity measurement is in need for experimental inputs to progress, this only motivates this analysis further.

References

- [1] Carsten Jensen. *Controversy and consensus: nuclear beta decay 1911 1934*. Springer Basel, 1999.
- [2] Carlo Giunti and Chung Wook. Kim. *Fundamentals of neutrino physics and astrophysics*. Oxford University Press, 2016.
- [3] Laurie M. Brown. The idea of the neutrino. *Physics Today*, 31(9):23–28, 1978.
- [4] Wolfgang Pauli. Pauli letter collection: letter to Lise Meitner. <https://cds.cern.ch/record/83282>.
- [5] Frederick Reines and Clyde L. Cowan. The neutrino. *Nature*, 178(4531):446–449, 1956.
- [6] 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] K. Kodama et al. Observation of tau neutrino interactions. *Physics Letters B*, 504(3):218 – 224, 2001.
- [8] Richard L. Garwin, Leon M. Lederman, and Marcel Weinrich. Observations of the failure of conservation of parity and charge conjugation in meson decays: the magnetic moment of the free muon. *Phys. Rev.*, 105:1415–1417, Feb 1957.
- [9] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105:1413–1415, Feb 1957.
- [10] M. Goldhaber, L. Grodzins, and A. W. Sunyar. Helicity of neutrinos. *Phys. Rev.*, 109:1015–1017, Feb 1958.
- [11] Wikipedia contributors. File:standard model of elementary particles.svg, 2020. [Online; accessed 22-March-2020].
- [12] John N. Bahcall et al. New solar opacities, abundances, helioseismology, and neutrino fluxes. *The Astrophysical Journal*, 621(1):L85–L88, jan 2005.
- [13] John N. Bahcall. Solar neutrinos. i. theoretical. *Phys. Rev. Lett.*, 12:300–302, Mar 1964.

-
- [14] Raymond Davis. Solar neutrinos. ii. experimental. *Phys. Rev. Lett.*, 12:303–305, Mar 1964.
- [15] Solving the mystery of the missing neutrinos. <https://www.nobelprize.org/prizes/themes/solving-the-mystery-of-the-missing-neutrinos>.
- [16] K. Abe et al. Solar neutrino measurements in super-kamiokande-iv. *Phys. Rev. D*, 94:052010, Sep 2016.
- [17] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, Aug 1998.
- [18] T Kajita, E Kearns, and M Shiozawa. Establishing atmospheric neutrino oscillations with Super-Kamiokande. *Nuclear Physics B*, 908:14–29, 2016.
- [19] A. Yu. Smirnov. Solar neutrinos: Oscillations or No-oscillations? 2016.
- [20] Q. R. Ahmad et al. Direct evidence for neutrino flavor transformation from neutral-current interactions in the sudbury neutrino observatory. *Phys. Rev. Lett.*, 89:011301, Jun 2002.
- [21] T. Araki et al. Measurement of neutrino oscillation with kamland: Evidence of spectral distortion. *Phys. Rev. Lett.*, 94:081801, Mar 2005.
- [22] S. Abe et al. Precision measurement of neutrino oscillation parameters with kamland. *Phys. Rev. Lett.*, 100:221803, Jun 2008.
- [23] J. K. Ahn et al. Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment. *Phys. Rev. Lett.*, 108:191802, 2012.
- [24] J. K. Ahn et al. RENO: An Experiment for Neutrino Oscillation Parameter θ_{13} Using Reactor Neutrinos at Yonggwang. 2010.
- [25] Y. Abe et al. Improved measurements of the neutrino mixing angle θ_{13} with the Double Chooz detector. *JHEP*, 10:086, 2014. [Erratum: *JHEP*02,074(2015)].
- [26] T. Araki et al. Measurement of neutrino oscillation with KamLAND: Evidence of spectral distortion. *Phys. Rev. Lett.*, 94:081801, 2005.
- [27] S. Abe et al. Precision Measurement of Neutrino Oscillation Parameters with KamLAND. *Phys. Rev. Lett.*, 100:221803, 2008.
- [28] P. Adamson et al. Combined analysis of ν_{μ} disappearance and $\nu_{\mu} \rightarrow \nu_e$ appearance in MINOS using accelerator and atmospheric neutrinos. *Phys. Rev. Lett.*, 112:191801, 2014.
- [29] K. Abe et al. Combined analysis of neutrino and antineutrino oscillations at t2k. *Phys. Rev. Lett.*, 118:151801, Apr 2017.

