

Robust Signal Extraction Methods and Monte Carlo
Sensitivity Studies for the Sudbury Neutrino Observatory and
SNO+ Experiments

by

Alexander Joseph Wright

A thesis submitted to the Department of Physics, Engineering Physics and Astronomy

in conformity with the requirements

for the degree of Doctor of Philosophy

Queen's University

Kingston, Ontario, Canada

September 2009

Abstract

The third and final phase of the Sudbury Neutrino Observatory (SNO) experiment utilized a series of ^3He proportional counters called Neutral Current Detectors (NCDs) to detect the neutrons produced by the neutral current interactions of solar neutrinos in the detector. The number of neutrons detected by the NCDs, and hence the total flux of ^8B solar neutrinos, has been determined using two novel signal extraction techniques which were designed to be robust against potential unexpected behaviour in the NCD background. These techniques yield total ^8B solar neutrino flux measurements of $5.04_{-0.40}^{+0.42}(\text{stat}) \pm 0.28(\text{syst}) \times 10^6 \text{cm}^{-2}\text{s}^{-1}$ and $(4.40 - 6.43) \times 10^6 \text{cm}^{-2}\text{s}^{-1}$, which are in good agreement with previous SNO results and with solar model predictions, and which confirm that previous NCD analyses were not unduly affected by unexpected background behaviour.

The majority of the hardware from the now-completed SNO experiment will be reused to create a new liquid scintillator based neutrino experiment called SNO+. An important part of the SNO+ physics program will be a search for neutrinoless double beta decay, carried out by dissolving ^{150}Nd into the scintillator. The sensitivity of the SNO+ experiment to neutrinoless double beta decay has been evaluated. If loaded at 0.1% (w/w) with natural neodymium, after 1 kT·a of data taking SNO+ would have a 90% C.L. sensitivity of $T_{1/2}^{0\nu} > 8.0 \times 10^{24}$ a or better 50% of the time; if the experiment were run with neodymium enriched to 50% in ^{150}Nd this limit improves to 57×10^{24} a. Under a reasonable choice for the ^{150}Nd neutrinoless double beta decay matrix element, these half lives correspond to upper limits on the effective Majorana neutrino mass of 112 meV and 42 meV, respectively. These limits are competitive with those expected from all other near-term neutrinoless double beta decay experiments.

Acknowledgements

The author would like to extend his thanks to all of those who have provided him with guidance, support and friendship over the past years.

The Particle Astrophysics Group at Queen's University provided a wonderfully supportive environment during my tenure as a graduate student. The faculty, including the emeritus faculty, were friendly, knowledgeable and approachable, and always took the time to assist students. The postdoctoral researchers and technical staff were extremely capable and genuinely interested in helping the students with their projects. My fellow graduate students proved to be extremely valuable collaborators, as well as great people to be around. In my experience, the Queen's Particle Astrophysics Group was composed entirely of excellent physicists who were also good people, which made working as a part of the group extremely rewarding on both the professional and the personal level.

The faculty at Queen's University, particularly Ian Towner, provided excellent coursework instruction which gave me the theoretical basis necessary to understand the experiments upon which I worked.

The SNO and SNO+ collaborations provided me with outstanding examples of how large collaborations can do excellent science while maintaining amicable and mutually respectful relationships among their members. I understand that this is not always the case in collaborations of this size, and it has made working on these experiments both fruitful and enjoyable.

While individual thanks could, and perhaps should, be given to each person with whom I have worked, I will restrict myself to a further four. First, special mention is due to Mark Chen for his excellent supervision, teaching and guidance. My frequent long discussions with Mark taught me more about physics and about being a good physicist than did any other aspect of my graduate

experience. Second, I wish to acknowledge Aksel Hallin who, during the time when our tenures at Queen's overlapped, was effectively my second supervisor. I have tried in many ways to model myself as an experimentalist after Aksel's excellent example. Third, I wish to thank Ken Clark for many helpful and enjoyable discussions about both physics and non-physics topics. Finally, I would like to thank Chuck Hearns for his practical advice regarding just about everything. I consider each of these gentlemen to be a close personal friend, and knowing them has enriched me personally as much as it has professionally.

Finally, and most importantly, I would like to thank my family for their support and encouragement over the years. Without the strong foundation which they provided I would not have been able to achieve what I have.

This work was supported financially by the Natural Sciences and Engineering Research Council through the P.G.S. and C.G.S. programs, and the Government of Ontario through the O.G.S.S.T. program.

Statement of Originality

As is always the case in experiments involving large collaborations, achieving the results described in this thesis required the effort of many individuals. As is also generally the case, however, the thesis focuses on those areas towards which the author made particularly important contributions.

The robust signal extraction techniques developed for use on the SNO NCD data, which are described in Chapter 4, were entirely the work of the author. They were developed in the context of discussions within the SNO collaboration about how best to treat the poorly understood NCD background in signal extraction, in which the author played a significant role. The author also participated in the development of the Monte Carlo model of the NCD pulse shapes and energy spectra and made important contributions to the development of a pulse shape analysis technique (which is not described in this thesis) which provided some ability to distinguish alpha and neutron events in the NCDs. In addition, the author played a central role in the development of the semi-empirical method of scope live fraction determination described in Section 5.3.3, and was responsible for defining and evaluating the systematic uncertainties (described in Section 5.2) associated with the use of an empirical neutron PDF in the signal extraction.

On the SNO+ experiment, the author was involved in almost every aspect of the development of the project. The author made important contributions to the development of the LAB liquid scintillator, playing a central role in the development of the LAB - acrylic chemical compatibility tests and the characterization of the LAB optical properties. The author was also responsible for modifying the SNOMAN Monte Carlo code to allow it to simulate the optical properties of a liquid scintillator and for using the modified Monte Carlo to study the light output of SNO+ under different optical conditions (including the addition of the neodymium double beta decay isotope). These

efforts are described in Section 6.2.1. The author was also responsible for the initial investigations of the feasibility (both optical and chemical) of loading double beta decay isotope into the LAB via nanoparticle suspension, and assisted in the development of the ultimate organometallic loading technique, both of which are described in Section 6.1.5. At the same time, the author created the physics sensitivity Monte Carlo code described in Section 7.4.1, and used it to investigate the sensitivity of the SNO+ experiment to the low energy solar neutrinos and to demonstrate that it was possible for SNO+ to carry out a competitive search for neutrinoless double beta decay, as discussed in Chapter 7. These studies formed an important part of the physics basis for the ultimately successful SNO+ funding proposals.

Contents

| | |
|--|-----|
| Abstract | ii |
| Statement of Originality | v |
| Table of Contents | vii |
| List of Tables | xii |
| List of Figures | xiv |
| Chapter 1. Introduction | 1 |
| Chapter 2. Solar Neutrinos and the Sudbury Neutrino Observatory | 4 |
| 2.1 Solar Neutrinos and the Solar Neutrino Problem | 4 |
| 2.2 The Sudbury Neutrino Observatory | 7 |
| 2.3 The Published SNO Solar Neutrino Results | 11 |
| 2.4 Neutrino Oscillations | 13 |
| 2.4.1 Vacuum Oscillations | 13 |
| 2.4.2 Matter Oscillation | 15 |
| 2.4.3 Current Constraints on the Solar Neutrino Mixing Parameters | 19 |
| Chapter 3. The SNO Neutral Current Detector Phase | 20 |
| 3.1 The Neutral Current Detectors | 20 |
| 3.2 Pulse Production in the NCDs | 22 |
| 3.3 NCD Data Acquisition | 27 |

| | | |
|-------------------|---|-----------|
| 3.4 | NCD Signals and Backgrounds | 28 |
| 3.4.1 | Neutron Signal | 29 |
| 3.4.2 | Alpha Backgrounds | 30 |
| 3.4.3 | Other Physics Backgrounds | 31 |
| 3.5 | Non-Physics Backgrounds and Data Cleaning | 32 |
| 3.5.1 | Data Cleaning | 32 |
| 3.5.2 | NNNAs | 34 |
| 3.6 | NCD Calibration | 38 |
| 3.6.1 | Electronics Calibration | 39 |
| 3.6.2 | Physics Calibration | 39 |
| 3.7 | Abnormal NCD Behaviour | 40 |
| 3.8 | Understanding the Alpha Background | 41 |
| 3.8.1 | Alpha Monte Carlo | 43 |
| 3.9 | Published NCD Results | 45 |
| Chapter 4. | Robust Determination of the NCD Neutron Number | 49 |
| 4.1 | The “Polynomial Method” | 49 |
| 4.1.1 | Polynomial Order | 50 |
| 4.1.2 | Binning | 57 |
| 4.1.3 | Fit Range | 57 |
| 4.2 | The “Overlap Method” | 58 |
| 4.2.1 | Method Description | 58 |
| 4.2.2 | Fit Range | 64 |
| 4.2.3 | Binning | 64 |
| 4.2.4 | Assumed Alpha Shape | 65 |
| 4.2.5 | Defining the Overlap Limits | 67 |
| 4.3 | Testing the Methods | 72 |

| | | |
|-------------------|--|-----------|
| 4.4 | Robust Signal Extraction Summary | 87 |
| Chapter 5. | Robust NCD Measurement of the Total Solar Neutrino Flux | 90 |
| 5.1 | The NCD Neutron Number | 90 |
| 5.2 | Neutron Number Systematic Uncertainties | 92 |
| 5.2.1 | Energy Resolution | 92 |
| 5.2.2 | Energy Scale | 95 |
| 5.2.3 | MUX Threshold Effects | 96 |
| 5.2.4 | ^{24}Na Statistics | 97 |
| 5.2.5 | ^{24}Na Non-Uniformity | 97 |
| 5.2.6 | Backgrounds in the ^{24}Na | 98 |
| 5.2.7 | Neutron PDF Systematic Uncertainties Summary | 99 |
| 5.2.8 | Verification with 2006 ^{24}Na Data | 101 |
| 5.3 | Live Fraction and Detection Efficiency | 101 |
| 5.3.1 | Shaper Live Fraction | 102 |
| 5.3.2 | MUX Live Fraction | 102 |
| 5.3.3 | Scope Live Fraction | 102 |
| 5.3.4 | Shaper Threshold Efficiency | 104 |
| 5.3.5 | MUX Threshold Efficiency | 104 |
| 5.3.6 | Data Cleaning Efficiency | 104 |
| 5.3.7 | Efficiency and Live Fraction Summary | 106 |
| 5.4 | Neutron Backgrounds in the NCDs | 106 |
| 5.5 | Live Time and Capture Efficiency | 110 |
| 5.5.1 | Live Time | 110 |
| 5.5.2 | Capture Efficiency | 110 |
| 5.5.3 | Live Fraction and Detection Efficiency Summary | 111 |
| 5.6 | NCD ^8B Flux Result | 111 |

| | | |
|--|--|------------|
| 5.6.1 | Neutron Rate to Flux Conversion | 111 |
| 5.6.2 | The NCD ^8B Flux Measurement | 112 |
| 5.7 | SNO Conclusion | 114 |
| Chapter 6. The SNO+ Experiment | | 115 |
| 6.1 | SNO+ Physics | 115 |
| 6.1.1 | Solar Neutrinos | 116 |
| 6.1.2 | Reactor Anti-neutrinos | 119 |
| 6.1.3 | Geo-Neutrinos | 121 |
| 6.1.4 | Supernova Neutrinos | 123 |
| 6.1.5 | Neutrinoless Double Beta Decay | 124 |
| 6.1.6 | Physics Summary | 126 |
| 6.2 | Experimental Design and Development | 126 |
| 6.2.1 | Scintillator Development | 126 |
| 6.2.2 | Acrylic Vessel Hold-Down | 133 |
| 6.2.3 | SNO+ Current Status and Outlook | 135 |
| Chapter 7. The SNO+ Neutrinoless Double Beta Decay Search | | 136 |
| 7.1 | Introduction to Neutrinoless Double Beta Decay | 136 |
| 7.1.1 | The Neutrinoless Double Beta Decay Rate | 137 |
| 7.1.2 | Implications of Detecting Neutrinoless Double Beta Decay | 139 |
| 7.1.3 | Calculation of Double Beta Decay Matrix Elements | 141 |
| 7.2 | Neutrinoless Double Beta Decay in SNO+ | 143 |
| 7.3 | Development of Neodymium Loaded Liquid Scintillator | 144 |
| 7.4 | Simulations to Estimate the SNO+ Sensitivity to Neutrinoless Double Beta Decay | 148 |
| 7.4.1 | Simulation Technique | 148 |
| 7.4.2 | Backgrounds Simulated | 150 |
| 7.4.3 | Fit Range | 152 |

| | | |
|-------|--|-----|
| 7.4.4 | Quantifying the $0\nu\beta\beta$ Sensitivity | 153 |
| 7.5 | Simulated SNO+ $0\nu\beta\beta$ Sensitivity | 154 |
| 7.5.1 | Comparison with Other $0\nu\beta\beta$ Experiments | 157 |
| 7.6 | Impact of Detector Parameters on the SNO+ $0\nu\beta\beta$ Sensitivity | 160 |
| 7.6.1 | Background Levels | 160 |
| 7.6.2 | Light Output/Energy Resolution | 162 |
| 7.6.3 | Neodymium Loading Levels | 162 |
| 7.6.4 | Energy Resolution Systematics | 163 |
| 7.6.5 | Pileup Backgrounds and Excited State Decays | 170 |
| 7.7 | Summary of the SNO+ Neutrinoless Double Beta Decay Search | 178 |
| 7.8 | SNO+ Conclusion | 179 |
| | References | 180 |

