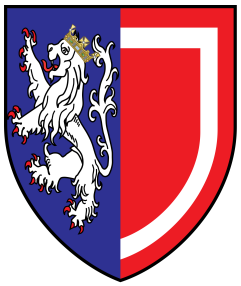
# Rapid Identification of Supernovae at SNO+

Luca A. Cavalli

Balliol College, Oxford



Thesis submitted in partial fulfilment of the requirements for the degree of Master of Science by Research at the University of Oxford

Michaelmas Term, 2015

# Rapid Identification of Supernovae at SNO+

Luca A. Cavalli
Balliol College, Oxford

Michaelmas Term, 2015

## Abstract

Core-collapse supernova are some of the most cataclysmic events known to occur in the universe. During the brief time of collapse and explosion, a supernova releases as much energy as the progenitor star does over its entire lifetime. 99% of this energy is thought to be emitted as neutrinos, detectable by SNO+ via charged and neutral current interactions. This thesis explores the use of distinctive supernova neutrino interactions, specifically inverse beta decay ($\bar{\nu}_e + p \rightarrow n + e^+$) and $^{12}$C de-excitation $\left(\nu + {}^{12}C \rightarrow {}^{12}C^* + \nu' \rightarrow {}^{12}C + \gamma\right)$, in order to rapidly tag supernova candidates on an event-by-event basis using a processor running near-line to the main SNO+ computing architecture. This was achieved via the design and construction of a prototype trigger system, known as the L3, running within the main SNO+ analysis software package RAT. Preliminary results over simulated supernova data at 10, 30 and 100 kpc are presented, alongside tests over early SNO+ dark-run data. Finally, L3 efficiency calculations are performed over Monte Carlo data in scintillator, with the prototype trigger demonstrating greater than 95% efficiency at tagging galactic supernova out to approximately 20 kpc.

## Preface

The aim of this work was the exploration of the use of distinctive supernova neutrino interactions in the SNO+ detector in order to rapidly tag supernova candidates in near real-time. This involved the creation and testing of a prototype C++ program, named the *L3 Processor*, designed to run as a component of the main SNO+ analysis software package RAT. This particular component type is known as an *event processor* within RAT, and may observe or modify event information as it is passed through the processor. As a collaborative effort, RAT contains a variety of modules and tools which may be utilised to run within RAT event processors. The L3 uses a number of these tools to carry out its checks, particularly the *Segmentor* stand-alone data utility class, written by Jack Dunger at Oxford, which is used as a core component of the isotropy checks conducted on water data.

As a work that explores the identification of supernova neutrinos, this work also discusses the theory and experimental aspects of supernova neutrino detection within the SNO+ detector in depth. In order to do this effectively, this work utilises a number of plots created by Belina von Krosigk for her own PhD thesis, *Measurement of Proton and Alpha Particle Quenching in LAB based Scintillators and Determination of Spectral Sensitivities to Supernova Neutrinos in the SNO+ Detector*. These plots greatly assist the reader in understanding the spectral aspects of neutrino detection at SNO+, key in determining which interactions may be selected for use in the L3 Processor.

Luca Cavalli
*June 2016*

# Contents


**1 Introduction** **1**
1.1 Motivation . . . 1
1.1.1 Project Objectives . . . 2
1.2 A Brief History of the Neutrino . . . 2
1.3 A Brief History of Supernovae . . . 6
1.4 Summary . . . 8

**2 Theory** **10**
2.1 Supernovae . . . 10
2.1.1 Classification . . . 10
2.1.1.1 Type Ia . . . 11
2.1.1.2 Type Ib/c . . . 11
2.1.1.3 Type II . . . 11
2.1.2 Core Collapse Supernovae . . . 13
2.1.2.1 Overview . . . 13
2.1.2.2 Final Stages of Core Collapse . . . 14
2.1.2.3 Neutrino Energy Loss Mechanisms and Spectra Formation . . . 15
2.1.2.4 Current Modelling . . . 18
2.2 Neutrino Properties . . . 20
2.2.1 Neutrino Oscillations . . . 20
2.2.1.1 Two Neutrino Approximation . . . 21
2.2.1.2 Three Neutrino Mixing . . . 22
2.2.2 Matter Effects . . . 24
2.2.3 Neutrino Mass and the Seesaw Mechanism . . . 25
2.2.4 Sterile Neutrinos . . . 28
2.3 Supernova Neutrino Interactions at SNO+ . . . 29
2.3.1 Inverse Beta Decay . . . 31
2.3.2 Carbon Nuclear Interactions . . . 32
2.3.2.1 Exclusive Transitions . . . 32
2.3.2.2 Inclusive Knockout Transitions . . . 34
2.3.3 Neutrino-Proton Elastic Scattering . . . 35
2.3.4 Neutrino-Electron Elastic Scattering . . . 35
2.3.5 Summary of Interaction Channels . . . 36

**3 The SNO+ Detector** **37**
3.1 Introduction . . . 37
3.2 Detector Design . . . 38
3.2.1 Structure . . . 38
3.2.2 Detection Principle . . . 38
3.2.3 Experimental Phases . . . 40
3.2.3.1 Water Phase . . . 40
3.2.3.2 Scintillator Phase . . . 40

3.2.3.3 Loaded Scintillator Phase . . . 41
3.3 Hardware and Software . . . 41
3.3.1 Overview . . . 41
3.3.2 PMTs . . . 42
3.3.3 Electronics . . . 42
3.3.4 Calibration Systems . . . 43
3.3.5 Software . . . 45
3.4 Detector Physics Goals . . . 45
3.4.1 Neutrinoless Double Beta Decay . . . 45
3.4.2 Supernova Neutrinos . . . 48
3.4.3 Low Energy Solar Neutrinos . . . 48
3.4.4 Geoneutrinos and Reactor Neutrinos . . . 48
3.4.5 Exotics: Nucleon Decay and Axions . . . 49
3.5 Detector Backgrounds . . . 49
3.5.1 Intrinsic Detector Radioactivity . . . 49
3.5.2 Cosmogenic Muons . . . 52
3.5.3 Solar and Atmospheric Neutrinos . . . 54

**4 The L3 Supernova Trigger** **56**
4.1 Overview . . . 56
4.2 Supernova Triggers at SNO+ . . . 57
4.2.1 The Event Builder and the Level 2 Trigger System . . . 58
4.2.2 The L3 System Concept . . . 59
4.3 The L3 Trigger Algorithm . . . 61
4.3.1 Components . . . 62
4.3.1.1 Data Cleaning and Non-Physics Triggers . . . 62
4.3.1.2 Timing Checks . . . 62
4.3.1.3 Isotropy Checks . . . 63
4.3.1.4 Event Reconstruction . . . 65
4.3.1.5 Supernova-Like Event Counters . . . 67
4.3.2 Output and SNEWS Alert . . . 68
4.4 L3 Prototype Performance . . . 69
4.4.1 L3 Prototyping and Simulation Testing . . . 69
4.4.2 Timing Checks . . . 70
4.4.2.1 Simulation Data . . . 70
4.4.2.2 MDC Data . . . 71
4.4.3 Isotropy Checks . . . 71
4.4.3.1 Simulation Data . . . 72
4.4.3.2 MDC Data . . . 73
4.4.4 Energy Checks . . . 74
4.4.4.1 Fitter Performance . . . 75
4.4.4.2 Energy Fit Results . . . 76
4.4.5 L3 Efficiency in Scintillator . . . 79
4.4.5.1 Poisson Estimation of L3 Efficiency . . . 80
4.5 Issues . . . 85

**5 Conclusion and Outlook** **87**

# Chapter 1

# Introduction

## 1.1 Motivation

Whilst the advent of large scale neutrino detection experiments have provided an excellent opportunity to study the neutrino, they have also paved the way for a completely new field of astrophysics: neutrino astronomy. The elusive nature of neutrinos, initially forming such a barrier to their detection, now provides us with a unique way of observing the universe. Their infinitesimal cross section of interaction enables us to spot point sources of neutrino production, with little to no directional interference as they pass through even the densest of matter. It is this characteristic which makes neutrinos invaluable for the study of core-collapse supernovae on a galactic scale: during core-collapse and the resulting explosion, neutrinos and antineutrinos carry away $\sim 99\%$ of the gravitational binding energy of the progenitor star at energies of several MeV [1], corresponding to $\sim 10^{53}$ ergs[1], and escape almost unimpeded from the dying star.

Similarly, the extreme astrophysical environment in which the neutrinos are produced provides an opportunity to study neutrino properties unseen elsewhere; the interior of a supernova is so dense as to be completely opaque to neutrinos for a short period of time, allowing unusual neutrino mixing effects to arise [2]. Additionally, these conditions force neutrinos of different flavours to decouple from the collapsing star at different radii and times, creating distinctive timing and energy spectra which encode substantial information as to the intrinsic properties of the neutrinos themselves [3]. With sufficient statistics, the observation of supernova neutrinos is expected to reveal the neutrino

[1]The erg is a unit of energy corresponding to $10^{-7}$ Joules, originating in the CGS system of units. The erg is commonly used in supernova studies for historical reasons, and is often coupled with the Foe, corresponding to $10^{51}$ ergs or $10^{44}$ J. For comparative purposes, at its current luminosity, the Sun will release approximately 1.2 Foe over the entire course of its lifetime.

mass hierarchy, which itself then has consequences for the determination of any potential neutrino Majorana mass [4]. Finally, the highly dense and turbulent conditions within supernovae create neutrino oscillation signatures which provide information as to the structural features of the supernova itself [5].

As a multipurpose, kiloton-scale scintillation detector, SNO+ is well positioned to detect these neutrinos via a number of interactions at relatively low energy thresholds. For any given supernova explosion, the neutrino flux will be our first indication that a supernova has occurred, allowing an alert to be issued to the astronomical community before the explosion is visible hours or days later. These short time-scales make quick detection and tagging imperative to ensure maximum physics coverage. This work explores the use of two major interaction signals, working in concert with a simple data processor to monitor for supernova bursts in real time. Ideally, once fully implemented, this system could augment current near line processing on SNO+ to provide a robust supernova alert system via the Supernova Early Warning Network, SNEWS.

### 1.1.1 Project Objectives

This work is an initial study for real time monitoring using the current SNO+ software package. The primary objective was to construct a prototype processor, known as the Level 3 Trigger or *L3*, acting as an extension to current monitoring tools, which incorporates a small number of key checks over real-time data. This prototype could be further augmented in future work to provide additional and more robust functionality. Even in its prototype form however, a simple L3 should fulfill the role of basic supernova monitoring; galactic supernovae are only predicted to occur 1-3 times per century [6], so this is an almost essential requirement given the inherent unpredictability of the events themselves.

## 1.2 A Brief History of the Neutrino

The roots of the discovery of the neutrino lay in an attempt to explain one of the most baffling mysteries of atomic physics of the early 20th century: why did electrons emitted via beta decay, one of the few known types of radioactivity at time, appear to violate conservation of energy? By 1911 it was well known that in beta decay, the atom would release an electron and transmute to a different elemental form, yet the efforts of Lise Meitner and Otto Hahn showed the first hints that the energy spectra of beta particles was continuous, in a seeming violation of the conservation of energy. This was later confirmed in 1914 by James Chadwick using one of the first Gieger counters.

Unlike alpha and gamma radiation, in which the energy of the emitted particles was simply equal to the difference in energy of the initial and final nuclear states, the beta decay spectra appeared to span the range from almost zero to the maximum energy permitted by the decay.

It wasn't until 1931 when Wolfgang Pauli, in his famous 'Dear Radioactive Ladies and Gentlemen' letter to the Federal Institute of Technology in Zurich, almost apologetically postulated an answer [7]. The solution was to create a new, extremely lightweight, hitherto undetected particle, which would carry away the remaining energy of the decay and explain the phenomenon. In 1934 Enrico Fermi then used this hypothetical particle to construct the most successful model of beta decay to date [8], the Fermi Interaction, and introduced the world to the Neutrino [9].

It wasn't until 1956 that Cowan and Reines at Savannah River, South Carolina, would prove beyond doubt the existence of the neutrino. By capitalising on advances in scintillator technology and utilising the vast numbers of neutrinos produced in fission reactions, they were able to observe the process $\bar{\nu}_e + p \rightarrow e^+ + n$, known as inverse beta decay (IBD), occurring on protons in water with an observed cross section corresponding to the predicted value of 6 x $10^{-44}$ $cm^2$ [10]. Over the decades, their experiment provided a blueprint for many neutrino experiments that would come after: a large tank of target material and scintillator surrounded by photomultiplier tubes. In this way, the minuscule cross section of neutrino interactions could be almost completely mitigated, with any interaction signals amplified by scintillator light and picked up by sensitive PMTs. Crucially, Cowan and Reines also made use of a unique property of the IBD interaction; the light from the instantaneous annihilation of the positron is followed microseconds later by the light emitted from the neutron capturing on hydrogen, permitting a distinctive 'delayed coincidence' tag to be applied to these events. Fittingly, the delayed coincidence tag is still used today by the majority of neutrino experiments measuring electron antineutrino reactions due to its reliability and consistency, and has been instrumental in the determination of many neutrino properties.

In the 1960's, work at the Brookhaven National Laboratory showed conclusively that muons could be produced from neutrinos originating from pion decay. It was well known that pion decay only produced muons via the decay $\pi \rightarrow \mu^- + \nu$, with the muon subsequently decaying via $\mu \rightarrow e^- + \bar{\nu}$. If all neutrinos were identical, neutrino pair annihilation should take place some portion of the time, and a photon should appear alongside the electron as a final state particle as one possible channel. Yet, Pontecorvo and Hinks, alongside several other experiments, showed that this never occurs [11]. There was a solution to this paradox: the neutrino created alongside the muon could be associated with that particle, a muon-neutrino $\nu_\mu$, likewise the neutrino produced alongside the electron could be of a different type, the electron-neutrino $\nu_e$. By creating muons, the Brookhaven

team had showed that muon neutrinos interact via inverse beta decay just like electron neutrinos, producing a muon rather than an electron. In doing so, they conclusively identified this second type of neutrino. Later work at SLAC in 1974 confirmed the existence of a third lepton, the tau, whilst precise measurements of the Z boson decay width at LEP in the 80's conclusively demonstrated that only three active flavours of neutrino could exist [12].

Observations of the oscillation of neutral kaon states in the 1950s, along with the inherent difficulties encountered in attempting to measure the tiny mass of the neutrino, prompted speculation, led by Bruno Pontecorvo, that the neutrino may itself oscillate between different flavour eigenstates. If this could be conclusively shown, then it would follow that neutrinos do indeed possess mass, and a long standing question of particle physics could be resolved. Whilst long baseline measurements could not be pursued until later, it was postulated that neutrino oscillation should manifest itself via beta decay energy spectra: if the spectra themselves depended on the mass of the neutrino, then any mixing of neutrino mass eigenstates should reveal itself as a mixture of energy spectra, producing deviations in the final spectrum. During the 1980s seemingly precise measurements of tritium and sulphur revealed a tiny, yet significant, distortion of spectra corresponding to a 17 keV neutrino. Further studies of the beta decays of $^{13}$C and $^{63}$Ni also demonstrated the effect, along with a number of other independent experiments, though it took until the early 1990s for this to be shown to be incorrect and the product of a combination of experimental error and backscatter off the experimental apparatus [13]. Whilst ultimately proved false, the controversy prompted a substantial amount of debate into the nature of neutrino oscillations and neutrino mass, which otherwise may not have taken place, and thus allowed much further development of the theoretical underpinnings of the concept of the neutrino.

Lending weight to the oscillation hypothesis was a problem that had plagued the astrophysical community since the 1960s. Neutrinos were theorised to be produced in prodigious quantities by the fusion reactions powering the Sun, yet all experiments to date had recorded a solar neutrino rate far lower than expected [14]. This meant that either the Standard Model of particle physics, or the Standard Solar Model, or both, were very wrong. Given the massive progress made in both for the previous half century however, the Solar Neutrino Problem (as it became known) was a deeply worrying development. Efforts by Ray Davis Jr. at the Homestake Mine in South Dakota progressed for decades, yet consistently demonstrated a deficit between the amount of electron neutrinos produced in the Sun and the number recorded on Earth [15]. A summary of the Homestake Experiment results between 1970 and 1979 is shown in 1.1.

In 1999, the Super Kamiokande collaboration showed a deficit in the number of atmospheric neu-

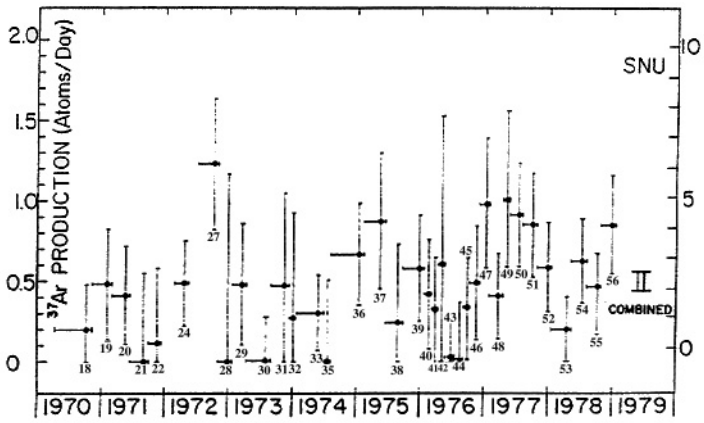



Figure 1.1: Summary of results showing the individual Homestake Experiment runs 18 to 56 undertaken between 1970 and 1979. Contemporary theoretical models as of 1980 expected a solar neutrino flux of 7.8 Solar Neutrino Units [SNU]. As Ray Davis Jr. and John Bahcall write in [17]: "The current difference between theory and observation using the best available estimates for the parameters is about a factor of three".

trinos passing through the Earth compared to those from directly overhead, indicating conclusively that neutrinos changed flavours during their longer passage through the Earth [16]. This provided strong evidence for non-zero neutrino mass and oscillation, but still failed to solve the solar neutrino problem. The resolution would require a different method to measure neutrinos, crucially one that was independent of neutrino flavour.

Two years later the Sudbury Neutrino Observatory published results, shown in Figure 1.2, conclusively demonstrating that the total number of neutrinos reaching the Earth was consistent with the standard solar model, and that the electron neutrinos produced in the Sun were in fact oscillating into different flavour states on their way to Earth [18]. This was achieved by observing neutrino interactions originating from solar $^{8}$B decays on heavy water, which allowed simultaneous measurement of three independent interaction channels: charged current, neutral current, and elastic scattering. As only the charged-current interaction is sensitive to electron neutrinos, the comparison of neutral-current to charged-current interactions provided an excellent means of determining the existence of neutrino flavour change. Likewise, the ability to measure the total flux of $^{8}$B solar neutrinos provided a good test of contemporary models of neutrino production in the Sun. In one fell swoop, the success of the SNO experiment proved the existence of both neutrino oscillations and the neutrino mass, without which the oscillations themselves would be impossible, and ushered in a new age of neutrino physics.

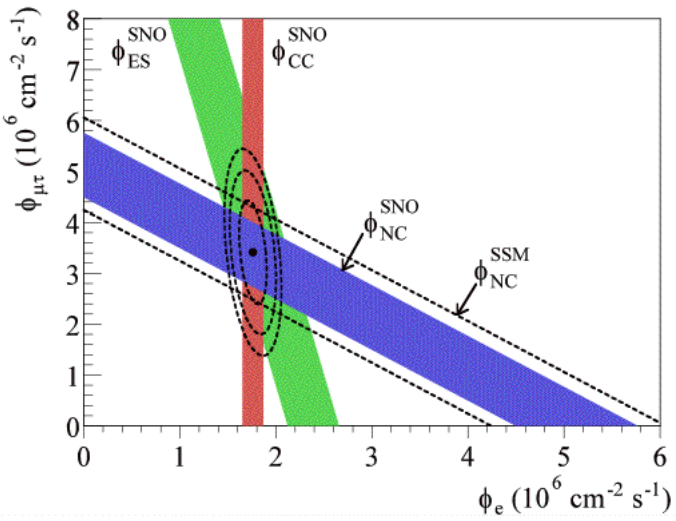



Figure 1.2: Results from the Sudbury Neutrino Observatory, showing the fluxes as determined by the SNO charged current (CC), neutral current (NC), and elastic scattering (ES) interactions[18]. The flux for the neutral current interaction, sensitive to all flavours, matches beautifully with the theoretical prediction of the standard solar model (SSM) as indicated by the dashed lines. The errors represented are $\pm$ 1$\sigma$, whilst the best fit values for the electron neutrino flux $\phi_e$ and the nonelectron neutrino flux $\phi_{\mu,\tau}$ are shown.

In 2012, researchers at Daya Bay, building on results from Double Chooz [19], MINOS [20], T2K [21], and RENO [22], definitively determined a non-zero result for one of the last unmeasured parameters of neutrino oscillation, the mixing angle $\theta_{13}$[2] [23]. In doing so, these experiments have opened the door to studies of CP violation in the neutrino sector, potentially presenting an answer to current theories of leptogenesis in the very early Universe.

## 1.3 A Brief History of Supernovae

Mankind have been aware of supernovae since at least 150 AD, when Chinese astronomers first recorded the appearance of a 'guest star' in the night sky. Over the following centuries, numerous recorded accounts exist of these 'new' stars appearing in the night sky, although it was not until Tycho Brahe's observations of SN1572 that any real empirical studies began in earnest. Indeed, the root of the name 'supernova' comes from the title of Brahe's work on these 'nova stellae' or new stars [24].

[2]Along with the three squared mass differences $\Delta m_{ij}$ and the phase difference $\delta$, the three neutrino mixing angles $\theta_{12}$, $\theta_{23}$, and $\theta_{13}$ parametrise the PMNS mixing matrix and relate the extent of mixing of one neutrino mass state to another. Further detail on mixing angles and the PMNS matrix is given in Section 2.2.1.

As technology advanced, astronomers began to notice existing stars undergoing periodic fluctuations in light output. Christened 'novae', these objects began to be studied in detail in the 19th century, with the first spectroscopic observations taken in 1866 by Huggins yielding highly unusual results, displaying "a spectrum unlike that of any celestial body we have previously examined" [25]. Significantly, Huggins was the first to suggest that the variability in light output of nova could be due to a series of cataclysmic explosions. Once the distance to the Andromeda galaxy was determined at the beginning of the 20th century, it became clear that a number of bright nova events lying along the line of sight of the galaxy were actually occurring within the galaxy itself. Give the magnitude of these nova, it was calculated that these objects were releasing vastly more energy than was first thought. Postulating that these events were caused by the gravitational collapse of massive stars into neutron stars and becoming the first to name them 'super-novae', Fritz Zwicky ushered in the modern era of supernovae theory, and through further searches in the late 1930s identified a number of supernova candidates in distant galaxies [26].

Quickly, it was realised by Walter Baade that these new supernova events could be responsible for the rapid expansion of certain nebula, whilst the efforts of Rudolph Minkowski categorised supernova events in types I and II based on the presence of hydrogen in their emission spectra. Rounding off this earlier work, Fred Hoyle and others proposed the first theory of supernova nucleosynthesis as a precursor to the gravitational collapse of massive stars, leading to the establishment of the modern theory of supernova.

Further progress was made in the later half of the 20th century, particularly in the development of supernovae modelling and observation techniques, and the number of cataloged supernova has risen exponentially. Finally, the observation of SN 1987A[3] by Kamiokande-II [27], Baksan [28], and IMB [29] marked the first detection of supernova via neutrino emission and resulted in the first supernova event to be observed across all bands of the electromagnetic spectrum, ushering in the modern era of neutrino astronomy and the supernova early warning system, SNEWS. The detection of the neutrino burst from SN1987A conclusively demonstrated that neutrinos play a leading role in driving supernovae explosions, and is widely taken as a confirmation of our understanding of the major features of core-collapse supernovae. It additionally allowed major constraints to be placed on exotic cooling mechanisms arising from any new particles emitted from the supernova core, in particular right-handed neutrinos and axions, and kickstarted serious attempts to further model

[3]SN1987A is the label given to the collapse of Sanduleak-$69^0$ 202 in the Large Magellanic Cloud, an approximate distance of 51 kpc from the Earth.

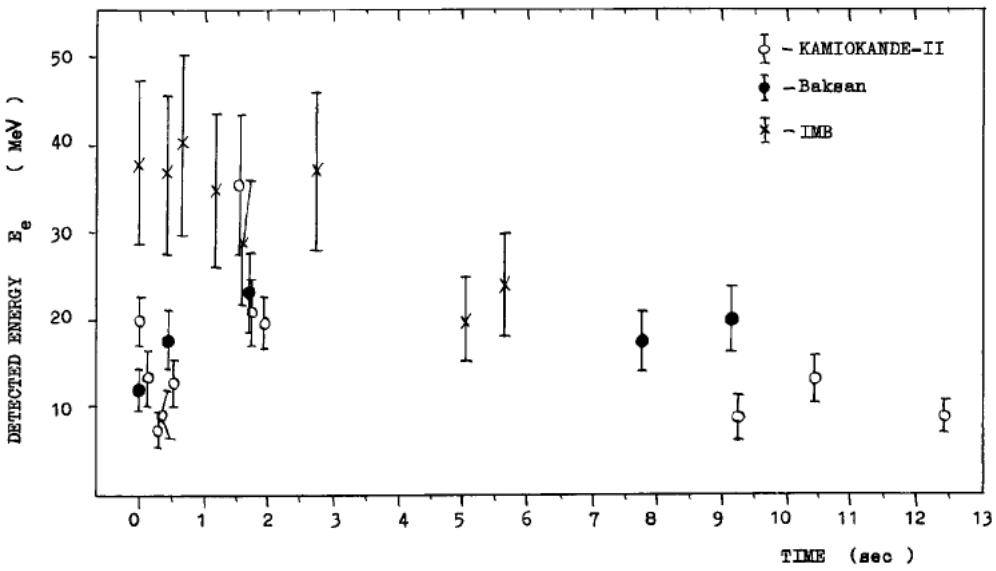



Figure 1.3: Energy of all detected SN1987A neutrino events at 0735 UTC on February 23, 1987. Time t = 0.0 corresponds to the time of the first detected neutrino event for each signal observed.

core-collapse supernovae explosions as a predominantly neutrino-driven effect [30]. A summary of all detected events from SN1987A is shown in Figure 1.3.

## 1.4 Summary

As one of the few known astrophysical processes to straddle the physical disciplines of particle physics, astrophysics, and nuclear physics, the interest in supernovae from the perspective of a neutrino experiment should be clear. The richness of physical information encoded in supernova signals is almost unrivalled in the study of neutrino physics, offering the opportunity to potentially resolve crucial uncertainties in our understanding of fundamental neutrino properties and interactions. With sufficiently high statistics, kilotonne-scale neutrino detectors such as SNO+ have the potential to study rare self-induced oscillation effects, determine the influence of matter effects, and even establish the true neutrino mass hierarchy. Thus, given the timescales involved, the ability of SNO+ to quickly and correctly identify core-collapse supernovae is of paramount importance in supernova data collection and in maximising the physics potential of the events themselves.

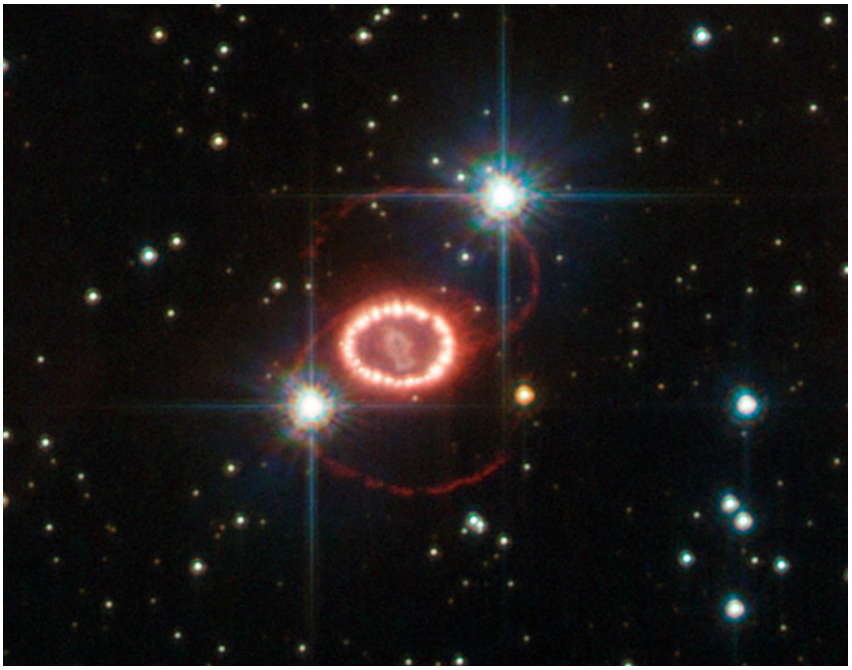


Figure 1.4: Hubble Space Telescope image of the remnant of core collapse supernova SN 1987A in the Large Magellanic Cloud. The three bright rings around the remnant are composed of material ejected via the stellar wind of the progenitor star, Sanduleak -69° 202, prior to explosion.

# Chapter 2

# Theory

## 2.1 Supernovae

Core collapse supernovae (CCSNe) are some of the most cataclysmic events in our universe. During the brief period of gravitational collapse and resulting explosion, the supernova releases the same amount of energy as the star radiates over its entire lifetime [5]. Uniquely, in a core collapse supernova, over 99% of the initial gravitational binding energy is emitted as neutrinos, around 10% of the initial mass of the star. This link between neutrino physics and astrophysics provides a unique opportunity to study both the intrinsic properties of the neutrino and the mechanism of core collapse.

### 2.1.1 Classification

For historical reasons, supernovae (SNe) are broadly divided into Type I and Type II categories based on their spectral output and light curves via the Minkowski-Zwicky scheme, shown in Figure 2.1. Type II SNe posses lines of hydrogen in their absorption spectra, whilst Type I do not. These types are then further subdivided based on their individual spectral characteristics, with each of these subdivisions essentially representing a different mechanism for stellar explosion. As this work is focused on the detection of cosmic neutrinos emitted from core collapse supernovae, i.e. SNe Types Ib/c and II, the remainder of the physics discussion will focus on this mechanism, but a brief outline of the entire supernova taxonomic system is given below for completeness.

#### 2.1.1.1 Type Ia

Type I supernovae can be further subdivided into Types Ia, Ib, and Ic depending on the presence of singly ionised silicon in their absorption spectra; Type Ia presents this line, at 615nm, whilst Type Ib/c do not. Type Ib and Ic and then differentiated by the presence or lack of helium in their absorption spectra respectively.