-
- [30] P. Adamson et al. Constraints on Oscillation Parameters from ν_e Appearance and ν_μ Disappearance in NOvA. *Phys. Rev. Lett.*, 118(23):231801, 2017.
- [31] K. Abe et al. Evidence for the Appearance of Atmospheric Tau Neutrinos in Super-Kamiokande. *Phys. Rev. Lett.*, 110(18):181802, 2013.
- [32] R. Wendell et al. Atmospheric neutrino oscillation analysis with sub-leading effects in Super-Kamiokande I, II, and III. *Phys. Rev.*, D81:092004, 2010.
- [33] M. G. Aartsen et al. Measurement of Atmospheric Neutrino Oscillations at 6–56 GeV with IceCube DeepCore. *Phys. Rev. Lett.*, 120(7):071801, 2018.
- [34] R. Abbasi et al. The Design and Performance of IceCube DeepCore. *Astropart. Phys.*, 35:615–624, 2012.
- [35] S. Adrian-Martinez et al. Measurement of the atmospheric ν_μ energy spectrum from 100 GeV to 200 TeV with the ANTARES telescope. *Eur. Phys. J.*, C73(10):2606, 2013. [Eur. Phys. J.C73,2606(2013)].
- [36] B. Aharmim et al. Combined Analysis of all Three Phases of Solar Neutrino Data from the Sudbury Neutrino Observatory. *Phys. Rev.*, C88:025501, 2013.
- [37] K. Abe et al. Solar Neutrino Measurements in Super-Kamiokande-IV. *Phys. Rev.*, D94(5):052010, 2016.
- [38] G. Bellini et al. First evidence of *pep* solar neutrinos by direct detection in borexino. *Phys. Rev. Lett.*, 108:051302, Feb 2012.
- [39] M. Cribier. Results of the whole GALLEX experiment. *Nucl. Phys. Proc. Suppl.*, 70:284–291, 1999. [284(1999)].
- [40] J. N. Abdurashitov et al. Solar neutrino flux measurements by the Soviet-American Gallium Experiment (SAGE) for half the 22 year solar cycle. *J. Exp. Theor. Phys.*, 95:181–193, 2002. [Zh. Eksp. Teor. Fiz.122,211(2002)].
- [41] Raymond Davis, Don S. Harmer, and Kenneth C. Hoffman. Search for neutrinos from the sun. *Phys. Rev. Lett.*, 20:1205–1209, May 1968.
- [42] John N. Bahcall. Neutrinos from the sun. *Scientific American*, 221(1):28–37, 1969.
- [43] G. Bellini et al. Precision measurement of the ${}^7\text{Be}$ solar neutrino interaction rate in borexino. *Phys. Rev. Lett.*, 107:141302, Sep 2011.
- [44] G. Bellini et al. Measurement of the solar ${}^8\text{B}$ neutrino rate with a liquid scintillator target and 3 mev energy threshold in the borexino detector. *Phys. Rev. D*, 82:033006, Aug 2010.