List of Tables

| | | |
|-----|--|-----|
| 4.1 | The trials penalty in polynomial selection. | 53 |
| 4.2 | The effect of the $\Delta\chi^2$ target on superfluous parameter acceptances. | 54 |
| 4.3 | The effect of the $\Delta\chi^2$ target on false parameter rejections. | 55 |
| 4.4 | The overlap of the NCD PRL PDFs. | 71 |
| 4.5 | The fraction of correct overlap neutron confidence intervals in Monte Carlo tests. . . | 81 |
| 4.6 | Polynomial order selection in the presence of hypothetical NNNAs. | 82 |
| 4.7 | The effect of polynomial order on fits to the blind data set. | 83 |
| 4.8 | Overlap limit tests on the blind data set. | 86 |
| 4.9 | Tests adding known NNNAs and alphas to the blind data. | 87 |
| 5.1 | The effect of polynomial order on fits to the full NCD data set. | 91 |
| 5.2 | Summary of the ^{24}Na systematic uncertainties | 100 |
| 5.3 | NCD live fractions and threshold efficiencies. | 106 |
| 5.4 | Summary of the neutron backgrounds | 109 |
| 5.5 | Summary of the systematic uncertainties in the ^8B flux measurement. | 112 |
| 6.1 | The expected SNO+ supernova neutrino signal. | 125 |
| 7.1 | ^{150}Nd nuclear matrix element calculations. | 143 |
| 7.2 | Baseline SNO+ $0\nu\beta\beta$ sensitivity with natural neodymium. | 157 |
| 7.3 | Baseline SNO+ $0\nu\beta\beta$ sensitivity with enriched neodymium. | 158 |
| 7.4 | Pileup event fraction as a function of background rate. | 175 |
| 7.5 | Pileup based limits on the tolerable levels of low energy radioisotopes in SNO+. . . . | 176 |

7.6 The results of neodymium γ -ray counting. 177

List of Figures

| | | |
|------|--|----|
| 2.1 | The solar fusion reactions | 4 |
| 2.2 | The solar neutrino energy spectra | 5 |
| 2.3 | The discrepancy between solar model predictions and experimental results. | 6 |
| 2.4 | The SNO detector. | 8 |
| 2.5 | The SNO Phase I energy fit. | 10 |
| 2.6 | The SNO Phase II isotropy fit. | 11 |
| 2.7 | The SNO neutral current, charged current, and elastic scattering flux results. | 12 |
| 2.8 | The energy dependence of the solar neutrino survival probability. | 18 |
| 2.9 | The current experimental constraints on the solar neutrino mixing parameters. | 19 |
| 3.1 | A schematic diagram of a Neutral Current Detector string. | 21 |
| 3.2 | The positions of the NCDs within SNO. | 22 |
| 3.3 | Example neutron pulse shapes. | 25 |
| 3.4 | Example alpha pulse shapes. | 26 |
| 3.5 | The neutron energy spectrum. | 29 |
| 3.6 | The alpha energy spectra. | 32 |
| 3.7 | Examples of spurious NCD events. | 33 |
| 3.8 | The effect of the data cleaning cuts on the NCD energy spectrum. | 35 |
| 3.9 | The energy spectra of some of the identified NNNAs. | 37 |
| 3.10 | The ^4He alpha spectra. | 42 |
| 3.11 | The Monte Carlo alpha spectrum. | 46 |
| 3.12 | The NCD energy fit published in PRL. | 48 |

| | | |
|------|---|-----|
| 4.1 | Tests of the polynomial choice algorithm. | 53 |
| 4.2 | The correlation between pull and $\Delta\chi^2$ | 56 |
| 4.3 | The effect of bin size on the overlap fits. | 66 |
| 4.4 | The background shape assumed in the overlap fits. | 67 |
| 4.5 | A comparison of the Monte Carlo and ^4He alpha spectra. | 68 |
| 4.6 | The overlap of different possible background shapes. | 70 |
| 4.7 | The overlap spread admitted by Monte Carlo alpha spectrum systematics. | 72 |
| 4.8 | Monte Carlo ensemble tests of the robust signal extraction methods. | 74 |
| 4.9 | Comparison of the effect of potential NNNAs on polynomial and standard fit biases. | 75 |
| 4.10 | A comparison of the effect of NNNAs on the overlap and standard fit biases. | 76 |
| 4.11 | A comparison of the effect of NNNAs on the overlap and polynomial fit biases. | 77 |
| 4.12 | A comparison of the effect of NNNAs on the polynomial and standard fit pulls. | 78 |
| 4.13 | A comparison of the effect of NNNAs on the overlap and standard fit pulls. | 79 |
| 4.14 | A comparison of the effect of NNNAs on the overlap and polynomial fit pulls. | 80 |
| 4.15 | Fits to the blind data set. | 84 |
| 4.16 | Fits to the blind data set with differing overlap limits. | 85 |
| 5.1 | Final fits to the full NCD data set. | 91 |
| 5.2 | The changes in the NCD energy calibration constants over time. | 94 |
| 5.3 | The ^{24}Na energy spectrum after systematic smearing. | 95 |
| 5.4 | The effect of the MUX threshold on the neutron spectrum. | 97 |
| 5.5 | The scope live fraction as a function of time since last event. | 105 |
| 5.6 | A comparison of the current neutral current flux results with the previous SNO results. | 113 |
| 6.1 | The solar electron neutrino survival probability curve with non-standard interactions. | 117 |
| 6.2 | Reactor anti-neutrino oscillation effects in SNO+. | 121 |
| 6.3 | The geo-neutrino energy spectrum. | 122 |
| 6.4 | The results of the Series I LAB - acrylic compatibility tests. | 129 |

| | | |
|------|---|-----|
| 6.5 | The current results of the Series II LAB - acrylic compatibility tests. | 130 |
| 6.6 | A candidate design for the SNO+ acrylic vessel hold-down. | 134 |
| 7.1 | The limits on the effective Majorana neutrino mass as a function of the minimum neutrino mass. | 140 |
| 7.2 | The optical attenuation of 0.6% $\text{Nd}_2\text{O}_3(\text{np})$ suspension in pseudocumene. | 146 |
| 7.3 | The simulated light output in SNO+ as a function of neodymium loading. | 147 |
| 7.4 | The simulated SNO+ energy spectrum including low energy backgrounds. | 151 |
| 7.5 | The effect of fit range on the SNO+ $0\nu\beta\beta$ fits. | 152 |
| 7.6 | The baseline SNO+ $0\nu\beta\beta$ sensitivity. | 156 |
| 7.7 | The SNO+ $0\nu\beta\beta$ signal at $\langle m_\nu \rangle = 270$ meV. | 156 |
| 7.8 | The effect of background levels on the SNO+ $0\nu\beta\beta$ limits. | 161 |
| 7.9 | The effect of energy resolution on the SNO+ $0\nu\beta\beta$ limits. | 162 |
| 7.10 | The effect of neodymium loading on the SNO+ $0\nu\beta\beta$ limits. | 163 |
| 7.11 | The effect of energy resolution uncertainty on the SNO+ $0\nu\beta\beta$ limits. | 164 |
| 7.12 | The addition of exponential tails to the simulated SNO+ energy response. | 165 |
| 7.13 | The effect of energy resolution tails on the SNO+ $0\nu\beta\beta$ limits. | 166 |
| 7.14 | The effect of energy resolution tails projected onto the spillover axis. | 168 |
| 7.15 | SNO+ false $0\nu\beta\beta$ “detections.” | 169 |
| 7.16 | Effect of excited state decays on the $2\nu\beta\beta$ spectrum. | 172 |
| 7.17 | The distortion of the $2\nu\beta\beta$ spectrum by event pileup. | 174 |

Chapter 1. Introduction

In recent years, our understanding of the fundamental properties of neutrinos has been significantly altered. The discovery of the non-zero neutrino mass, which was first suspected as a solution to the Solar Neutrino Problem and which has now been conclusively demonstrated through the observation of neutrino oscillations by a number of experiments, represented the first significant change in the Standard Model of Particle Physics in many years. The existence of a small but non-zero neutrino mass poses an interesting problem for theorists and experimentalists alike; for theorists, the challenge is to explain how such a small mass might arise, while the challenge for experimentalists is to precisely describe the additional neutrino properties and behaviours that are made possible by the non-zero mass. In particular, there is a great deal of interest in determining the absolute scale and ordering of the neutrino masses and in discovering whether neutrinos are Majorana or Dirac particles. Determining these properties, which would fill some of the last gaps in our knowledge of the Standard Model particles, will help us to understand the role that neutrinos played in the evolution of the universe, and perhaps even help to explain the origin of the matter-antimatter asymmetry.

The Sudbury Neutrino Observatory was one of the experiments which first demonstrated the existence of a non-zero neutrino mass by showing conclusively that the solar neutrinos undergo flavour change. After this initial discovery, SNO continued to operate, collecting data to confirm its initial results and to measure more precisely the parameters which describe solar neutrino oscillations. The second chapter of this thesis describes the SNO experiment in the context of its contribution to the solution of the Solar Neutrino Problem. Neutrino oscillations and the formalism describing them are also introduced. Chapter 3 gives further detail about the third and final phase of SNO. During this

phase, the neutrons produced by neutrino interactions within the detector were primarily detected using an array of proportional counters. The design and operation of these counters is described, and the results of the phase are summarized.

In determining the number of neutrons captured by the proportional counters, it was necessary to perform a signal extraction in a regime where the behaviour of the neutron signal was understood very well but the behaviour of the background was more difficult to predict. This ill-defined uncertainty in the description of the NCD background posed a serious potential problem to traditional χ^2 and maximum likelihood signal extraction techniques wherein an incorrect description of the background shape would in general result in a bias in the extracted number of neutrons. Chapter 4 describes the development and validation of two novel signal extraction techniques which were created to help address this problem by minimizing their dependence on prior knowledge of the background distributions. In the first method, a polynomial is used to describe the background shape, with the optimal polynomial order necessary to describe the data selected *a posteriori* by comparing, using a goodness-of-fit based selection algorithm, fits to the data using different polynomial orders. In the second method, a neutron number confidence interval is determined using only the standard assumption that it is physically unreasonable for the background shape to mimic the behaviour of the neutron signal too closely. In both methods, lessening the dependence of the signal extraction on prior knowledge of the background spectrum makes the method more robust than traditional signal extraction techniques when unexpected background spectra are present in the data, at the cost of a decrease in fit precision. In Chapter 5, these new signal extraction techniques are used to determine the number of neutrons detected by the proportional counter array and hence to produce a new, robust, measurement of the total flux of ^8B solar neutrinos.