Type Ia supernovae explosions are generated by the accretion of material onto a pre-existing white dwarf core, resulting in violent, runaway re-ignition of nuclear fusion, known as carbon detonation. As the the white dwarf accretes material, it comes to within 1% of the Chandrasekhar limit[1] before carbon fusion is ignited, preventing collapse into a black hole. As white dwarfs are supported against collapse solely by electron degeneracy pressure rather than thermal pressure, fusion reactions cannot be regulated by a period of cooling expansion as with main sequence stars. As a consequence, once fusion is reignited, it becomes a runaway reaction and the white dwarf 'detonates'. Because of the uniform mass of white dwarfs that explode via this mechanism, Type Ia SNe have a consistent peak luminosity, and are routinely used as standard candles by the astronomical community [31].

#### 2.1.1.2 Type Ib/c

Whilst Type Ib/c are included in the Type I class due to their lack of hydrogen, in practice they are far closer to Type II SNe than Type Ia in that they are triggered by core collapse rather than accretion. They are thought to be the consequence of the collapse of massive stars which have exhausted, or been stripped of, their entire reserve of hydrogen [32]. In the case of Type Ic SNe, the majority of their helium reserves have also been depleted before collapse. The progenitors of these supernova types are thought to be very massive stars at late stages of evolution which have shed their outer layers of light elements to extremely powerful stellar winds. Due to the fact that Type Ib/c SNe only occur in these particular stars, they are exceedingly rare. Other than this, the mechanism of collapse is very similar to Type II SNe, and consequently are included in the more general classification of CCSNe.

#### 2.1.1.3 Type II

Type II supernovae are all triggered by the core collapse of a single massive star after the exhaustion of available fuel. Depending on the initial mass and metallicity of the star, the resulting implosion

[1]The Chandrasekhar limit is defined as the maximum mass of a white dwarf star able to be sustained by electron degeneracy pressure alone. Above this limit, the white dwarf undergoes catastrophic collapse into a neutron star. The currently accepted value for the Chandrasekhar mass limit is 1.44 solar masses.

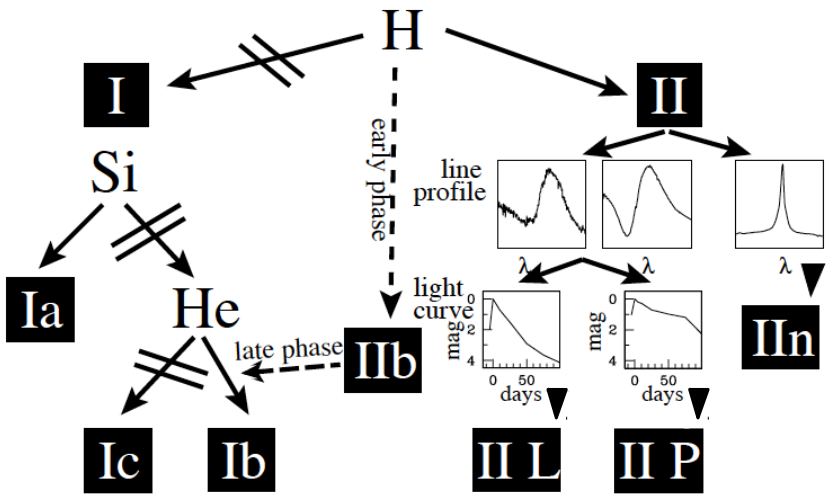



Figure 2.1: The Minkowski-Zwicky supernova classification scheme, displayed as a simple flowchart. Detailed supernova classification requires identification of spectral features, coupled with analysis of line profiles, luminosity, and spectral evolution [33].

then either results in the violent expulsion of stellar material, or in the case of complete gravitational collapse, a black hole. Type II SNe are further categorised based on the the light curve of the resulting explosion [33]. Type II-L demonstrate a linear decline in luminosity over time whilst Type II-P show evidence of a plateau. This difference is thought to be caused by the presence of the hydrogen envelope - Type II-L have already shed their outer layer of hydrogen, whilst in Type II-P SNe the shock front ionises the hydrogen shell, resulting in a significant increase in opacity.

Type II SNe can then be further subdivided into categories based on their spectral characteristics. Type IIn SNe show the presence of a narrow hydrogen peak in their spectra, associated with very high mass loss progenitor stars, where Type IIb have a weak hydrogen emission line in their initial spectra which vanishes over time, before presenting a spectrum similar to a Type Ib SNe. It is therefore thought that Type IIb SNe may have progenitors that undergo mass loss via Roche-lobe overflow[2] in a binary system. Indeed, post explosion observations for SN 1993J have revealed a compact blue companion alongside the fading supernovae, and similar results have been claimed for SN 2011dh [34].

[2]The Roche lobe describes the region around a single star in a binary star system within which objects are gravitationally bound to that star. Roche lobe overflow occurs when stellar material overflows this boundary and begins to fall toward the companion star.

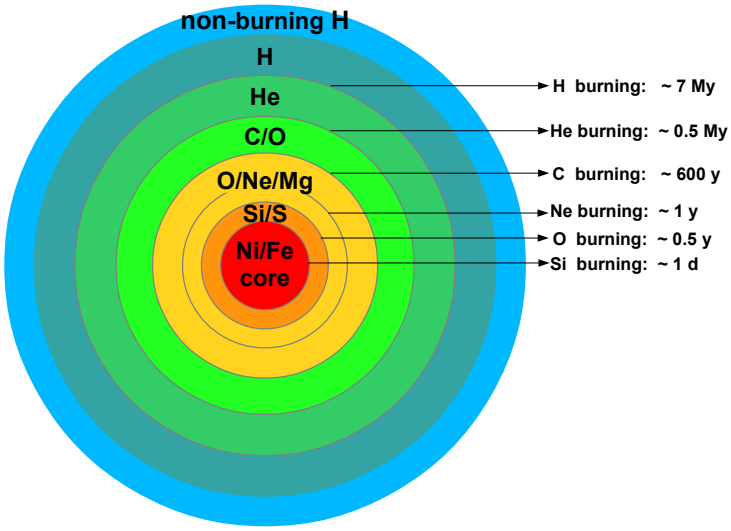



Figure 2.2: Shell structure schematic of a Type II core-collapse supernova progenitor star just prior to collapse, after formation of the iron core. Typical timescales for the burning of different elements during nucleosynthesis are given, and calculated for a star with a progenitor mass of $M_P = 25M_\odot$. The radial thickness of the burning layers is not shown to scale [35].

### 2.1.2 Core Collapse Supernovae

#### 2.1.2.1 Overview

Fundamentally, the collapse of any massive star results from the inability of the star to sustain its own mass via the pressure generated from nuclear fusion reactions in the core. This collapse occurs in stages as the star transitions from burning hydrogen, to helium, to progressively heavier elements depending on the initial mass of the star. As the fusion reactions of each stage become exhausted, the star collapses until the gravitational pressure generated is sufficient to reignite fusion in progressively heavier elements, producing a characteristic "onion-shell" structure, shown schematically in Figure 2.2. Each step produces progressively heavier nuclei, which produce progressively less energy via fusion. This process continues until the inner core of the star consists almost entirely of iron, which no longer releases energy via fusion. At this point, the core is under enormous gravitational pressure and is only supported against collapse by electron degeneracy pressure. Once the mass of the core exceeds the Chandrasekhar limit, degeneracy pressure is no longer sufficient to support the core from collapsing in on itself and the star begins to implode. It is this final phase of collapse that is responsible for the vast majority of neutrinos produced during a supernova, and the mechanism is detailed below.

#### 2.1.2.2 Final Stages of Core Collapse

As shell burning progresses, silicon burning deposits increasing quantities of iron 'ash' in the core of the progenitor star. Once the mass of this core exceeds the Chandrasekhar mass of 1.44 $M_{\odot}$, electron degeneracy pressure is no longer capable of supporting the star against collapse, and the core begins to fall in on itself. At this stage, the rate of collapse is highly sensitive to the total number of electrons present in the core; this is dictated by the competing weak processes of electron capture and beta decay. Partial photodisintegration of the iron nuclei of the core serves to further accelerate the pace of collapse, removing energy from the core and decreasing the number of electrons available to slow collapse. At densities of around $10^{12}$ g cm$^3$ the core begins to 'trap' neutrinos as their scattering time becomes larger than the time taken for core collapse [3].

Eventually the inner core reaches densities of $10^{14}$ g cm$^3$, comparable to those of the atomic nucleus. At this stage, the inner core experiences an abrupt decrease in compressibility, forcing the in-falling material to decelerate and rebound, creating a shock front which rapidly drives through the outer core, dissociating heavy elements into their constituent nuclei. This dissociation causes a massive increase in electron capture events on free protons, resulting in a burst of electron neutrinos and furthering the neutronisation of the core. The shock then weakens, preventing runaway expansion and allowing material to continue to fall inward toward the core remnant.

At this point the core remnant is supported solely by neutron degeneracy pressure, and has essentially become a proto-neutron star. Due to the pressures and temperatures present, the remnant can now only cool via neutrino pair-production and diffusive loss of neutrinos of all flavours. After around 10s, the core becomes fully transparent to neutrinos, rapidly radiating a large fraction of its energy away in this form. The combination of these effects release a large fraction of the total energy of the neutron core, around $10^{46}$ J, in a ten second burst of neutrinos. These neutrinos then re-energise the stalled shock via heating of the high-density shock front, causing rapid expansion and driving the final explosion of the star.

The final outcome of a core-collapse supernova is dependent solely on the progenitor mass. Below around 20 $M_{\odot}$, the collapse simply results in a neutron star remnant. Above this, the remnant eventually collapses to a black hole, although with an associated explosion. For a progenitor star above 50 $M_{\odot}$, a star is thought to collapse directly into a black hole with no associated explosion, although this upper limit is highly uncertain and model-dependent [36].

#### 2.1.2.3 Neutrino Energy Loss Mechanisms and Spectra Formation

As we now know, the vast majority of the energy loss undergone during stellar collapse is radiated away via neutrinos, with only a negligible fraction lost through electromagnetic radiation and kinetic energy. As the star collapses, it experiences a number of stages in which it undergoes this energy loss to neutrinos of different flavours, and this in turn dictates the final spectra of neutrinos observable to SNO+ and other neutrino detectors.

During infall, neutrino production is dominated by electron capture on heavy nuclei, including $^{56}$Fe, and free protons. Prior to neutrino trapping at densities of around $10^{11} - 10^{12}$g cm$^3$, these neutrinos radiate away freely with increasing luminosity as high energy $\nu_e$, reaching a total energy output of around $10^{51}$ ergs. [37]

At the point of core bounce, a huge burst of electron capture events occur on the remaining $^{56}$Fe nuclei, due to the lower Q value of the interaction compared to that of free protons. At densities of $10^{14}$ g cm$^3$, this produces a wave of trapped $\nu_e$ immediately behind the travelling shock front [38]. As the shock is opaque to $\nu_e$ initially these neutrinos build up until the shock reaches densities of $10^{11}$g cm$^3$, at which point the shock front becomes transparent to $\nu_e$ and the neutrinosphere forms, finally releasing this build up as an enormous burst of neutrinos known as the deleptonisation burst [3].

This energy loss, around $10^{51}$ ergs, severely weakens the shock and causes it to stall at a post-bounce distance of 200-300 km, around 100 ms from the time of shock formation. As the shock stalls, matter in the outer core continues to fall toward the centre, but is rapidly decelerated by its passage through the shock front, resulting in quasi-stationary accretion onto the core remnant. Unless this mass accumulation is sufficient to result in black hole formation, the shock front is revived around 500 ms later, ejecting the outer stellar material and causing the final, delayed explosion of the star [39].

After core bounce and during the accretion period, neutrinos of all flavours are formed thermally through a number of interactions in addition to those $\nu_e$ produced via electron capture, as shown in Table 2.1. The behaviour of these neutrinos then dictates the shape of the final neutrino energy spectrum of the supernova.

The passage of neutrinos through the supernova depends entirely on their flavour at production; $\nu_x$[3] can only interact via neutral-current, and therefore exhibit very different behaviour from $\nu_e/\bar{\nu}_e$

[3] $\nu_x$ describes neutrino flavours $\nu_\mu$ and $\nu_\tau$, both of which may only undergo neutral current interactions and are therefore considered indistinguishable. Additionally, $\nu_x$ is taken to refer to both $\nu_x/\bar{\nu}_x$ due to their identicality for the purposes of this discussion, in contrast to $\nu_e/\bar{\nu}_e$ which behave very differently within supernova.

| Interaction | Mechanism | Flavours produced |
|---|---|---|
| $e^+ + n \rightarrow p + \bar{\nu}_e$ | Positron capture on neutrons | $\nu_e$ |
| $N + N \rightarrow N' + N' + \nu + \bar{\nu}$ | Nucleon-nucleon bremsstrahlung | All flavours |
| $e^\pm + N \rightarrow e^{\pm'} + N' + \nu + \bar{\nu}$ | Electron-nucleon bremsstrahlung | All flavours |
| $e^+ + e^- \rightarrow \nu + \bar{\nu}$ | Electron-positron bremstrahlung | All flavours |
| $\gamma \rightarrow \nu + \bar{\nu}$ | Plasmon decay | All flavours |
| $\gamma + e^\pm \rightarrow e^\pm + \nu + \bar{\nu}$ | Photoannihilation | All flavours |
| $\nu_e + \bar{\nu_e} \rightarrow \nu_{\mu,\tau} + \bar{\nu}_{\mu,\tau}$ | $\nu_e$-$\bar{\nu}_e$ annihilation | $\nu_x$ |

Table 2.1: Table of thermal neutrino production interactions in core collapse supernovae, taken from [40]. All neutrinos above have energies ∼40 MeV as determined by the surrounding temperatures within the supernova itself. Where $\nu_x$ can only react via neutral-current interactions, $\nu_e$ and $\bar{\nu}_e$ can also undergo charged-current interactions.

which can also interact via charged-current interactions. This has a number of consequences, but of particular importance is the formation of neutrinospheres at different radii for different neutrino flavours, as determined by the variation in production rates and interaction cross sections of each flavour. Neutrino-neutrino coherent scattering is a major contributor to neutrino opacity, and therefore the neutrino density of each flavour greatly affects the final position of the neutrinosphere for any given flavour [41].

$\nu_e/\bar{\nu}_e$ are kept in thermal equilibrium by charged current processes until they decouple from the neutrinosphere, exhibiting a thermal spectrum that corresponds to the matter temperature at that radius. As the density of protons is much higher than that of neutrons, the rate of electron capture far outreaches that of positron capture, leading to a $\nu_e$ neutrinosphere at a greater radius, and at a lower temperature, than that of $\bar{\nu}_e$. The neutrino cross sections of all flavours increase with $E^2$, and as the different energy groups decouple at different temperature and radii, the neutrinospheres actually have some finite thickness and therefore their neutrino spectra are actually averaged over several temperatures and exhibit only quasi-thermal behaviour.

$\nu_x$ are kept in thermal equilibrium by neutral current processes, in a similar manner to $\nu_e/\bar{\nu}_e$ and charged current interactions, until decoupling at some characteristic radius and temperature. This position is much closer to the core than the decoupling point of $\nu_e/\bar{\nu}_e$ due to fewer sources of interaction, and therefore $\nu_x$ decouple at higher energies. These $\nu_x$ do not free stream from this position as $\nu_e/\bar{\nu}_e$ do however, and instead undergo elastic scattering interactions with heavy nucleons in the outer core at a point known as the energysphere [42]. The average energies of these $\nu_x$ however, around 10 MeV, mean that these scattering processes are inefficient at energy transfer and therefore allow the $\nu_x$ to diffuse outwards until they reach thermal equilibrium with colder matter at the outer edge of this transport sphere. Thus, the energy spectra of $\nu_x$ is smeared out and it's mean energy $\langle E_{\nu_x} \rangle$ is shifted slightly downwards compared to the quasi-blackbody spectra of $\nu_e/\bar{\nu}_e$. This results

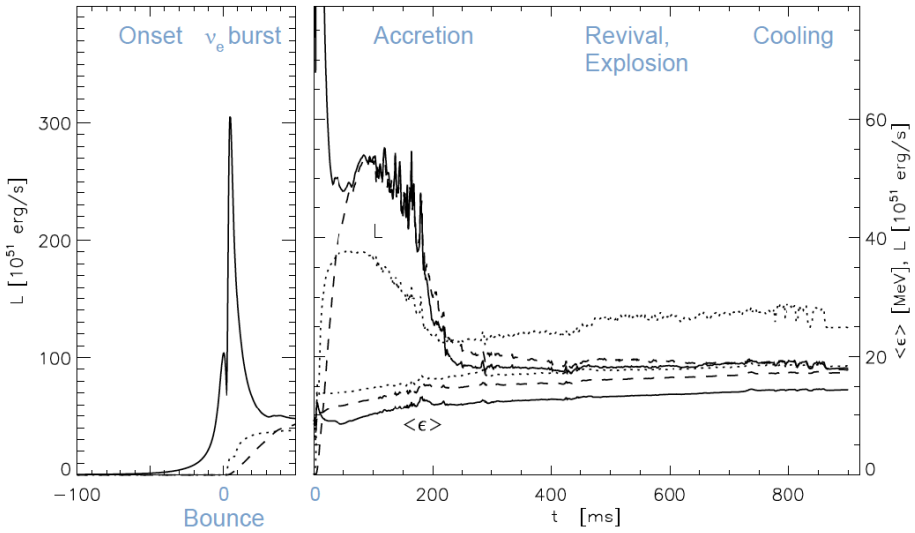



Figure 2.3: Sample time evolution of the luminosities of the different neutrino flavours $\nu_e$ (solid line), $\bar{\nu}_e$(dashed line), and $\nu_x$ (dotted line) after simulation of the explosion of a progenitor star with mass $M_P = 15M_\odot$, with the explosion obtained after manipulation [43]. The post bounce phase is shown on the right side of the plot, and also shows the mean energies of the emitted neutrinos $\langle E \rangle$ .The cooling phase is not shown, and lasts approximately 10 seconds with roughly exponentially decreasing luminosity.

in a 'pinched', narrower spectral width for $\nu_x$ than its $\nu_e/\bar{\nu}_e$ cousins [42].

From this, it can be seen than the mean energies of $\nu_e/\bar{\nu}_e/\nu_x$ follow the hierarchy:

$$\langle E_{\nu_x} \rangle \gtrsim \langle E_{\bar{\nu}_e} \rangle > \langle E_{\nu_e} \rangle \tag{2.1}$$

The position of each neutrinosphere, $R_\nu$, are expected between around 50 km and 100 km from the core, smaller than the radius at which the shock stalls. This creates a luminosity 'bump' for each flavour of neutrino, where the neutrinos that have built up behind the shock wave are released as the shock passes each respective neutrinosphere [37]. The only exception to this exists for higher energy neutrinos, which by nature exhibit a higher interaction rate and therefore decouple later due to higher opacity. In addition, neutrinos produced deep in the core of the supernova take hundreds of milliseconds to reach each respective neutrinosphere, and continue to stream during shock stagnation. The time evolution of neutrino luminosities for a core collapse supernova of progenitor mass $15M_\odot$ is shown in Figure 2.3.

#### 2.1.2.4 Current Modelling

It is now known that the energy of the shock front is insufficient to lead to runaway explosion after the energy removal caused by the deleptonisation burst for progenitors above ∼8-10 $M_{\odot}$. Whilst neutrino heating of the stalled shock front is known to play a large part in driving final stellar dissociation, 1D spherically symmetric models fail to adequately explain the phenomenon. The shock is therefore expected to be revived by some multi-dimensional combination of mechanisms, currently poorly understood.

Some fraction of neutrinos emitted from the various neutrinospheres are known to deposit energy in the dense post-shock region, mostly through $\nu_e$ and $\bar{\nu}_e$ captures on free nucleons. This heating then competes with neutrino-cooling effects from diffusive losses detailed in Table 2.1. Close to the neutrinospheres, neutrino cooling dominates and creates a cooling layer, though this cooling rate declines faster than the heating rate with distance from the core. Thus, there exists a point known as the gain radius $R_g$, at which heating begins to dominate. If sufficient energy can be deposited between the gain layer and the stalled shock position then, in theory, the shock can be revived and the final explosion explained. Whilst neutrino heating alone is known not to be sufficient to drive the final expansion, it is thought that some combination of neutrino heating, convective processes, and hydrodynamic instabilities in the post-shock layer may impart enough energy to cause explosion. Though 1D spherically symmetric models have been shown to exhibit major problems in adequately explaining supernova, 2 dimensional axisymmetric models have been far more successful [44]. Axisymmetric models exhibit both convective heating behaviour and anisotropic effects in shock formation[4] , permitting a reduction in the critical neutrino luminosity required for explosion compared to spherically symmetric models, and providing a basic confirmation of the neutrino-driven explosion theory [47].

The critical luminosity, $L_{\nu,c}(\dot{M})$, described above provides a parameter for the runaway conditions required of supernova, and acts as the threshold above which runaway expansion must occur. It depends on the mass accretion rate of the shock $\dot{M}$: steady accretion cannot be maintained for $L_{\nu} > L_{\nu,c}(\dot{M})$. $L_{\nu,c}(\dot{M})$ increases with $\dot{M}$ as heating must overcome the pressure from in-falling material, whilst higher progenitor masses also increase $L_{\nu,c}(\dot{M})$ in line with the higher gravity of

[4]Recent simulations have demonstrated a new anisotropic effect named SASI, or Standing Accretion Shock Instability, which is thought to help explain the explosion mechanism of heavier progenitors - standing waves cause distortions in the shock front, leading to non-uniform heating, allowing some areas of the front to expand further than others and thus increasing the time available for heating. These instabilities result in deformed explosions, additionally helping to explain both asymmetries in observed supernovae ejecta, and the measured proper motion and rotation of young pulsars [45][46].

the proto-neutron star. For runaway expansion to occur, post shock material must be heated faster than it can be accreted downward into the cooling layer, imposing a timescale requirement on the process: the timescale for neutrino heating must be less than the timescale of advection through the gain layer, $t_\nu < t_{\text{advection}}$. The critical luminosity can then be determined as follows:

$$L_{\nu,c}(\dot{M}) \propto \beta^{-\frac{2}{5}} \dot{M}^{\frac{2}{5}} M_{NS}^{\frac{4}{5}} \tag{2.2}$$

Where $\beta$ is the ratio of the post-shock and pre-shock densities, and $M_{NS}$ is the mass of the proto-neutron star[40].

In 2D simulations, $L_{\nu,c}(\dot{M})$ is greatly reduced by convective overturn effects within the post-shock layer [43]. These effects describe non-radial, partially turbulent mass motions which increase the dwell time of material in the the gain layer. Convective downdrafts carry post-shock material to the vicinity of $R_g$, whilst the outward rise and expansion-cooling of buoyant neutrino-heated bubbles reduce neutrino-cooling by remission and push the shock farther out, further prolonging the dwell time of material in the gain layer. In addition, non-radial oscillations and shock 'slosh', also have favourable effects on neutrino driven explosive conditions by reducing $L_{\nu,c}(\dot{M})$. The combination of these effects lift core conditions very close to the critical threshold required for explosion [44].

3D models confirm the observations seen in 2D models, showing shock expansion to be quasi-stationary, and driven by the inflation of neutrino-heated bubbles. These bubbles then fragment via turbulence effects, and lead to potentially large shock deformation, pushing the shock outward in all directions more continuously then the effects shown in 2D simulations. That said, so far 3D models have been inconclusive; higher order shock effects have been been shown to occur, though not at sufficient levels to account for explosion, whilst some new effects seen actually increase $L_{\nu,c}(\dot{M})$. It is possible that some 3D models exclude strong low-order modes of post-shock mass motions, potentially exceeding threshold, and that neutrino driven convective effects actually destroy SASI coherence. Likewise, it is also possible that hydrodynamic instabilities which are highly dependent on initial progenitor conditions and asymmetries during infall [48], excite oscillations in the post-bounce regime and are responsible for the onset of explosion, though this remains unclear at best.

Ultimately, while neutrino luminosities, average energies, and neutrino energy spectra are all greatly affected by the the approximation of progenitor properties, and are therefore heavily model dependent, the fundamental features of supernova neutrino emission are generic and their overall trends comparable [44]. Thus, provided the time evolution of parameters agree with theory qual-

itatively, detector sensitivities remain valid across all models. Consequently, for the purposes of supernova monitoring at SNO+, and therefore for the purposes of this thesis, the absence of a conclusive theory of neutrino-driven supernova explosions is of little effect.

## 2.2 Neutrino Properties

Neutrinos are neutral spin-$\frac{1}{2}$ fundamental particles that interact solely through the weak interaction, and are therefore classified as leptons. Initially thought to be massless, recent discoveries have confirmed that neutrinos posses infinitesimal but definite mass, and therefore also interact gravitationally. Lacking any comprehensive quantum theory of gravity however, and considering the extremely weak effect of gravity on a subatomic scale, this aspect can safely be neglected when discussing neutrino interactions. The discovery of a definite neutrino mass does have a number of interesting consequences however, presented in detail below.

### 2.2.1 Neutrino Oscillations

Neutrinos are unique in that they are the only known fundamental particle to undergo oscillation. After the first observation of the quantum mechanical oscillation of neutral kaon states in the 1960s, neutrino oscillations were first conceived as simple transitions between neutrino and antineutrino states. The later discovery of higher lepton flavours expanded the hypothesis to include three neutrino flavour and mass eigenstates. Whilst all neutrinos are produced with some definite flavour eigenstate $\mid \nu_\alpha >$, this 'weak eigenstate' is in fact a coherent superposition of the mass eigenstates $\nu_i$:

$$\mid \nu_\alpha >= \sum_i U_{\alpha i} \mid \nu_i > \tag{2.3}$$

Where U is some mixing matrix. In the case of the electron neutrino $\nu_e$ for example:

$$\mid \nu_e >= U_{e1} \mid \nu_1 > + U_{e2} \mid \nu_2 > + U_{e3} \mid \nu_3 > \tag{2.4}$$

In the general case of three neutrino oscillation, transformation from a general neutrino flavour basis $\nu_\alpha$ to a general mass basis $\nu_i$ is then conducted as follows:

$$\begin{bmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{bmatrix} = \begin{bmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{bmatrix} \begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{bmatrix} \tag{2.5}$$

As the neutrino propagates, the mass eigenstates each propagate at a different velocity, resulting in differing superpositions of mass eigenstates over time as the relative phases between each eigenstate shift. These coherent superpositions then manifest themselves over a baseline L as different flavour eigenstates, responsible for the detection and production of the particle we identify as the neutrino via the weak interaction [49]. The predecessor to SNO+, the Sudbury Neutrino Observatory, confirmed the existence of neutrino oscillations in 2001 by observing events sensitive to all species, thereby solving the famous solar neutrino problem.

#### 2.2.1.1 Two Neutrino Approximation

Whilst a detailed treatment of three-neutrino oscillations will not be given, as the mixing terms rapidly become cumbersome, it is possible to derive the neutrino oscillation probability $P(\nu_\alpha \rightarrow \nu_\beta)$ by approximating the mixing in terms of two neutrinos Using the 2x2 mixing matrix U:

$$U = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \tag{2.6}$$

The mixing of the flavour eigenstates $\nu_\alpha$, $\nu_\beta$ with mass eigenstates $\nu_1$, $\nu_2$ is then given via the following relation:

$$\begin{pmatrix} \nu_\alpha \\ \nu_\beta \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} \tag{2.7}$$

Using the plane wave solution to the Schrodinger equation, the time evolution of neutrino states at some point (x, t) is then given by the following:

$$\begin{pmatrix} \mid \nu_\alpha(x,t) > \\ \mid \nu_\beta(x,t) > \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} e^{-i\phi_1} & 0 \\ 0 & e^{-i\phi_2} \end{pmatrix} \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} \mid \nu_\alpha(0,0) > \\ \mid \nu_\beta(0,0) > \end{pmatrix} \tag{2.8}$$

The probability of finding $\nu_\beta$ at some distance from a source of pure $\nu_\alpha$ is then simply given by the square of the amplitude:

$$P(\nu_\alpha \to \nu_\beta) = |< \nu_\beta(x,t) \mid \nu_\alpha(0,0) >|^2 \tag{2.9}$$

Which, after simplification, yields:

$$P(\nu_\alpha \to \nu_\beta) = \sin^2(2\theta)\sin^2(\frac{\Delta m^2 L}{4E_\nu}) \tag{2.10}$$

A full derivation can be found in [50].

Using the two neutrino approximation above, the probability of detecting a neutrino of flavour $\beta$ produced as flavour $\alpha$, $P(\nu_\alpha \to \nu_\beta)$, is then:

$$P(\nu_\alpha \to \nu_\beta) = \sin^2(2\theta)\sin^2(1.27\Delta m^2\,[\mathrm{eV}^2]\frac{L\,[\mathrm{km}]}{E\,[\mathrm{GeV}]}) \tag{2.11}$$

The angle $\theta$ denotes the mixing angle, and relates the flavour states to the mass states. A mixing angle of zero would result in no oscillation from the source neutrino as the flavour states would be identical to the mass states, whilst a maximal mixing angle of $\frac{\pi}{4}$ would ensure that all $\alpha$ neutrinos would oscillate to flavour $\beta$ at some distance L.

The mass-squared difference term ensures that the neutrinos must have mass, and that the masses must be different, for without it the oscillation probability would yield zero. It also dictates the limitations of neutrino oscillation experiments; they can only determine the differences of the masses, rather than the absolute masses of each species. Additionally, the term gives no information as to the mass hierarchy: $m_1$ may be larger than $m_2$, but $\Delta m^2 \to -\Delta m^2$ yields the same oscillation probability.

Whilst the last term, $\frac{L}{E}$, can be fully manipulated by experiment, for a fixed production source such as supernova this term dictates the range of ($\theta$,$\Delta m^2$) over which we may observe neutrinos as the oscillation probabilities for other combinations become vanishingly low.

#### 2.2.1.2 Three Neutrino Mixing

Currently, the single best description of the observed data on neutrino oscillations is given by assuming three-flavour neutrino mixing in vacuum. For three generations, neutrino oscillations are entirely governed by 7 parameters: three mixing angles $\Delta m^2_{ij}$, three squared mass differences $\theta_{13}$, and one phase $\delta$. Using these, we can then parameterise the above mixing matrix U as follows:

$$U = \begin{bmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{bmatrix} \tag{2.12}$$

Where $s_{ij} = sin\theta_{ij}$ and $c_{ij} = cos\theta_{ij}$, $\theta_{ij} = [0, \pi/2]$ ,$\delta = [0, 2\pi]$

This is the Pontecorvo-Maki-Nakagawa-Sakata mixing matrix, or PMNS matrix, and can be thought of as the leptonic counterpart to the CKM mixing matrix for quarks. As of 2015, all mixing parameters have been observed and measured other than the sign of $|\Delta m_{23}|$ and $\delta$. The phase $\delta$ is responsible for any potential CP violation observed in the neutrino sector, and is non-zero only for this case. As this is yet to be shown, it leads to complications in the accurate determination of $\theta_{13}$ due to the presence of $\delta$ in the matrix element $U_{e3}$, despite the recent confirmation of a non-zero $\theta_{13}$ result from the Daya Bay experiment [23].