-
- [45] O. Yu. Smirnov et al. Measurement of neutrino flux from the primary proton–proton fusion process in the Sun with Borexino detector. *Phys. Part. Nucl.*, 47(6):995–1002, 2016.
- [46] Mass hierarchy cartoon. <http://ignatz.phys.rpi.edu/site/index.php/the-physics-2/>, Accessed: 25/08/19.
- [47] P.F. de Salas, D.V. Forero, C.A. Ternes, M. Tórtola, and J.W.F. Valle. Status of neutrino oscillations 2018: 3σ hint for normal mass ordering and improved cp sensitivity. *Physics Letters B*, 782:633 – 640, 2018.
- [48] K. Abe et al. Search for cp violation in neutrino and antineutrino oscillations by the t2k experiment with 2.2×10^{21} protons on target. *Phys. Rev. Lett.*, 121:171802, Oct 2018.
- [49] M. A. Acero et al. New constraints on oscillation parameters from ν_e appearance and ν_μ disappearance in the nova experiment. *Phys. Rev. D*, 98:032012, Aug 2018.
- [50] M. A. Acero et al. First Measurement of Neutrino Oscillation Parameters using Neutrinos and Antineutrinos by NOvA. 2019.
- [51] James Strait et al. Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE). 2016.
- [52] Justo Martín-Albo. Sensitivity of DUNE to long-baseline neutrino oscillation physics. *PoS, EPS-HEP2017:122*, 2017.
- [53] Stephanie Bron. Hyper-Kamiokande: Towards a measurement of CP violation in lepton sector. In *Proceedings, Prospects in Neutrino Physics (NuPhys2015): London, UK, December 16-18, 2015*, 2016.
- [54] M. Tanabashi et al. Review of particle physics. *Phys. Rev. D*, 98:030001, Aug 2018.
- [55] Nathalie Deruelle and Jean-Philippe Uzan. *Relativity in Modern Physics*. Oxford University Press, 2018.
- [56] F. Couchot et al. Cosmological constraints on the neutrino mass including systematic uncertainties. *Astronomy & Astrophysics*, 606, 2017.
- [57] Beate Bornschein, Uwe Besserer, Markus Steidl, Michael Sturm, Kathrin Valerius, Jürgen Wendel, and KATRIN Collaboration. The five phases to standard tritium operation of katrin. *Fusion Science and Technology*, 71(3):231–235, 2017.
- [58] Ali Ashtari Esfahani et al. Determining the neutrino mass with cyclotron radiation emission spectroscopy—project 8. *Journal of Physics G: Nuclear and Particle Physics*, 44(5):054004, mar 2017.
- [59] L. Gastaldo et al. The Electron Capture ^{163}Ho Experiment ECHO: an overview. *J. Low Temp. Phys.*, 176(5-6):876–884, 2014.

-
- [60] C. Alduino et al. Measurement of the two-neutrino double-beta decay half-life of ^{130}Te with the cuore-0 experiment. *The European Physical Journal C*, 77(1):13, Jan 2017.
- [61] J. Schechter and J. W. F. Valle. Neutrinoless double- β decay in $\text{su}(2)\times\text{u}(1)$ theories. *Phys. Rev. D*, 25:2951–2954, Jun 1982.
- [62] Michael Duerr, Manfred Lindner, and Alexander Merle. On the quantitative impact of the schechter-valle theorem. *Journal of High Energy Physics*, 2011(6):91, Jun 2011.
- [63] Jonathan Engel and Javier Menéndez. Status and future of nuclear matrix elements for neutrinoless double-beta decay: a review. *Reports on Progress in Physics*, 80(4):046301, mar 2017.
- [64] A. Gando et al. Search for majorana neutrinos near the inverted mass hierarchy region with Kamland-Zen. *Phys. Rev. Lett.*, 117:082503, Aug 2016.
- [65] A. Batista et al. Open questions in cosmic-ray research at ultrahigh energies. *Frontiers in Astronomy and Space Sciences*, 6:23, 2019.
- [66] A. Aab et al. Observation of a large-scale anisotropy in the arrival directions of cosmic rays above 8×10^{18} ev. *Science*, 357(6357):1266–1270, 2017.
- [67] 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, 2003.
- [68] Shirley Weishi Li and John F. Beacom. First calculation of cosmic-ray muon spallation backgrounds for MeV astrophysical neutrino signals in super-kamiokande. *Phys. Rev. C*, 89:045801, Apr 2014.
- [69] Shirley Weishi Li and John F. Beacom. Spallation backgrounds in super-kamiokande are made in muon-induced showers. *Phys. Rev. D*, 91:105005, May 2015.
- [70] R.P. Kokoulin. Uncertainties in underground muon flux calculations. *Nuclear Physics B - Proceedings Supplements*, 70(1):475 – 479, 1999. Proceedings of the Fifth International Workshop on topics in Astroparticle and Underground Physics.
- [71] National Research Council. *Neutrinos and Beyond: New Windows on Nature*. The National Academies Press, Washington, DC, 2003.
- [72] John F. Beacom et al. Physics prospects of the J inping neutrino experiment. *Chinese physics C*, 41(2):23002, 2017.
- [73] F. P. An et al. Cosmogenic neutron production at D aya B ay. *Phys. Rev. D*, 97:052009, Mar 2018.