The SNO experiment has now been completed, but investigations of neutrino properties at SNO-LAB will continue with the new SNO+ experiment. SNO+ is a liquid scintillator experiment that is currently being constructed from the SNO detector. An introduction to the SNO+ experiment, including its history and development and an overview of its physics potential is given in Chapter

6. SNO+ will pursue a broad program of neutrino research, including investigation of the low energy solar neutrinos, geo-neutrinos, reactor anti-neutrinos, and supernova neutrinos. In addition, by loading the scintillator with neodymium SNO+ will also stage a search for neutrinoless double beta decay. Detection of neutrinoless double beta decay would demonstrate that neutrinos are Majorana particles and provide information about the absolute neutrino mass. The search for neutrinoless double beta thus has the potential to provide answers to two of the most important outstanding questions about the properties of the neutrino. In Chapter 7, the double beta decay process is introduced and the SNO+ double beta decay search is described. The sensitivity of the SNO+ double beta decay search is then estimated using Monte Carlo simulations and the effect of different experimental factors on the sensitivity of the experiment is explored. A comparison of the SNO+ double beta decay sensitivity with that of other near- and medium-term double beta decay experiments is also given.

In summary, this thesis deals with the study of fundamental neutrino properties in two separate but related experiments. The analysis of the data from the final phase of the Sudbury Neutrino Observatory confirms again the observation that neutrinos have mass and helps to further constrain the solar neutrino oscillation parameters. The development of the SNO+ experiment, on the other hand, sets the stage for further investigations of neutrino behaviour, including the staging of a competitive search for neutrinoless double beta decay. Thus, the thesis describes how an experiment which helped to demonstrate that neutrinos are massive particles will, in a new form, remain at the forefront of efforts to learn more about the nature of that mass and will, as a result, continue to expand our knowledge of these particles at the most fundamental level.

Chapter 2. Solar Neutrinos and the Sudbury Neutrino Observatory

2.1 Solar Neutrinos and the Solar Neutrino Problem

Since the early part of the 20th century it has been understood that the energy emitted by the Sun is produced through a series of nuclear fusion reactions [1]. These reactions are now understood to occur in two series, the “*pp* chain” and the “CNO cycle,” both of which ultimately result in the fusion of protons into ⁴He nuclei. These reaction series are illustrated in Figure 2.1. It is interesting to note that the early papers on solar fusion make no reference to neutrinos; the solar neutrinos were first discussed in 1948, about 10 years later [2].

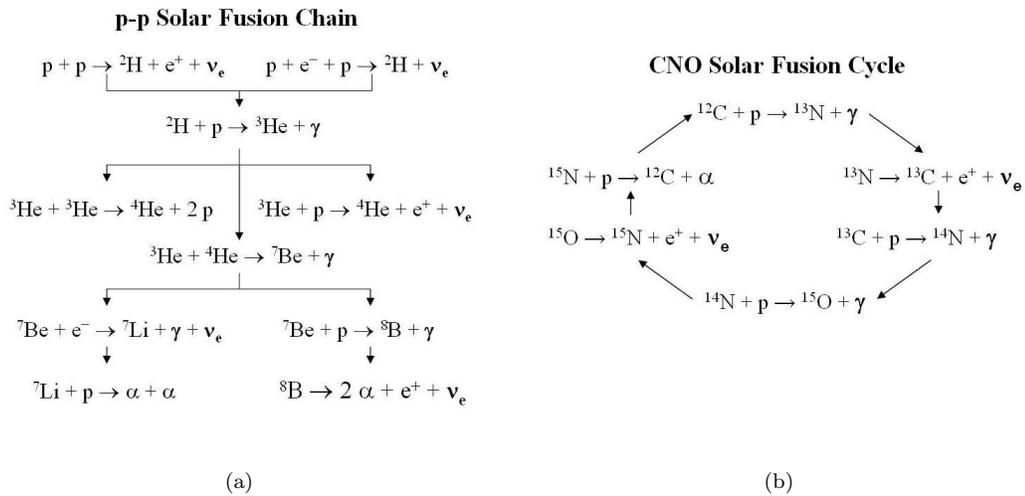


Figure 2.1: The nuclear fusion reactions that produce the energy emitted by the Sun. Of these, the *pp* chain is responsible for the bulk of the solar energy generation. Only the dominant reactions in the CNO cycle are shown.

By the early 1960’s, the cross-sections for most of the important solar fusion reactions had been

measured well enough, and models of the solar interior were developed to a sufficient extent that the fluxes and energy spectra of the solar neutrinos could be predicted with some reliability (a modern prediction of the solar neutrino spectra is shown in Figure 2.2). These calculations predicted that the flux of higher energy solar neutrinos was large enough to be detectable by terrestrial experiments; as a result, a chlorine-based solar neutrino experiment was constructed in the Homestake Mine in South Dakota [2].

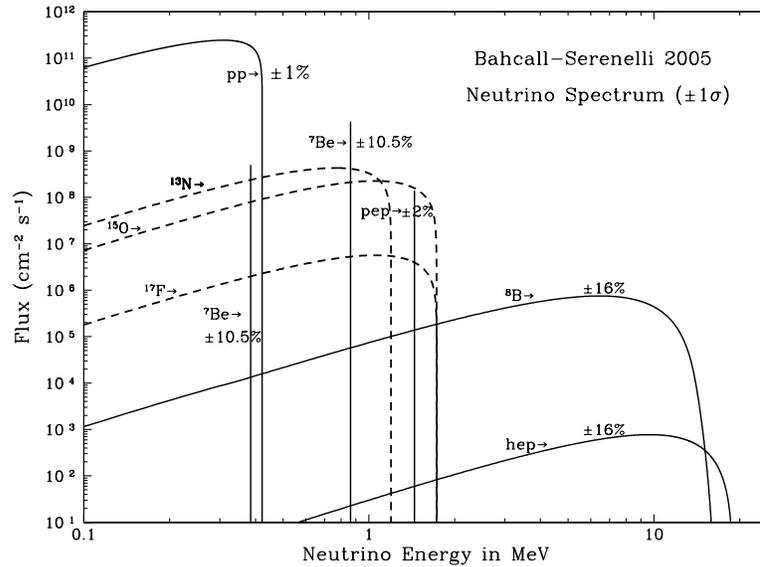


Figure 2.2: The energy spectra of the solar neutrinos as predicted by solar modelling. The pp chain neutrino fluxes are represented by solid lines, while the CNO neutrino fluxes are represented by dashed lines. Figure from [3].

The flux of solar neutrinos measured by this “Homestake Chlorine Experiment” [4] and subsequent radiochemical [5, 6, 7] and real-time [8, 9] experiments was surprising; in every case, the measured neutrino flux was significantly lower than the solar model predictions. As it became evident that this discrepancy could be ascribed neither to errors in the solar model calculations nor to experimental mistakes, the difference between the measured and predicted solar neutrino fluxes became known as the “Solar Neutrino Problem.” Figure 2.3 shows a modern summary of the discrepancy between the solar neutrino flux measured by the radiochemical and light water experiments and the predictions of the standard Solar Model.

Total Rates: Standard Model vs. Experiment
Bahcall–Serenelli 2005 [BS05(OP)]

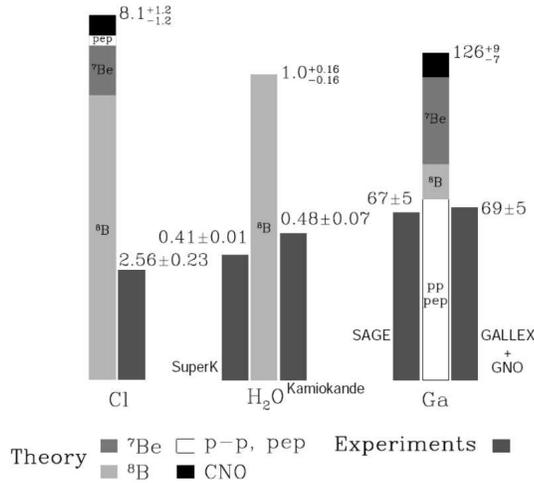


Figure 2.3: The discrepancy between the predictions of the Standard Solar Model and the solar neutrino fluxes measured by the radiochemical (Ga and Cl) and elastic scattering (H_2O) experiments that formed the basis of the Solar Neutrino Problem. Figure adapted from [3].

Numerous potential solutions to the solar neutrino problem were discussed over the subsequent years [10]. These could be divided into two groups; those which suggested an incomplete understanding of the Sun, in which case fewer neutrinos than initially expected might be produced, and those which predicted a new behaviour of the neutrino, by which some fraction of the neutrinos produced in the Sun could evade detection by the Earth-based experiments. In the former category were non-standard solar models, which considered, among other ideas, non-standard isotopic abundances in the solar interior, possible effects of magnetic fields and rotations of the solar interior, and thermal dis-equilibrium between the solar surface and interior. Possible resonances in the cross-sections for the solar fusion reactions were also discussed. The latter category included the possibility of neutrino decay to a light scalar or pseudo-scalar particles and another neutrino flavour [11], the transition of neutrinos to (sterile) right-handed neutrinos as the result of magnetic precession [12], and oscillations between neutrino flavours similar to those that had already been observed in the

quark sector [13, 14, 15]¹.

In 1985 it was realized [16] that a heavy water ($^2\text{H}_2\text{O}$ or “D₂O”) based solar neutrino experiment would be able to distinguish between the reduced neutrino production and flavour change scenarios. In such an experiment, neutrinos would be detected primarily through two nuclear interactions with deuterium: the charged current interaction,

$$\nu_e + d \rightarrow p + p + e^-, \quad (2.1)$$

which is sensitive only to electron neutrinos and the neutral current interaction,

$$\nu_{e,\mu,\tau} + d \rightarrow \nu_{e,\mu,\tau} + n + p, \quad (2.2)$$

which is equally sensitive to all three neutrino flavours. A heavy water experiment would thus be able to measure the total active solar neutrino flux using the neutral current interaction and, by comparing that flux with the electron neutrino flux measured by the charged current interaction, determine directly whether or not the solar neutrinos undergo flavour change.

2.2 The Sudbury Neutrino Observatory

The Sudbury Neutrino Observatory (SNO) was a heavy water based solar neutrino detector built to resolve the Solar Neutrino Problem. The detector, shown in Figure 2.4, consisted of 1000 T of heavy water in a 12 m diameter acrylic sphere suspended in 7400 T of ultra-pure light (i.e. isotopically normal) water to shield the experiment from radioactive backgrounds. The active volume was surrounded by an array of about 9500 photomultiplier tubes (PMTs) which were used to detect the Čerenkov light given off by the energetic electrons produced, directly or indirectly, by the neutrino interactions in the heavy water. To shield against background from cosmic radiation, the experiment was located 6800 feet underground in VALE INCO’s Creighton mine near Sudbury, ON. The design of the SNO experiment is described in detail in [17].

¹ Neutrino oscillations could account for the observed deficit of solar neutrinos because the early neutrino detectors were either only sensitive to, or primarily sensitive to, electron flavour neutrinos, like those produced in the solar fusion reactions. Neutrino oscillations would cause some of these solar electron neutrinos to appear as muon or tau neutrinos on Earth, thereby lowering the number of neutrinos detected by these experiments.

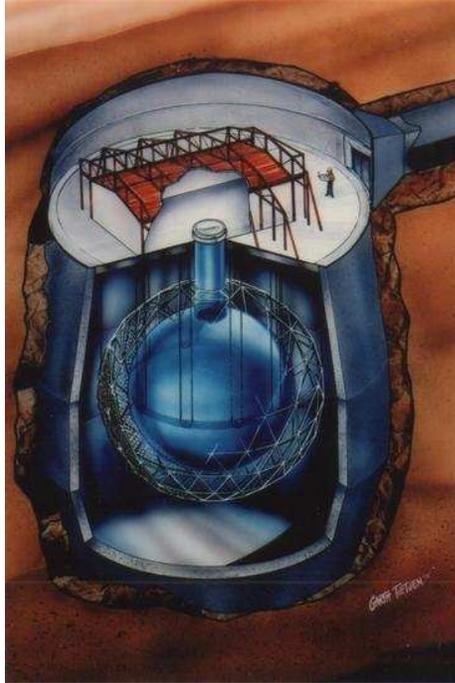


Figure 2.4: An artist's depiction of the SNO detector.