While the absolute magnitude of $|\Delta m_{23}|$ has been determined through measurement of atmospheric neutrino oscillations $\nu_\mu \to \nu_\tau$, without the presence of matter effects the absolute magnitude of the mass difference is impossible to determine, and therefore the actual neutrino mass hierarchy remains unclear. In contrast, the absolute value of $\Delta m_{12}$ has been determined via measurements of the oscillation of solar neutrinos, which are subject to matter effects as they propagate through the Sun, thus allowing us to resolve the mass ordering $m_1 < m_2$. More details on this mechanism are given in the following subsection.

Currently, the PMNS matrix is thought to be unitary, leading to a number of useful relations through the unitary matrix properties:

$$U^\dagger U = I \to U^\dagger = U^{-1} = (U^*)^T$$

Hence:

$$\begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{bmatrix} = \begin{bmatrix} U_{e1}^* & U_{\mu 1}^* & U_{\tau 1}^* \\ U_{e2}^* & U_{\mu 2}^* & U_{\tau 2}^* \\ U_{e3}^* & U_{\mu 2}^* & U_{\tau 3}^* \end{bmatrix} \begin{bmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{bmatrix} \tag{2.13}$$

And:

$$\begin{bmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{bmatrix} \begin{bmatrix} U^*_{e1} & U^*_{\mu 1} & U^*_{\tau 1} \\ U^*_{e2} & U^*_{\mu 2} & U^*_{\tau 2} \\ U^*_{e3} & U^*_{\mu 2} & U^*_{\tau 3} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{2.14}$$

Should the PMNS matrix be shown not to be unitary through experiment, then either a sterile neutrino or new physics would be required. Any sterile neutrino could also exist as an extension to the 3 x 3 mixing matrix, extending it to 4 x 4 or even 5 x 5 depending on the number of sterile neutrinos. If the neutrino is a Majorana particle[5] [51], two additional phases are required, $\alpha, \beta$ which influence the rate of any neutrinoless double beta decay process [49].

### 2.2.2 Matter Effects

As neutrinos propagate through a dense material such as the Sun, they experience interactions in matter which modify the mixing angles and mass differences from their vacuum values [52][53]. This occurs due to the presence of electrons in dense matter which interact only with electron neutrinos at typical energies, increasing the effective mass of $\nu_e$ and resulting in the addition of a potential term $H_{int}$ to the Hamiltonian of the $\nu$ system in matter, $H_m = H_0 + H_{int}$. Where oscillations in vacuum occur due to the phase differences arising from differences in mass of the neutrino eigenstates, and therefore differences in their velocities, oscillations in matter are caused by the difference in effective potential between different flavour eigenstates. Using a two neutrino approximation, $\nu_\alpha \rightarrow \nu_\beta$, this results in a modification to the mixing angle in matter $\theta_m$ given by [54]:

$$\sin 2\theta_m = \frac{\sin 2\theta}{\sqrt{(\frac{\Delta V}{\Delta m^2} - \cos 2\theta)^2 + \sin^2 2\theta}} \tag{2.15}$$

Where $\theta$ is the mixing angle in vacuum, and $\Delta m^2$ is the square of the mass differences in vacuum. $\Delta V = V_\alpha - V_\beta$, the difference in potential experienced by each neutrino flavour.

This expression yields a number of interesting consequences. Firstly, when $\Delta V \rightarrow \infty$, $\sin 2\theta_m \rightarrow 0$, and therefore oscillations cannot occur. This suggests that in extremely dense matter, such as in stellar cores, no oscillation occurs. Additionally, there also exists a point where the value of the electron number density, and therefore the potential, forces the oscillation probability to 100%. This is known as the MSW resonance, and occurs when the ratio between the potential difference and the square of the mass difference $\frac{\Delta V}{\Delta m^2} = \cos 2\theta$, oscillating all neutrinos from one flavour to the other.

[5] A Majorana particle is defined as any fermion that is its own antiparticle. Majorana neutrinos are explored in depth in Section 2.2.3.

The most useful consequence of this relationship however is the sign dependence of $\Delta m^2$: swapping $\Delta m^2 \rightarrow -\Delta m^2$ yields a different result for the effective mixing angle $\theta_m$. This relationship was used in solar neutrino oscillation measurements to determine the mass ordering of $m_1$ and $m_2$, though it is not yet known if $m_1 < m_2 < m_3,$, or $m_3 < m_1 < m_2$, and can also be employed in analyses of supernova data, thus potentially resolving the orientation of the neutrino mass hierarchy.

Whilst it is theoretically possible to determine the mass hierarchy through measurement of $|\Delta m^2_{13}| \approx |\Delta m^2_{23}|$ via long baseline experiments, it is definitely possible to determine the hierarchy through supernova neutrino fluxes, either through direct measurement or through the observation of Earth matter effects [55]. The latter method has the advantage of being completely model-independent, and can take advantage of the SNEWS network to search for signal differences in detectors positioned on opposite sides of the planet, for example IceCube and SNO+. Provided that one detector directly measures the signal, the second detector will be shadowed by the Earth to some degree, and differences in the neutrino flux between the shadowed and unshadowed detectors should point conclusively to either the normal or inverted mass hierarchy.

Direct measurements of neutrino flux from supernova should present features characteristic of mixing, and therefore can also be used to determine the mass hierarchy albeit in a model-dependent way. The spectral characteristics of a supernova neutrino burst are key to this, with the initial deleptonisation/neutronisation burst exhibiting different flavour compositions at Earth depending on the mass hierarchy: normal hierarchy will show a pure $\nu_x$ peak with a composite spectrum of $\bar{\nu}_e$, whilst inverted would show a mixed flavour $(\nu_e, \nu_x)$ neutronisation peak with a composite spectrum of $\bar{\nu}_e$. In addition, the Earth matter effects should only be present in a specific channel for a given hierarchy: $\bar{\nu}_e$ for the normal hierarchy, and $\nu_e$ for the inverted hierarchy. Finally, the mean energy ordering of the supernova neutrinos observed on earth will be an excellent indicator of neutrino mixing itself. We know, due to the energy dependence of neutrinosphere-decoupling, that neutrinos of different flavours will be created from a supernova at mean energies inversely proportional to their respective neutrinosphere radius: $\langle E_{\nu_e} \rangle < \langle E_{\bar{\nu}_e} \rangle \lesssim \langle E_{\nu_x} \rangle$. Consequently, the degree to which the mean energies as measured on Earth differ from this ordering can tell us substantial amounts about the mixing process.

### 2.2.3 Neutrino Mass and the Seesaw Mechanism

The extremely small mass of the neutrino, less than a millionth of the electron mass, is as of yet unexplained. One potential solution lies with the so-called 'seesaw' mechanism, in which an

extremely heavy, right-handed neutrino exists in addition to the lightweight, left-handed, standard model neutrino. The large mass of the right-handed neutrinos would explain the very low mass of the neutrino, as the mass of the known, left-handed neutrino would be proportional to the inverse of the mass of the hypothetical right-handed neutrino.

This mechanism arises naturally when attempting to expand the Standard Model to include the massive neutrinos that we see experimentally [56]. For other massive fermions, the Dirac mass term of the Lagrangian includes some Yukawa coupling $f_q$ to the vacuum expectation value of the Higgs field $\langle\varphi\rangle_0$ to give the particle mass $f_q \langle\varphi\rangle_0$. For some quark $q^0$:

$$L_{D_q} = f_q \langle\varphi\rangle_0 (\bar{q_L^0}) q_R^0 + h.c. \tag{2.16}$$

So, for our massive neutrino, we clearly need to include a Yukawa coupling $f_\nu$ and a right-handed neutrino field $\nu_R^0$ in the Dirac mass term[6]:

$$L_{D_\nu} = f_\nu \langle\varphi\rangle_0 (\bar{\nu_L^0}) \nu_R^0 + h.c. \tag{2.17}$$

However, using a Higgs field VEV $\langle\varphi\rangle_0$ of 246 GeV and a neutrino mass $< 0.05$ eV requires a Yukawa coupling constant of $\sim 10^{-13}$, more than six orders of magnitude lower than that of the electron. We are therefore quite well motivated in searching for alternate explanations for this low mass, and it can be pursued elegantly through the addition of a right-handed Majorana mass term $L_{M_R}$. In the above, the addition of a right-handed neutrino field is required to generate the Dirac mass term, despite our understanding that right-handed fields do not take part in the weak interaction. This therefore ensures that there is no reason not to include an additional right-handed Majorana mass term $L_{M_R}$ which violates neither conservation of energy nor conservation of weak isospin:

$$L_{m_R} = -\frac{m_R}{2} (\bar{\nu}_R^0)^c \nu_R^0 + h.c. \tag{2.18}$$

Giving us a total neutrino mass term, including Dirac and Majorana components, of:

$$L_{m_\nu} = -m_D \bar{\nu}_R^0 \nu_L - \frac{m_R}{2} (\bar{\nu}_R^0)^c \nu_R^0 + h.c. \tag{2.19}$$

[6]In the following derivation, h.c. denotes the hermitian conjugate of the expression, $q^0$ denotes some real particle with left and right handed component fields $q_L^0$ and $q_R^0$ respectively, and $(q^0)^c$ denotes the complex conjugate of $q^0$.

$$= \frac{1}{2}\begin{bmatrix}(\bar{\nu}_R^0)^c & (\bar{\nu}_R^0)\end{bmatrix}\begin{bmatrix}0 & m_D\\ m_D & m_R\end{bmatrix}\begin{bmatrix}\nu_L^0\\ (\nu_R^0)^c\end{bmatrix} + h.c. \tag{2.20}$$

Where $m_D$ and $m_R$ are the Dirac and Majorana mass parameters, related in the second expression as the neutrino mass matrix $M_\nu$. [If we suppose that the Dirac mass parameter $m_D$ is comparable in order of magnitude to other fermions, and we also suppose that $m_R \gg m_D$.]-move. We can then diagonalise the mass matrix $M_\nu$ with the following transformation, using some unitary matrix $Z$:

$$Z^T M_\nu Z = D_\nu \tag{2.21}$$

$$D_\nu = \begin{bmatrix}m_1 & 0\\ 0 & m_2\end{bmatrix} \tag{2.22}$$

We then obtain the diagonal matrix $D_\nu$ whose diagonal elements are the positive-definite eigenvalues of $M_\nu$, $m_1$ and $m_2$. To first order, Z corresponds to:

$$Z = \begin{bmatrix}1 & \frac{m_D}{m_R}\\ -\frac{m_D}{m_R} & 1\end{bmatrix}\begin{bmatrix}i & 0\\ 0 & 1\end{bmatrix} \tag{2.23}$$

Hence,

$$D_\nu = \begin{bmatrix}\frac{m_D^2}{m_R} & 0\\ 0 & m_R\end{bmatrix} \tag{2.24}$$

Giving us our eigenvalues for $M_\nu$:

$$m_1 \simeq \frac{m_D^2}{m_R} \qquad m_2 \simeq m_R. \tag{2.25}$$

Recasting $L_{m_\nu}$ in terms of the mass eigenfields, we define the chirally left-handed two component column vector $\nu_L$:

$$\nu_L \equiv Z^{-1}\begin{bmatrix}\nu_L^0\\ (\nu_R^0)^c\end{bmatrix} \tag{2.26}$$

We can then define the two component field $\nu$with components $\nu_1$ and $\nu_2$ as:

$$\nu \equiv \nu_L + (\nu_L)^c \equiv \begin{bmatrix} \nu_1 \\ \\ \nu_2 \end{bmatrix} \tag{2.27}$$

Finally, we may use the fact that scalar covariant combinations of fermion fields can connect only fields of opposite chirality, we can rewrite $L_{m_\nu}$ as follows:

$$L_{m_\nu} = -\sum_{i=1}^{2} \frac{m_i}{2} \bar{\nu}_i \nu_i \tag{2.28}$$

Assuming the Dirac mass parameter $m_D$ is on the order of typical fermion masses, and that the right-handed Majorana mass term $m_R$ is much heavier than $m_D$, we obtain a elegant explanation for the diminutive neutrino mass: neutrino masses are so small because the right-handed neutrino mass is so large. This is known as the Seesaw Mechanism. Each light neutrino $\nu_1$ would also have a heavy neutral partner $\nu_2$, reflected by the mass eigenvalue $m_2$ above, which exists on some high mass scale and is hypothesised to exist as a very heavy neutral lepton, potentially responsible for leptogenesis in a number of theories [56].

We may also see that Equation 2.27 is identical under charge conjugation. Any particle whose field possesses this property is identical to its antiparticle, known as a Majorana particle, and therefore we see that the eigenstates of a combined Dirac-Majorana mass term $L_{m_\nu}$are in fact Majorana neutrinos. Thus, the seesaw mechanism potentially provides not only a explanation for the neutrino mass scale, but also the leptonic asymmetry of the universe and the mass nature of the neutrino.

Whilst the above treatment is only for one neutrino, it is readily possible to expand the derivation to include three leptonic generations and mixing. This has the net consequence of creating the three light 'classical' neutrinos $\nu_{1,2,3}^{light}$ and three heavy partner neutrinos $\nu_{1,2,3}^{heavy}$. We know from experiment that the heavy neutrinos cannot exist below 80 GeV, further increasing the motivation for the seesaw mechanism by increasing the likelihood of heavy neutral leptons at the GUT scale. Finally, using the above combined Dirac-Majorana mass term and assuming only three observable generations of light neutrino, we allow an approximately unitary PMNS matrix to arise naturally [56].

### 2.2.4 Sterile Neutrinos

It is also possible that should any right-handed neutrinos exist, they would not couple to the weak force. These are known as sterile neutrinos, and while no evidence of their existence has been found to date, they are not inconsistent with the concept of the seesaw mechanism and are theoretically well-motivated. It has been hypothesised that sterile neutrinos may be produced via neutrino mixing

between active and sterile components, and would therefore exist as an additional mass eigenstate in neutrino oscillations. If they exist, they would have the potential to explain a number of unexplained phenomena in physics, including dark matter, baryogenesis, and dark radiation [57].

A number of active experiments are currently searching for the sterile neutrino, and whilst no clear observations have yet been made to date [58], two independent reactor experiments have reported anomalies in the rate of antineutrino production potentially consistent with the existence of a fourth neutrino species [59]. Additionally, recent analyses of CMB data using the Planck satellite have suggested the presence of higher relativistic degrees of freedom potentially consistent with a (3+1) sterile neutrino [60].

## 2.3 Supernova Neutrino Interactions at SNO+

Type II supernovae release ∼99% of the gravitational binding energy of the progenitor star as neutrinos, mostly via pair production of all flavours as detailed in Section 1.1.2. At the densities of core collapse, the stellar interior is opaque to neutrinos which then scatter and thermalise as they become the primary energy transport agents of the stellar core. As the core cools, the neutrinos decouple at different temperatures depending on their flavour composition: heavier neutrino flavours decouple at higher temperatures as they only undergo neutral-current interactions with matter. This results in an approximately thermal spectrum of $\nu_e/\bar{\nu}_e$, and a modified[7] thermal spectrum for $\nu_x$, with the mean energy hierarchy given by Equation 2.1. Putting aside individual detector capabilities, the total event rate on Earth is then solely dependent on distance to the supernova, via a simple $\frac{1}{r^2}$ relationship. Whilst we expect future megaton-scale neutrino detectors such as HyperK to be capable of detecting supernovae out to the Andromeda galaxy [61], currently we are limited to galactic scales of ∼50 kpc.

These neutrino interactions can then be subdivided into neutral current and charged current reactions, depending on the identity of the weak vector boson involved. In principle, neutrinos are detectable via both reaction types, though in practice detector energy thresholds and materials dictate which is dominant. Historically, the dominant channel for all neutrino detection has tended to be the charged current inverse beta decay reaction, or IBD ($\bar{\nu}_e + p \rightarrow e^+ + n$). This is due mainly to its clear interaction signature: the characteristic time difference between the near-instantaneous positron annihilation and the later neutron capture provide a robust delayed-coincidence tag.

[7]The spectrum of $\nu_x$ is modified by the presence of an energysphere prior to the true $\nu_x$ neutrinosphere, within which $\nu_x$ undergo elastic scattering interactions with heavy nuclei prior to free streaming from the surface of the neutrinosphere. This results in a down-shifted, 'pinched' modification of the original thermal spectrum of $\nu_x$. Further details are given in Section 1.1.2.3.

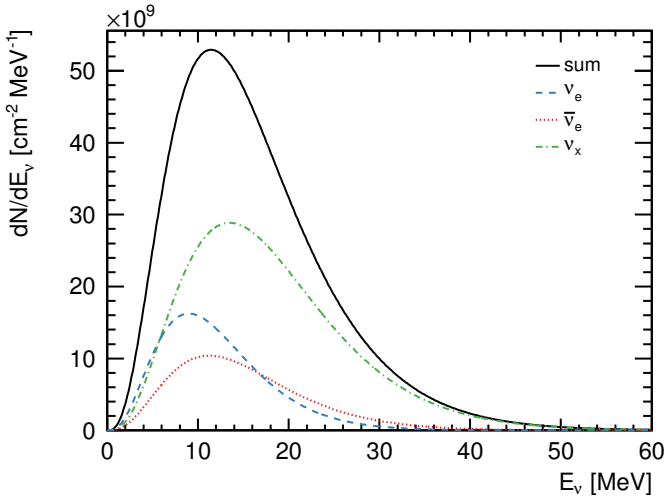



Figure 2.4: Sample neutrino fluence $\frac{dN}{dE_\nu}$ as a function of neutrino energy $E_\nu$ for a supernova at 10 kpc and total energy 300 foe. [40]

There exist a multitude of interactions aside from IBD of course, however previously technological limitations have prevented experiments from observing them reliably. For example, all 24 detected events from SN1987A arose from inverse beta decay [62]. Whilst an excellent channel for observing supernova, IBD is limited in that it is only sensitive to $\bar{\nu_e}$ and is therefore not capable of investigations of heavy-flavour neutrino interactions, and consequently is not sensitive to the total supernova neutrino flux. Liquid scintillator experiments provide an opportunity to diversify the number of visible interactions whilst remaining highly sensitive to IBD events. In addition, recent studies have shown a new interaction channel thought to be highly sensitive to supernova neutrinos: neutrino-proton elastic scattering, or PES [63]. Whilst these events are expected to form the majority of SN interactions at SNO+, the proton recoil spectrum is soft, and is therefore difficult to separate from other low energy backgrounds.

In the case of SNO+, and liquid scintillator detectors more generally, the low energy thresholds and large mass of homogeneous target material presents an opportunity to observe some of the rarer neutrino interactions. The variety of interaction channels open to liquid scintillator experiments provides SNO+ with the ability to observe a far greater range of the overall supernova neutrino flux, and therefore potentially probe BSM physics. For example, there exist a number of neutrino interactions on the nucleus of $^{12}$C, both charged and neutral current, which may prove fruitful in the search for the neutrino mass hierarchy. In particular, the neutral current excitation of $^{12}$C is dominated by heavier flavour neutrino species, and therefore allows these events to be closely identified with $\nu_\mu$ and $\nu_\tau$ interactions [4].

This section provides an overview of the range of supernova neutrino interactions observable at SNO+.

### 2.3.1 Inverse Beta Decay

The dominant charged current supernova interaction at all neutrino observatories is inverse $\beta$ decay, known as IBD, which describes the process:

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

It has a large and clean theoretical cross section, with a threshold of ~ 1.8 MeV. The visible energy of the reaction is directly related to the incoming neutrino, with the total energy encapsulated in the energies of the annihilating positron and thermalised neutron, respectively. The prompt energy observable in the detector is then simply the sum of the positron kinetic energy and the annihilation energy, and is given by $E_{prompt} = E_{pos} + 2m_e$. The energy of the interacting antineutrino, $E_{\bar{\nu}}$, is then closely related to the prompt energy of the IBD positron as follows [64]:

$$\begin{aligned} E_{\text{prompt}} &\simeq E_{\bar{\nu}} + (m_p - m_n - m_e) + 2m_e \\ &= E_{\bar{\nu}} - 0.78 \text{ MeV} \end{aligned}$$

Where $m_p$ denotes the mass of the proton, $m_n$ the mass of the neutron, and $m_e$ the electron mass.

The IBD signal can be characterised by the presence of two light peaks, known as the prompt and delayed peaks respectively. The first, prompt, signal arises from the near-instantaneous annihilation of the positron on an electron in the detector medium, whilst the delayed signal arises from the capture of the thermalised neutron hundreds of microseconds later. In SNO+, this predominantly occurs on hydrogen 200-260 microseconds later than the prompt light signal, and is marked by a de-excitation gamma ray with characteristic energy 2.223 MeV. Whilst 99% of thermalised neutrons capture on hydrogen, a small minority capture on $^{12}$C and release a de-excitation gamma of 4.945 MeV [40].

In addition to the above, the IBD signal can be clearly tagged using the coincidence between the above events, and defining a characteristic $\Delta t$ where this may occur. Commonly, the coincidence tag for IBD sets a simple time and volume limit, defined by the mean free path of the neutron,

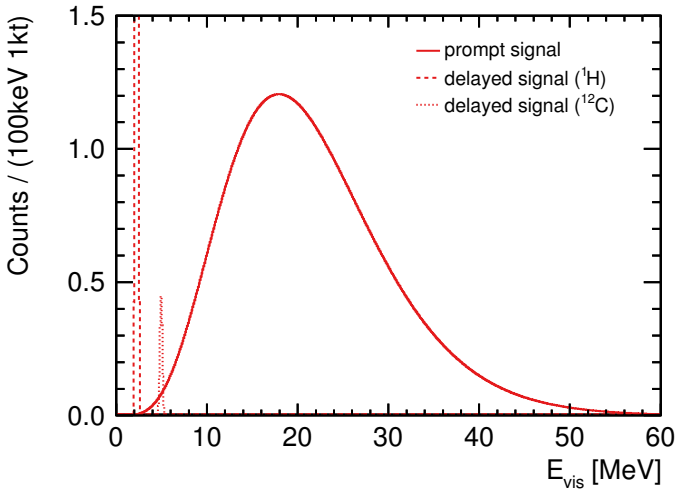



Figure 2.5: Visible inverse beta decay spectrum in LAB, assuming an incoming neutrino fluence of Figure2.4 [40].

around the location and time of the prompt light and looks for a neutron capture event within that window. When implemented alongside energy cuts, this tagging method is highly effective, and leads to a detection efficiency of ∼0.94 in similar liquid scintillator experiments such as KamLAND [65]. Interestingly however, for supernova the possibility arises that the rate of IBD events in scintillator is simply too great, and the effectiveness of the tagging algorithm drops as multiple IBD events occur in the same time and volume windows. Supernovae closer than ∼350 pc are expected to exceed this limit, provided that half of the neutrinos produced are emitted within the first second as projected [40].

### 2.3.2 Carbon Nuclear Interactions

#### 2.3.2.1 Exclusive Transitions

The large quantities of $^{12}$C available in liquid scintillator allows unique interactions to occur in this detector medium that are unobservable in water. The $^{12}$C nucleus provides a triplet of super-allowed, exclusive transitions from the ground state of $(0+, 0)$ to a triad of $(1+, 1)$ states: $^{12}N_{g.s.}$, $^{12}B_{g.s.}$, and $^{12}C^{*}(15.11$ MeV) [66]. Supernova neutrinos interact with the nucleus in three ways: charged-current capture of $\overline{\nu}_e$, charged-current capture of $\nu_e$, and neutral-current inelastic scattering of $\nu_x/\overline{\nu}_x$. All three have distinctive detector signatures, with the charged-current events undergoing delayed-coincidence beta decay following the interaction, and the neutral-current interaction events releasing a mono-energetic gamma ray at 15.11 MeV. The interactions are listed below along with

their threshold energies $E_T$:

$$\begin{aligned} \bar{\nu}_e + {}^{12}C &\rightarrow {}^{12}B_{g.s.} + e^+ \; (E_T = 14.39 \text{ MeV}) \\ &\downarrow \\ {}^{12}B_{g.s.} &\rightarrow {}^{12}C + e^- + \overline{\nu}_e \end{aligned} \tag{2.29}$$

$$\begin{aligned} \nu_e + {}^{12}C &\rightarrow {}^{12}N_{g.s.} + e^- \; (E_T = 17.34 \text{ MeV}) \\ &\downarrow \\ {}^{12}N &\rightarrow {}^{12}C + e^+ + \nu_e \end{aligned} \tag{2.30}$$

$$\begin{aligned} \nu + {}^{12}C &\rightarrow {}^{12}C^* + \nu' \; (E_T = 15.11 \text{ MeV}) \\ &\downarrow \\ {}^{12}C^* &\rightarrow {}^{12}C + \gamma \, (15.11 \text{ MeV}) \end{aligned} \tag{2.31}$$

All three reactions possess high energy thresholds, and are therefore only sensitive to the high energy fraction of the $\nu_e/\bar{\nu}_e$ spectrum. The charged current reactions are distinctive in that their short lived daughter particles are observable using a combination of the delayed signal from the beta decay of $^{12}$N/$^{12}$B, together with the prompt signal from the $e^{\mp}$ event in a manner very similar to the delayed coincidence tag used in IBD interactions above. It is projected that future, multi-kilotonne detectors may use an efficient form of this technique to greatly increase sensitivity to all supernova flux components $\nu_e, \bar{\nu}_e$ and $\nu_x$ [67].

As the neutral-current reaction is sensitive to all neutrino species it possesses the largest cross section of the above at SNO+ by a considerable margin, with the number of 15.11 MeV gamma events expected to resolve well above the supernova-induced background, producing an observable spectral feature, shown in Figure 2.6. The expected event rates for the charged current interactions seen in Equations 2.29 and 2.30 are far lower, and are non-visible against the IBD positron spectrum background without use of the delayed coincidence tag mentioned above.

#### 2.3.2.2 Inclusive Knockout Transitions

In addition to the above exclusive transitions of $^{12}$C , there exist a number of so-called 'inclusive' transitions which include an excited-state daughter nucleus alongside a knockout nucleon [66]. These interactions possess far lower cross sections than those of the exclusive transitions above and are therefore not prominent in the overall supernova visible energy spectra. However, the knockout nucleons observed in these inclusive transitions create additional backgrounds to both PES and IBD events, and thus are one of the leading contributors to overall supernova-induced background rates [40].

The primary knockout channels from neutrino interactions on $^{12}$C are detailed below, along with their energy thresholds $E_T$. Equations 2.32 and 2.33 detail the charged current interactions, whilst Equations 2.34 and 2.35 give the neutral current interactions:

$$\begin{aligned} \nu_e + {}^{12}C &\rightarrow {}^{12}N^* + e^- \ (E_T = 17.95 \text{ MeV}) \\ &\downarrow \\ &{}^{11}C + p \end{aligned} \tag{2.32}$$

$$\begin{aligned} \bar{\nu}_e + {}^{12}C &\rightarrow {}^{12}B^* + e^+ \ (E_T = 17.26 \text{ MeV}) \\ &\downarrow \\ &{}^{11}B + n \end{aligned} \tag{2.33}$$

$$\begin{aligned} \nu + {}^{12}C &\rightarrow {}^{12}C^* + \nu' \ (E_T = 15.46 \text{ MeV}) \\ &\downarrow \\ &{}^{11}B + p \end{aligned} \tag{2.34}$$

$$\begin{aligned} \nu + {}^{12}C &\rightarrow {}^{12}C^* + \nu' \ (E_T = 18.73 \text{ MeV}) \\ &\downarrow \\ &{}^{11}C + n \end{aligned} \tag{2.35}$$

As with the exclusive interactions above, the knockout interactions all possess high energy thresholds and are therefore only sensitive to the high energy fraction of the supernova neutrino spectrum. These high energy interactions on the $^{12}$C nucleus produce excited states above the nucleon separation energy, causing the nucleus to decay by particle emission and producing a knockout nucleon with high mean energy, averaging $E_N \gtrsim 2$ MeV. Knockout protons and recoil protons induced by knockout neutrons from a major background to PES events, whilst any knockout neutrons will readily thermalise and capture on $^{1}$H or $^{12}$C, imitating IBD events. Additionally, the final state nuclei cannot be rejected via coincidence tagging: $^{11}$B is stable, whilst $^{11}$C has a long lifetime of $\tau \sim 29$ minutes. Knockout interactions on the $^{12}$C nucleus therefore create an irreducible background to both PES and IBD events from supernova at SNO+.

### 2.3.3 Neutrino-Proton Elastic Scattering

The process $\nu + p \rightarrow \nu + p$ has only recently been proposed as a viable means of supernova detection at liquid scintillator experiments [63]. As all flavours of neutrinos and antineutrinos can undergo elastic scattering with no kinematic energy threshold, proton elastic scattering is thought to comprise the single largest interaction channel for supernova at SNO+. Additionally, as the proton recoil spectrum is proportional to the neutrino spectrum, it allows for the potential measurement of separate energies and temperatures of the heavier neutrino species.

However, proton elastic scattering events typically have a very soft proton recoil spectrum due to the mass of the proton, with the relationship $T_p \simeq \frac{2E_\nu^2}{M_p}$, where $T_p$ is the kinetic energy of the recoiling proton. In addition, the scintillation light output from the recoiling proton is heavily quenched, shifting the bulk of observed events below 500 keV electron equivalent. This energy regime is problematic due to high background levels, lowering the observed event spectrum for this channel to a third of the total event spectrum assuming a detector threshold of ∼200 keV. Whilst it is theoretically possible for SNO+ to lower its threshold on detection of a supernova, in practice this yields far lower statistical precision and is therefore unfavoured.