-
- [74] T. Hagner et al. Muon-induced production of radioactive isotopes in scintillation detectors. *Astroparticle Physics*, 14:33–47, August 2000. Provided by the SAO/NASA Astrophysics Data System.
- [75] G. Bellini et al. Cosmogenic Backgrounds in Borexino at 3800 m water-equivalent depth. *Journal of Cosmology and Astroparticle Physics*, 8:049, August 2013.
- [76] S. Abe et al. Production of radioactive isotopes through cosmic muon spallation in KamLAND. *Physical Review C*, 81(2):025807, February 2010.
- [77] M. Aglietta, G. Badino, G. Bologna, C. Castagnoli, and A. Castellina. Neutron flux generated by cosmic-ray muons at 5200 hg/sq CM s.r. underground - Depth-neutron intensity curve. *Nuovo Cimento C Geophysics Space Physics C*, 12:467–477, August 1989.
- [78] R. Hertenberger, M. Chen, and B.L. Dougherty. Muon-induced neutron and pion production in an organic liquid scintillator at a shallow depth. *Physical Review D*, 52:3449–3459, December 1995.
- [79] F. Boehm et al. Neutron production by cosmic-ray muons at shallow depth. *Physical Review D*, 62(9):092005, November 2000.
- [80] S.C. Blyth et al. Measurement of cosmic-ray muons and muon-induced neutrons in the Aberdeen Tunnel Underground Laboratory. *Physical Review D*, 93(7):072005, April 2016.
- [81] H. M. Araújo et al. Measurements of neutrons produced by high-energy muons at the Boulby Underground Laboratory. *Astroparticle Physics*, 29:471–481, July 2008. Provided by the SAO/NASA Astrophysics Data System.
- [82] L. Reichhart et al. Measurement and simulation of the muon-induced neutron yield in lead. *Astroparticle Physics*, 47:67–76, July 2013. Provided by the SAO/NASA Astrophysics Data System.
- [83] S. Andringa et al. Current Status and Future Prospects of the SNO+ Experiment. *Adv. High Energy Phys.*, 2016:6194250, 2016.
- [84] B. Aharmim et al. Measurement of the cosmic ray and neutrino-induced muon flux at the sudbury neutrino observatory. *Phys. Rev. D*, 80:012001, Jul 2009.
- [85] Benjamin Schmidt. Latest results from the cuore $0\nu\beta\beta$ search and future non-ge solid detectors (cupid, supernemo etc.). Nuphys, 2018.
- [86] N. D. Scielzo et al. Double-beta decay Q values of Te-13-, Te-128, and Te-120. *Phys. Rev.*, C80:025501, 2009.
- [87] Jack Dunger. Cosmogenic and externals analysis. SNO+-docDB 4187, 2017.