In addition to the charged current and neutral current reactions discussed earlier, SNO also detected neutrinos through the neutrino-electron elastic scattering interaction,

$$\nu_{e,\mu,\tau} + e^- \rightarrow \nu_{e,\mu,\tau} + e^-. \quad (2.3)$$

This interaction (which is also the neutrino detection channel used in the light water neutrino experiments) is sensitive to all three neutrino flavours, although the cross section is about six times higher for interactions with electron neutrinos.

The charged current and elastic scattering interactions in SNO were detected by observing the Čerenkov light from the directly produced energetic electrons in the PMT array. Because Čerenkov light production is both prompt and highly directional, it was possible to use the relative positions, distribution in time, and number of photons collected by the PMT array to reconstruct the position, direction and energy of the primary electron. In the case of the elastic scattering interaction, the direction of the electron is strongly correlated with the direction of the neutrino, while the direction of the charged current electrons is weakly anti-correlated with the neutrino direction. The direction

of an event with respect to the position of the Sun thus provided a powerful means of separating elastic scattering events from the charged current and (isotropic) neutral current signals.

Separation of the charged current and neutral current interactions was more subtle, as it involved separating the Čerenkov signals of the directly produced charged current electrons from the secondary signals produced by the neutral current neutrons. To ensure that the charged current - neutral current separation was done correctly and accurately, the SNO experiment was carried out in three phases, each with a different primary means of performing this separation.

In the first phase of the experiment, which ran from November 1999 to May 2001, the active volume of the detector contained pure heavy water. In this configuration, the neutrons produced by the neutral current interaction were primarily detected through their capture on deuterium nuclei, which resulted in the production of single 6.25 MeV γ -rays. These γ -rays subsequently Compton scattered in the heavy water, each of them typically producing a single electron that emitted Čerenkov light. In Phase I, then, the neutral current and charged current interactions both resulted in the creation of single Čerenkov electrons, meaning that the signals could not be differentiated on an event-by-event basis. However, the energy of the electron produced in the charged current reaction is strongly correlated with energy of the incident neutrino, resulting in a spectrum of electron energies that is quite different from the electron energy spectrum produced from neutron captures on deuterium (which is independent of neutrino energy). In Phase I, then, the neutral current and charged current signals were primarily separated by fitting the observed electron energy spectrum to the expected charged current and neutral current spectra [18]. This fit is shown in Figure 2.5.

The main disadvantage to the use of an energy fit to separate the charged current and neutral current signals was that the ^8B solar neutrino spectral shape must be assumed in order to calculate the expected charged current spectrum. This was not ideal because (as will be discussed later) neutrino oscillations, if they occurred, were expected to be energy dependant, and would therefore result in distortions of the expected charged current spectrum. Thus, in the second phase of the experiment, which ran from July 2001 to October 2003, about two tonnes of NaCl was dissolved

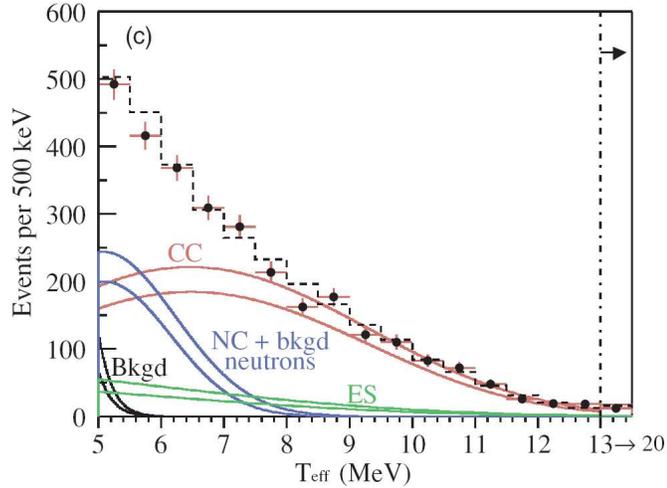


Figure 2.5: The energy fit used in Phase I as the primary charged current - neutral current separation variable. The charged current spectrum, which reflects the ${}^8\text{B}$ solar neutrino energy spectrum, extends to higher energy than the neutral current spectrum, which is simply the energy spectrum produced by neutron captures on ${}^2\text{H}$. Figure from [19].

into the heavy water. As ${}^{35}\text{Cl}$ has a higher neutron capture cross section than deuterium, most of the neutral current neutrons detected in Phase II were captured by ${}^{35}\text{Cl}$. Such captures released a total 8.6 MeV distributed among several γ -rays (2.5 on average), each of which could subsequently Compton scatter. Therefore, in Phase II each neutron capture resulted in the production of multiple Čerenkov electrons. Having multiple Čerenkov electrons meant that the light produced by a neutral current event during Phase II was more isotropic on average than the light produced by a (single electron) charged current event. It was therefore possible to use an isotropy parameter, called β_{14} , to statistically separate the charged current and neutral current signals [20], as shown in Figure 2.6. This made it possible to separate and measure the charged current and neutral current fluxes without relying on their energies. In fact, the separation of the charged current and neutral current fluxes in Phase II was done independently in the different energy bins, which allowed SNO to make a measurement of the charged current spectrum and, through it, the ${}^8\text{B}$ solar neutrino spectrum.

Finally, in the third phase of SNO, which ran from November 2004 to November 2006, the neutral current neutrons were primarily detected through an array of ${}^3\text{He}$ proportional counters that was deployed in the heavy water volume. These “Neutral Current Detectors” or “NCDs” provided a

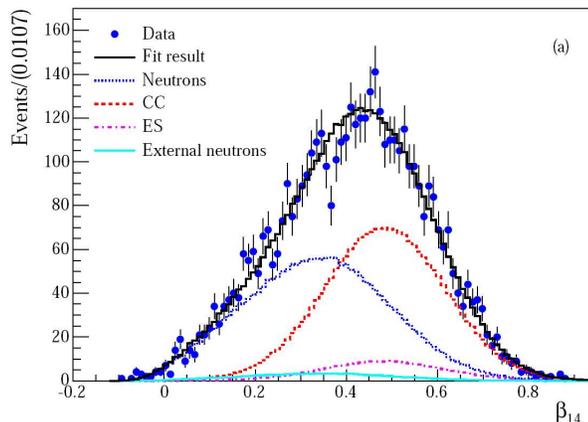


Figure 2.6: The β_{14} isotropy parameter fit used in Phase II as the primary charged current - neutral current separation variable. Events which produced light more isotropically in the detector, like neutron captures on ^{35}Cl which subsequently produced multiple γ -rays, had lower values of β_{14} , while less isotropic events like the charged current and elastic scattering interactions which produced only single Čerenkov electrons had higher β_{14} values. Figure from [20].

completely independent detection channel for the neutral current interactions, and hence reduced the correlation between the charged current and neutral current flux measurements². The NCD measurement of the total solar neutrino flux is the topic of the SNO portion of this thesis.

2.3 The Published SNO Solar Neutrino Results

The SNO experiment showed conclusively that the solar neutrinos undergo flavour change [18, 20, 21]. The total flux of solar neutrinos, as measured by the neutral current reaction in each phase of the experiment, is shown in Figure 2.7(a). Each of these measurements is in good agreement with the predictions of the Standard Solar Model ($(5.69 \pm 0.91) \times 10^6 \text{cm}^{-2}\text{s}^{-1}$ [3]). However, as seen in Figure 2.7 (b) and (c), the charged current and elastic scattering fluxes are significantly lower, indicating that approximately 66% of the electron neutrinos produced in the Sun appear as muon or tau neutrinos on Earth.

The SNO results, in combination with those of the other radiochemical and real-time solar

² In Phases I and II, the fits were essentially background free in the relevant energy range, and the elastic scattering signal was strongly constrained by its directionality with respect to the direction of the Sun. Thus, increasing the neutral current flux in the fit tended to decrease the charged current flux and vice versa, meaning that the uncertainties on the charged current and neutral current fluxes were highly correlated. Reducing this correlation reduces the uncertainty on the charged current / neutral current flux ratio (the most important number in the neutrino flavour change investigations), and hence increases the sensitivity of the SNO experiment to neutrino oscillations.

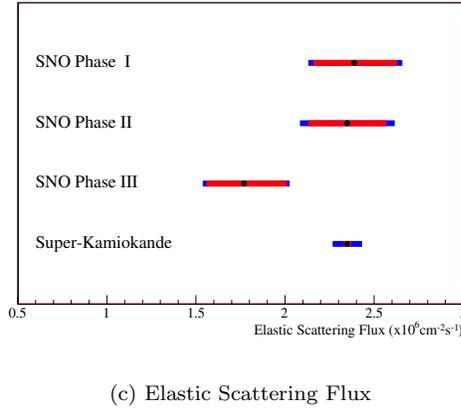
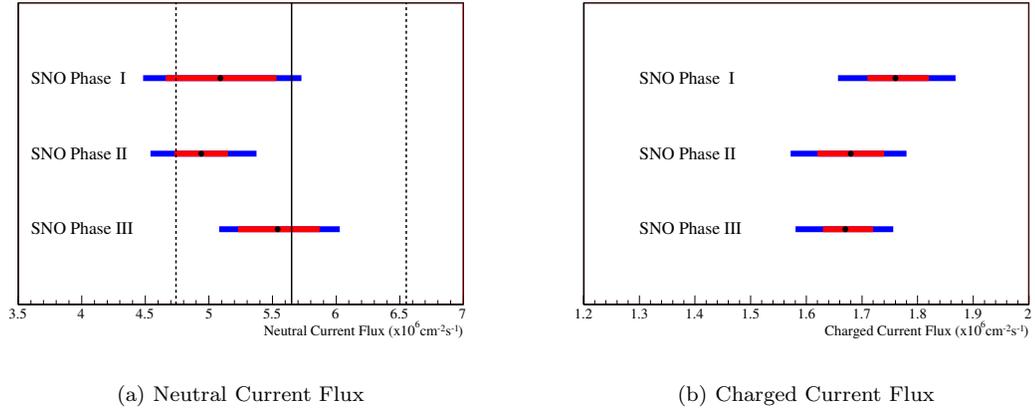


Figure 2.7: The flux results from the three phases of SNO [19, 20, 21]. The total solar neutrino flux inferred from the rates of neutral current interactions observed in SNO is shown in (a). (b) and (c) show the total neutrino fluxes inferred from the charged current and elastic scattering event rates under the assumption that the solar neutrino flux is purely ν_e . For each phase, the best fit point is shown in black, the inner red band indicates the range of the statistical uncertainties and the outer blue band shows the range of the total uncertainty (the quadratic sum of the statistical and systematic uncertainties). There is clearly a significant difference between the neutral current, charged current and elastic scattering fluxes, indicating that the solar neutrinos undergo flavour change (note that the x-axis scales are different in the different plots). The prediction of the Standard Solar Model (BS05, [3]) is shown as a vertical band in (a) for comparison with the SNO result - the solid line shows the central value prediction and the dashed lines represent the $\pm 1\sigma$ uncertainties. The high statistics measurement of the elastic scattering flux from the 1496-day data set of the Super-Kamiokande experiment [9] is shown in (c) for comparison. The apparently low elastic scattering number from the SNO Phase III result has been investigated and appears to be entirely consistent with a statistical fluctuation which seems to have occurred in a single energy bin.

neutrino experiments discussed earlier [4, 5, 6, 7, 8, 9, 22] and the KamLAND reactor anti-neutrino experiment [23], show that the observed neutrino flavour change is well described by matter enhanced “MSW”-type neutrino oscillations (described in the next section) but not by neutrino decay or de-coherence models [24]. The Solar Neutrino Problem has thus been resolved and neutrino oscillations are now accepted as a fundamental aspect of neutrino physics.

2.4 Neutrino Oscillations

Excellent reviews of the neutrino oscillation theory relevant to solar neutrino oscillations are given in Chapter 9 of [25] and in [26]. The experimentally important oscillation phenomena are outlined below, with the reader referred to the aforementioned references for more detailed derivations.