### 2.3.4 Neutrino-Electron Elastic Scattering

As above, neutrino scattering off electrons is possible for all six neutrino types, and as a purely leptonic process its cross section is known to a high degree of precision. However, it is difficult to distinguish electron elastic scattering from proton elastic scattering due to the lack of reliable tagging methods for either. Whilst in principle pulse shape discrimination could be used, the low efficiency

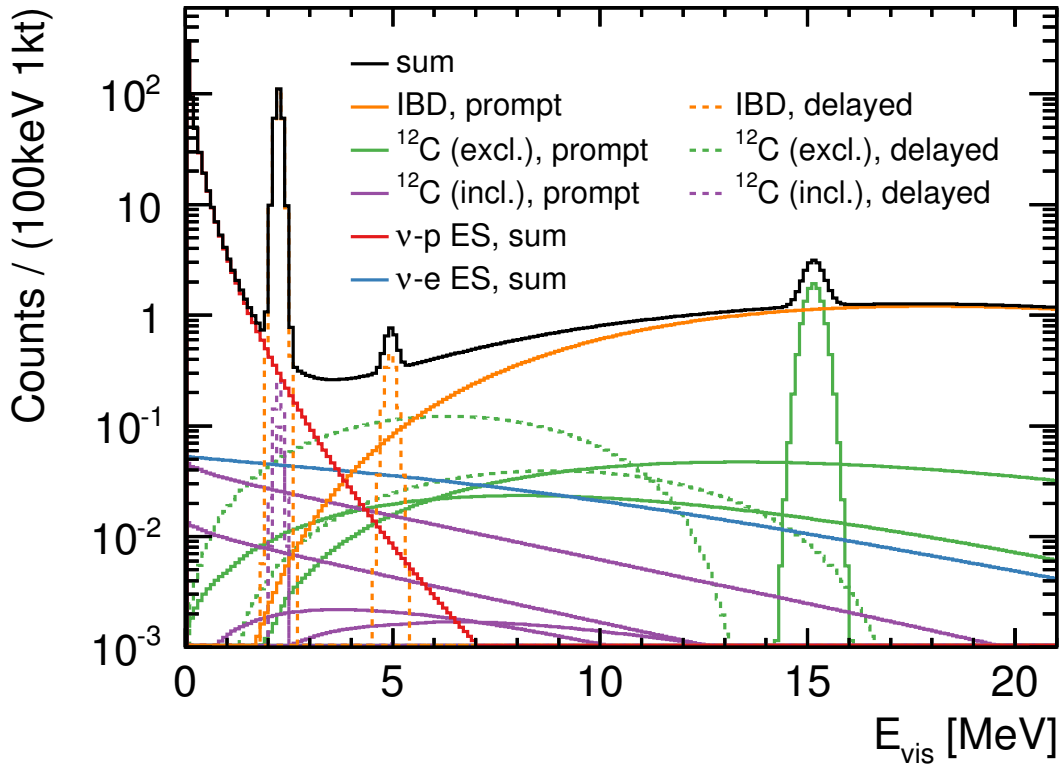



Figure 2.6: Visible energy spectra for all supernova interaction channels at SNO+, assuming a neutrino fluence of Figure 2.4, obtained from [40].

of this method in liquid scintillator results in electron elastic scattering becoming a irreducible background for proton elastic scattering events at SNO+ [40].

### 2.3.5 Summary of Interaction Channels

All visible supernova interaction channels at SNO+ are shown in Figure 2.6, using the reference 10 kpc supernova fluence given in Figure 2.4. The spectral features of the main supernova interactions of IBD and PES can be readily observed, with the prompt light from IBD events dominating above 5 MeV, and PES events dominating below 2 MeV. The IBD neutron capture peaks on hydrogen and carbon are also readily apparent at 2.2 MeV and 5 MeV respectively. Note the delayed component of the inclusive knockout transitions on $^{12}$C imitating the delayed IBD peak at 2.2 MeV. At higher energies, the neutral current de-excitation of $^{12}$C is clearly visible at 15.1 MeV, resolving well above the background of IBD prompt light.

# Chapter 3

# The SNO+ Detector

## 3.1 Introduction

SNO+ is the successor experiment to SNO, the Sudbury Neutrino Observatory, and is situated in Vale's Creighton Mine #9 in Sudbury, Canada [68]. SNO functioned as a multipurpose heavy water Cerenkov detector and rose to prominence in the physics community in 2001 for the key role it played in the resolution of the solar neutrino problem. SNO+ seeks to repurpose the original structure of the SNO experiment to search for neutrinoless double beta decay via conversion into a kilotonne-scale liquid scintillator detector, whilst remaining sensitive to a number of other processes and phenomena. This will be achieved by replacing the heavy water of the old SNO detector with LAB-PPO liquid scintillator, along with upgrades of the detector hardware and materials [69].

The detector itself is housed in the 5000 $m^2$ SNOLAB complex, an expansion of the original SNO experiment facility, alongside a variety of additional experiments such as HALO and DEAP. The location of the detector 2 kilometers underground, an equivalent of 6010 meters of water overburden, allows almost unrivalled protection from the cosmic ray background and results in a cosmogenic muon intensity orders of magnitude lower than at similar facilities such as Gran Sasso and the Soudan mine (see Figure 3.8). As such, it is an ideal site from which to conduct highly sensitive neutrino experiments.

## 3.2 Detector Design

### 3.2.1 Structure

The SNO+ detector structure, shown in Figure 3.1, consists of a central acrylic vessel, or *AV*, 12m in diameter, surrounded by a stainless steel geodesic support structure known as the *PSUP,* housing ∼9300 8-inch photomultiplier tubes. The PSUP extends the diameter of the detector assembly to 17.8 m, the entirety of which is submerged in approximately 7 kt of ultra-pure water in a cylindrical, urylon-lined cavity 34m deep [69]. Above the detector assembly lies the *deck*, which houses the detector electronics, and connects the experiment to the main SNOLAB cleanrooms. The AV is connected to the experiment deck by a cylindrical extension 6.8 m long known as the neck, permitting access to the AV for inspection, cleaning, and filling purposes. The top of the neck then houses the cover gas system, designed to prevent the ingress of radon-contaminated mine air into the AV [70], and is affixed to the deck area by means of the universal interface or UI. The main body of the AV is secured to the deck and floor by a system of hold-up and hold-down ropes, the combination of which is necessitated by the density differences between water and liquid scintillator [70].

Once fully operational, the AV will contain approximately 780 tons of scintillator loaded with the double beta decay isotope $^{130}$Te [69]. By replacing the original $D_2O$ with LAB-PPO liquid scintillator, the number of photons emitted per unit particle energy will be an order of magnitude greater than that of the Cerenkov process, greatly reducing the energy threshold for event detection and therefore enabling the search for neutrinoless double beta decay.

### 3.2.2 Detection Principle

Fundamentally, SNO+ is a light detector: it is designed to detect and record the energy and position of an interaction via measurement of its photon output. Initially, SNO+ will function as a water Cerenkov detector in much the same manner as the original SNO experiment; charged particles produced by neutrino interactions transit the detector at relativistic velocities, producing a Cerenkov light cone in the detector medium. The light in this cone can then be detected by the surrounding PMTs.

Later experimental phases will make use of liquid scintillator as the detector medium, utilising a different mechanism to detect particle interactions: charged particles excite the organic scintillator molecule, light is released as the molecule de-excites, and the light is then observed by the PMTs. This lowers the energy threshold of detection considerably compared to Cerenkov light, which ne-

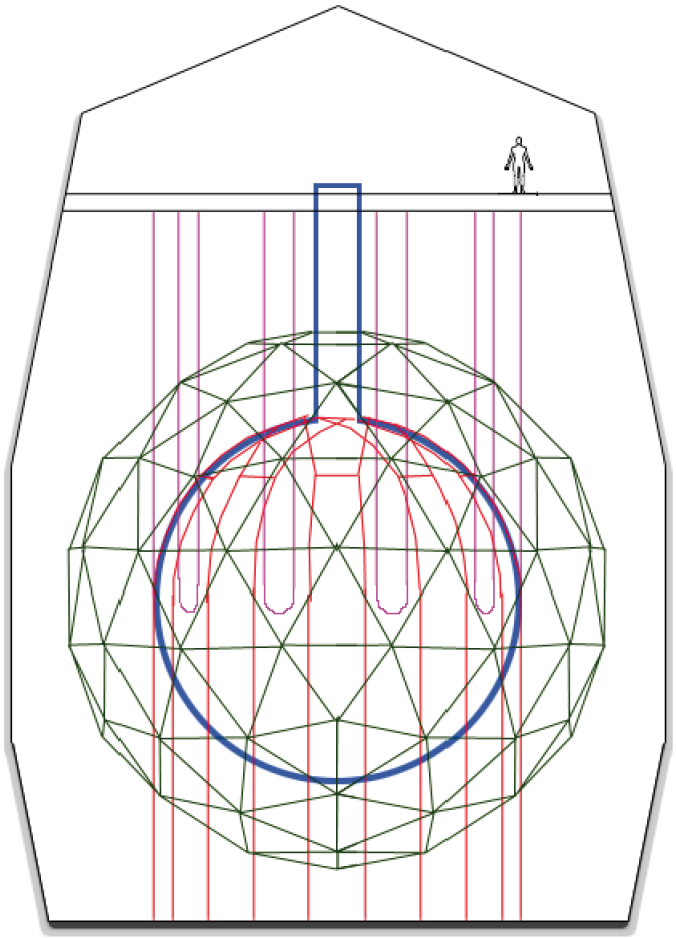


Figure 3.1: The SNO+ detector and cavity, shown to scale, from [71]. The electronics deck is positioned at the top of the diagram, and is connected to the AV (blue) via the universal interface and neck structure. The AV is secured to the deck by the hold up ropes (purple), and secured to the cavity floor by hold down ropes (red). The PSUP (green) surrounds the AV with $\sim$ 9300 PMTs contained in a geodesic support structure. During detector running, the cavity will be filled with ultra-pure water shielding, and the AV will be filled with water or scintillator.

cessarily requires high energy particles in order to produce the characteristic Cerenkov cone.

The Tellurium-loaded phase differs from pure scintillator only in that it aims to principally detect neutrinos originating from the decay of $^{130}$Te within the detector medium rather than neutrinos from external sources, but remains capable of both. It is only in this low energy regime facilitated by the scintillator where neutrinoless double beta decay may be observed; for a Majorana neutrino mass of $m_{\beta\beta} = 200$ meV, we would expect a signal at $\sim$2.5 MeV [69].

### 3.2.3 Experimental Phases

After commissioning, SNO+ plans to run over a series of consecutive experimental phases, each designed to progressively ramp up the physics functionality of the detector whilst providing a robust means of measuring actual experimental backgrounds. In some cases, these phases allow secondary physics goals to be pursued[1]. The phases of SNO+ are named according to the detector medium in the AV.

#### 3.2.3.1 Water Phase

The initial phase of SNO+ involves filling the AV with ultra-pure light water in order to conduct calibration tasks and initial background measurements. Due to its very well understood properties, water will be used to test all detector subsystems and calibration instruments under live conditions. In addition, the presence of water in the AV permits measurement of *external* and *surface* backgrounds, which are defined respectively as background sources originating in the surrounding rock and detector materials, and background originating on the surface of the acrylic vessel itself. This allows accurate measurement of irreducible backgrounds in the cleanest environment possible, with the added bonus of further reducing background via the leaching of lead from the AV surface into the water, which will then be drained before later phases. Further discussion of detector backgrounds is given later in the chapter.

During this phase SNO+ will function effectively as a low-resolution water Cerenkov detector, and thus physics goals will be primarily limited to nucleon decay searches, solar background measurements, and supernova monitoring.

#### 3.2.3.2 Scintillator Phase

During the scintillator phase of SNO+ the AV will be filled with the main liquid scintillator mixture: a combination of the aromatic hydrocarbon solvent *linear alkylbenzene* (LAB), which is excited by interactions within the detector but does not re-emit easily, and a 2 g/L concentration of the fluor *2,5-diphenyloxazole* (PPO) to form *LAB-PPO*. Exitations of the LAB molecule are transferred non-radiatively to the PPO fluor, which then emits this energy as light visible to the detector PMTs. The LAB and PPO are separately purified underground at a purpose-built processing plant, which then mixes the liquids via a fluid handling and distribution system, outputting the mixture directly into the AV. This facility will purify the mixture initially, as the AV is filled, then continuously over the

[1]For the purposes of this thesis, SNO+ should be capable of supernova monitoring immediately upon commission.

course of running in order to remove contaminants which may leak from the AV to the scintillator [72].

A number of key measurements are planned for this phase, including internal scintillator background measurements, light yield measurements, and optical scattering and attenuation measurements using the instruments calibrated in the water phase. All of these are essential before the primary physics programme can begin. A number of higher-level physics analyses are also possible during this phase, including measurements of reactor neutrinos, geoneutrinos, and low energy solar neutrinos.

#### 3.2.3.3 Loaded Scintillator Phase

The main phase of SNO+ is that of the loaded phase, in which the liquid scintillator mixture will be 'loaded' with the double beta decay isotope $^{130}$Te. Initially, the scintillator is to be loaded to 0.3% by mass, permitting preliminary searches for neutrinoless double beta decay. After a period of stable running, incremental increases of $^{130}$Te can then be added in order to increase sensitivity. Currently, loading levels of up to 3% are being explored, with the potential of probing deeply into the inverted mass hierarchy parameter space. The physics of the loaded scintillator phase, including the benefits of $^{130}$Te loading and detector limits, are discussed in greater depth in section 3.4.1.

## 3.3 Hardware and Software

### 3.3.1 Overview

The deck houses the bulk of the electronics hardware used in SNO+, along with the AV universal interface, calibration systems and subsystems, and general access hatches. The PMT cables run up through the deck and are connected to the electronics architecture, organised into a system of cards and crates. The crates then connect to the CAEN analogue-digital converter, which along with the master trigger card, comprises the timing rack. The timing rack is responsible for defining the timeframe of an event, before sending the information to the Event Builder. The Event Builder receives event information from a number of sources in addition to the timing rack, and collates the information to define an 'event', before sending data to the SNO+ DAQ system for processing, storage, and analysis.

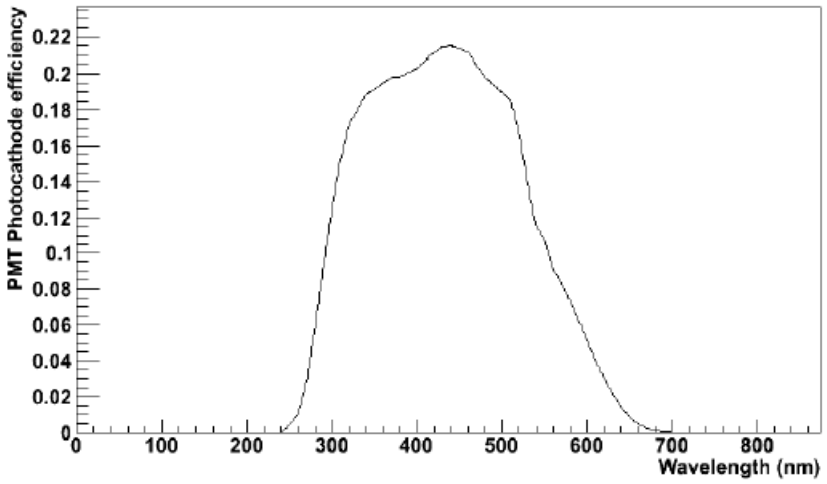



Figure 3.2: Quantum efficiency of the SNO+ Hamamatsu R1408 photomultiplier tubes against photon wavelength [71].

### 3.3.2 PMTs

SNO+ utilises ∼9300 8-inch diameter Hamamatsu R1408 photomultiplier tubes from the original SNO experiment [68], each one coupled to a 27 cm diameter reflective concentrator. With the addition of the concentrators, this extends the total detector coverage of the PMTs to over 50%, greatly enhancing the sensitivity of the experiment [73]. Despite thorough cleaning before deployment, the PMTs are one of the major contributors to external background in the experiment due to radioactive contaminants within the PMT components, though this is significantly mitigated by the presence of the pure water shielding between the PSUP and the AV.

The PPO fluor in the LAB-PPO cocktail emits light at wavelengths between 350 and 400 nm, whereas the maximum quantum efficiency of the SNO+ PMTs occurs at ∼440 nm, necessitating the use of wavelength shifters within the scintillator mixture to maximise detector efficiency. The quantum efficiency of the SNO+ PMTs is shown in Figure 3.2. Two wavelength shifters are currently being explored for use in the experiment; *bisMSB*, which emits between 380 and 450 nm, and *perylene*, which emits between 440 and 500 nm. Both match well with the peak quantum efficiency of the SNO+ PMTs.

### 3.3.3 Electronics

Once light is registered by an individual PMT, it sends a signal to a PMT interface card, or PMTIC, which acts as a junction for 32 individual PMT channels. This signal is then forwarded to a daughter

board, or DB, which carries out preliminary measurements and signal discrimination/integration. The daughter board then differentiates trigger-dependent and trigger-independent quantities from the signal, whilst creating a signal pulse for various pre-defined triggers if above threshold. All information from the daughter board is then sent on to a front end card, or FEC, each of which handles 32 channel inputs. These FECs are organised into 19 'crates', each containing 16 FECs, which together read out all PMT channels.

The information from a single FEC in a each crate is then forwarded to the respective crate's trigger card, or CTC, which combines the channel signals from each FEC within that crate into a single summed quantity for each defined trigger type. All non-trigger information is sent to an XL3 card in the same crate, whilst the per-crate trigger signals are sent on to the Master Trigger Card - Analogue, or MTC/A, which combines all trigger signals into a single, detector-wide trigger signal for each trigger type. These detector-wide trigger type signals are finally forwarded to the Master Trigger Card - Digital, or MTC/D, which compares the signals against a pre-defined threshold for each type, and issues a global trigger if the thresholds are met within a given time window. Once a global trigger is issued, the XL3 card is allowed to send its information on to the event builder and data acquisition systems.

The information received by these systems is divided into four base quantities by the daughter boards, each of which is essentially the raw output from the PMTs measured against some parameter. Three of these relate the measured charge of the PMT to a gain setting at some defined integration time (long or short), whilst the fourth simply measures the total integrated signal time. These quantities are then the fundamental data that SNO+ uses for all its analyses, alongside a simple count of the number of PMTs hit, and permit calculation of derived quantities such as reconstructed energy and position [74]. A simple schematic of the above system is shown in Figure 3.3.

The Event Builder and its successive event processors are key stages in supernova monitoring and will be discussed in detail in the next chapter.

### 3.3.4 Calibration Systems

Calibration systems at SNO+ include both optical and radioactive sources, each designed to test a different aspect of detector response. The primary optical calibration system is *ELLIE* - an externally mounted system of fibre optic cables connected to LED and laser inputs. Subsystems of ELLIE are in place to monitor and measure specific aspects of the response; *TELLIE* will check the timing response of the PMTs via pulsed LEDs, *SMELLIE* will conduct tests of the scattering properties of

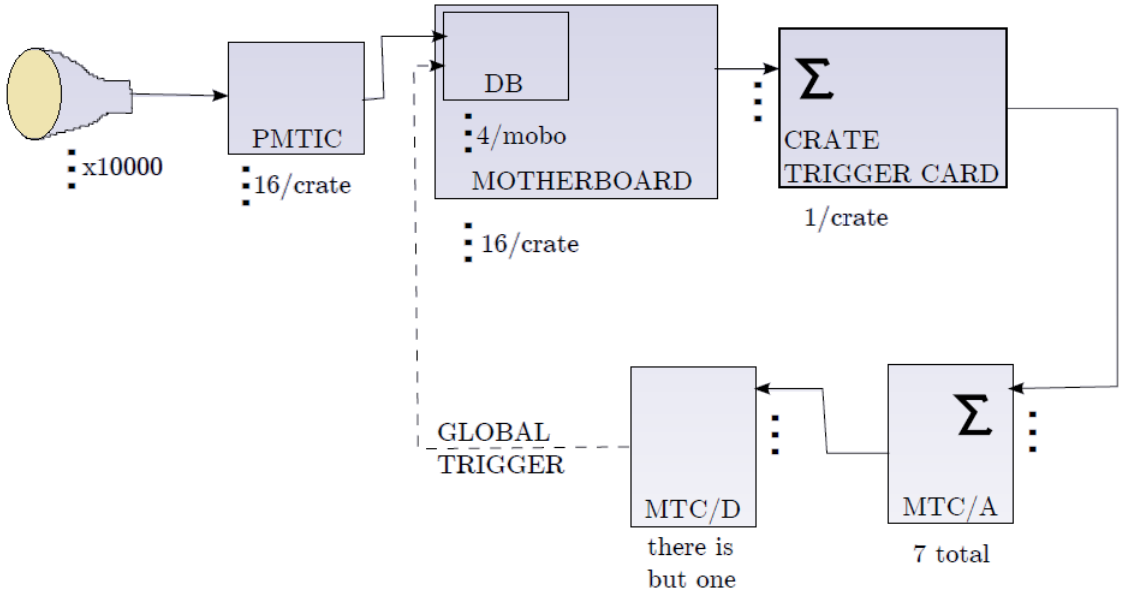



Figure 3.3: Simple schematic of the electronics and global trigger system at SNO+. The XL3s, DAQ system, and Event Builder are not shown. [71]

the liquid scintillator medium via a series of milliwatt lasers, and *AMELLIE* will monitor for changes in the attenuation properties of the scintillator mixture over time using wide-angle LEDs. ELLIE therefore enables SNO+ to conduct frequent calibrations of the detector PMTs without the need for the insertion of internal sources, thereby minimising the risk of contamination and increasing efficiency by minimising detector downtime.

Whilst a number of other secondary optical calibration sources exist at SNO+, the most relevant to this thesis is that of the supernova calibration source, or SNC. The SNC is designed to simulate supernova neutrino emissions using a pulsed laser diode, which is able to output very high-rate bursts of photons isotropically within the detector. This will primarily stress-test the data acquisition systems and electronics architecture of SNO+, whilst studying the data capture performance of the detector at various simulated supernova distances. As the possibility of card, crate, and even full detector breakdown is possible at the exceptionally high rates of nearby supernovae, stress-testing is vital. The SNC is capable of simulating supernovae as close as 500 pc.

Radioactive sources will be inserted into the detector via a system of umbilical ropes, designed to hold the source at varying positions within the AV. The radioactive sources will then be used to check the energy scale and resolution of the detector, whilst uncovering any systematic uncertainties and determining reconstruction efficiencies. In addition, SNO+ will deploy underwater cameras to monitor the position of the AV, hold-up and hold-down ropes, and calibration sources within the

detector.

### 3.3.5 Software

The Reactor Analysis Tool, or RAT, functions as the primary Monte Carlo software used in the SNO+ experiment, and is additionally responsible for conducting processing tasks on both simulated and real data. In this sense, it is ubiquitous at SNO+ and is the primary detector simulation software used throughout the collaboration. RAT utilises up-to-date GEANT4 libraries for its physics process simulations, whilst the large-scale data framework software ROOT is employed to handle data management tasks.

Alongside RAT and ROOT, the software architecture contains a number of standalone programs to conduct near-line data processing and on-line monitoring. It is formed by a hierarchy of processing software, each of which conduct increasingly complex checks on the data received, and store it accordingly. These are known as the L1 and L2 trigger systems, and as the creation of the L3 system is the primary concern of this thesis, this aspect of the software is discussed in depth in Chapter 4.

## 3.4 Detector Physics Goals

### 3.4.1 Neutrinoless Double Beta Decay

The observation of neutrinoless double beta decay is the primary aim of the SNO+ experiment, and would herald a major discovery of the fundamental nature of the neutrino. Neutrinoless double beta decay is only possible if the neutrino is its own antiparticle, referred to as a Majorana particle, and would demonstrate clearly an example of lepton number violation. In doing so, it would provide strong support for the current theory of leptogenesis and lepton-antilepton asymmetry.

If the neutrino is indeed a Majorana particle, then a proportion of normal double beta decay should be neutrinoless. In standard double beta decay, a nucleus emits two electrons and two electron antineutrinos, with the energy of the interaction shared between the particles. As only the electrons are readily observable, the neutrinos appear to "steal" a portion of the decay energy, and the kinetic energy spectrum of the electrons produces the continuous spectrum characteristic of beta decay. Under neutrinoless double beta decay, the two neutrinos interact as virtual particles, effectively annihilating one another. This results in a highly distinctive final state of two back-to-back electrons who share the entirety of the decay energy between them, resulting in a monoenergetic

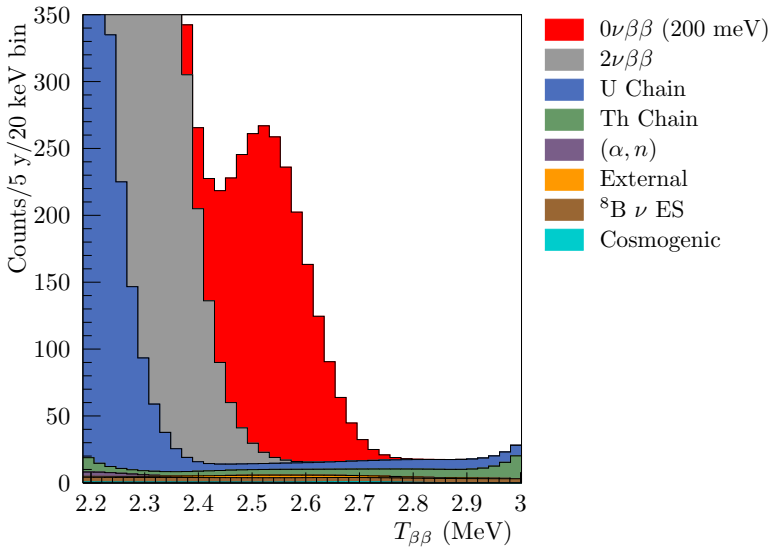



Figure 3.4: Plot of the expected $0\nu\beta\beta$ spectrum at SNO+ at 0.3% $^{130}$Te loading, assuming a Majorana neutrino mass of 200 meV and showing all major backgrounds. The x-axis shows the combined final-state electron energies. [69].

peak at the Q value of the reaction. For $^{130}$Te, this peak is centred at ~2.5 MeV. If the process is observed, SNO+ can then proceed to measure the half life of the decay, and in doing so obtain the effective Majorana mass of the neutrino. Current theory suggests a Majorana mass $m_{\beta\beta}$ of 200 - 400 meV for a half life on the order of $10^{25}$ y [69].

As the effective Majorana mass is dependant on the neutrino mass hierarchy, SNO+ is only sensitive to $0\nu\beta\beta$ should the neutrino mass ordering be inverted, i.e. $m_3 < m_1 \lesssim m_2$. This is shown in Figure 3.5. The uncertainty on $m_{min}$ is a consequence of our lack of knowledge of the absolute scale of the neutrino masses, explored in Section 2.2. The width of the bands reflect uncertainties in our knowledge of phases, mixing angles, and squared mass differences. The allowed $m_{\beta\beta}$ range is dependent on the mass hierarchy: NS denotes normal spectrum, and IS denotes the inverted spectrum. The two bands converge in the degenerate region, where $m_1 \simeq m_2 \simeq m_3$.

The method employed by SNO+ has a number of advantages over its competitor experiments in the search for neutrinoless double beta decay [75]. Due to the sheer size of the liquid scintillator volume, external backgrounds can be minimised through fiducialisation of the detector medium, whilst the spatial distribution of any internal background isotopes will be uniform. The process of phasing the experiment over time additionally allows the internal and external backgrounds to be measured before and after deployment of the isotope, as well as enabling the detector response to be throughly tested prior to deployment. Similarly, the liquid scintillator cocktail can be removed

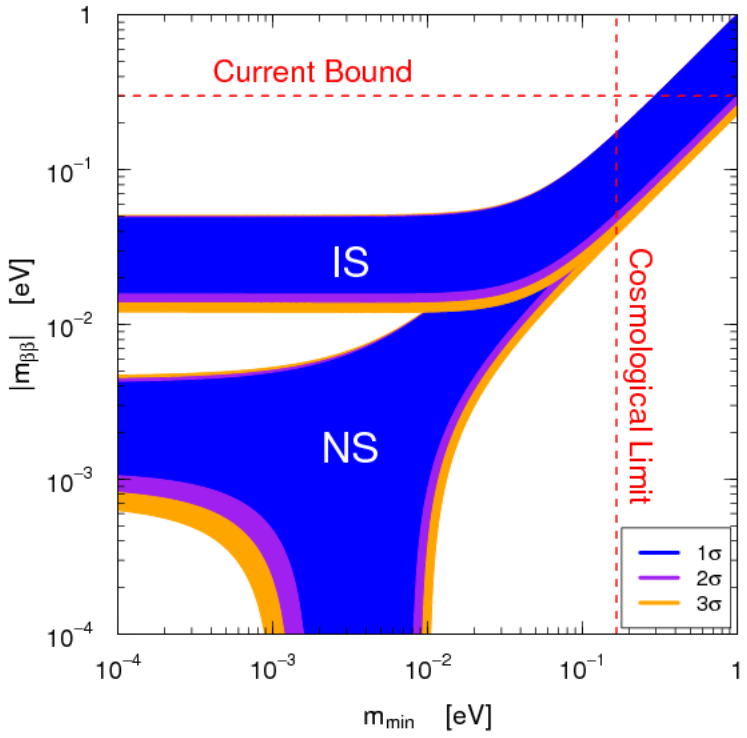



Figure 3.5: Plot of the range of allowed values for the effective Majorana neutrino mass $m_{\beta\beta}$ in terms of the mass of the lightest neutrino $m_{min}$.

and repurified at any time should internal background levels become too high. Finally, the method is scaleable: Tellurium loading can be easily and relatively affordably scaled up in order to increase any signal against the major backgrounds.

Tellurium was chosen as the $0\nu\beta\beta$ isotope of choice for multiple reasons. There is a large natural abundance[2] of $^{130}$Te, at $\sim$34%, enabling high loading of the liquid scintillator volume without any need for enrichment whilst remaining relatively inexpensive. Additionally, $^{130}$Te possesses one of the longest $2\nu\beta\beta$ half lives of all $0\nu\beta\beta$ isotopes, at $\sim 7\times10^{20}$ yr, reducing the background from $2\nu\beta\beta$ significantly. Lastly, $^{130}$Te retains high levels of light transmission, minimal scattering, and acceptable light yields even up to high loading levels of $\sim$5%. A plot of the expected $0\nu\beta\beta$ spectrum at SNO+ at 0.3% loading, assuming a Majorana neutrino mass of 200 meV, is shown in Figure 3.4. As can be seen, at this mass any peak from $0\nu\beta\beta$ is expected to resolve well above the dominant $2\nu\beta\beta$ background given sufficient detector running.

[2]In this case, "abundance" refers to the natural abundance, or total amount of a given isotope available on Earth relative to the total of all isotopes of that element combined.

### 3.4.2 Supernova Neutrinos

Core collapse supernovae are a prodigious source of neutrinos, with an estimated 99% of the gravitational binding energy of the star being converted into neutrinos of all flavours upon collapse [3]. These neutrinos give us a unique insight into the fundamental mechanisms involved in neutrino-driven supernovae explosions, whilst also supplying us with an excellent means of determining some of the lesser known properties of the neutrinos themselves.