-
- [88] Tanner Kaptanoglu and others. Double beta official plots. using te+dda optics. SNO+-docDB 5681, 2019.
- [89] SNO+ Collaboration. Search for invisible modes of nucleon decay in water with the sno+ detector. *Phys. Rev. D*, 99:032008, Feb 2019.
- [90] SNO+ Collaboration. Measurement of the ^8B solar neutrino flux in SNO+ with very low backgrounds. *Phys. Rev. D*, 99:012012, Jan 2019.
- [91] K. Abe et al. Solar neutrino results in Super-Kamiokande-III. *Phys. Rev.*, D83:052010, 2011.
- [92] B. Aharmim et al. Measurement of the ν_e and total ^8B solar neutrino fluxes with the sudbury neutrino observatory phase-iii data set. *Phys. Rev. C*, 87:015502, Jan 2013.
- [93] G. Bellini et al. Search for solar axions emitted in the m1-transition of $^7\text{Li}^*$ with borexino ctf. *The European Physical Journal C*, 54:61–72, 01 2008.
- [94] Mark Stringer. *Sensitivity of SNO+ to supernova neutrinos*. PhD thesis, University of Sussex, January 2019.
- [95] Ruben Darío Guerrero. Making jablonsky diagrams using tikz. Accessed: 05/06/19, 2016.
- [96] J B Birks. Scintillations from organic crystals: Specific fluorescence and relative response to different radiations. *Proceedings of the Physical Society. Section A*, 64(10):874–877, oct 1951.
- [97] James Waterfield. *Optical calibration system for SNO+ and sensitivity to neutrinoless double-beta decay*. PhD thesis, University of Sussex, April 2017.
- [98] Ian Coulter. *Modelling and reconstruction of events in SNO+ related to future searches for lepton and baryon number violation*. PhD thesis, Oxford University, 2013.
- [99] M. J. D. Powell. A tolerant algorithm for linearly constrained optimization calculations. *Mathematical Programming*, 45(1):547–566, Aug 1989.
- [100] Nuno Barros et al. Water phase unidoc. SNO+-docDB 5421, 2018.
- [101] Edward Leming. Drive correction. SNO+-docDB 4592, 2017.
- [102] Anthony LaTorre. Measurement of the trigger efficiency in sno+. SNO+-docDB 4723, 2017.
- [103] Richard Bonventre. Guide to ecals and board testing. internal note, 2015.

-
- [104] Evelina Arushanova. *Development of an ^{90}Y calibration source and rejection of pileup backgrounds in the SNO+ experiment*. PhD thesis, Queen Mary University of London, 2018.
- [105] Robert Stainforth. *Characterising the Optical Response of the SNO+ Detector*. PhD thesis, University of Liverpool, 2016.
- [106] L Lebanowski et al. Reconstructed 16n spectra and fits with energy calibration. SNO+-docDB 5407, 2019.
- [107] Mark Anderson et al. Ambe neutron analysis review. SNO+-docDB 5701, 2019.
- [108] Thorium and uranium chains. http://www.world-nuclear.org/uploadedImages/org/info/radioactive_decay_series_th_u.png, Accessed: 25/07/19.
- [109] S. Agostinelli et al. GEANT4: A Simulation toolkit. *Nucl. Instrum. Meth.*, A506:250–303, 2003.
- [110] Glg4sim package. <http://neutrino.phys.ksu.edu/~GLG4sim/>.
- [111] M.D. Lay and M.J. Lyon. An experimental and monte carlo investigation of the r1408 hamamatsu 8-inch photomultiplier tube and associated concentrator to be used in the sudbury neutrino observatory. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 383(2):495 – 505, 1996.
- [112] V.A. Kudryavtsev. Muon simulation codes music and musun for underground physics. *Computer Physics Communications*, 180(3):339 – 346, 2009.
- [113] Christopher Kyba. Measurement of the atmospheric neutrino induced muon flux at the sudbury neutrino observatory. *Dissertations available from ProQuest*, 01 2006.
- [114] H. Zhang et al. Supernova relic neutrino search with neutron tagging at super-kamiokande-iv. *Astroparticle Physics*, 60:41 – 46, 2015.
- [115] M. Kleinschmidt et al. Production of positive pions by neutron-proton collisions. *Zeitschrift fur Physik A Atoms and Nuclei*, 298(4):253–266, Dec 1980.