2.4.1 Vacuum Oscillations

Standard or “vacuum” neutrino oscillations of the type first proposed by Gribov and Pontecorvo [13] arise, as in the quark sector, from imperfect alignment between the neutrino mass eigenstates (ν_1 , ν_2 , and ν_3) and flavour eigenstates (ν_e , ν_μ , and ν_τ). The relationship between the mass and flavour eigenstates is given by the so-called Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, \hat{U}_{PMNS} , via

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \hat{U}_{PMNS} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}. \quad (2.4)$$

As the PMNS matrix is believed to be unitary, it can be expressed in its most general form in terms of three rotation angles and one complex phase. For neutrino oscillations, the mixing angles are conveniently chosen to be θ_{12} , θ_{23} , and θ_{13} , and the complex phase is usually labelled δ . The general

form of the PMNS matrix can then be written

$$\hat{U} \equiv \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \quad (2.5)$$

$$= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \times \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \times \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (2.6)$$

$$= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \quad (2.7)$$

where $s_{ij} \equiv \sin \theta_{ij}$ and $c_{ij} \equiv \cos \theta_{ij}$.

Quantum mechanics dictates that as a superposition of mass eigenstates (e.g. a neutrino produced in a pure flavour eigenstate) propagates, the probability of that particle being detected in any one of the flavour eigenstates changes with time. In fact, these probabilities oscillate, with the frequency of oscillation determined by the difference between the energies of the propagating states. With the assumptions that the neutrino rest masses are small compared to their energies and that the different mass eigenstates propagate with equal momenta, this energy splitting can be expressed in terms of the differences between the mass eigenstates. For a neutrino produced in the electron flavour eigenstate, the probability, P_{ee} , that the neutrino is detected in that same eigenstate after propagating for a time t is given by

$$\begin{aligned} P_{ee} &= 1 - \frac{1}{2} \cos^4(\theta_{13}) \sin^2(2\theta_{12}) \left(1 - \cos\left(\frac{\Delta^2 m_{21} t}{2E\hbar}\right) \right) \\ &\quad - \frac{1}{2} \sin^2(2\theta_{13}) \left(1 - \cos\left(\frac{\Delta^2 m_{31} t}{2E\hbar}\right) \right) \\ &\quad - \frac{1}{2} s_{12}^2 \sin^2(2\theta_{13}) \left(\cos\left(\frac{\Delta^2 m_{31} t}{2E\hbar}\right) - \cos\left(\frac{\Delta^2 m_{21} + \Delta^2 m_{31} t}{2E\hbar}\right) \right), \end{aligned} \quad (2.8)$$

where E is the total energy of the propagating neutrino and $\Delta^2 m_{ij} = m_i^2 - m_j^2$ is the difference in the squares of the masses of the neutrino mass eigenstates. Thus, for neutrino oscillations there are three important mass splittings, Δm_{12} , Δm_{23} , and Δm_{13} , although only two of these are independent

once their ordering is known.

Because the size of the neutrino production region in the Sun is large compared to neutrino oscillation lengths, terrestrial solar neutrino experiments are sensitive only to the phase-averaged survival probability (i.e. the average over t in Equation 2.8):

$$P_{ee} = 1 - \frac{1}{2} \cos^4(\theta_{13}) \sin^2(2\theta_{12}) - \frac{1}{2} \sin^2(2\theta_{13}). \quad (2.9)$$

Thus we can see that if neutrino oscillations were restricted to vacuum phenomena, the solar neutrino survival probability would be independent of neutrino energy and depend only on two mixing angles

The situation is in fact even simpler than this. Although three flavour neutrino oscillations are in general described by 6 parameters (three mixing angles, two mass splittings and one complex phase), it turns out that, because θ_{13} is small and Δm_{12} is significantly less than Δm_{23} [27], the PMNS matrix factors naturally, both physically and mathematically, into the form shown in Equation 2.6. The different terms in this equation are experimentally accessible through different experiments; the θ_{13} term by short baseline reactor anti-neutrino experiments, the θ_{23} term by atmospheric neutrino experiments and the θ_{12} term by solar neutrino experiments and by long-baseline reactor experiments [26]. For solar neutrino experiments, therefore, the two-flavour approximation, in which θ_{13} is taken to be zero and hence in which only θ_{12} and Δm_{12}^2 appear, is sufficiently accurate to describe the observed oscillation phenomena. Under this approximation, the phase-averaged electron neutrino survival probability in vacuum is simply

$$P_{ee} = 1 - \frac{1}{2} \sin^2(2\theta_{12}). \quad (2.10)$$

2.4.2 Matter Oscillation

As Wolfenstein [14] and then Mikheyev and Smirnov [15] pointed out, there is an additional oscillation effect when neutrinos propagate through matter. As neutrinos propagate through matter, the neutrino-matter interaction adds an “effective mass” term to the neutrino mass (Wolfenstein used the analogy of matter having an “index of refraction” for neutrino propagation). If all neutrinos experienced identical interactions with matter this would have no effect on oscillation phenomena.

However, electron neutrinos do interact differently with matter than the other neutrino flavours due to the possibility of charged current interactions between electrons and electron neutrinos. This means that electron neutrinos gain a different “effective mass” from matter interactions than do the other neutrino flavours, and this difference in effective mass can affect neutrino oscillations.

In the two-flavour approximation, the survival probability for a neutrino propagating in uniform matter is given by

$$P_{ee} = 1 - \frac{1}{2} \sin^2(2\theta_m) \left(1 - \cos \left(\frac{2\pi L \Delta M}{4\pi E \hbar c} \right) \right), \quad (2.11)$$

where the matter mixing angle, θ_m , is given by

$$\sin(2\theta_m) = \frac{\Delta m^2}{\Delta M} \sin(2\theta) \quad (2.12)$$

or by

$$\tan(2\theta_m) = \frac{\Delta m^2 \sin(2\theta)}{\Delta m^2 \cos(2\theta) - 2\sqrt{2}G_F n_e E}, \quad (2.13)$$

and the difference in the matter energy eigenstates, $\Delta M = E_2 - E_1$, is given by

$$\Delta M = \sqrt{\left(\Delta m^2 \cos(2\theta) - 2\sqrt{2}G_F n_e E \right)^2 + \left(\Delta m^2 \sin(2\theta) \right)^2}. \quad (2.14)$$

In these expressions, G_F is the Fermi constant, n_e is the electron number density, θ and Δm^2 are the vacuum mixing angle and the rest mass splitting, respectively (there is only one of each in the two flavour approximation), and E_2 is defined to be larger than E_1 . It can now be seen that, in addition to the mixing angle, the survival probability for neutrinos propagating in uniform matter depends on the neutrino energy and mass splitting, as well as the matter density and the strength of the neutrino-matter interaction. One important feature of matter mixing is that, as can be seen from Equation 2.13 or from Equations 2.12 and 2.14, the matter mixing is maximal when

$$\cos(2\theta) = \frac{2\sqrt{2}G_F n_e E}{\Delta m^2}, \quad (2.15)$$

regardless of the size of the vacuum mixing angle.

For neutrinos propagating in non-uniform matter densities (like the neutrinos propagating out of the Sun), determining the survival probability is more challenging and involves continual “re-mapping” of the neutrino mass eigenstates as the density changes. Conceptually, the matter effect

can be thought of as changing the relationship between the mass and flavour eigenstates. This means that an electron neutrino produced in the very high matter density of the solar core has a different admixture of mass eigenstates than an electron neutrino produced in vacuum. As the neutrinos produced in the core of the Sun propagate through the smoothly decreasing solar density to the surrounding vacuum, they change adiabatically (i.e. undergo equilibrium re-mappings of the mass eigenstates, but not jumping between mass eigenstates) from the high-density mass-to-flavour-eigenstate relations to the low density relations. Thus, for example, it is possible that ν_e is primarily ν_1 in vacuum and primarily ν_2 in the dense matter of the solar core. So when the (primarily) ν_2 neutrinos produced as ν_e in the Sun propagate to the (vacuum-like) Earth (where ν_2 has a much smaller ν_e component) they appear to be primarily non- ν_e . This “flavour swap” occurs when the matter density at the point of neutrino production is greater than, or equal to, the density given by the resonant mixing condition (Equation 2.15).

Analytic calculation of the solar survival probability can be achieved for simple density profiles, for example in situations where the solar density changes linearly or exponentially with radius. The linear density approximation is illustrative: the phase-averaged survival probability for electron neutrinos produced on the “high density” side of the resonant region and detected on the “low density” side is in this case given by

$$P_{ee} = \frac{1}{2} + \left(\frac{1}{2} - P_j\right) \cos(2\theta_m) \cos(2\theta), \quad (2.16)$$

where θ_m corresponds to the matter density at the production point,

$$P_j = \exp\left(\frac{-\pi\Delta m^2 \sin^2(2\theta)}{4E \cos(2\theta)} \left(\frac{n_e}{|dn_e/dr|}\right)_{\text{res}}\right), \quad (2.17)$$

and “res” refers to the resonant density region. Here P_j gives the probability of a jump between flavour eigenstates as a result of the change in relationship between the flavour and mass eigenstates³. Note that the survival probability still depends on the neutrino energy, the mixing angle and the neutrino mass splitting.

³ In the adiabatic approximation, the neutrinos do not change mass eigenstates, so any change in the relationship between the mass and flavour eigenstates must appear in the flavour basis.

Numerical calculations which include realistic solar density profiles, which account for inexactness in the adiabatic approximation⁴ and which properly treat neutrinos of lower energy for which a resonant transition does not occur, have been carried out by the SNO collaboration, among others. Figure 2.8 shows the solar neutrino survival probability over the relevant energy range for solar mixing parameters near the current global best-fit values [21]. As can be seen in the Figure, at high energies, the resonant transition is “fully on.” The survival probability is low, changes slowly with energy, and to a large extent is determined only by the mixing angle (i.e. is fairly independent of neutrino energy). At very low energies, the solar electron density is not high enough for the matter effect to be significant and neutrino oscillations are essentially vacuum oscillations. The transition between these extremes is determined by the turn-on in the resonant condition; the position of this “survival upturn” is determined by the solar matter density and the neutrino mass splitting.

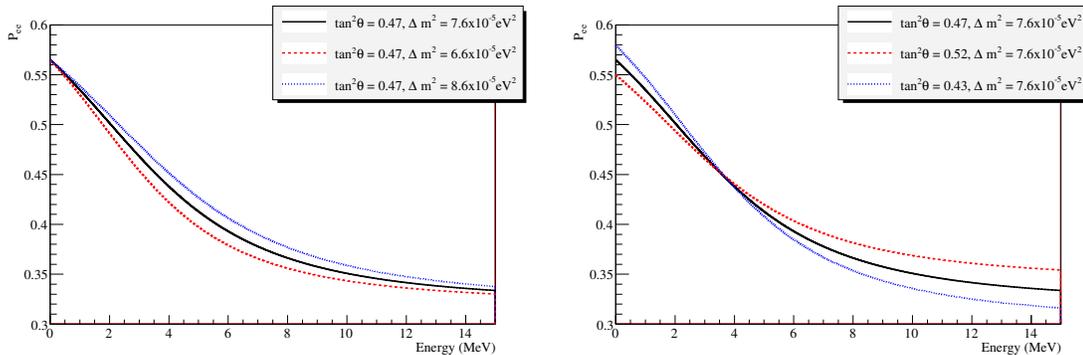


Figure 2.8: The solar electron neutrino survival probability as a function of energy. As shown on the left, changing the neutrino mass splitting (the changes shown here are about five times larger than the uncertainty on the global solar + KamLAND best fit uncertainty) changes the survival probability in the transition region, but not at the high and low energy asymptotes. Changing the mixing angle, on the other hand, changes the survival probability in the high and low energy regions, as shown on the right (these curves show the $\pm 1\sigma$ variation about the current global best fit mixing angle from [21]). In all cases, the survival probabilities shown are averaged over the solar neutrino production region.

⁴ For the adiabatic condition to be perfect, the propagation distance through the region of the Sun at which the resonance condition applies must be long compared to the oscillation wavelength at that matter density.