The Baksan underground telescope have placed an upper bound on the mean frequency of galactic core-collapse supernovae at $f_{\text{collapse}} = 0.13\ \text{yr}^{-1}$ at 90% C.L., meaning that we are currently overdue [6]. During the last observed supernova event, SN1987A, a total of 24 events were observed at a distance of 50 kpc. At SNO+, a supernova at 10 kpc is expected to produce at least 652 events[3] in the water phase, and over 757 events during the scintillator phase across all flavours [1], although higher energy thresholds during the water phase will greatly reduce sensitivity to these events compared to the scintillator phase. As the focus of this paper, a detailed treatment of supernova physics and supernova neutrino interactions at SNO+ is given in Chapter 2.

### 3.4.3 Low Energy Solar Neutrinos

Due to its location over 2 km underground, SNO+ is uniquely positioned to measure low energy neutrinos from the Sun. At this depth, the background from cosmogenically activated isotopes is low enough to offer SNO+ high sensitivity, whilst the low energy threshold of the liquid scintillator detector medium enables probes of the solar neutrino spectrum at energies orders of magnitude lower than those examined by SNO. This opportunity is at its height during the unloaded scintillator phase of detector running, which will provide excellent sensitivity to CNO, pep, and low energy $^{8}$B solar neutrinos, potentially enabling accurate measurement of the metallicity of the Sun and precise studies of the MSW effect [69]. The solar neutrino spectrum for a range of interactions is shown in Section 4.5.3 in Figure 3.10.

### 3.4.4 Geoneutrinos and Reactor Neutrinos

SNO+ is able to observe the antineutrinos emitted from both nuclear reactors and naturally occurring radioactive decays in the Earth's crust. Reactor neutrinos are expected to contribute around 90 events per year at SNO+, and thus enable studies of neutrino oscillation parameters due to the

[3] An 'event' at SNO+ is taken to mean any physics interaction occuring in the detector which generates sufficient light for the detector to issue a global trigger and record data. The system by which SNO+ records data is dealt with in depth in Section 3.3.

known distance of the reactors themselves and their well-understood energy spectra.

Geoneutrinos originate from naturally occurring Uranium and Thorium decay chains in the Earths crust and mantle, and thus also present a well known energy spectra. Measurement of these geoneutrinos unifies particle and geophysics, and potentially permits the accurate measurement of the radiogenic heat flow of the Earth. SNO+ will collaborate with a number of similar experiments located around the world to create a global analysis of mantle and crust contributions to the overall geoneutrino spectrum, a first in geophysics [69].

### 3.4.5 Exotics: Nucleon Decay and Axions

SNO+'s location deep underground presents a real opportunity for exotic physics searches. In particular, the search for invisible nucleon decays and axion-like particles can be enabled due to low background levels.

Invisible nucleon decay is the hypothetical process by which nucleons decay into invisible daughter particles, the search for which is planned in the water phase. SNO+ plans to search for the invisible decay of oxygen nucleons, which, in the case of either proton or neutron decay, would present a de-excitation gamma signature at around 6 MeV. Potentially, SNO+ can increase the current limits on the half life of invisible nucleon decay by a factor of two in just six months of running time [76].

Axion-like particles are hypothetical elementary particles produced in the Sun as an extension to the QCD Lagrangian, and should be observable at SNO+ via two main mechanisms. Both Compton conversion of the axion and the axioelectric effect should produce $\sim$5 MeV of energy deposition in the detector, and searches for this signal are planned in both the water and loaded scintillator phases. The loaded scintillator phase in particular offers the opportunity to improve the current limits on the axioelectric effect by many orders of magnitude due the the large mass of the $^{130}$Te nucleus and the effect's dependence on nuclear charge [69].

## 3.5 Detector Backgrounds

### 3.5.1 Intrinsic Detector Radioactivity

Other than the cosmogenic muon background, there exists another, more problematic background intrinsic to the detector materials themselves. All detector materials exhibit some form of radioactivity, and in a detector as sensitive as SNO+ the resultant decay products are a major cause of experimental background. Intrinsic radioactivity can be divided into three main categories, de-

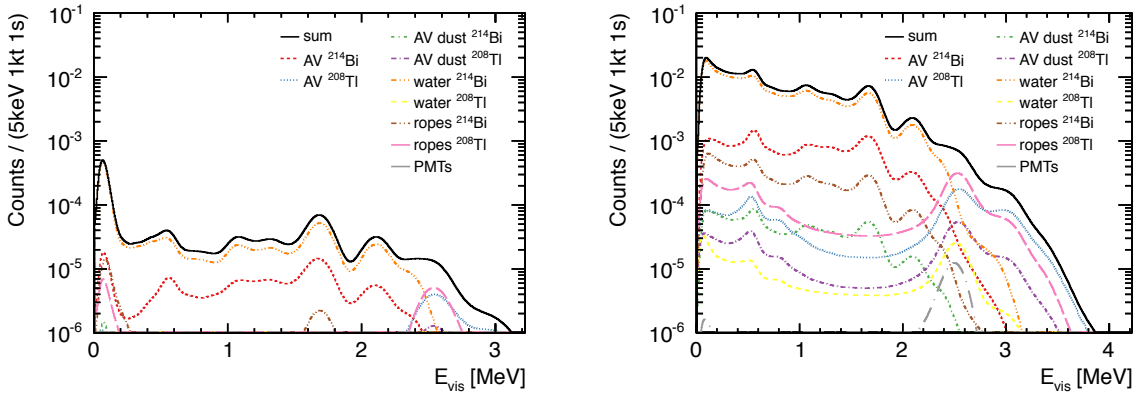



Figure 3.6: Comparison of external background rates in the detector with (left) and without (right) a fiducial volume cut at $R_{fv} = 5$m.[40]

pending on their position relative to the liquid scintillator volume: *external*, *surface*, and *internal* backgrounds. External background events originate from outside the AV, for example in the PMTs, cavity rocks, or hold-down ropes. Surface background events originate from radioactive decays on the inner surface of the acrylic vessel, whilst internal backgrounds are caused by radioactive contaminants within the liquid scintillator itself. Internal backgrounds are the largest contributor of the three to intrinsic detector background at SNO+ [77].

Surface and external backgrounds can be strongly reduced by introducing a fiducial cut to the scintillator volume. By convention, the fiducial volume at SNO+ is defined as $R_{fv} \leq 5$ m, yielding a fiducial mass of 0.45 kt, although this is different for different analyses - $0\nu\beta\beta$ searches use a fiducial volume cut of $R_{fv} \leq 3.5$ m. External backgrounds in particular can be greatly reduced by introducing a fiducial volume cut. A comparison between event rates with and without the fiducial volume cut is shown in Figure 3.6.

The rate of surface and internal backgrounds at SNO+ is detector-specific, and mostly due to the history of usage of the AV in the previous SNO experiment. Due to a fundamentally different detection principle, SNO had a far higher event threshold, and therefore concern for intrinsic backgrounds in the AV was minimal due to the lower typical energies of such events. As the detector was eventually decommissioned in preparation for SNO+ and drained after a highly successful physics program, the acrylic vessel was exposed to raw mine dust before insulation. Mine dust is heavily contaminated with $^{222}$Rn, which decays to $^{210}$Pb. $^{210}$Pb readily absorbs into materials such as the surface of the AV, and the legacy of this decommissioning is present in the modern experiment via the radioactive daughters of $^{210}$Pb, $^{210}$Po and $^{210}$Bi, which now contaminate the surface of the AV. These isotopes then gradually leach into the scintillator volume over the lifetime of the detector,

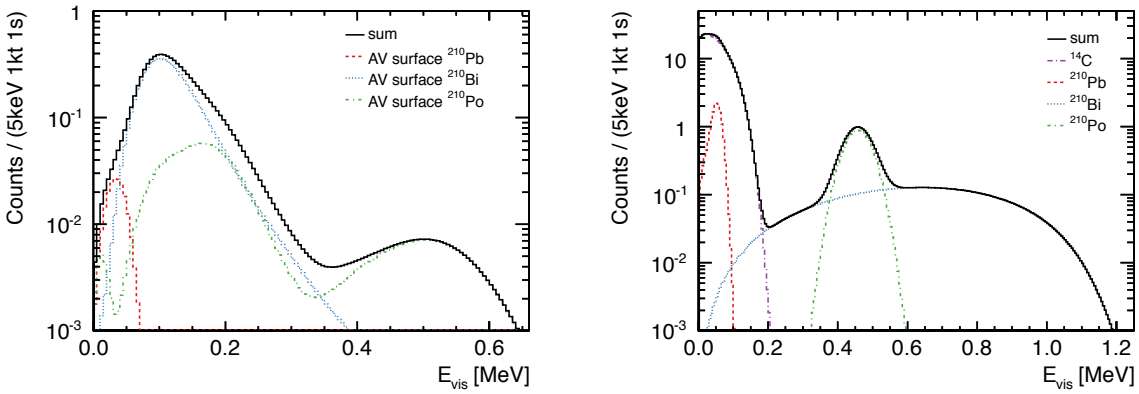



Figure 3.7: Simulated detector surface backgrounds (left) and internal backgrounds (right), including the effects of leaching.[40]

increasing the internal background commensurately. Whilst some of these contaminants will be eliminated via the draining of the AV at the end of the initial water phase, a significant number will remain, contributing to overall background rates.

An additional significant internal background process to those mentioned above is that of the decay of $^{14}$C via beta emission. $^{14}$C is constantly produced through cosmogenic neutron interactions, and is generally present in all organic materials at a roughly fixed trace level. $^{14}$C has a relatively short lifetime despite its low abundance, and is not removable via purification due to its similarity with $^{12}$C. Hence, $^{14}$C is by far and away the largest internal background at SNO+ - an order of magnitude larger than other processes at 391.46 kt$^{-1}$ s$^{-1}$ [40]. It is this background, due to its fundamentally irreducible nature, that is responsible for setting the trigger threshold of SNO+ at 200 keV. The spectra and rates of these surface and internal events are shown in Figure 3.7, with the substantial $^{14}$C background evident below 0.2 MeV. Whilst a significant background for a number of physics searches at SNO+, the low energy of these events ensure they pose a minimal threat to the supernovae searches considered in this thesis.

Although a number of other radioactive processes contribute both to the internal and surface background event rates, these processes are essentially eclipsed by BiPo and $^{14}$C events, and are therefore neglected in this work. A table of the most significant internal backgrounds at SNO+ is given in 3.1.

| Isotope | $\tau$ | Decay Type (BR [%]) | Energy [keV] | Rate [$kt^{-1}$ $s^{-1}$] |
|---|---|---|---|---|
| $^{14}C$ | $8.27\times10^3$ y | $\beta$ (100) | 156.5 | 391.46 |
| $^{210}Po$ | 200 d | $\alpha$ (99) | 5.305 | 16.49 |
| $^{210}Bi$ | 7.23 d | $\beta$ (99) | 1161 | 15.68 |
| $^{210}Pb$ | 32.2 y | $\beta$ (81, 19) | 17, 63 | 15.68 |
| Total | | | | 439.31 |

Table 3.1: Table of internal background isotopes with their decay types, branching ratios, lifetimes, and energies. Isotopes are ordered by decreasing contribution to overall background rate. Alpha decay energies correspond to particle energy, whilst beta decay energies indicate the endpoint of the spectrum. The Pb/Bi/Po rates include leaching effects into the scintillator volume. These values assume 9 months of ultra-pure water in the AV, followed by surface cleansing, prior to scintillator deployment. [40]

### 3.5.2 Cosmogenic Muons

Cosmic rays are extremely high energy particles incident on the Earth from deep space, very often originating outside our solar system. They are composed of a wide range of particles, predominantly ionised nuclei, with around 90% of these being hydrogen ions - simple protons. These high energy protons produce spallation and a shower of secondary particles when they collide with the Earth's atmosphere, resulting in a cascade of lighter particles. The most common secondary particles are mesons, pions and kaons, which subsequently decay into relativistic muons. These muons then penetrate deep into the Earth's crust, and form a major background even at the depths of the SNO+ experiment. Despite possessing one of the deepest experimental sites in the world, the sensitivity of the detector ensures SNO+ still has to contend with a considerable muon flux background, both in the form of direct muon transit of the cavity and second order effects from muon spallation in the surrounding rock. For comparative purposes, the total muon flux for a range of underground laboratories is shown in Figure 3.8.

In scintillator, the high energy of these muons manifests itself as a very large number of hit PMTs, often numbering in the high thousands. Whilst distinctive, this is also the signature of a number of other high energy events including those produced by supernova neutrinos. Fortunately, the predictable path of these muons through the detector often results in a "through and through" event with clear directionality from the double Cerenkov events in the surrounding water, permitting discrimination and tagging via the SNO+ classifiers and data cleaning processors. Additionally, there exist so-called "grazing" muons, which enter the cavity and the water, but do not pass through the acrylic vessel. These pose an additional challenge for tagging purposes, but are generally well understood. The detector muon tracks and their classification is shown in Figure 3.9.

A further source of cosmogenic detector backgrounds originates from muon interactions in the rock surrounding the cavity. These spallation interactions create large quantities of nucleons, partic-

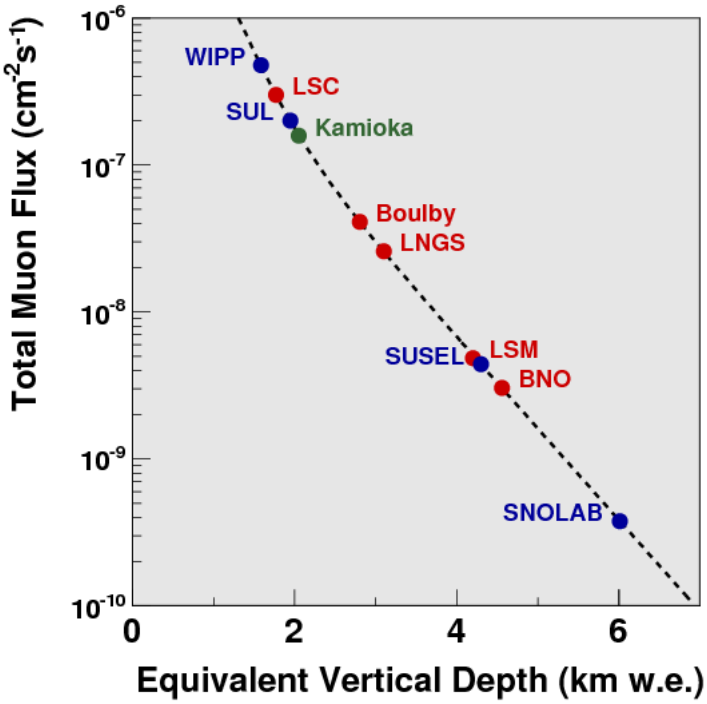



Figure 3.8: Total muon flux vs. equivalent vertical depth for a range of underground laboratories. As of 2015, SNOLAB's depth of 6km m.w.e is only bettered by the Jinping Underground Laboratory in China (not shown), positioned at a depth of 6720 meters water equivalent

ularly fast neutrons. Whilst a neutron background is present at all sites due to natural radioactivity in the rock, these muon-induced neutrons possess a very hard energy spectrum, allowing them to penetrate much further into the detector. As with the direct muon background, these events are external to the detector, and therefore in principle excludable with cuts on dimensionality and position fitting.

As can be seen in figure 3.8, the muon flux at SNO+ is $\Phi_\mu$ =(3.31±0.09) × $10^{-10}$ cm$^{-2}$ s$^{-1}$. We therefore expect (3.74±1.02) × $10^{-4}$ muon s$^{-1}$ in the total scintillator volume of 780 t, equivalent to (4.79±1.31) × $10^{-4}$ kt$^{-1}$s$^{-1}$[40]. Compared to the burst time of a typical supernova, these figures are negligible, and therefore the cosmogenic muon background can be safely disregarded for the purposes of this thesis.

The muon induced neutron background rate at SNO+ in the liquid scintillator volume is necessarily lower than the the cosmogenic muon rate, whilst the spallation neutron rate from the surrounding rock is $< 10^{-7}$ kt$^{-1}$s$^{-1}$. Together with the spallation and internal cosmogenic neutron backgrounds, the combined neutron background from radioactive sources in the AV, photodisintegration, and $\alpha$ particle interactions totals $<10^{-3}$ kt$^{-1}$s$^{-1}$ [40]. Once again, this figure is far lower than the expected

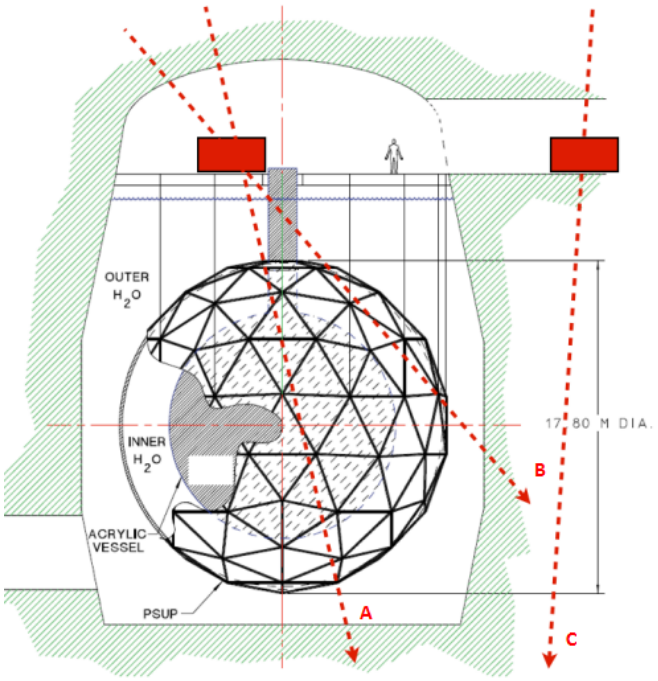



Figure 3.9: Diagram showing the various tracks of muons through the detector. Track A shows a typical 'through and through' muon event, whilst track B shows the path of a so called 'grazing' muon. Track C avoids the detector entirely, but may result in neutron spallation background due to muonic attenuation in the surrounding rock. Original diagram from [69], modified to illustrate possible muon paths.

rate of supernova interactions at SNO+ and can also be neglected for the purposes of this thesis.

### 3.5.3 Solar and Atmospheric Neutrinos

The background rate of solar neutrino interactions at SNO+ is largely constant with respect to time, and is dependent solely on the rate of fusion reactions in the Sun. Solar neutrinos possess a wide energy spectra given the variety of reactions present in the solar interior, though the majority above ~1.5 MeV predominantly arise from the decay of $^{8}$B. In this sense, it is irreducible, but very well understood.

Fortunately for supernova studies, the rate of solar neutrino interactions compared to the rate of prospective supernova events at SNO+ is low, at $< 0.03$ kt$^{-1}$s$^{-1}$ [69]. Atmospheric neutrinos, originating from cosmic ray interactions in the upper atmosphere, have a low enough rate of interaction at SNO+ to be fully negligible at $\lesssim 10^{-5}$ kt$^{-1}$ s$^{-1}$ [40]. The solar neutrino background spectrum below 3 MeV is shown in Figure 3.10. The dominance of $^{8}$B above ~1.5 MeV is clearly seen, whilst below this point the solar background spectrum increases progressively with reducing energy due to the influence of pep, $^{7}$Be and pp chain neutrino interactions respectively. All super-

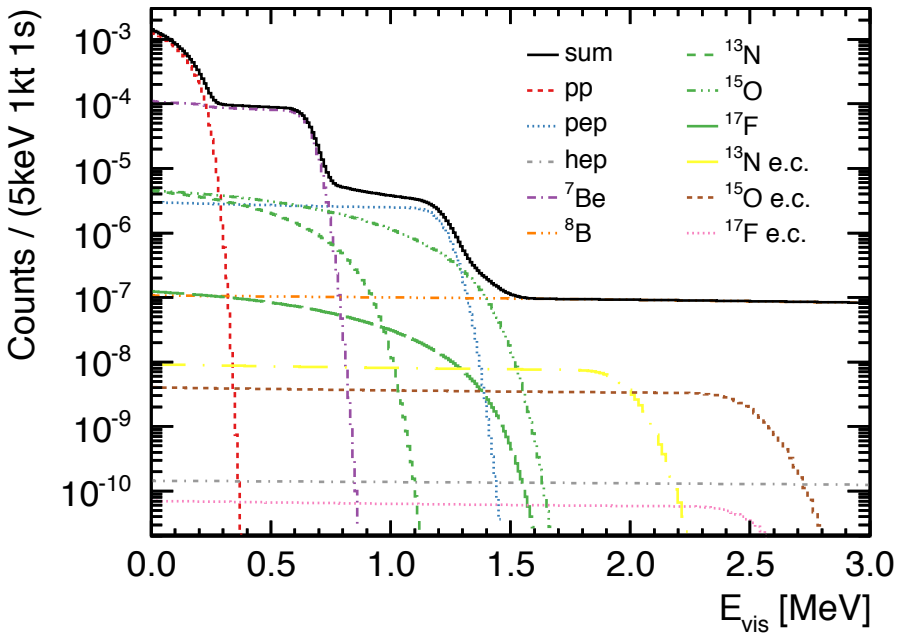



Figure 3.10: Solar background spectrum below 3 MeV for unoscillated electron neutrinos in scintillator at an energy resolution of $^{6\%}/\sqrt{E_{vis}}[MeV]$. [40].

nova neutrino interactions of interest to this thesis take place above 1.5 MeV, and therefore remain largely unaffected by the solar neutrino background.

# Chapter 4

# The L3 Supernova Trigger

## 4.1 Overview

Supernovae are prized events in the study of neutrinos. Very few other sources have the potential of outputting neutrinos in the same quantities, nor be as sensitive to the intrinsic properties of the particles themselves. Determination of mass mixing effects, the neutrino mass hierarchy, and neutrino oscillation parameters such as $\theta_{13}$ are all potentially possible with sufficient statistics [55][78]. All supernovae produce a characteristic burst of neutrinos, on the order of 10 seconds long, in which large quantities of information about the progenitor star and the neutrinos themselves are encoded in the form of time and energy spectra. While only 24 neutrino events were detected from SN 1987A, analyses of these few events have produced hundreds of papers and shed a large amount of light upon the mechanism of core collapse itself.

As these bursts are so rare, and last only a few seconds, it is imperative that experimental trigger systems are fast enough to detect them, whilst remaining detailed enough to tag them as supernova candidates. As neutrino experiments have the potential to provide such an effective early warning system for the astronomical community, it is additionally important that the latency between data acquisition and signal confirmation be kept to the bare minimum.

Whilst the current trigger system at SNO+ is eminently capable of recording a supernova, it lacks the ability to perform detailed analysis and therefore differentiate between real data and instrumental background on short timescales. Any data received would be categorised simply as a "burst", placing any genuine supernova signal alongside unrelated physics events, flashers, and even detector breakdown. It is then clearly important that an additional stage of processing is applied to

recorded bursts, both for internal analyses and also to provide the capability for any report to the SNEWS network[1]. The L3, or Level 3 Trigger, is this additional stage.

For any potential L3 system, a number of minimum goals must be met. Primarily, the system must be able to distinguish between a real supernova and background, with minimal false alerts and a fast reaction time on the order of minutes. To do this, it required a number of different checks to be run over the burstfile output of the L2, specifically checks on isotropy, data cleaning, and energy. A simple, quick position and energy fit conducted over the events can then potentially give us large amounts of information as to the nature of the burst.

## 4.2 Supernova Triggers at SNO+

In principle, a supernova burst should be unmistakeable. A huge flux of neutrino events pass through the detector in a very short timespan, offering a clear signature of the collapse of a massive star. In practice however, these same strengths can also be weaknesses. High Nhit bursts in the detector are not unique to supernovae, and are caused far more often by defects in the detector; flashing PMTs or simple electronic malfunctions can easily mimic a supernova. Similarly, the uncertainty in knowing at what distance a galactic supernova will occur results in large variability in the appearance of any signal in the detector, requiring a very low energy threshold to be run constantly. This poses challenges to data handling and storage.

Given that we only expect a galactic supernova perhaps 3 times per century on average [33], without robust monitoring there exists the risk of data misinterpretation and the potential for the loss of time-sensitive data. Mitigating these issues is therefore of the very highest priority in supernova searches, and may be achieved through constant software monitoring and near-line analysis.

SNO+ has a method in place to monitor the detector for burst events, namely the L2 trigger/filter system operating in concert with the *Stonehenge* program [79]. This system runs between the event builder and the writing of data to the Grid[2], with the purpose of running the detector at a low threshold in order to constantly watch for supernova candidates whilst not retaining the low energy data permanently. This system watches for events using a simple Nhit criterion, then monitors the

[1]SNEWS refers to the Supernovae Early Warning System network, a collaboration of international neutrino detectors sensitive to the neutrino signal from core collapse supernovae. The goal of SNEWS is to provide the astronomical community with a prompt alert of the occurrence of a galactic core-collapse event, and to optimise global sensitivity to core-collapse supernovae signals. SNEWS has been running since 2005, and currently encompasses seven neutrino experiments: Super Kamiokande (Japan), LVD (Italy), IceCube (Antarctica), KamLAND (Japan), Borexino (Italy), Daya Bay (China), and HALO (China).

[2]The Grid is a distributed set of computer storage and processing resources, accessible through a common interface.

time window following the initial burst event. If subsequent events occur within the time window expected of supernova, it is retained as a burstfile.

The L3 processor and trigger system is then responsible for taking the burstfiles produced by L2 and performing a far more rigorous, yet efficient, series of checks to verify whether these burstfiles contain supernova candidate events This allows a much higher level of confidence in the issuing of any global supernova alert than could be provided with burstfiles alone. As the primary focus of this thesis, the creation of this system, its implementation, preliminary results, performance, and issues will be considered with some detail for the remainder of the chapter.

### 4.2.1 The Event Builder and the Level 2 Trigger System

The event builder is an integral part of the SNO+ DAQ system, and plays a pivotal role in the integration of data from a variety of different DAQ sources into a single, readable file. The builder reads a single stream of data from the master trigger card (MTC), the PMTs, and the CAEN ADC via an Ethernet connection to ORCA, before reorganising and integrating this data into physical events using the global trigger IDs, or GTIDs, associated to each[3]. This is then all written into ZDAB[4] files to be further analysed by the monitoring system, processed, and ultimately passed to the Grid [74].

The Level 2 trigger system at SNO+ consists of a number of ZDAB utilities coupled to the Stonehenge program, alongside a number of bash scripts to handle file movements and launching the various jobs. Stonehenge receives ZDAB data directly from the event builder via a virtual filesystem. It can then perform a number of basic analysis tasks, such as simple cuts, the recording of event rates, and burst monitoring. Essentially, it is a system by which SNO+ can retain lower Nhit data than normal for a brief time period without altering the final data being written to the Grid. As the concern of this thesis is in employing the output of the L2 for supernova detection, the method by which a burst is detected is key.

The Stonehenge burst detection algorithm defines a burst in a highly specific way: the burst 'begins' when a sufficiently large number of high Nhit events occur within a particular time interval. It then 'ends' when the rate of these high Nhit events drops below some specified threshold. The algorithm therefore identifies four parameters that define a burst event: initial Nhit threshold, the burst rate, the burst length, and the event drop rate. These parameters can be altered to change

[3] A global trigger ID is assigned to every incidence of a global trigger occurring in the SNO+ detector. The GTID can then be used as a unique identifier of each event. The process of issuing a global trigger is detailed in Section 3.3.3.

[4] The ZDAB legacy file format was originally used on SNO and is still used for storing all raw data from SNO+ events. RAT reads these files natively and may also convert ZDAB files into ROOT files for wider use.

the behaviour of the burst monitor using a configuration file. Once a burstfile is produced, it is then converted to a ROOT file by a script within the L2 system and deposited within a directory with other supernova candidates. Any events that exceed the initial threshold for burst detection, without meeting the full burst criteria, are kept as *L1* files and are retained on an underground RAID array for one week [79].

### 4.2.2 The L3 System Concept

The L3 system has a seemingly simple task to fulfil, but due to the constraints of the environment in which it is applied, it must also be highly efficient and streamlined. Ultimately, it must be capable of telling operators how likely it is that any given event is a supernova, and it must do this quickly. It therefore straddles the demands of any basic near-line analysis (speed, data handling, robustness) with the challenges of a physics analysis where accuracy is crucial. It was decided that the best compromise on this front was to use a simple series of checks which can be applied quickly, extending the work done by L2, yet still yielding good physics.

By definition, these checks must be able to handle data on an event-by-event basis, denying the trigger access to high statistics analyses. At first glance this is problematic for supernova monitoring - the strongest observable signal of supernovae in current generation detectors arises from the delayed coincidence tag on antineutrino IBD events. This method requires running over multiple events to spot the coincidence tag however, and is therefore unusable on an event-by-event basis. Similarly, whilst neutrino-proton elastic scattering is thought to be the single largest supernova signal at SNO+, the soft scintillation light released by the recoiling proton is non-identifiable on an event-by-event basis.

Fortunately, there exist other supernova-specific interactions which should be readily apparent in any supernova neutrino burst, and are therefore potentially identifiable by any prospective L3. The first of these is the neutron-capture component of the IBD interaction - the thermalised neutron produces a monoenergetic peak at 2.22 MeV when captured on hydrogen, and as IBD forms the largest visible signal from supernovae at SNO+, this peak should be extremely distinctive. In addition, the nuclear excitation of $^{12}$C via neutral current neutrino interactions produces a monoenergetic 15.11 MeV de-excitation gamma for each event. Due to the substantial number of carbon nuclei in the scintillator volume, these events are thought to comprise the third-largest supernova signal at SNO+ during the scintillator phase. The combination of these interaction signatures therefore provide one of the most viable means of detecting supernova on an event-by-event basis at high speed.

When coupling these tags with a number of other, more generic, checks it is possible to provide a much more reliable estimation of the likelihood of any one burstfile containing a supernova. In particular, checks on event isotropy and timing can be used to eliminate a large fraction of common detector backgrounds and retrigger events, thereby increasing confidence of detection.

Supernova interactions are expected to occur with no preferential direction at SNO+ despite the inherent directionality of a supernova event: scintillation light smears out any directional information until it is unrecognisable, sacrificing directional detail for much lower energy thresholds than large water Cerenkov detectors. However, the structure of the SNO+ detector itself opens the possibility of using the pure water shielding around the AV as a means of determining event directionality during the scintillator phase. The PSUP volume contains around twice the volume of the AV itself, potentially creating sensitivity to supernova direction via the use of Cerenkov light cones and event timing. For the simple case of events taking place along the line-of-sight to the supernova, any events taking place within the water shield should manifest themselves as prompt light within the detector, as the Cerenkov light will reach the PSUP prior to any scintillation light. We would then also expect the majority of these events to occur on the opposite side of the detector to the direction of the supernova due to the directionality of the Cerenkov light cone itself. However, this analysis is beyond the scope of this thesis, and thus for the remainder of the work supernova events that occur during the scintillator phase will be assumed to be fully isotropic.