2.4.3 Current Constraints on the Solar Neutrino Mixing Parameters

As described in the previous section, the neutrino oscillations of interest to solar neutrino experiments can be described by two mixing parameters, θ_{12} and Δm_{12}^2 . The current experimental constraints on these parameters, as derived from the SNO results alone, from the combination of the data from all solar neutrino experiments, and from the combination of the data from all of the solar neutrino experiments plus the KamLAND long-baseline reactor neutrino experiments, are shown in Figure 2.9. As can be seen, to a large extent the global constraint in Δm_{12}^2 comes from the reactor anti-neutrino measurement, while the constraint on θ_{12} comes from the SNO results.

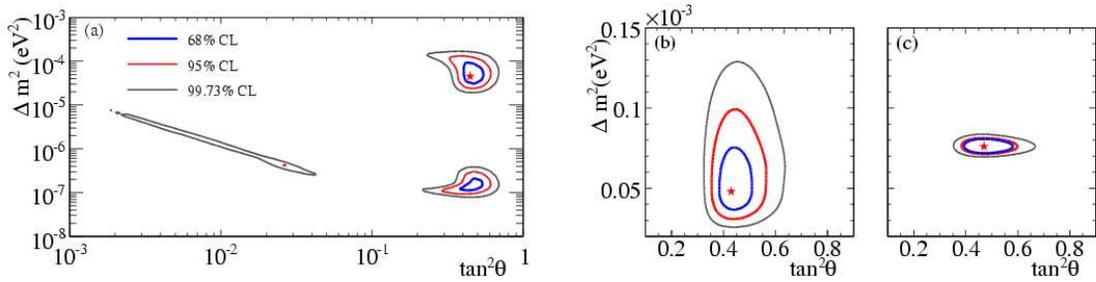


Figure 2.9: The current experimental constraints on the solar neutrino mixing parameters from (a) SNO alone, (b) all solar neutrino experiments, and (c) solar neutrino experiments + KamLAND. Figure from [21]. Note that the axis ranges are different in (a) than in (b) and (c). In (c) alone, the inner blue contour is drawn at the 90% C.L. rather than the 68% C.L. indicated in the legend. This figure includes the information from the NCD results described in Section 3.9.

Chapter 3. The SNO Neutral Current Detector Phase

During Phase III of the SNO experiment, the neutrons produced in the heavy water volume by neutral current neutrino interactions were primarily detected by an array of proportional counters called the “Neutral Current Detectors” or the “NCDs.” In this Chapter, the NCDs and their operation during the third phase of SNO are described, and the previously published NCD results outlined.

3.1 The Neutral Current Detectors

The Neutral Current Detectors were ^3He proportional counters which were deployed in the SNO heavy water volume during the third phase of the experiment. The NCDs detected neutrons through the $n+^3\text{He} \rightarrow p+^3\text{H}+764 \text{ keV}$ neutron capture interaction, which has an extremely large (5333 b) cross section. This large cross section meant that even when deployed in a relatively sparse array the NCDs provided the dominant neutron detection channel during Phase III of SNO.

The construction, deployment, and operation of the NCD array and the analysis of the NCD results are described in detail in [17], [21], and especially in [28]. The treatment here largely follows from those works.

As shown in Figure 3.1, a SNO NCD consisted of a 2-3 m long, 5.08 cm diameter thin walled nickel tube which was filled to 2.5 atm with an 85:15 (by pressure) mixture of ^3He and CF_4 . A $50\mu\text{m}$ diameter copper anode wire to collect the ionization electrons produced by the passage of charged particles through the gas ran down the centre of each NCD, passing out through the endcaps via fused-silica high-voltage feed-throughs. At the time of their deployment into the SNO detector, 3 to 4 of these “counters” were laser welded end-to-end, producing 9-11m NCD “strings.” An open-ended

inductive delay line adding ~ 90 ns to the transit time of reflected pulses was added to the bottom of each string. Below the delay line, a vectran line attached the NCD to two acrylic spheres which were used to attach the NCDs to pre-installed acrylic anchor positions on the bottom of the SNO acrylic vessel. NCD read-out and high-voltage supply were provided via single cables which ran from the top of each string along the upper inside surface of the SNO acrylic vessel and out the neck.

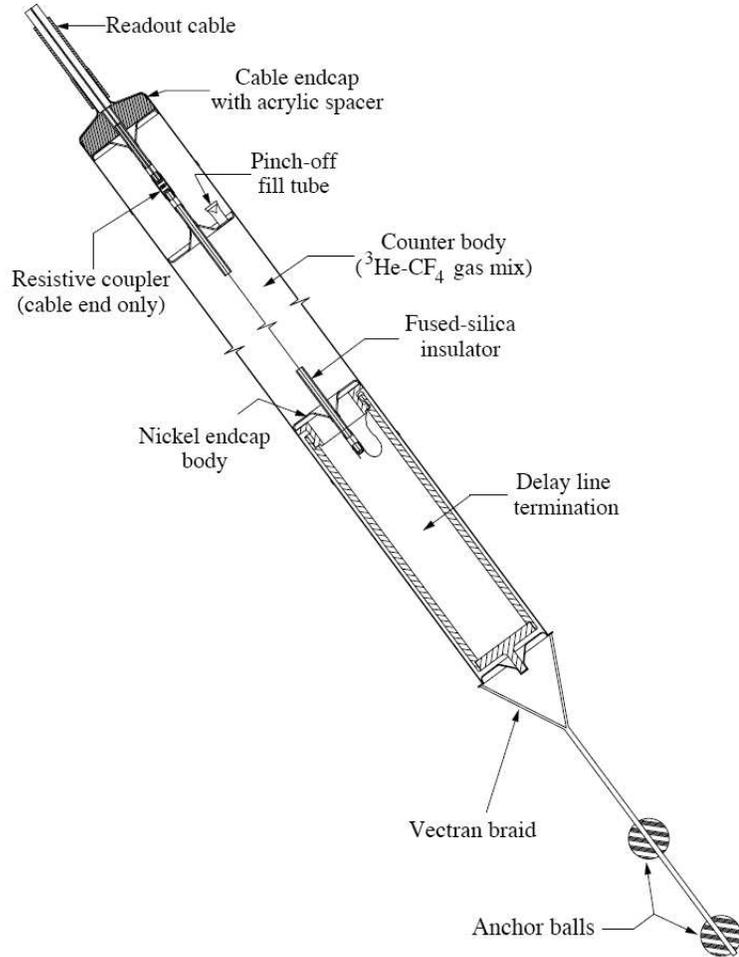


Figure 3.1: A schematic view of an NCD string. In this image the body of the NCD is greatly compressed. Figure from [17].

A total of 36 ^3He NCDs totalling 398m were deployed on a 1m grid within the SNO detector. The positions and names of the NCD strings are shown in Figure 3.2. In addition, 4 NCDs filled with ^4He in place of the ^3He were deployed to allow the study of non-neutron backgrounds.

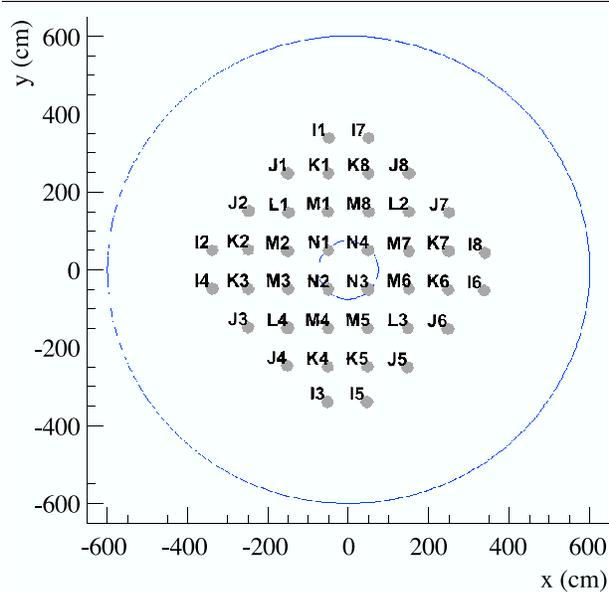


Figure 3.2: A cross section of the SNO detector at the equator showing the position of each NCD string. The ^4He strings were I2, I3, I6, and I7. The smaller circle at the centre shows the position of the neck of the acrylic vessel. Figure from [28].

3.2 Pulse Production in the NCDs

In the NCD analysis, two main observables were important: the total charge collected by the NCD for each event, and the “pulse shape” of each event (which was an oscilloscope trace of the charge leaving the NCD as a function of time). In order to understand these observables, it is necessary to understand the basic ideas behind pulse production in the NCDs.

As ionizing particles like the energetic proton and triton (^3H nucleus) produced from neutron capture on ^3He passed through the NCD gas, they ionized the gas to produce electron-ion pairs. These “primary” electron-ion pairs were distributed along the track of the primary particle and, on the timescales relevant to pulse shape creation, were produced essentially instantaneously. When the NCDs were live, the anode wires were held at 1950V relative to the NCD body, so that the electrons drifted towards the anode wire while the positive ions drifted away from it. The time required for a primary electron to drift to the anode was dependant on the radius at which the electron was produced; it required about $3.5 \mu\text{s}$ for an electron to drift across the entire radius of the NCD, and

this time decreased non-linearly with radius. Thus, the time profile of the arrival of charge at the anode wire gave a mapping of the radial profile of the primary ionization density.

The arrival of a charge signal at the anode wire resulted in identical pulses propagating up and down the anode wire. The downward-going component reflected from the bottom of the delay line and arrived at the top of the NCD delayed by ~ 100 ns relative to the upward going pulse. For very narrow pulses, the prompt and delayed signals were easily distinguished; for longer pulses the prompt and delayed signals substantially overlapped. The time delay between the prompt and reflected pulse varied with the additional distance travelled by the reflected pulse and could thus, in principle, be used to extract the z-position of an event within an NCD.

In order to record and make use of this pulse shape information, it was necessary that the charge read-out of the NCDs be fast compared to the spread in the electron times-of-arrival. In the NCDs, the physical limit on the readout time is set by the ion drift speed in the gas. This is because the majority of the charge collected by the anode was trapped as “image charge” by the positive ions produced close to the wire by the multiplication avalanches (described below). The image charge was released (and could hence be recorded) only as the ions drifted away from the anode. The total drift time for an avalanche ion to reach the outer cathode wall was several milliseconds, resulting in a decaying “ion tail” on each of the recorded pulses. However, the amount of trapped image charge decreased relatively quickly during the first part of the ion drift ($\sim 50\%$ was released in the first $2 \mu\text{s}$ [28]), so the radial charge deposition information was accessible to the analysis in the pulse shape.

A final important factor in pulse production by a physics event in an NCD was the space charge effect. In order to produce a detectable signal from the relatively small number of primary electrons created by a typical physics event in the NCD gas, the NCDs were run at high enough voltage such that within about $100 \mu\text{m}$ of the anode wire the electric field was strong enough that a drifting electron obtained sufficient energy between subsequent collisions with gas molecules to produce secondary ionization. As each of these secondary electrons could also produce additional ionizations, an “avalanche” of ionizations occurred; as a result, each primary electron resulted in the creation

of ~ 200 secondary electrons, and hence a MeV-level primary ionization of the gas (resulting in a few tens of thousands of primary electron-ion pairs) could produce a detectable charge on the anode wire. As discussed earlier, the cloud of positively charged ions produced by an electron avalanche drifted away from the anode wire quite slowly and beneath the cloud, the electrons trapped as image charge resulted in a reduced electric field.

Depending on the geometry of a physics event in the NCD, it was possible that primary electrons from different parts of the ionization track impinged on the same part of the anode wire at different times. In this case, the electrons which arrived at the anode later passed through the clouds of “space charge” ions produced by the avalanches of the earlier electrons and experienced the reduced electric field beneath. This meant that the later primary electrons, and any secondary avalanche electrons that they produced, gained less energy between collisions and were less likely to produce additional ionizations. As a result, the “gas gain” (i.e. the average number of secondary electrons produced per primary electron) of the later primary electrons was reduced by the space charge effect. This had two major effects on the analysis. First, it meant that the amount of charge collected by the NCD was not exactly proportional to the energy deposited in the gas; for alpha particle tracks perpendicular to the anode wire (which maximizes the space charge effect), for example, the amount of collected charge was approximately 50% less than the amount collected for an equivalent alpha travelling parallel to the wire. Second, the pulse shapes of high space charge events were noticeably affected, with the amplitude of the later part of the pulse shape significantly suppressed.

Taking all of these factors together, a typical neutron capture in the NCDs resulted in the collection by the anode wire of approximately 0.5 pC of charge over the period of a few microseconds. This signal is small but nonetheless detectable. To further illustrate the information available in the NCD pulse shapes, some representative neutron pulse shapes are shown in Figure 3.3, and some example alpha pulse shapes are shown in Figure 3.4.

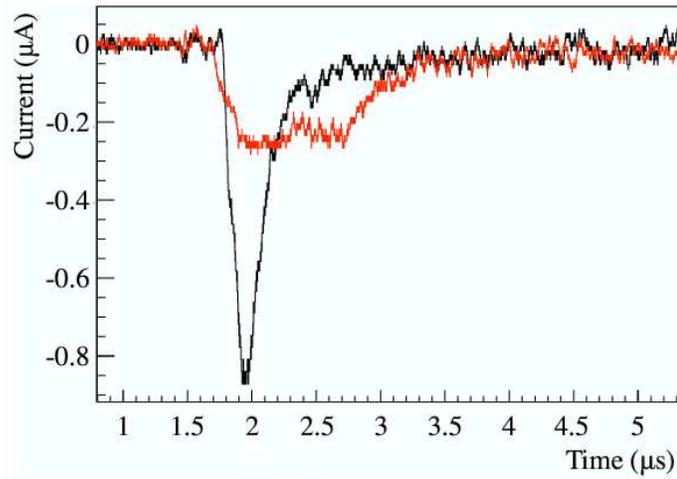


Figure 3.3: Example neutron pulse shapes taken from a neutron calibration run. The narrower pulse was produced by a neutron capture which resulted in the proton and triton tracks begin nearly parallel to the anode wire, so all of the charge was collected in a relatively short time. The wider pulse resulted from an interaction where the (back-to-back) proton and triton tracks were nearly perpendicular to the anode wire. The ionization signal arrived at the anode over the period of a few microseconds, and two bumps corresponding to the Bragg peaks (which are increases in ionization density which occurs near the end of charged particle tracks) of the proton and triton energy loss profiles are visible. Figure from [28]

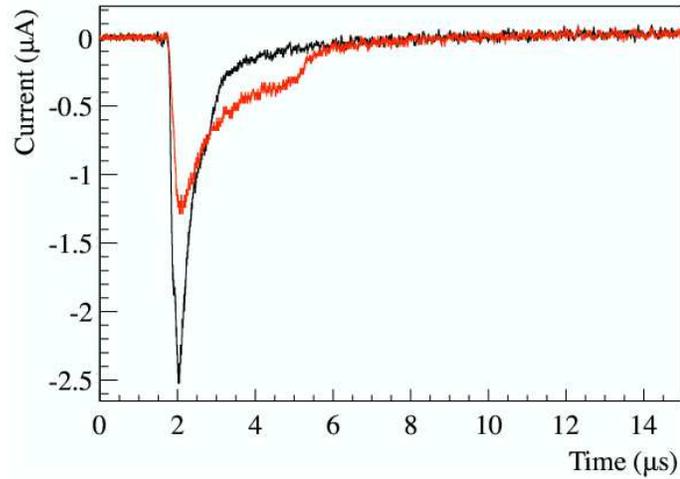


Figure 3.4: Example alpha pulse shapes taken from a ^4He string. The narrower pulse is a pulse produced by a track which was essentially parallel to the wire, so all of the charge was collected in a relatively short time. The slow decay after $3\mu\text{s}$ is the “ion tail” resulting from the slow release of electrons trapped as “image charge” on the anode wire as the positive ions drift away from the anode. The wider pulse is a pulse where the alpha track was essentially perpendicular to the anode. As the alpha originates from the inside surface of the nickel tube, the Bragg peak occurs nearest the wire. Therefore the charge from the Bragg peak arrives at the anode first, resulting in a “forward-peaked” pulse shape; the charge from the rest of the pulse continues to arrive over the next $3\text{-}3.5\mu\text{s}$. In fact, the effect of the Bragg peak is exaggerated because the derivative of electron drift time with respect to radius increases with radius; thus even a flat ionization profile would result in a “forward peaked” pulse shape, although to a lesser extent than what is seen here. Figure from [28].

3.3 NCD Data Acquisition

In order to make use of the pulse shape information described above while still being able to record events at sufficient rates to allow high-rate calibrations (and to be sensitive to possible neutrino bursts from Galactic supernovae), the NCD read-out electronics system was divided into two streams. The NCD signals passed through current pre-amplifiers, which converted the current signal into a voltage, and were then split into two parallel streams: the “shaper” stream and the “MUX” stream.

The “shaper-ADC system” was a system of shaping amplifiers, one for each NCD string, that measured the total charge collected by the NCD for each event. The shaper-ADC system maintained a rolling integral of the voltage from each pre-amplifier. This integral was then differentiated, with zero crossings of the derivative used to define maxima in the integrated charge and hence candidate events. A level-crossing discriminator accepted or rejected candidate events based on their total maximum integrated charge (i.e. the value of the charge integral at the time of zero-crossing of the derivative) and the maximum of the charge integral for accepted events was read-out by a 12 bit ADC. A “shaper event” was then produced in the SNO data stream. The shaper system was designed to be fast enough (with a dead-time of approximately $235 \mu\text{s}$ per event) to record data from higher-rate calibration sources and to provide sensitivity to possible supernova neutrino bursts.

The “multiplexer” or “MUX” stream recorded the event pulse shapes. Each NCD channel entering the MUX system was monitored by a level-crossing discriminator. Signals from channels that crossed the threshold were passed, with an added voltage offset, to one of four logarithmic amplifiers. Up to 12 strings were connected to each logarithmic amplifier (the signal from a given string always went through the same amplifier); if multiple strings on the same amplifier had simultaneous MUX triggers, the signals were summed before entering the amplifier. The output from the logarithmic amplifier was then passed to one of two digital oscilloscopes (the choice between the scopes was based on a toggle bit which in principle alternated events between the scopes). At the scope, the event, from about 1500 ns before the trigger time to a total of $15 \mu\text{s}$ was digitized at 1 GHz, and read out through a GPIB interface. The MUX trigger and scope information then became part of

a “MUX event” in the SNO data stream (with a time stamp based on the MUX trigger time). The readout of the oscilloscopes was relatively slow (~ 0.6 s), limiting the throughput of the two-scope system to about 1.2 Hz. This was adequate for regular running, but too slow to record pulse shapes for all events during calibration running. In the event that the scope toward which the toggle bit directed a log-amp signal was busy, a “partial MUX event,” containing the MUX trigger information but no oscilloscope trace was produced. As the MUX trigger had a much smaller dead time than the scope (~ 1 ms), these partial MUX events allowed MUX-shaper trigger coincidences to be used as a method of data cleaning even on higher rate data.

The dual shaper-MUX system meant that each physics event produced two events in the SNO data stream. These events were associated with one another offline based on the relative timing of the events (the MUX trigger occurred 8-15 μ s before the shaper trigger due to the shaper integration time). Only “correlated events” with both MUX and shaper triggers were considered in the analysis; for neutrino data it was further required that each MUX event have an associated scope trace.

3.4 NCD Signals and Backgrounds

The NCDs were built to detect the passage through the NCD gas of the energetic proton and triton produced by neutron captures on the ^3He ; they were, however, sensitive to the passage of any ionizing particle. This meant that background events, in particular alpha particles produced by the decay of radioactive isotopes in the NCD components, were also present in the NCD signal. It was originally hoped that the neutron and alpha events would be distinguishable based on differences in their pulse shapes. A good deal of effort has gone into realizing this separation, and some promising results have been obtained. This pulse shape separation will not, however, be discussed here. Instead, as in the previously published NCD analysis ([21], discussed below), the neutron and alpha signals will be separated based on their energy spectra. Therefore, it is the energy spectra of the neutrons, alphas, and other backgrounds which are of principal interest to the current discussion.

3.4.1 Neutron Signal

The capture of a neutron by a ${}^3\text{He}$ nucleus produces a back-to-back proton-triton pair with a total kinetic energy of 764 keV,



In the NCDs, the proton and triton resulting from a neutron capture propagated through the NCD gas depositing, in many cases, the entire 764 keV into the gas. It was also possible, however, that either the proton or the triton (or, with a tiny geometric probability, both), impacted the nickel body or the anode wire of the NCD before depositing all of their energy¹, and hence did not deposit their full energy in the gas. This geometric effect resulted in a rather unique energy spectrum, shown in Figure 3.5, which extended downward from 764 keV.

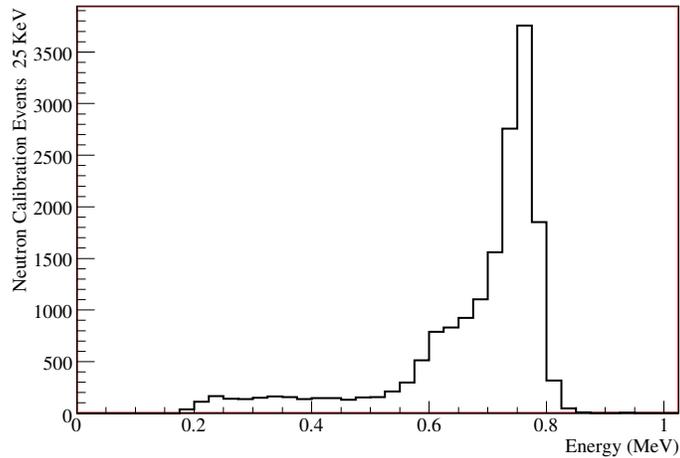


Figure 3.5: The energy spectrum produced in the NCDs by a source of neutrons uniformly distributed throughout the SNO heavy water volume. The “full energy peak,” where the full energy of both the proton and triton are deposited in the gas, is visible at 764 keV. The total absorption of the proton’s energy through collision with the NCD wall results in the 191 keV shoulder, while the total absorption of the triton’s energy results in the shoulder at 573 keV. This higher energy shoulder is significantly enhanced by the downward smearing of the full energy peak by the space charge effect.

¹ The propagation distance of the proton in the NCD gas was ~ 1 cm, while the triton travelled ~ 3 mm.

3.4.2 Alpha Backgrounds

The majority of the physics signals detected by the NCDs were alpha particles emitted by radioactive isotopes in and on the NCD components. The energy of a typical alpha particle was several MeV, and hence most of these alphas were significantly higher in energy than the neutron peak. However, some alpha particles were degraded by passing through nickel or copper before entering the NCD gas, while others passed through the NCD gas on a chord and struck the nickel wall before depositing all of their energy. Both of these effects could cause alpha particles to appear in the neutron region of the energy spectrum (in the final data there were about five times as many alpha events as neutron events in the neutron energy window). The alpha backgrounds in the NCDs can be grouped into four main categories:

1. **Bulk Alphas:** ^{238}U and ^{232}Th chain decays in the bulk of the nickel NCD bodies produced alphas with energies as high as 8.8 MeV. Due to degradation of the alphas as they propagated out of the bulk of the nickel and into the NCD gas, and to space charge effects, a significant fraction of these events fell into the neutron energy region.
2. **Surface Alphas:** During NCD construction, radioactive daughters of ^{222}Rn in the laboratory air plated out on the inside surface of the NCD nickel. The ^{222}Rn decay chain includes ^{210}Po which emits a 5.3 MeV alpha and which is supported by ^{210}Pb with a 22-year half-life. In order for these surface ^{210}Po alphas to appear in the neutron energy window, it was necessary for them to re-enter the nickel wall after traversing a short chord through the NCD gas. This relatively unlikely track geometry meant that a much smaller fraction of surface events than bulk events appeared in the neutron energy region. However, a higher overall number of surface events made the total contributions of the two classes in the neutron energy region similar.
3. **Wire Alphas:** Surface and bulk events originating from the anode wire were also detected by the NCDs and had noticeably different pulse shapes than alphas originating from the nickel. Even surface wire alphas could be degraded by passing on a chord through the (very thin)

anode wire before entering the gas, and wire alphas were much more likely to experience significant space charge effects than were alphas from the nickel. Wire alphas are thought to have been responsible for approximately 2% of the alpha events recorded by the array [29]. However, a lack of “chord tracks” through the NCD gas meant that a much smaller fraction of wire alphas appeared in the neutron energy region compared to the other types of alpha.