In the water stage, electron elastic scattering interactions do potentially contain directional information, but due to the low event counts of these interactions compared with other events in the water stage, any anisotropies from elastic scattering are expected to be invisible against the background of other, isotropic, supernova events. Only larger mass water Cerenkov detectors, such as Super Kamiokande with a total volume of 32 kt, are expected to be able to resolve supernova directions [80].

Electron elastic scattering interactions are only expected to produce $\sim$8 events kt$^{-1}$ in water for a 10 kpc supernova, whilst the major interaction signal in water, IBD, which does possess an almost isotropic distribution, produces approximately 265 events kt$^{-1}$ at 10 kpc [80]. All other supernova interactions at SNO+ in the water phase are expected to be fully isotropic. Accordingly, supernova events at SNO+ in the water phase should occur with an approximately isotropic distribution across the detector volume, while events in the scintillator phase should be fully isotropic. As a large number of instrumental backgrounds are anisotropic or localised, this check on event isotropy is expected to be a powerful discriminator of burstfile data in water.

The second simple check we can do is one of timing. For high-frequency bursts with high average Nhits, detector retriggers are a common issue. Retriggers are caused by late light that has lingered long enough in the AV to exceed the trigger lockout time of 440 ns and pass into the next event GTID window, thereby *retriggering* the detector by lighting up one or more PMTs. These can seriously affect any analysis of the data by introducing large numbers of "false" low-Nhit events, and are best cut for this reason. This cut can be achieved effectively with minimal loss of signal by ensuring the cut time lies suitably close to the GTID window. Given a 10 kpc supernova, SNO+ expects approximately 888 events in scintillator over 10 seconds, or an event rate of ∼89 Hz [40]. As the GTID interval at SNO+ is set by the detector lockout time of 440 ns, a 500 ns cut should eliminate the bulk of retrigger events whilst having a negligible impact on data gathering. A cut too close to this lockout window will miss late retrigger events, thus 500 ns was selected as the best balance between eliminating the greatest number of retriggers and minimising the impact on data collection.

Finally, simple data cleaning checks should act as an additional discriminator of instrumental background. SNO+ utilises a number of data-cleaning processors for analysis, including a comprehensive masking system for identified instrumental backgrounds specific to the SNO+ hardware. By integrating these pre-existing utilities into any L3 processor, it should be possible to further reduce the probability of misidentifying supernova events.

## 4.3 The L3 Trigger Algorithm

The checks and cuts detailed above are then implemented via a RAT processor, L3Proc, with future versions utilising a control shell script that recognises when burstfiles have been passed into the directory and automatically runs L3Proc over the data. The current prototype is limited to manual running. A simple algorithm is used within L3Proc to analyse the burstfiles passed to it by the L2, with the events within each burstfile undergoing a succession of checks before finally being subject to event reconstruction. The L3 also adjusts each stage of the algorithm depending on the material inside the detector. Each stage of the algorithm is detailed in Section 4.3.1 in order of application.

As well as conducting these checks, the L3 also outputs a number of plots and detailed information about the burst, and contains the structure necessary to output a SNEWS alert signal. This aspect of the L3 is detailed in Section 4.3.2.

### 4.3.1 Components

#### 4.3.1.1 Data Cleaning and Non-Physics Triggers

Data cleaning processors are the main method by which instrumental backgrounds are identified and tagged in SNO+ data, and are mostly concerned with spotting electronic cross-talk and PMT malfunctions. In total, the "datacleaning" processor has 15 background tags which it can implement over a dataset. These are then added to the event tree for that dataset, allowing later processors to manipulate or remove specific backgrounds, or even the entire set of 15. The L3 will eventually use this functionality to ensure the burstfiles are as clean as possible, particularly important given the increased likelihood of detector malfunction during the very high rates expected for supernova bursts.

In addition to data cleaning tags, the SNO+ detector has a number of triggers used solely for calibration and testing purposes. As these runs are quite clearly incompatible with supernova monitoring, a cut is in place within the L3 to remove any events identified with a non-physics trigger bit. In practice, this means all triggers are passed other than PulseGT, Pedestal, EXT8, and EXTASYNC. The first two denote diagnostic and electronic calibration runs, whilst the latter two are used for the optical calibration runs of ELLIE and its subsystems. There of course exists a very small probability of supernova detection during these runs, though the risk of this is deemed negligible given supernova timescales.

#### 4.3.1.2 Timing Checks

The primary check on timing done by the L3 prototype is that of the retrigger cut. As mentioned previously, this cut is in place to prevent the inclusion of large numbers of "false" low Nhit events, generated by residual late light in the detector. This becomes particularly problematic during high energy, high-rate events such as close supernova bursts, when very large numbers of retriggers can drown out signals at low energies and therefore disrupt near-line supernova checks.

The L3 also measures the interval between the initial and final events observed within the burstfile to determine an estimate of the length of the burst. The accuracy of this measurement is largely dependent on the specific L2 burst criteria, which may initiate or complete a burst measurement prematurely, or even after the true burst has begun. However, as the margin of error on the L2's burst determination is on the order of milliseconds, and the characteristic timescale of supernova neutrino bursts is on the order of 10 seconds, short of L2 breakdown this uncertainty should have only a marginal impact on the final outcome in the case of a supernova event. The initial neutron-

isation signal, containing the vast majority of supernova neutrino events, should reliably initiate the burst processor by satisfying the rate requirements. The rate difference between this initial signal and the later cooling period may potentially create issues with premature burst termination, though this ought to be readily corrected and optimised prior to full detector operation via the supernova calibration source, or even high-rate TELLIE runs.

Additional timing checks carried out by the L3 include cuts on orphan events. *Orphan* events are commonly associated with detector malfunctions during high-rate data taking, and whilst used somewhat as a catch-all term at SNO+, its use in this thesis specifically refers to events which cannot be associated with a global trigger by the event builder. These events are then "orphaned" and left without a GTID. As these events are essentially noise, they are cut by the L3 using a check on the GTID timing - any event with a GTID clock count of zero must be an orphan, and is therefore removed. Whilst the data cleaning processors should in theory remove orphans prior to L3 running, their removal is important enough to warrant contingency checks within the L3 processor itself.

Finally, the L3 carries out calculations of the time residual of an event, shown in Figure 4.1. Time residuals are used in SNO+ to measure the time of light emission with respect to the event, rather than with respect to detector triggering. The time residual per PMT for each event is calculated via the following equation:

$$t^i_{\text{res}} = t^i_{\text{hit}} - t^i_{\text{flight}} - t_{\text{reco}}$$

Where $t^i_{\text{res}}$ represents the time residual, $t^i_{\text{hit}}$ is the $i$th PMT's hit time, $t^i_{\text{flight}}$ is the time of flight of the photon from the reconstructed position to the $i$th PMT, and $t_{\text{reco}}$ is the reconstructed event time. We then expect all interactions in scintillator to display some standard form, split between the prompt light of the singlet transitions and the late light of the triplet transitions.

Predominantly used in optical calibration analyses, the L3 outputs time residual plots as an additional tool for operators monitoring for supernovae as they offer valuable timing and positional information that may prove useful in later analysis. As such, their calculation is therefore anticipated to be actively utilised in future, more sophisticated iterations of the L3.

#### 4.3.1.3 Isotropy Checks

Supernova events at SNO+ should produce light that is distributed almost isotropically within the detector medium. It should be noted however that during the scintillator phase, any event causing

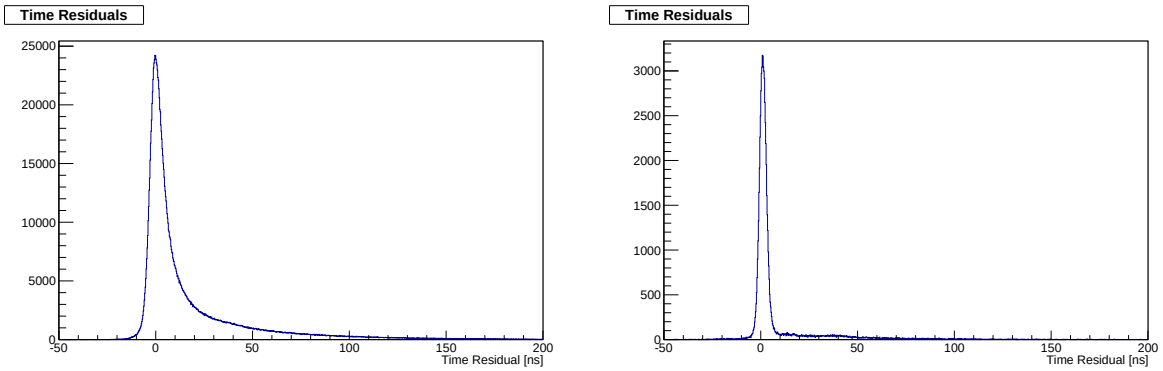



Figure 4.1: Sample time residual output of the L3 for scintillator (left) and water (right) for a simulated supernova at 10 kpc. Note the broad scintillator peak compared to the water peak, evidence of the delay between singlet and triplet transitions in LAB.

scintillation will appear to be isotropic given the nature of the mechanism. Thus, all isotropy checks discussed below are currently envisioned to be used only during the water phase. SNO+ does provide an isotropy classifier at the macro level, though the additional delay in processing burstfile data this way is unwelcome, and due to time constraints, unused.

To ensure that the L3 accurately measures the degree of isotropy in the detector, the L3 processor utilises a tool within RAT known as the segmentor. The segmentor divides the PSUP into a user-specified number of segments $n_{\mathrm{div}}^2$ of approximately equal area in $\cos\theta$ and $\varphi$, with $n_{\mathrm{div}} = 10$ by default, before allocating a number of PMTs to each segment [81]. It is then possible to trace back any hit PMT to its given segment, and use this information accordingly. As backup, the L3 additionally outputs a map of all PMTs hit over the course of the burstfile in $\cos\theta$ and $\varphi$, allowing a visual interpretation of isotropy information to be built up for each burst.

The L3 uses the segmentor by examining the PMT hit locations for each event in the burstfile, and checking which segment contains the most Nhits for any given event. This segment is then assigned *primary* status. Each segment can only be assigned primary status once over the course of the burst, and the L3 keeps track of each unique primary segment assigned. Once the burst is complete, the L3 compares the number of unique primary segments hit to the total number of segments, and outputs this information both explicitly and in the form of percentage detector coverage. Isotropic bursts will typically show a very high detector coverage percentage of 70-100%, whereas anisotropic bursts typically produce low detector coverage regardless of the energy regime in which they occur. As the bulk of instrumental backgrounds produce anisotropic bursts, the segmentor can therefore be a powerful tool in discriminating supernova events in water.

The advantages of the segmentor method of isotropy checking is apparent compared to the alternatives: it is relatively quick, efficient, accurate, and largely independent of the energy regime of the burst. The latter is useful when considering bursts from supernovae at distances greater than 10 kpc. However, the method does have its drawbacks. Without a threshold, segments can be assigned primary status despite very low Nhits, with very low statistical significance. Similarly, segments very high in Nhits may also be ignored should the burst be of a high enough energy, when these segments may contain significant isotropy information. These issues may be mitigated by introducing a low energy threshold in the former case, and altering the resolution of the segmentor check in the latter case. In general however, the influence of these effects on the accuracy of the check itself is expected to be low.

#### 4.3.1.4 Event Reconstruction

For any event, SNO+ must translate the raw charge and time data of the PMTs to real physical information: the position, time, and energy of a given interaction. SNO+ has a wide variety of fitters to accomplish this task, each used for slightly different purposes, though all designed to reconstruct events as accurately as possible from raw PMT information. Typically, the most robust of these fitters are time consuming; in effect, proper fitter utilisation involves a fine balance between the accuracy of the fit and the speed by which the fit is obtained. For the vast bulk of analysis purposes, speed is a secondary concern. Supernova monitoring requires a different approach however.

Strictly speaking, event reconstruction is not necessary to spot a simple burst in the detector. Given the rarity of supernova however, it becomes vitally important to be able to properly identify the origin of a burst. Energy reconstruction in particular allows the observation of supernova "fingerprints" which would otherwise be unobservable: the peaks of the IBD interaction at 2.2 MeV, and of $^{12}$C de-excitation at 15.11 MeV. Position and time reconstruction permit better energy fits, and are therefore also vital for this purpose. Unfortunately however, the typical timescales of supernova neutrino bursts are much smaller than the processing time of the best fitters. This is problematic when one of the primary purposes of supernova monitoring is to alert the wider astronomical community promptly.

Consequently, the L3 must use the most efficient fitters possible, without sacrificing accuracy. Fortunately, the idiosyncratic nature of supernova neutrino bursts ensures that the accuracy of the fitters need not be as high as those used in detailed physics analysis, such as $0\nu\beta\beta$ searches, so the L3 may trade accuracy for speed without a great loss of statistical significance in the final result.

Likewise, for a typical analysis, a number of additional classifiers and optimisers are run to tailor the analysis to the interaction being studied. In the case of the L3, these are largely unnecessary and simply delay the processing time further.

For position and time fitting, the L3 will use the *quad* position fitter [82]. The quad fitter fits point-like events by using multiple combinations of four PMTs and calculating an analytical solution for the origin of the light source for each combination, assuming a constant speed of light. Each solution of x, y, z, and t is then combined, and the median returned to yield a position and time fit for a given event. Due to its simplicity, the quad fitter is primarily used in SNO+ to provide a starting point or *seed* for more complex fitters, but provides enough accuracy to be used for fast position reconstruction purposes such as those required by the L3.

For energy fitting, the L3 will use the *energyFunctional* fitter [83] at the macro level, using the result from the quad fitter as a seed. Previously, the standard method for energy reconstruction at SNO+ involved the use of an Nhits lookup table, whereby the Nhits of an event would be compared to a standard table of Nhits for varying energies and radii to "reconstruct" the original information. Practically this was not ideal as it failed to account for the asymmetries of the detector, whilst the production of sufficiently accurate tables for each detector medium was inefficient and time consuming. The energy functional method instead deconstructs the problem by mapping the Nhits of an event to a number of independent functions. In principle, the Nhits of any event at SNO+ should depend on three main variables: the energy of the event, the radius of the event, and the z-coordinate of the event. The latter two are included due to the intrinsic asymmetry of the detector; events positioned near the edge of the AV or the neck can result in highly non-linear effects on Nhit counts. As a simple treatment, the energy functional then determines the energy of an event by using the following equation:

$$\text{Nhits} = f(E) \times f(r) \times f(z)$$

This yields more accurate energy reconstruction than the older lookup method with no loss of speed, and is now used as standard in the general purpose *scintFitter* combination fitter and others.

The L3 combines the quad and energyFunctional fitters into the *fastfitter*[5], which is then implemented via macro over burstfiles generated in both water and scintillator data. The performance of this fitter can be seen in Figure 4.4.4.1.

[5]For the purposes of this thesis, this combination of quad and energyFunctional fitters is referred to as the fastfitter, although this term is not used in SNO+ and is for convenience purposes only.

#### 4.3.1.5 Supernova-Like Event Counters

One of the key responsibilities of the L3 is to inform the SNO+ collaboration of the likelihood of a supernova event having occurred in a given burstfile. Whilst later iterations of the L3 are expected to be quite sophisticated in this regard - outputting SN likelihood information online for each burst and eventually notifying the SNEWS network automatically - initially the L3 is constrained to simply give fast and accurate information to the user. One of the most effective methods for this task is via event reconstruction of the burstfile, and in particular, energy reconstruction. As detailed in previous chapters, supernovae produce a highly distinctive range of interactions at SNO+ in all experimental phases. Two of the most significant of these are IBD and $^{12}$C de-excitation.

Any major physics analysis of IBD will invariably use the delayed-coincidence tag between positron annihilation and neutron capture due to its high efficacy and reliability. However, for this method to be effective, events must be studied en-mass post data-collection, making it ill-suited to the demands of near-line supernova monitoring which, by definition, must process burstfiles event-by-event. Fortunately, the sheer number of events from a supernova burst ensures an IBD rate much higher than those of typical reactor antineutrino and geoneutrino events. This manifests itself as a large monoenergetic peak at 2.2 MeV in the reconstructed energy of the burst, representing the gamma released by thermalised-neutron capture on hydrogen. Thus, in its initial iteration, the L3 counts events around this peak within a $3\sigma$ confidence interval, as determined from large-scale simulations[6] of IBD events, and labels them *IBD-like*. Whilst not as accurate as a delayed-coincidence tag at spotting IBD events, typical supernova event rates ensure it is nevertheless an effective tool for supernova monitoring purposes.

The de-excitation of $^{12}$C is anticipated to be the third largest supernova signal at SNO+. It possesses a very clear signature through its monoenergetic 15.1 MeV gamma, and is expected to yield enough events to be above the supernova-induced background for the high energy regime. Accordingly, the L3 takes advantage of this monoenergetic peak, and will initially count these events as $^{12}$*C-like* within a $3\sigma$ confidence interval. Along with the monoenergetic peak, a further signal of this interaction at SNO+ comes from low-Nhit Cerenkov events occurring in the water shielding between the AV and PSUP: the long attenuation length of the 15.1 MeV gamma allows photons produced near the AV boundary to escape outwards without causing scintillation. According to

[6]These limits were determined by simulating 25000 randomly-distributed IBD events in scintillator, then reconstructed using the same combination of fitters used by the L3. The outcome of these simulations for both the scintfitter and fastfitter can be seen in the fitter comparison plot in Figure 4.5.

simulation, this occurs in around 10% of $^{12}$C de-excitation events, and shows up as a consistent low-Nhit peak. Whilst potentially useful in any later version of the L3, the difficulty in spotting these events against the low-energy background discourages their use initially.

Taken in isolation, the above "event-like" counters are both somewhat simplistic, yet together they are greater than the sum of their parts - short of any software malfunction, any burstfile which passes the other L3 checks, whilst still yielding high counts for both $^{12}$C-like and IBD-like events has a very high likelihood indeed of containing a supernova. The occurrence of these two peaks simultaneously in any background is negligible, whilst the speed of implementation within the processor itself ensures minimal processor overhead once event reconstruction is completed.

### 4.3.2 Output and SNEWS Alert

One the main roles of the L3 is to output burst information to the user in the form of plots and logs. This enables later analysis of the L3 processed burst file, especially important should an automated L3 issue a SNEWS alert, or present borderline data. Currently the L3 outputs a number of plots to a ROOT file, and all relevant non-graphical burst information to terminal, though this is later envisaged to output to a log file alongside the ROOT file. The plots contain a range of event information including energy, position, and isotropy distributions, as well as PMT hitmaps and timing information. The non-graphical burst information contains general details such as burst length, total number of events, and average Nhits, along with counts of common signs of instrumental background such as orphans[7].

In its current prototype form, the L3 contains only the framework for any potential supernova alert, with the intention of this being greatly expanded at a later date. Currently, the framework looks for a number of supernova indicators and issues a dummy alert to terminal should a specified combination of these exceed threshold. These indicators differ for scintillator and water phases, and involve a combination of event criteria, with a focus on energy information, timing information, and, in the case of water, isotropy information. As the current dummy SNEWS alert is only in place for testing purposes, the criteria required to issue an alert are not rigorous. Later work on this aspect of the L3 will be subject to stringent requirements from the central SNEWS group at Brookhaven National Laboratory, and therefore the actual criteria used will be decided at the collaboration level of SNO+.

[7]For the purposes of this thesis, an orphan event is defined as an event with no associated GTID. A full definition is given in Section 4.4.2.

## 4.4 L3 Prototype Performance

The performance of the L3 can be scrutinised by examining its behaviour in three key areas: timing, isotropy, and energy reconstruction. These areas each encompass a number of independent checks conducted by the L3 itself, including water-specific isotropy checks and scintillator-specific energy reconstruction checks, which may be analysed independently of each other. This allows us to determine the overall efficiency of the L3 prototype, whilst highlighting any shortcomings in order to inform future development.

### 4.4.1 L3 Prototyping and Simulation Testing

The development of the L3 was facilitated mainly through the use of supernova simulation burstfiles. Supernova neutrino bursts were simulated using RAT at 10 kpc, 30 kpc, and 100 kpc in both water and scintillator, using files originally generated by the SNGEN supernova simulator [84]. The L3 was then run over these simulated bursts in order both to test the functionality of the code and fine-tune its features. For the majority of testing, the code was run over the 10 kpc simulated burst - as the "standard" supernova within neutrino astronomy, the vast majority of supernova neutrino research is carried out for a prospective supernova at this distance from the Earth. The L3 was therefore essentially calibrated over this burstfile - all previous studies of expected supernova event rates at SNO+ have been carried out at this distance, and therefore irregularities in the code are most easily spotted using this method with minimal potential for misinterpretation. Due to the very low background expected at rates and energies typical of supernova neutrino bursts, these simulations were conducted without any background processes enabled[8].

Unfortunately, for reasons detailed in Section 4.5, errors within the Monte Carlo simulated burstfiles used for the L3 analysis precluded any assessment of the use of the $^{12}$C de-excitation tag discussed earlier in the chapter. For this reason, the remainder of this thesis will focus on the use of IBD as the primary interaction used for supernova tagging, alongside all additional checks and cuts, although a preliminary $^{12}$C tag is currently functional within the L3 processor itself.

Once a functional L3 prototype was complete, it was then run over burstfiles generated by the SNO+ Mock Data Challenge, or *MDC*, undertaken in July 2015 [85]. The MDC was designed to test the full data processing chain of the SNO+ detector, from the raw data on the builder to the final calibrated files on the Grid. Necessarily, this involved the L2 and production of burstfiles that

[8] An analysis of common detector backgrounds, as well as their impact on supernova monitoring, is given in Chapter 2.

the L3 would be expected to analyse once fully operational. It was therefore an ideal first test for the L3 prototype under semi-real operating conditions. As the L3 is currently not capable of being run near-line, the MDC files were instead analysed locally, though this should have had little effect on the results. Some issues were present in the structure of the MDC data[9] which prevented full analysis of the L3's capabilities - namely the use of energy fitters - although timing and isotropy checks were enabled.

The MDC data files were selected from a number of previous data runs, taken during the partial-fill detector dark runs in December 2014. The run selection was intended to cover a wide range of run types, including bad runs, and was therefore ideal for L3 testing purposes. Each run type was useful in testing the different capabilities of the L3 prototype, and these are discussed in the performance section below.

### 4.4.2 Timing Checks

#### 4.4.2.1 Simulation Data

Figure 4.2 shows a comparison of supernova events in scintillator at low energy with and without the retrigger cut for a simulated supernova at 10 kpc, using the fastfitter. Clearly apparent is the large decrease in reconstructed low energy events in the retrigger-cut data, resulting in a far more prominent 2.2 MeV gamma peak from neutron capture. The retrigger cut therefore performs well over simulated data, permitting signal discrimination whilst reducing what is essentially detector noise.

Without removal, retriggers dominate the low energy regime, which has ramifications for the L3 despite being outside the main energy regions of interest. Retriggers add large numbers of "false" low energy events, and therefore have the potential to greatly affect calculations of L3 efficiency by imitating genuine PES events, which typically occur at energies less than 1 MeV. PES events are expect to form the majority of supernova neutrino interactions at SNO+ in both phases, and therefore an estimate of their total count is key to providing an estimate of supernova distance and L3 efficiency. This is explored in more depth in Section 4.4.5.

[9]Specifically, the MDC data inherited the incorrect run-types that were set during the December 2014 data-taking, which required manual intervention into the ORCA DB run documents used during the challenge. Simultaneously, issues with burstfile subrun ID setting were observed in the L2 prior to the MDC. It is thought that a combination of these may have led to the MDC burstfiles examined in this thesis to possess null run IDs, causing the fitters to malfunction. The correction of this error was deemed too time-consuming for use in this thesis, and considerably outside it's scope.

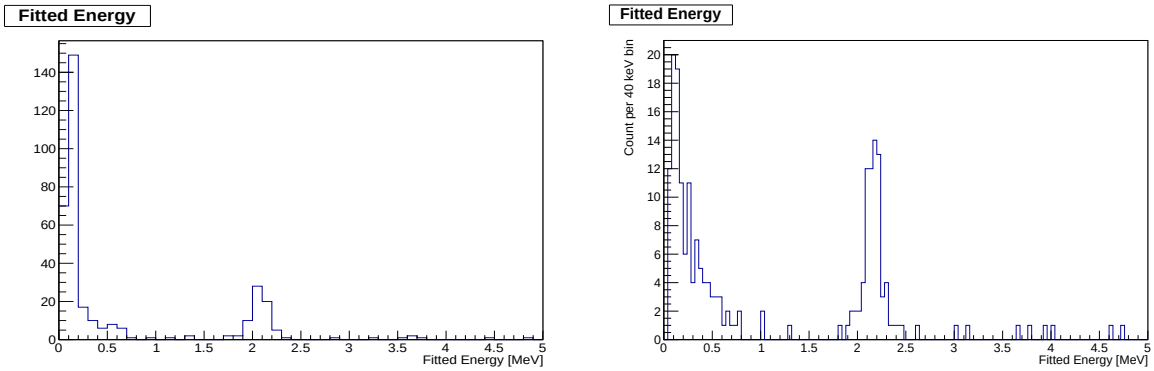



Figure 4.2: L3 output plots of the reconstructed event energy with (left) and without (right) retriggers from a 10 kpc simulated supernova in scintillator. Note the marked increase in low energy events with retrigger events included, as well as the clear peak from IBD neutron capture at 2.2 MeV. The low energy event count is reduced by over a factor of three after the cut is applied, revealing potential PES events.

#### 4.4.2.2 MDC Data

The MDC files proved useful in testing the L3's response to large numbers of orphan events, along with retrigger discrimination. Largely, the L3 was very effective at spotting orphans over a range of run types including high-rate, low-rate, and TSLP crate-by-crate[10] runs, as shown in Table 4.1. All events not tagged as orphans despite possessing an orphan trigger in the event information were labelled as zero-hit events by the L3 processor, and therefore still eliminated by the L3 trigger.

| MDC Run Type | Total Events | Total Orphans | L3 Orphan Count | L3 Zero-Hit Count |
|---|---|---|---|---|
| V. High Rate DQ Test | 845 | 840 | 839 | 1 |
| Low Rate DQ Test | 17 | 14 | 14 | 0 |
| TSLP Crate-by-Crate | 11 | 11 | 10 | 1 |

Table 4.1: Table of L3 orphan counts during MDC testing. The total orphan count was determined from the number of orphan triggers issued by the detector. During running, large quantities of orphans should be removed by the datacleaning processor using this trigger bit. The L3 conducts additional orphan contingency checks should the datacleaning processor fail.

### 4.4.3 Isotropy Checks

The main isotropy check undertaken by the L3 prototype involves the use of the segmentor mentioned in in Section 4.3.1 - by comparing the number of unique segments hit over the course of a burst we may determine a rough measure of the total isotropy of the event. We expect supernova events in both water and scintillator to be fully isotropic, though isotropy is only meaningful in the water case, and therefore all isotropy tests were conducted solely in water using the 10 kpc and 30 kpc

[10]TSLP crate-by-crate refers to timing slope runs, conducted for electronic calibration purposes, in which the channels in each crate are checked sequentially for their response to delays in the global trigger.

simulated supernova data sets.

The segmentor detector coverage percentage given by the L3 should be high for real supernova at all distances, although the total number of unique primary segments will scale with the number of events passed by the L3 should the number of events passed be of the same order as the number of segments, which defaults to 100 though may be set by the user. Routine detector backgrounds such as flashers typically show highly localised events, and therefore their detector segmentor coverage should be low. Short of total detector breakdown, it is expected that the segmentor check should identify anisotropies in the detector with some confidence. The MDC therefore offers the chance to test the segmentor check against what can be considered background patterns, particularly when using localised run types such as crate-by-crate TSLP runs, whilst the supernova simulations enable checks to be made of segmentor performance with realistic supernova event counts.

In addition to the primary segmentor isotropy check, the L3 also outputs PMT hitmaps of each burstfile, as well as reconstructed positional information. These offer valuable redundancy in isotropy checking, and allow operators the ability to visually inspect burst patterns to supplement segmentor checks. In addition to the burst maps, the L3 also outputs a 2D histogram of the position fit in $\cos\theta$ / $\phi$, permitting an additional visual isotropy check. In principle a 2D position fit should offer a very good estimate of the isotropy of a burst, although in reality the histogram is subject to large errors from the results of the water fitters, and thus should be considered in concert with the fitter-independent segmentor results and PMT hitmaps in order to form an accurate assessment of isotropy.

#### 4.4.3.1 Simulation Data

Water simulations for both 10 kpc and 30 kpc data sets showed high degrees of isotropy, detailed in Table 4.2. The segmentor checks on the 10 kpc dataset exhibited total isotropy with a segmentor coverage of 100%. The 30 kpc data showed lower segmentor coverage of 52%, though this is mostly an effect of the resolution - the passed event count of 74 is less than the resolution of the segmentor, and therefore the number of primary segments hit relative to the passed event count is a better indicator of isotropy in this case, corresponding to approximately 70%. Examination of the PMT hitmaps for both datasets confirms high degrees of isotropy, shown in Figure 4.3. The simulated 100 kpc dataset was also examined, though due to low event counts for this distance it is not possible to draw meaningful conclusions as to the performance of the L3 in this regard.

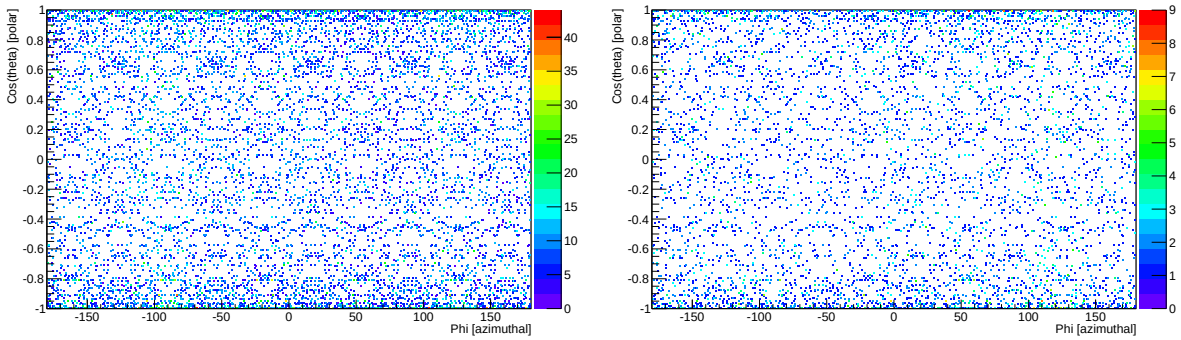



Figure 4.3: L3 PMT hitmap output in $\cos\theta$, $\phi$ for a simulated supernova at 10 kpc (left) and 30 kpc (right) in the water phase. Note the highly isotropic distribution of hit PMTs in both datasets. Whilst the information contained in these plots alone is insufficient for supernova tagging, it permits operators to discriminate common burst background via pattern recognition (see Figure 4.4)

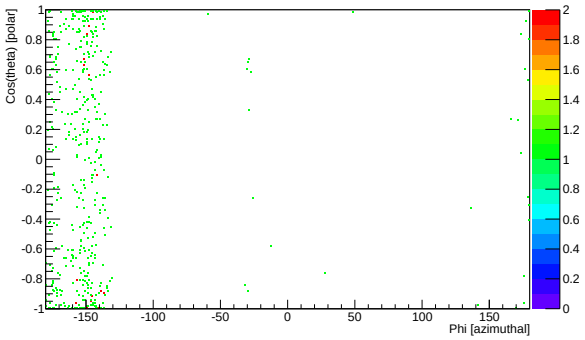



Figure 4.4: L3 PMT hitmap output in $\cos\theta$, $\phi$ for the MDC TSLP crate-by-crate run. Note the highly anisotropic distribution of hit PMTs - TSLP crate-by-crate runs are a useful for the L3 in that they imitate common detector background events such as flashers, which are commonly highly localised.

| Supernova Distance [kpc] | Passed Event Count | Segmentor Coverage [%] |
|---|---|---|
| 10 | 590 | 100 |
| 30 | 74 | 52 |
| 100 | 6 | 6 |

Table 4.2: Table of segmentor performance for simulated data sets at 10 kpc, 30 kpc, and 100 kpc during the water phase. Note the decrease in segmentor performance with event count - as the segmentor coverage is decided by the number of unique primary segments hit, events with low segmentor coverage and low passed event counts may actually demonstrate high degrees of isotropy, as seen in the 30 kpc dataset. This may be mitigated via alteration of the resolution of the segmentor from its default value.

#### 4.4.3.2 MDC Data

The MDC burstfile data displayed a mixture of results - all showed low segmentor coverage, though a number of these runs had a total number of events much lower than the segmentor resolution of 100,

| MDC Run Type | Total Event Count | Segmentor Coverage [%] |
|---|---|---|
| V. High Rate DQ Test | 845 | 5 |
| Low Rate DQ Test | 32 | 14 |
| 2h DQ Detector Run | 32 | 23 |
| TSLP Crate-by-Crate | 12 | 1 |

Table 4.3: Table of segmentor performance for a selection of MDC run types. All MDC burstfile data originates from the partial-fill detector darkruns from December 2014.

as can be seen in Table 4.3. Thus, some run types, particularly general detector runs and variable rate data quality runs, actually display higher isotropy than is evident in the detector coverage percentage.

Of all the MDC data, the most useful for the purposes of testing L3 performance was given by the TSLP crate-by-crate run, which exhibited prominent anisotropic characteristics. As the major backgrounds for supernova detection are detector backgrounds, themselves highly anisotropic, it is significant that the L3 was able to detect this clearly during processing. Figure 4.4 shows the PMT hitmap for the TSLP run, in which the anisotropic nature of the burstfile is evident.

### 4.4.4 Energy Checks

The energy checks conducted by the L3 were perhaps the most important, as they offered a chance to test the concept behind the entire trigger: certain supernova interactions should be readily apparent in quick analyses of supernova data, and therefore employable as the primary means of supernova discrimination.

As shown below, in Figure 4.4.4.1, this is demonstrated to be justified - the 2.2 MeV neutron capture gamma is self-evident in scintillator even in 30 kpc data. This result, combined with the high average energies of events, is a vindication for the use of fast energy reconstruction over burstfile data. Given that 30 kpc is approximately the diameter of the Milky Way galaxy, it is fair to say that SNO+ can be confident of galactic supernova detection in the scintillator phase using this method, provided the L3 is itself operational.

The $^{12}$C de-excitation tag was not shown to be an effective discriminator of supernova bursts in scintillator during testing, though this is primarily due to issues identified within the supernova simulation files themselves, which will be discussed in detail in Section 4.5.

Unfortunately, due to the limitations of the fitters, the efficacy of tagging supernova through energy reconstruction is greatly reduced during the water phase. The reconstruction of events in water is challenging, and exceptionally difficult at lower energies. This is primarily because the only means of triggering the detector during this phase is via Cerenkov radiation, which necessarily

requires high energy charged particles in order to be emitted. The water fitters rely on Nhit counts and reconstructed positions in order to calculate energies, and therefore face a challenge at low Nhit counts which can potentially reconstruct as a wide range of possible energies. This is an artefact of the large poisson fluctuations inherent at low Nhits, which result in a large energy uncertainty on low-Nhit events.

Thus, in effect, the inaccuracy of the water fitters at low energy can be thought of as a binning error: the number of Nhits is a poor measure of a continuous variable when the range of possible Nhit values corresponding to that variable is high. In contrast, due to the much higher Nhit counts of scintillator events, the range of possible Nhit values corresponding to a particular energy value is low, and therefore brute-force Nhit lookup methods are far more effective at reconstructing low energy events.

Ultimately, this means that energy reconstruction is not an efficient means of conclusive supernova detection during the water phase. The major supernova interactions expected during this phase, IBD and proton elastic scattering, or PES, both pose large challenges to the L3 with only light water in the AV. Proton elastic scattering is not visible, as mentioned earlier, whilst IBD can only be detected via it's positron spectrum. The 2.2 MeV neutron capture signal, essential to the L3 during the scintillator phase, is also non-viable as a means of supernova detection during the water phase due to the energy limitations of the water fitters. The water fitters do still reconstruct high energy events effectively however, and therefore offer some sensitivity to the thermal $\nu_e/\bar{\nu}_e$ supernova spectra via the positron spectrum of IBD. This can be seen in Figure 4.8 below.

#### 4.4.4.1 Fitter Performance

In order to test the performance of the fastfitter combination, 25000 randomly distributed IBD events were simulated in scintillator using RAT. The events were Poisson distributed in time with a mean event rate of 10 $s^{-1}$, and the results were fitted with a simple Gaussian. The results of the fit can be seen in Figure 4.4.4.1. The use of the fastfitter reduces the latency of the L3 by a factor of 30, reducing the run time to the order of minutes in the case of 10 kpc scintillator data, with little impact on accuracy. The mean of the gaussian fit of the energy distribution using fastfitter reconstruction occurs at $\mu_{\text{fast}} = 2.172$, whilst the mean of the scintfitter distribution occurs at $\mu_{\text{scint}} = 2.221$. The standard deviation of the fastfitter distribution was $\sigma_{\text{fast}} = 0.099$, whilst the scintfitter distribution yielded $\sigma_{\text{scint}} = 0.095$.

In addition to comparing the performance of the fastfitter, these events were also used to obtain the $\pm 3\sigma$ limits on the IBD-like event counter used by the L3. The L3 therefore counts events as

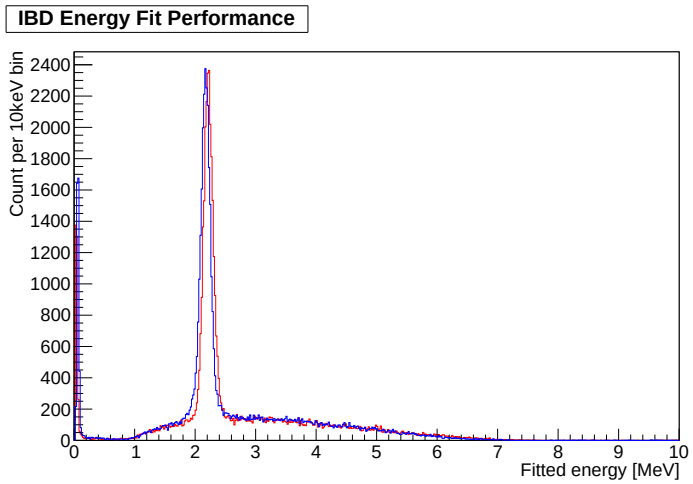



Figure 4.5: Comparison plot of the performance of the scintfitter (red) and the fastfitter (blue) over 25,000 IBD events in scintillator.

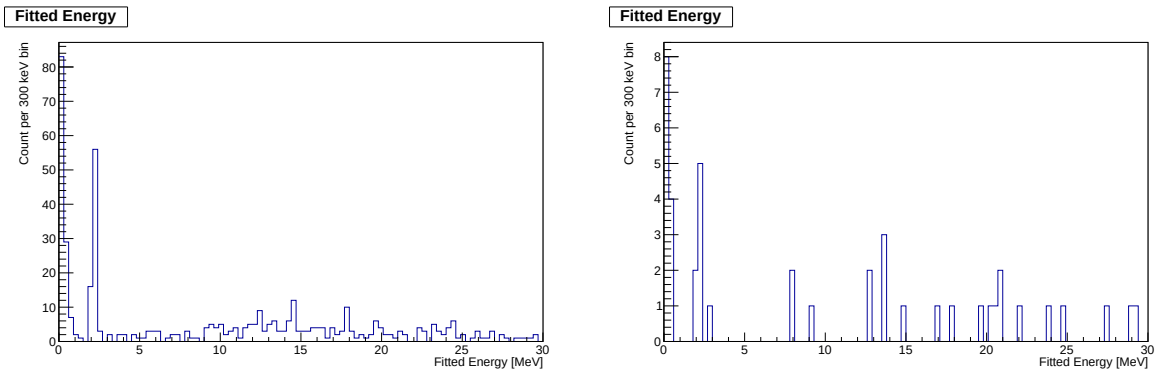



Figure 4.6: Fitted energy distributions for simulated supernova at 10 kpc (left) and 30 kpc (right) in scintillator using the fastfitter combination. Note the clear peak at 2.2 MeV in both datasets, and the high energy neutrino spectrum. Along with the distinctive IBD neutron capture peak at 2.2 MeV, the 10 kpc data shows a low, broad peak at high energy, characteristic of the thermal spectrum of supernova neutrinos (see Chapter 2).

IBD-like in the reconstructed energy range $1.87 < E_{\text{reco}} < 2.47$ MeV.

#### 4.4.4.2 Energy Fit Results

The L3 processor was run over both 10 kpc and 30 kpc simulated supernovae in scintillator in order to examine the reconstructed energy performance and efficiency of the L3 IBD tag. As can be seen in Figure 4.6, the peak from IBD neutron capture on hydrogen at 2.2 MeV is clear in both 10 kpc and 30 kpc data, whilst the high energy neutrino spectrum is also readily apparent at 10 kpc. Due to issues with the simulation, the event rates are reduced from their theoretical values, and the 15.11 MeV carbon de-excitation signal is consequently not visible in the data. The 30 kpc data is greatly

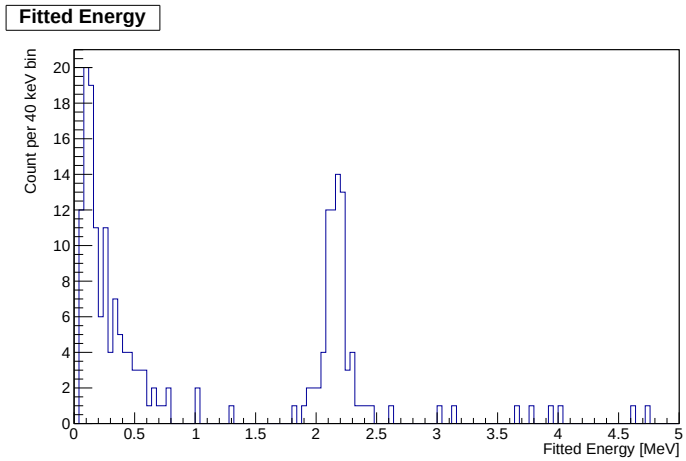



Figure 4.7: Fitted energy distribution for the 10 kpc simulated supernova data in scintillator at energies below 5 MeV. The neutron capture peak is clear here, permitting a relatively simple supernova tag. PES events dominate at energies below ∼1 MeV, though in practice the high background rates at low energy, combined with the detector trigger threshold of 200 keV, will eliminate a large proportion of these events from the L3's analysis.

reduced in resolution due to the reduction in event rates, although similar patterns are evident.

As shown in Table 4.4, the 10 kpc data returned 73 IBD-tagged events for 458 total events passed by the L3, whilst the 30 kpc returned 7 IBD-tagged events for 42 passed events. Initially, the L3 received more than double these events for each data set, though most events were discarded due to failed energy fits and were therefore not available for analysis. This issue will be elaborated on later in the chapter.

Due to the lack of background in these plots, PES events dominate the low energy regime at $< 1$ MeV, but the majority of these events are not expected to be visible in the detector due to a combination of high backgrounds at low energies, and the consequent detector trigger threshold of 200 keV. It is possible that any future L3 could temporarily set lower detector thresholds on burst detection, increasing the sensitivity to PES events, though this would be a complex undertaking and is therefore beyond the scope of this thesis.

The high energy distribution shows a faint peak between 10-15 MeV in the 10 kpc data; as the vast majority of events in this energy regime arise from IBD positron annihilation, we would expect the high energy regime of the reconstructed energy to be closely related to the neutrino fluence spectrum shown in Figure 2.4 in Chapter 2. The mean energy of $\bar{\nu}_e$ emitted from supernovae is expected at $\langle \bar{\nu}_e \rangle \simeq 12$ MeV [3], so for a limited data set to reveal this relationship is encouraging. The 30 kpc data is limited by much lower statistics, and therefore hints at neutrino distributions are not readily apparent, though a large proportion of high energy and IBD-tagged events still ensure

a supernova tag is possible. The L3 was also run over 100 kpc data in both water and scintillator, though the low event rate in both media would be unlikely to trigger the L2 burst monitor, and therefore unlikely to be processed by the L3. The results are included in both Table 4.4 and Table 4.5 for completeness however. The large decrease in passed events relative to detected events can be accounted for by errors in the hepevt conversion method used in scintillator, and is discussed in detail in Section 4.5. For a real supernova, these counts are expected to be almost equivalent, depending on orphan counts and retrigger rates.

| Distance [kpc] | Events Simulated | Events Detected | Events Passed | IBD-Like Event Count |
|---|---|---|---|---|
| 10 | 1290 | 926 | 458 | 73 |
| 30 | 133 | 101 | 42 | 7 |
| 100 | 11 | 9 | 5 | 1 |

Table 4.4: Table of event counts given by the L3 processor for a range of simulated supernova datasets in scintillator. "Events Simulated" shows the number of events in the original hepevt file used to generate the data, "Events Detected" refers to the number of simulated events observed in the detector, whilst "Events Passed" refers to the number of these events that successfully passed all L3 checks.

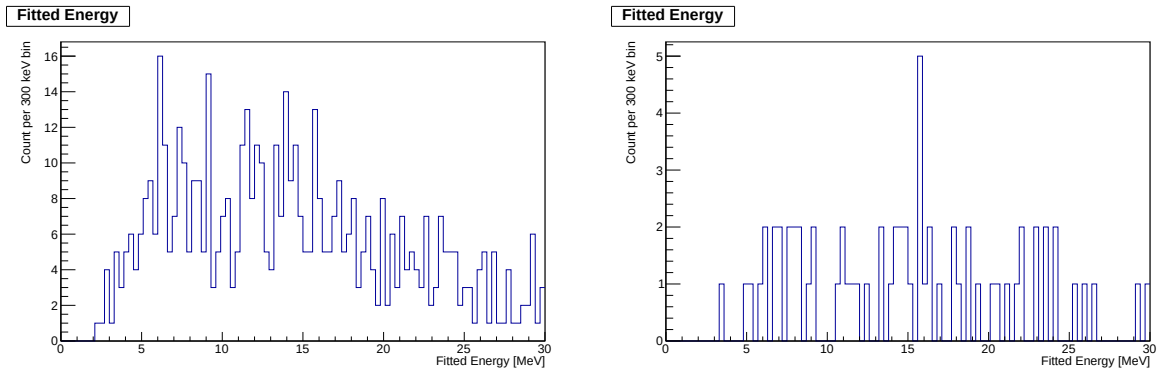



Figure 4.8: Fitted energy distribution for a simulated supernova at 10 kpc (left) and 30 kpc (right) in water using the fastfitter combination. Note the absence of the 2.2 MeV IBD signal in water along with far lower event rates. The high average neutrino energy is still apparent however, and therefore energy fitting is still of use during the water phase.

The L3 processor was also run over simulated supernovae in water at both 10 kpc and 30 kpc, and the results shown in Figure 4.8. As mentioned, these plots offer much lower amounts of information on first glance compared to the scintillator plots, though a high energy neutrino spectrum is still apparent in both datasets. Consequently, the L3 still retains the capability to assess supernova likelihood, though at far lower efficiencies than those of scintillator. An automated SNEWS trigger is therefore unlikely to be implemented during the water phase without considerable improvements to

| Distance [kpc] | Events Simulated | Events Detected | Events Passed | High Energy Event Count |
|---|---|---|---|---|
| 10 | 1290 | 591 | 590 | 556 |
| 30 | 133 | 74 | 74 | 73 |
| 100 | 11 | 6 | 6 | 5 |

Table 4.5: Table of events counts given by the L3 processor for a range of simulated supernova datasets in water. All column titles are equivalent to those given in Table 4.4 above, with the exception of the high energy event count, defined as $E_{\text{reco}} > 5$ MeV. The low orphan counts and retrigger rates in water ensure that the passed event count is almost equivalent to the detected event count. The issue affecting the overall scintillator event counts is not present in water, resulting in a large quantities of detected events.

the IBD tagging methods of the L3, or the introduction of a greatly improved water energy fitter capable of the speeds required by the L3. A full breakdown of event counts in water is given in Table 4.5.

### 4.4.5 L3 Efficiency in Scintillator

The efficiency of the prototype L3 may be assessed by examining the response of the L3 to a range of simulated supernova datasets, whilst employing the same criteria used by the SNEWS trigger framework within the processor. The SNEWS trigger requires the burst to have timing characteristics typical of a supernova burst, and for a high percentage of passed events to exist relative to the total number of events detected in the burst. If these requirements are met, the SNEWS trigger then looks at the number of reconstructed high energy events in the burstfile, defined as $E_{\text{reco}} > 5$ MeV, before checking for the number of IBD-like and $^{12}$C-like tags. If sufficient numbers of events are of high energy, and a distinct peak of IBD-like events is seen, then the trigger issues a SNEWS dummy alert. Due to errors in the Monte Carlo data, $^{12}$C events are currently not visible, and are therefore not employed in this efficiency measure. For real data, the presence of a $^{12}$C de-excitation peak alongside the IBD neutron capture peak should provide confirmation of a supernova burst.

Whilst these criteria are rudimentary, and currently used to simply test the SNEWS trigger framework, the underlying principles allow us to define the expected efficiency of the L3 for supernovae of varying distances by examining the L3's response to decreasing event counts. If the above criteria are all met, and the number of "IBD-like" tagged events is large relative to the supernova-induced background at high energies, then we can be confident of supernova detection. In these conditions, the efficiency of the L3 can be defined as $\simeq$100%. As the event rate in the detector decreases, we would expect the number of tagged IBD-like events to decrease commensurately, before reaching a level where the IBD peak is no longer clearly defined against the high energy background.

At this point, we can no longer reasonably claim to have conclusively observed a supernova burst without further analysis. Whilst it may be possible for a future L3 to use high energy events alone to tag supernova, the current iteration of the L3 uses the conservative criteria that an IBD tag must be present before a SNEWS trigger may be sent.

Consequently, the efficiency of the L3, $\eta_{\mathrm{L3}}$, must then scale proportionally with the total event count passed by the L3, which will itself be proportionate to the inverse square of the distance to the observed supernova $d_{\mathrm{SN}}$, i.e.

$$\eta_{\mathrm{L3}} \propto n_{\mathrm{passed}} \propto \frac{1}{d_{\mathrm{SN}}^2} \tag{4.1}$$

This relationship allows a simple measure of the efficiency of the L3 to be determined. In the limit where $d_{\mathrm{SN}} \to \infty$, we will reach a point where the IBD event count will be too low to be statistically meaningful, and the L3 will no longer be effective in tagging supernova. By combining the above, we arrive at the following relationships:

$$\begin{aligned} \eta_{\mathrm{L3}} &\to 1 \quad as \quad d_{\mathrm{SN}} \to 0 \\ \eta_{\mathrm{L3}} &\to 0 \quad as \quad d_{\mathrm{SN}} \to \infty \end{aligned}$$

#### 4.4.5.1 Poisson Estimation of L3 Efficiency

As a first order approximation, we may begin by expecting the total number of tagged IBD-like events to be equal to the sum of the IBD-like signal and background events, $n_{\mathrm{passed}}^{\mathrm{IBD}} = (S_{\mathrm{IBD}} + B_{\mathrm{IBD}})$ where $S_{\mathrm{IBD}}$ denotes IBD-like signal events, and $B_{\mathrm{IBD}}$ denotes IBD-like background events.

In order to define our lower event threshold, $n_{\mathrm{low}}^{\mathrm{IBD}}$, we must first examine the employment of the IBD-like tag. In the circumstances that the IBD-like tag alone is being used to tag a supernova burst and send a SNEWS trigger from SNO+, then our limit should be stringent. In the condition that the IBD-like tag is part of a number of tags used to define a supernova, then our limit may be more relaxed and our IBD sensitivity consequently improved. The latter is the planned scenario for the L3 at SNO+, however, as these additional tags are not examined in this thesis, for the following calculation we will assume that the L3 operates using the IBD-like tag alone.

The rate of background events to IBD-like events at 2.2 MeV include major contributions from supernova-induced interactions[11], predominantly $\nu - p$ elastic scattering and the inclusive $^{12}$C trans-

[11] It should be noted that for the purposes of this efficiency calculation, supernova induced processes other than IBD are considered as background although under normal circumstances they would be included as signal events. This is in order to estimate the ability of the L3 prototype to detect supernova *in its current form*, which operates using only the IBD tag. Thus, this efficiency calculation may be thought of as the efficiency of the L3 in tagging IBD events in the detector, which may then be converted to an estimate of supernova tag efficiency using Equation 4.1 and assuming the only process detected is that of IBD.

itions detailed in Chapter 2. In addition, during the loaded-scintillator phase there will be a contribution from double beta decay neutrinos, due to the $2\nu\beta\beta$ energy spectrum endpoint at 2.5 MeV. Finally, solar neutrinos from the decay of $^8$B and intrinsic detector radioactivity backgrounds also contribute, though their rates are negligible compared to other backgrounds. Thus, the total IBD-like background rate $B_{\text{tot}}^{\text{IBD}}$ can then be estimated by considering the contributions from each of the major backgrounds[12]:

$$\begin{aligned} B_{\text{tot}}^{\text{IBD}} &= B_{\nu-p}^{\text{IBD}} + B_{12C}^{\text{IBD}} + B_{2\nu\beta\beta}^{\text{IBD}} \\ &= (0.2 \pm 0.05)\ \text{kt}^{-1}\text{s}^{-1} + (0.1 \pm 0.05)\ \text{kt}^{-1}\text{s}^{-1} + (0.15 \pm 0.1)\ \text{kt}^{-1}\text{s}^{-1} \\ &= 0.55 \pm 0.12\,\text{kt}^{-1}\text{s}^{-1} \end{aligned}$$

Considering an average supernova burst length of ∼10s, and using the highly conservative assumption that supernova-induced backgrounds occur at a constant rate over the course of the burst[13], we may then determine our threshold $n_{\text{low}}^{\text{IBD}}$. Assuming burst triggering of the L2 10 times per day, an average burst length of 10 seconds, and an acceptable SNEWS false alert rate of 1 per 3 months as used in Daya Bay[14] [86], this yields a required limit of less than 1 false positive IBD-like tag in ∼913 bursts. This equates to the probability of any L3 IBD tagged burst containing a supernova as 99.890%.

Assuming the events follow a Poisson distribution, $n_{\text{low}}^{\text{IBD}}$ may then be found by numerically solving the following integral for $N_{99.89\%}$:

$$\int_0^{N_{99.89\%}} A \times \frac{(S_{\text{IBD}} + B_{\text{IBD}})^n e^{-(S_{\text{IBD}}+B_{\text{IBD}})}}{n!} = 0.9989 \tag{4.2}$$

Where A is some normalisation constant and n is the number of detected events. Rounding $N_{99.89\%}$ to the closest integer, we obtain $n_{\text{low}}^{\text{IBD}} = 14$ events.

[12]Background rates were calculated for the IBD energy ROI given in Section 4.4.4.1 using known $2\nu\beta\beta$ neutrino spectra at 0.3% loading. Known $\nu - p$ / $^{12}$C rates for the SNO+ detector are given in Figure 2.6 from [40], though these will fluctuate with supernova distance - the given rates refer to a supernova at 10 kpc. Instrumental backgrounds are also expected to contribute heavily to the overall L3 background rate, though an estimate of their rate is not possible until detector commissioning, and therefore is unused in this thesis.

[13]The vast majority of supernova events are projected to occur within the first few seconds of the burst - see Figure 2.3.

[14]The burst trigger rate of the L2 is difficult to ascertain before detector commissioning and without knowledge of true instrumental background rates, but given the high initial rate requirements this ought to be low. The figure of 10 bursts per day should thus be taken as an order of magnitude estimate. The SNEWS system uses a coincidence requirement on all supernova alerts - no single detector can trigger the system without confirmation from at least one additional detector. Thus, the acceptable false alert rate from each detector may be set lower to ensure higher sensitivity without compromising the integrity of the system.

We may then calculate the efficiency of the IBD-like tag by defining the efficiency of the tag, $\eta_{\mathrm{IBD}}$, as:

$$\eta_{\mathrm{IBD}} = 1 - P(n < n_{\mathrm{low}}^{\mathrm{IBD}}) \tag{4.3}$$

We may then assume that the efficiency of the IBD-like tag will follow a Poisson distribution as follows:

$$\begin{aligned} \eta_{\mathrm{IBD}} &= 1 - Poi(\lambda = S + B,\ n < n_{\mathrm{low}}^{\mathrm{IBD}}) \\ &= 1 - \sum_{n=0}^{n_{\mathrm{low}}^{\mathrm{IBD}}} Poi(\lambda = S + B,\ n) \\ &= 1 - (e^{-(S+B)} \sum_{n=0}^{n_{\mathrm{low}}^{\mathrm{IBD}}} \frac{(S+B)^n}{n!}) \end{aligned} \tag{4.4}$$

By considering this distribution over a range of passed events we may then plot the efficiency of the L3 in terms of IBD tagged events, given in Figure 4.9. It should be noted that the above calculations do not take instrumental background into account - instrumental background is likely to be the largest contributor to supernova backgrounds at SNO+, though an accurate estimate is difficult at present. Instrumental backgrounds such as flashers typically imitate high energy events, so their impact at 2.2 MeV is expected to be low however. The above calculation also assumes that the L3 will be run over burstfiles with no lower event count limit, though in practice this will be defined by the L2 burst parameters. Consequently, large uncertainties exist for the L3 efficiency at low IBD-like event counts.

The number of IBD tagged events may be converted to the number of passed events expected for a given supernova distance by considering the mean number of IBD events tagged by the L3 over the three simulated data sets, $\left\langle R_{\mathrm{IBD}}^{\mathrm{tag}} \right\rangle$, given by:

$$\left\langle R_{\mathrm{IBD}}^{\mathrm{tag}} \right\rangle = \frac{1}{3}\left( \frac{n_{\mathrm{IBD}}^{10\mathrm{kpc}}}{n_{\mathrm{pass}}^{10\mathrm{kpc}}} + \frac{n_{\mathrm{IBD}}^{30\mathrm{kpc}}}{n_{\mathrm{pass}}^{30\mathrm{kpc}}} + \frac{n_{\mathrm{IBD}}^{100\mathrm{kpc}}}{n_{\mathrm{pass}}^{100\mathrm{kpc}}} \right) = 0.175 \tag{4.5}$$

Hence, assuming that the relationship between passed events, $n_{\mathrm{pass}}$, and IBD tagged events $n_{\mathrm{tag}}^{\mathrm{IBD}}$ is linear, over the full ten second burst we would expect:

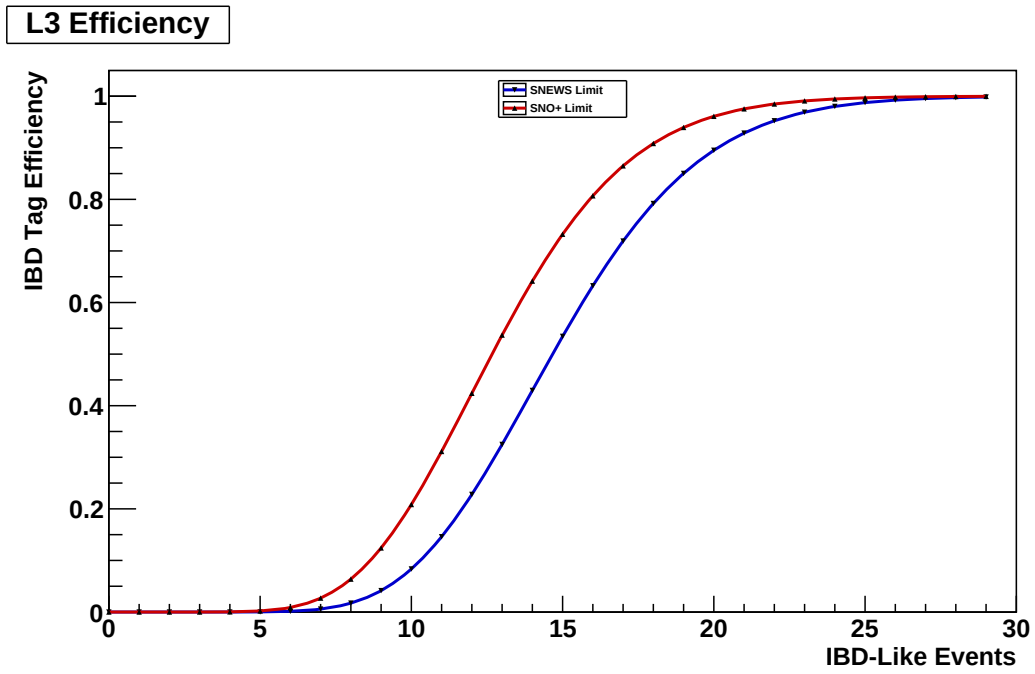



Figure 4.9: Plot of the L3 IBD tag efficiency with increasing IBD-like event counts $n^{\text{IBD}}_{\text{passed}} = n^{\text{IBD}}_{\text{tag}} = (S_{\text{IBD}} + B_{\text{IBD}})$, using Equation 4.4 above. Two limits are given for comparative purposes: the SNEWS limit of 1 false alert per 3 months, calculated using Equation 4.2, and an additional higher sensitivity limit of 1 false alert per week for use within SNO+.

$$n_{\text{pass}} = \frac{n^{\text{IBD}}_{\text{tag}}}{\left\langle R^{\text{tag}}_{\text{IBD}} \right\rangle} \tag{4.6}$$

Thus, using the relation given in Equation 4.1, and by using Equation 4.6 above, we may estimate the supernova distance in kiloparsecs $d_{\text{SN}}$. Using the number of passed events, $n_{\text{pass}}$, the simulated passed event count at 10 kpc[15] $n^{\text{10kpc}}_{\text{pass}}$ , and extrapolating from Table 4.4, we can approximate $d_{\text{SN}}$ as follows:

$$\begin{aligned} d_{\text{SN}} &= \sqrt{\frac{n^{\text{10kpc}}_{\text{pass}}}{n_{\text{pass}}}} \times 10\,\text{kpc} \\ &= \sqrt{\frac{n^{\text{10kpc}}_{\text{pass}} \left\langle R^{\text{tag}}_{\text{IBD}} \right\rangle}{n^{\text{IBD}}_{\text{tag}}}} \times 10\,\text{kpc} \end{aligned} \tag{4.7}$$

From Table 4.4, we see that the number of passed events at 10 kpc in scintillator using the current prototype of the L3 is $n^{\text{10kpc}}_{\text{pass}} = 458$. Using this as a measure of the current L3 efficiency, we

[15]This figure is obtained from Table 4.4. These counts are lower than expected due to errors in the Monte Carlo simulations. Consequently, the projected event count under real operating conditions is expected to be larger, with a consequently higher efficiency.

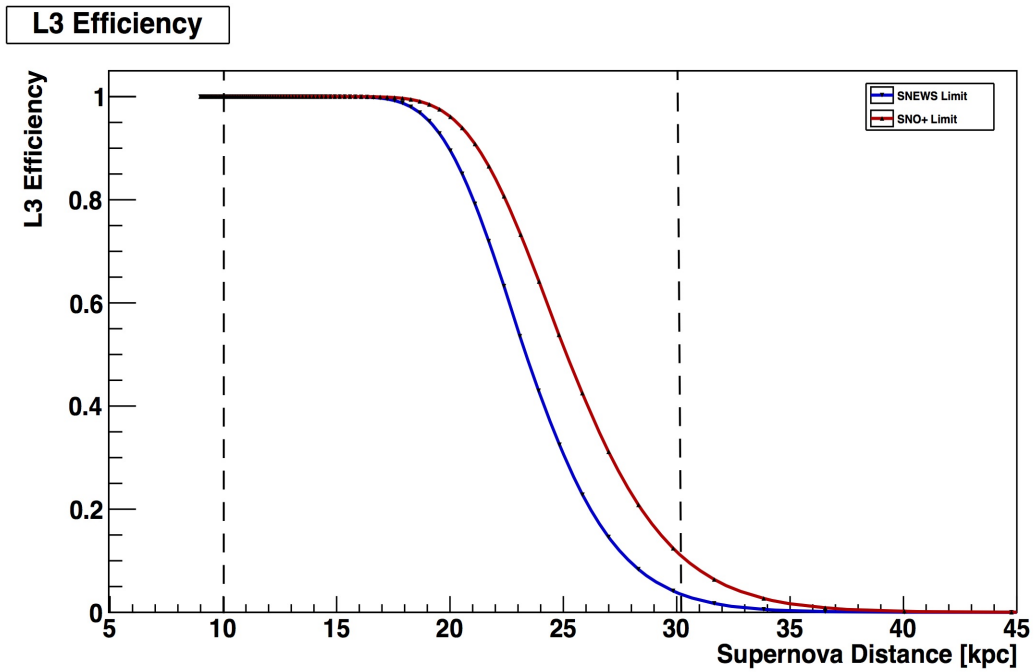



Figure 4.10: Plot of the overall L3 efficiency in scintillator with increasing supernova distance, using the passed event counts from the 10 kpc simulated supernova dataset, Equations 4.4 and 4.7, and assuming no detector trigger threshold.

may then plot the L3 efficiency $\eta_{\mathrm{L3}}$ against supernova distance $d_{\mathrm{SN}}$ using Equation 4.4 and Equation 4.7 above. This is shown in Figure 4.10. As with Figure 4.9, two limits are given for comparative purposes. For reasons discussed in Figure 4.9, the uncertainty on the L3 tagging efficiency increases substantially at lower event counts, and therefore increases commensurately for greater supernova distances. It is also expected that the efficiency of the L3 will be higher than the above in practice - this is due to higher rates of IBD than those seen in the Monte Carlo data, explained below, combined with the anticipated improvement from any applied $^{12}$C de-excitation tag. Similarly, this efficiency calculation catagorises all processes other than IBD events in the ROI as background, including supernova-induced processes. In reality, these events will be counted as signal events, increasing the L3 efficiency commensurately. Finally, it is expected that future iterations of the L3 will utilise the high energy event count to create an additional tag, further increasing sensitivity. For reference purposes, dashed lines show the distance of a "standard" galactic supernova at 10 kpc and the diameter of the Milky Way galaxy at $\sim 31$ kpc. The distance to SN1987A was $\sim 50$ kpc, whilst the Andromeda galaxy lies at $778 \pm 33$ kpc.

## 4.5 Issues

As mentioned above, a simple error was present in the simulation data for scintillator which is responsible for a number of problems seen in the above results. The error originated in the method of data generation: full supernova simulations are not currently supported in RAT, and so in order to generate them effectively other sources of data must be used. In this case, the source was a hepevt file, originally created by the SNGEN software [84], which listed all interaction information for a 10 kpc supernova at SNO+ as pure data: each interaction is assigned a daughter particle type(s), specific 4-momenta, and position. This was read into and converted by RAT, which then simulated the response of the detector to these events.

Unfortunately, due to an error within the supplied conversion script or original generating program SNGEN, the detector medium within the AV was taken to exist across the entire PSUP volume. Consequently, when these events were simulated within RAT, events which would only be possible on scintillator within the AV were also occurring in the water between the AV and PSUP. The SNO+ position and energy fitters only expect to reconstruct events within the AV volume for scintillator, so this bug has the net effect of reducing the number of events observed within the detector in relation to the number expected from theory, with the remainder classified as invalid fits. For the 10 kpc supernova simulation, 1290 events were calculated to occur within the detector by SNGEN with a range of interaction types, and were duly simulated in turn by RAT. Due to the error, only 416 events were reconstructed in scintillator from the total of 1290, a direct consequence of the ratio of the AV volume to that of the PSUP: $\frac{R^3_{\mathrm{AV}}}{R^3_{\mathrm{PSUP}}} \sim 0.3$.

The events within the original data source were calculated specifically to yield only events observable by the SNO+ detector for a supernova of a given distance. This means that of the finite number generated, all should be visible events within the detector. Due to the above error, a large fraction of these events occurred in RAT well outside their calculated interaction positions, scaling with $R^3$, and therefore resulted in a sizeable decrease in the total number of events observed and a large increase in the number of failed position and energy fits. Carbon de-excitation events were particularly affected; only observable within the scintillator volume, and predominantly taking place in the outer volume of the AV, the shifting of the 15.1 MeV $\gamma$ event positions into the water renders them invisible to the energy fitters, and therefore unusable by the L3. As only a modest number of these events are expected at SNO+ for an ideal supernova, this error has a particularly pronounced effect in the results, and is responsible for greatly reducing the efficacy of the carbon de-excitation

tag.

As IBD events are tagged within the L3 through the gamma produced via delayed neutron capture on hydrogen, the position shift of IBD events to water also has the effect of removing these events from the L3's grasp. 2.2 MeV events positioned outside the AV in water will likely fail to trigger the detector, whilst any successful reconstruction would likely be unrecognisable as an IBD event due to the minimal energy signature in water compared to scintillator. Fortunately, the effect of this bug on the isotropy and timing characteristics of supernova are negligible, and still permit accurate simulation for L3 development and testing purposes.

The issue most strongly affects the performance of the L3 trigger in scintillator, whilst in water it has little effect. This is due to the fact that during the water phase, the assumption of the detector medium extending through the PSUP volume is coincidentally correct. In scintillator however, event tagging is instrumental in the success of the L3 in its goal of supernova monitoring, and hence the error plays a major role in reducing the apparent efficiency of the L3 over supernova simulations in this thesis.

# Chapter 5

# Conclusion and Outlook

Early results from L3 prototype testing are promising, if still somewhat rudimentary. As a proof of concept, the L3 prototype is successful: it is capable of spotting supernova-specific neutrino interactions quickly and provides the basic structure necessary to provide SNEWS alerts at a later date. However, a lack of sophistication in its current iteration ensures it is not capable of such a leap before further development and optimisation, particularly during the water phase of SNO+. Results have also shown that the potential of the L3 is at its greatest during the scintillator phase, although future development could further increase its accuracy during the water phase. Inclusion of an effective tag for IBD positrons, an increase in the precision of the current isotropy check, and the inclusion of more sophisticated water-phase energy fitters could all potentially allow the L3 to be effective over a greater range of supernova distances in water. Similarly, whilst it is unlikely that the L3 will be capable of reliable SNEWS triggers during the water phase of SNO+ without substantial work, further development should still ensure that the L3 is capable of providing accurate burst parameters to the collaboration along with any supernova tag. Given that water phase running is anticipated to be far shorter than scintillator phase running, the lower efficacy of the L3 during this time is perhaps less problematic than it might be otherwise.

Consequently, whilst undoubtedly a crude tool in its current form, the prototype L3 processor and trigger has largely achieved the aims discussed earlier in this thesis, particularly for its use during the scintillator phase of detector running. By focusing on the core signatures of the burst, the L3 tags supernovae through simple analysis rather than burst detection, thereby ensuring that the L3 is both resilient to detector backgrounds and reliable in the tag itself. In providing a simple means of galactic supernova detection, it outlines the direction of travel for future near-line super-

nova monitoring at SNO+, and allows subsequent work to focus on refining and augmenting the trigger to allow greater sensitivity and confidence of detection. Whilst the L3 itself has yet to be subject to the demands of real near-line running, preliminary results from the prototype, run over local data, have shown that a small number of key parameter checks, administered quickly, can be considered an effective means of supernova tagging at SNO+.

The future work of the supernova group at SNO+ is then clearly defined: to further refine and develop the L3 into a trigger system capable of sending a SNEWS alert reliably, with minimal operator oversight. By honing the system to this point, the gains for the collaboration as a whole also increase – an L3 capable of sending a SNEWS signal with reliability must also be capable of determining supernova burst parameters with precision, therefore allowing the SNO+ collaboration access to the best possible statistics from the detector for further analysis. In addition, this L3 will offer operators indications as to potential areas of value within the data. For instance, an unusually high number of carbon de-excitation events would imply large numbers of neutral-current events occurring within the detector, and therefore would have strong implications for the strength of the $\nu_x$ spectrum [4]. In turn, this could then yield valuable insights into matter effects within the progenitor star, the formation of the neutrinospheres in the core, or even the nature of the neutrino-driven explosion itself.

An effective supernova monitoring infrastructure is vital if SNO+ is to achieve its physics aims in this area and contribute to the global supernova physics community. The potential physics yield of a single successful observation is so great that the implementation of viable monitoring tools should be considered a priority within the astrophysical sector of SNO+ research. Should SNO+ be successful in detecting a high-statistics galactic supernova neutrino signal, the sensitivity of the detection medium offers the ability to probe fundamental questions in neutrino physics, such as the sign of $\Delta m_{23}$, and therefore the nature of the neutrino mass ordering [55][78]. In turn, this has ramifications for SNO+'s primary physics goals: SNO+ is only sensitive to neutrinoless double beta decay should the mass hierarchy be inverted, even with high $^{130}$Te loading levels [69]. Additionally, successful observation of a high statistics supernova signal at SNO+ would offer the opportunity to study any divergences between current supernova models and the real physical signal, potentially resolving the mechanism driving the final stages of explosion. An efficient, sophisticated monitoring system is therefore key to maximising the supernova physics capabilities of the detector.

In this regard, the importance of the L3 is twofold: it ensures SNO+ is in the best position to record any supernova signal effectively, whilst also being responsible for the issuing of any global SNEWS alert from the SNO+ collaboration. The efficiency of the SNEWS system scales with

the number of neutrino detectors capable of sending a reliable alert, and therefore the inclusion of SNO+ provides valuable redundancy in the network. The impact of false alarms, routine detector downtime, detector calibration, and detector breakdowns are all heavily minimised with the inclusion of additional global detectors, and therefore the value of the system to the astrophysical community scales commensurately. In turn, the speed of the SNEWS alert dictates the response time of the astronomical community's measurement efforts, and therefore also maximises the global scientific gain from any single event. Given the nature of supernovae as a truly interdisciplinary research field, the potential benefit to astrophysics, particle physics, nuclear physics, and astronomy is substantial.

Ultimately, the prospects for supernova monitoring at SNO+ are greatly encouraging. The prototype L3 processor, whilst simple, has provided a direction of travel that can be pursued with potentially huge rewards should a galactic supernova occur within the lifetime of the detector.

# Bibliography


[1] M.A. Schumaker. Supernova detection with SNO+. *Nuc. Phys. B Proc. Supp.*, 00, 2010.

[2] A. Mezzacappa and O.E.B. Messer. Neutrino transport in core-collapse supernovae. *J. Comp. App. Math.*, 109:281–319, 1999.

[3] S. Woosley and T. Janka. The physics of core-collapse supernovae. *Nature Phys.*, 1:147–154, 2005.

[4] M.C. Chen L. Cadonati, F.P. Calaprice. Supernova neutrino detection in Borexino. *Astropart. Phys.*, 16:361–372, 2002.

[5] A. Mirizzi et al. Supernova neutrinos: Productions, oscillations and detection. *Il Nuovo Cimento*, 2015.

[6] E.N. Alexeyev and L.N. Alexeyeva. Twenty years of galactic observations in searching for bursts of collapse neutrinos with the Baksan underground scintillation telescope. *Jour. Exp. Theor. Phys*, 95, 2002.

[7] W. Pauli. "Dear Radioactive Ladies and Gentlemen". Letter to Tubingen, December 1930.

[8] F. L. Wilson. Fermi's theory of beta decay. *American Journal of Physics*, 36:12:1150–1160, 1968.

[9] H. Bethe and R. Peierls. The "Neutrino". *Nature*, page 532, 1934.

[10] Jr et al. C.L. Cowan. Detection of the free neutrino: a confirmation. *Science*, 124: 3212:103–124, 1956.

[11] G. Danby et al. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9:1:36–44, 1962.

[12] ALEPH Collaboration. A precise determination of the number of families with light neutrinos and of the Z boson partial widths. *Phys. Lett. B*, 235: 3, 4:399–411, 1990.

[13] F.E. Wietfeldt and E.B. Norman. The 17 kev neutrino. *Phys. Rep.*, 273:3:149–197, 1996.

[14] D. S. Harmer R. Davis and K. C. Hoffman. Search for neutrinos from the sun. *Phys. Rev. Lett.*, 20:1205–1209, 1968.

[15] R. Davis. A review of the Homestake solar neutrino experiment. *Prog. Part. Nucl. Phys.*, 32:13–32, 1994.

[16] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562, 1998.

[17] J. N. Bacall and R. Davis Jr. An account of the development of the solar neutrino problem. *Essays in Nuclear Astrophysics*, pages 243–285, 1982.

[18] Q. R. Ahmad et al. Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:1:011301, 2002.

[19] Y. Abe et al. Indication for the disappearance of reactor electron antineutrinos in the Double Chooz experiment. *Phys. Rev. Lett.*, 108:131801, 2012.

[20] P. Adamson et al. Improved search for muon-neutrino to electron neutrino oscillations in MINOS. *Phys. Rev. Lett.*, 107:181892, 2011.

[21] K. Abe et al. Indication of electron neutrino appearance from an accelerator-produced off-axis muon neutrino beam. *Phys. Rev. Lett.*, 107:041801, 2011.

[22] S. Kim et al. Observation of reactor electron antineutrino disappearance in the RENO experiment. *Phys. Rev. Lett.*, 108:191802, 2012.

[23] F. P. An et al. Observation of electron-antineutrino disappearance at Daya Bay. *Phys. Rev. Lett.*, 108:171803, 2012.

[24] T. Brahe. De nova et nullis aevi memoria prius visa stella, 1572.

[25] W. Huggins. On a new star. *Monthly Notices of the Royal Astronomical Society*, 26:275, 1866.

[26] W. Baade and F. Zwicky. On super-novae. *Proceedings of the National Academy of Sciences*, 20: 5:254–259, 1934.

[27] K. Hirata et al. Observation of a neutrino burst from the supernova SN1987A. *Phys. Rev. Lett*, 58:1490–1493, 1987.

[28] E.N.Alexeyev et al. Detection of the neutrino signal from SN1987A in the LMC using the INR Baksan underground scintillation telescope. *Phys. Lett. B*, 205:209–214, 1988.

[29] C. B. Bratton et al. Neutrinos from SN1987A in the IMB detector. *Nuc. Instrum. Methods Phys. Res., Sect A.*, 264:28–31, 1988.

[30] P. J. Kernan and L.M Krauss. Yet another paper on SN1987A: Large angle oscillations, and the electron neutrino mass. *arXiv*, 9410010v1[astro-ph], 1994.

[31] W. Hillebrandt and J.C. Niemeyer. Type Ia supernova explosion models. *Ann. Rev. Astron. Astrophys.*, 38:191–230, 2000.

[32] A. Heger et al. How massive single stars end their life. *Astrophys. J.*, 591:288–300, 2003.

[33] E. Cappellaro and M. Turatto. Supernova types and rates. In D. Vanbeveren, editor, *The Influence of Binaries on Stellar Population Studies*, volume 264, page 199. Brussels, 2001.

[34] P. J. Kernan and L.M Krauss. The type IIb supernova 2013df and it's cool supergiant progenitor. *Astron. Journ.*, 147:37, 2014.

[35] A.C. Phillips. *The Physics of Stars.* John Wiley & Sons, New York, 1999.

[36] J.F. Beacom, R.N. Boyd, and A. Mezzacappa. Black hole formation in core-collapse supernovae and time-of-flight measurements of the neutrino masses. *Phys. Rev. D*, 63: 073011, 2001.

[37] H. Janka et al. Theory of core collapse supernovae. *Phys. Rept.*, 442:38–74, 2007.

[38] L. Hüdepohl et al. Neutrino signal of electron-capture supernovae from core collapse to cooling. *Phys. Rev. Lett*, 104: 251101, 2010.

[39] H. Janka. Conditions for shock revival by neutrino heating in core-collapse supernovae. *Astron. Astrophys.*, 368:527–560, 2001.

[40] B. von Krosigk. *Measurement of proton and $\alpha$-particle quenching in LAB based scintillators and determination of spectral sensitivities to supernova neutrinos in the SNO+ detector.* PhD thesis, Technical University of Dresden, 2015.

[41] A. Burrows and J. Goshy. A theory of supernova explosions. *Astrophys. J.*, 416:L75–L78, 1993.

[42] G.G. Raffelt. Muon-neutrino and tau-neutrino spectra formation in supernovae. *Astrophys. J.*, 561:890–914, 2001.

[43] R. Buras et al. Two-dimensional hydrodynamic core-collapse supernova simulations with spectral neutrino transport. *Astron. Astrophys.*, 447:1049–1092, 2006.

[44] H. Janka et al. Core-collapse supernovae: Reflections and directions. *Prog. Theor. Exp. Phys.*, 01A309, 2012.

[45] J. Guilet R. Kazeroni and T. Foglizzo. New insights on the spin-up of a neutron star during core collapse. *MNRAS*, 456:126–135, 2016.

[46] J. M. Blondin and S. Shaw. Linear growth of spiral SASI modes in core-collapse supernovae. *Astrophys. J*, 656:366–371, 2006.

[47] S. Furusawa et al. The influence of inelastic neutrino reactions with light nuclei on the standing accretion shock instability in core-collapse supernovae. *Astrophys. J.*, 774:78, 2013.

[48] E. Müller A. Wongwathanarat and H. Janka. Three-dimensional simulations of core-collapse supernovae: From shock revival to shock breakout. *Astron. Astrophys.*, 577, 2015.

[49] B. Kayser. Neutrino oscillation physics. In G. Bellini and L. Ludhova, editors, *Proceedings of the International School on Astroparticle Physics*. IOS Press, Amsterdam, 2012.

[50] C. Giunti and C. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, Oxford, UK, 2007.

[51] E. Majorana. A symmetric theory of electrons and positrons. *Il Nuovo Cimento*, 14:171–184, 1937.

[52] L. Wolfenstein. Neutrino oscillations in matter. *Phys. Rev. D*, 17: 2369, 1978.

[53] S.P. Mikheev et al. Resonance amplification of oscillations in matter and spectroscopy of solar neutrinos. *Sov. J. Nucl. Phys.*, 42:913–917, 1985.

[54] S. Boyd. Neutrino Oscillations. Graduate Lecture Series, Warwick University, 2015.

[55] S. Dighe et al. Identifying the neutrino mass spectrum from the neutrino burst from a supernova. *Phys. Rev. D*, 033007, 2000.

[56] Boris Kayser. Neutrino mass, mixing, and flavor change. *arXiv*, [hep-ph]:0804.1497.

[57] K. N. Abazajian et al. Light sterile neutrinos: A white paper. *arXiv*, [hep-ph]:1204.5379.

[58] F.P. An et al. Search for a light sterile neutrino at Daya Bay. *Phys. Rev. Lett.*, 113: 141802, 2014.

[59] G. Mention et al. The reactor neutrino anomaly. *J. Phys.: Conf. Ser*, 408: 012025, 2013.

[60] A. Malchiorri E. Di Valentino and O. Mena. Dark radiation sterile neutrino candidates after Planck data. *JCAP*, 11:18, 2013.

[61] Hyper Kamiokande Working Group. The Hyper Kamiokande experiment.

[62] W.D. Arnett et al. Supernova 1987A. *Ann. Rev. Astron. Astrophys.*, 27:629–700, 1989.

[63] W.M. Farr J. F. Beacom and P. Vogel. Detection of supernova neutrinos by neutrino-proton elastic scattering. *Phys. Rev. D*, 66: 033001, 2002.

[64] F.P. An et al. New measurement of antineutrino oscillation with the full detector configuration at Daya Bay. *Phys. Rev. Lett.*, 115: 111802, 2015.

[65] S. Abe et al. Measurement of the $^{8}$b solar neutrino flux with the KamLAND liquid scintillator detector. *Phys. Rev. C*, 84: 035804, 2011.

[66] T. Yoshida et al. Neutrino-nucleus reaction cross sections for light element synthesis in supernova explosions. *Astrophys. Journ.*, 686:448–466, 2008.

[67] S. Perasso et al. Measurement of ortho-positronium properties in liquid scintillators. *e-print*, 1306.6001, 2013.

[68] J. Boger et al. The Sudbury Neutrino Observatory. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 449, 2000.

[69] S. Andringa et al. Current status and future prospects of the SNO+ experiment. *Adv. High En. Phys.*, 2015.

[70] N. Fatemighomi. Cover gas commissioning report. *SNO+ DocDB*, 2918:[Version 1], 2015.

[71] K. Majumdar. *On the Measurement of Optical Scattering and Studies of Background Rejection in the SNO+ detector*. PhD thesis, University of Oxford, 2015.

[72] L. Segui. Scintillator purification and development in the SNO+ experiment. (Presentation at the IoP HEPP & APP Conference 2015).

[73] G. Doucas et al. Light concentrators for the Sudbury Neutrino Observatory. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 370:579–596, 1996.

[74] A. Mastbaum. Electronics overview and status. *SNO+ DocDB*, 1339:[Version 1], 2012.

[75] S. Biller. SNO+ with tellurium. *Phys. Procedia*, 00:1–5, 2014.

[76] I. T. Coulter. *Modelling and reconstruction of events in SNO+ related to future searches for lepton and baryon number violation.* PhD thesis, University of Oxford, 2013.

[77] P. G. Jones. *Background Rejection for the Neutrinoless Double Beta Decay Experiment SNO+.* PhD thesis, 2011.

[78] P.D. Serpico et al. Probing the neutrino mass hierarchy with the rise time of a supernova burst. *Phys. Rev. D*, 85: 085031, 2012.

[79] K. Labe. *L2 User Guide.* SNO+ Collaboration, 2014.

[80] S. Ando and K. Sato. Determining the supernova direction by it's neutrinos. *Progress of Theoretical Physics*, 107:5:107–214, 2002.

[81] J. Dunger. Segmentor. *SNO+ DocDB*, 2783:[Version 4], 2014.

[82] I. T. Coulter. Quadfitter. *SNO+ DocDB*, 1031:[Version 1], 2011.

[83] K. Majumdar. Functional form energy reconstruction. *SNO+ DocDB*, 2782:[Version 1], 2014.

[84] S.P. Korte. Simulating supernova physics at SNO+. Master's thesis, 2009.

[85] J.R. Wilson et al. Mock data challenge - post mortem. *SNO+ DocDB*, 3436:[Version 1], 2015.

[86] H. Wei et al. Design, characterization, and sensitivity of the supernova trigger system at Daya Bay. *Astropart. Phys.*, 75:38–43, 2016.

# List of Figures


1.1 Summary of the individual Homestake Experiment runs, 1970-1979 . . . . . 5
1.2 Results from the Sudbury Neutrino Observatory . . . . . 6
1.3 Energy of all detected events from supernova SN1987A . . . . . 8
1.4 Hubble Space Telescope image of the SN1987 remnant in the Large Magellanic Cloud. 9

2.1 Minkowski-Zwicky supernova classification scheme . . . . . 12
2.2 Shell structure schematic of a Type II CCSN just prior to collapse . . . . . 13
2.3 Sample time evolution of the neutrino luminosities by neutrino flavour after core collapse 17
2.4 Sample neutrino fluence for a 10 kpc supernova . . . . . 30
2.5 Visible IBD spectrum in LAB assuming fluence of Figure 2.4 . . . . . 32
2.6 Summary of visible energy spectra for all SN interaction channels at SNO+, assuming a neutrino fluence of Figure 2.4 . . . . . 36

3.1 Scale diagram of the SNO+ Detector and cavity . . . . . 39
3.2 Quantum efficiency of the SNO+ photomultiplier tubes against photon wavelength . 42
3.3 Simple schematic of the electronics and global trigger system at SNO+ . . . . . 44
3.4 Plot of the expected $0\nu\beta\beta$ spectrum at SNO+ at 0.3% $^{130}$Te loading, assuming a Majorana neutrino mass of 200 meV . . . . . 46
3.5 Plot of the range of allowed values for the effective Majorana mass $m_{\beta\beta}$ in terms of the mass of the lightest neutrino $m_{min}$ . . . . . 47
3.6 Comparison of external detector background rates with and without a fiducial volume cut . . . . . 50
3.7 Detector surface and internal backgrounds including leaching effects . . . . . 51
3.8 Total muon flux against vertical depth for a range of underground laboratories . . . 53
3.9 Diagram of common muon paths through the SNO+ detector . . . . . 54
3.10 Spectrum of solar backgrounds below 3 MeV in the SNO+ detector . . . . . 55

4.1 Sample time residual output of the L3 in water and scintillator . . . . . 64
4.2 Comparison of the L3 reconstructed energy output with and without retriggers from a simulated supernova at 10 kpc in scintillator . . . . . 71
4.3 Comparison of the L3 PMT hitmap output in $cos\theta$, $\phi$ for simulated supernovae at 10 and 30 kpc in water . . . . . 73
4.4 L3 PMT hitmap output in $cos\theta$, $\phi$ for the MDC TSLP crate-by-crate run . . . . . 73
4.5 Comparison plot of scintfitter and fastfitter performance over 25,000 IBD events in scintillator . . . . . 76
4.6 Fitted energy distributions for simulated supernovae at 10 and 30 kpc in scintillator using the fastfitter . . . . . 76
4.7 Fitted energy distribution for a 10 kpc simulated supernova in scintillator below 5 MeV 77
4.8 Fitted energy distributions for simulated supernovae at 10 and 30 kpc in water using the fastfitter . . . . . 78
4.9 Plot of the L3 IBD tag efficiency against increasing IBD-like event counts in scintillator 83
4.10 Plot of the overall L3 efficiency in scintillator against increasing supernova distance . 84

# List of Tables

2.1 Summary of thermal neutrino production interactions in CCSNe . . . . . 16

3.1 Table of most common internal background isotopes . . . . . 52

4.1 Table of L3 orphan counts during MDC data testing . . . . . 71
4.2 Table of segmentor performance for a range of simulated supernovae distances in water 73
4.3 Table of segmentor performance for a selection of MDC run types . . . . . 74
4.4 Table of event counts given by the L3 processor for a range of simulated supernova datasets in scintillator . . . . . 78
4.5 Table of event counts given by the L3 processor for a range of simulated supernova datasets in water . . . . . 79

## Acknowledgements

First and foremost, I'd like to thank my supervisor Prof. Steve Biller for his guidance. Our discussions on supernova neutrinos first sparked this thesis, and later chats frequently focused and sharpened its direction and content. Similarly, I doubt most supervisors would have taken too kindly to one of their graduate students leaving to become a Navy pilot midway through a PhD, let alone providing the means for doing an MSc instead - your generosity will never be forgotten.

The second most important individual in the writing of this thesis is undoubtedly Dr. Belina von Krosigk, and her own doctoral work on supernova neutrinos at SNO+. Her work was frequently a guiding light for this thesis, and it's left its mark in a number of kindly donated plots that would have been far too time consuming to generate otherwise. Huge thanks to you Belina.

I'd also like to thank both Dr. Matt Strait and Kevin Labe at the University of Chicago for the enlightening discussions on near-line processing, and the technical side of actually implementing a new supernova monitoring processor. Again, this work couldn't have been completed without your input.

To my fellow SNO+ graduate students here at Oxford - a massive thank you. It's fair to say that our office is unique in the building for its cameraderie and generally brilliant atmosphere, so my thanks to all of you, new and old. I'd like to single out Krish Majumdar in particular for his tireless efforts in proofreading my chapters and eliminating a number of insidious errors. As for the rest of you - Chris, Jack, Jeff, Alex - our many conversations and discussions on all things under the Sun have enabled the acquisition of a large amount of the knowledge now contained in this work (and much that has no place anywhere near it). Thank you all!

I'd also like to thank my family and friends for their patience, and their remarkable tolerance for my seeming to disappear off the planet for long stretches of this thesis.

Finally, to Chiara: without you this would have all been impossible.

Luca Cavalli
*December 2015*