4. **Endcap Alphas:** Quartz feed-through tubes at either end of each NCD counter created, by design, 2.5 cm long dead regions from which charge could not be collected (this avoided complications due to the non-uniform electric fields at the ends of the counters). However, alpha tracks originating from within the dead endcap regions and propagating into the live region (or vice versa) had significantly different pulse shapes and energy spectra than “normal” alphas. Under the assumption that the alpha activity was uniformly distributed within the NCDs, endcap alphas were expected to form only a few percent of the total alpha population.

The alpha backgrounds observed in the NCDs were dominated by surface and bulk alpha events from the nickel, while the wire and endcap alphas were thought to be present at the few percent level [29]. The energy spectra of these dominant components, as simulated by the SNOMAN Monte Carlo software package (described in Section 3.8.1), are shown in Figure 3.6.

3.4.3 Other Physics Backgrounds

Apart from neutrons and alphas, the NCDs were also sensitive to high energy electrons, both those produced by beta decays and those produced by the Compton Scattering of γ -rays. Because electrons are relatively weakly ionizing, the path length required for an electron to deposit enough energy in the NCD gas to pass threshold was of the order of the length of the counters, which in turn meant that electrons were expected to provide a small component of the background at the very low end of the neutron energy region which would decrease quickly with energy. Due to the multiple changes in direction possible due to the electrons scattering both within the gas and off the nickel walls, however, electrons had the ability to produce a wide variety of pulse shapes.

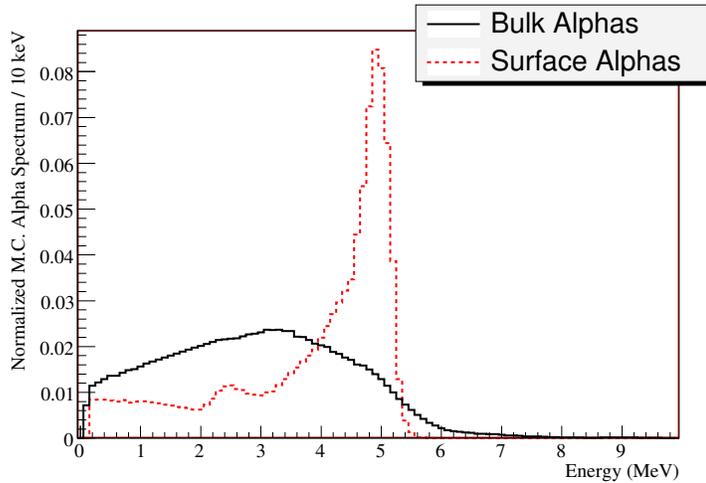


Figure 3.6: The energy spectra of bulk and surface alphas, as predicted by the SNOMAN Monte Carlo. The energy spectra of bulk ^{238}U and ^{232}Th chain decays are very similar in the neutron energy region, so these chains were not treated individually in most of the NCD analysis. The energy of the ^{210}Po peak is degraded slightly by space charge effects. The peak in the surface alpha spectrum around 2.5 MeV is composed of full energy polonium events with geometries that induced significant space charge effects.

3.5 Non-Physics Backgrounds and Data Cleaning

In addition to the physics backgrounds discussed above, the NCD system also recorded a large number of spurious events. These pulses took many forms and were generated by a variety of sources, not all of which are understood. Examples of these spurious pulses are shown in Figure 3.7.

3.5.1 Data Cleaning

In order to remove as many of these spurious events as possible from the analyzed data set, an extensive series of data cleaning cuts was developed.

The simplest data cleaning cut, which removed a large fraction of the spurious triggers, was the requirement that an analyzed event be “correlated,” (i.e. have both a shaper and a MUX trigger which occurred within the normal coincidence time window). This cut removed electronics noise which occurred upstream of the divergence of the MUX and shaper data streams, very narrow events which did not have enough integral charge to trigger the shaper, and very wide events which did not have sufficient amplitude to trigger the MUX.

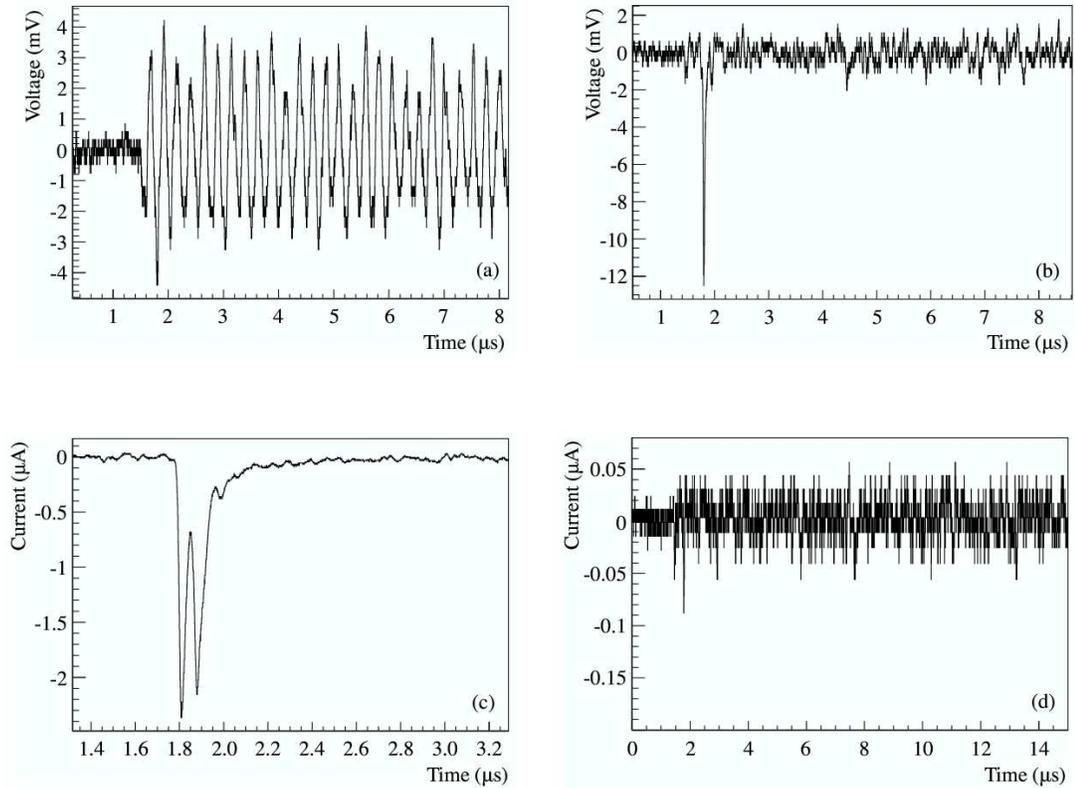


Figure 3.7: Examples of spurious NCD events: (a) shows oscillatory noise, likely from electrical pick-up in the NCD electronics, (b) shows what is likely a micro-discharge from within the pre-amplifier, (c) shows a breakdown event from the delay line, and (d) shows a “flat trace” trigger from baseline fluctuations. Signals thought to originate from the NCD are plotted in μA , while those which are thought to originate in or above the pre-amplifier are plotted in mV. Figures from [28].

The next level of data cleaning removed events that occurred as part of a burst (MUX channel bursts, shaper channel bursts, and NCD events in conjunction with PMT bursts were all removed), and those events which occurred in conjunction with the passage of a muon through the detector.

Finally, the oscilloscope traces of the remaining events² were passed through two complementary data cleaning paths, one of which examined the pulses in the time domain and one which operated in frequency domain. These data cleaning paths cut events which were too narrow to be physics events, oscillatory events with multiple zero crossings, and “flat trace” events, among others. The frequency-domain cuts also checked for the presence of the decaying ion tail in the latter part of the pulse, which should be a signature of all physics events.

The combination of all of the data cleaning cuts was very effective at removing the spurious instrumental backgrounds, with a neutron sacrifice that was measured to be less than 1%. The effect of the data cleaning cuts on the NCD energy spectrum is shown in Figure 3.8. Unfortunately, unlike the data cleaning cuts in the PMT data, it proved impossible to develop for the NCDs a method to set a numerical limit on the residual contamination of spurious events in the data after the data cleaning cuts were applied. Therefore, while it is believed that any residual contamination is negligible, this cannot be conclusively demonstrated.

3.5.2 NNNAs

In fact, there are several examples in the data of unexplained spurious events which pass the cleaning cuts in significant numbers. The NCDs upon which these events (which were colloquially called “Non-Neutron Non-Alphas” or “NNNAs”) were observed are listed below, along with brief descriptions of the NCD behaviour. The energy spectra of some of these strings are shown in Figure 3.9.

- **J3 (String 26):** This was the first string upon which NNNAs were identified, based on a significant distortion in the neutron-region energy spectrum (shown in Figure 3.9(a)). The events on this string were originally claimed to be “missing their ion tail,” and were in some

² For neutrino data it was also required that all events have a scope trace, so that the following checks could be performed. For high-rate calibration data, however, most events did not have scope traces due to the relatively long scope readout time. Therefore, in calibration runs the scope-based data cleaning checks were not required.

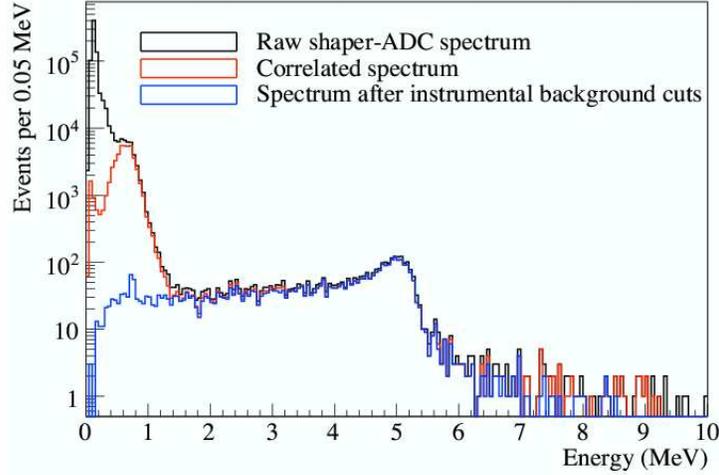


Figure 3.8: The effect of the data cleaning cuts on the NCD energy spectrum. Figure from [28]. The neutron full energy peak appears at 0.76 MeV. This Figure shows only a subset of the NCD data.

cases distinguishable by eye. Nonetheless the NNNA pulse shapes were similar enough to physics events that it proved impossible to reject these events on a pulse shape basis (in fact, pulse shape based analysis consistently classified these events as more similar to neutrons than to alphas). The rate of NNNA events on J3 was constant in time, raising the possibility that they may in fact have been produced by an as yet unidentified radioactive background.

- **N4 (String 0):** As in the case of J3, the N4 NNAs were identified based on a distortion in the energy spectrum of the string (see in Figure 3.9(b)). Unlike the J3 events however, the N4 NNAs are not uniform in time, occurring mostly in low rate bursts during a few runs.
- **K7 (String 8):** This string exhibited low rate bursts of events with pulse shapes similar to those seen on N4 [30]. As seen in Figure 3.9(c), however, the distortion in the K7 energy spectrum extends to higher energies than does the N4 distortion.
- **I7 (String 3):** This ^4He string exhibited high-rate bursts of shaper activity, some of which also resulted in MUX triggers. Some of these events had pulse shapes similar to the other NNAs that were identified. The energy spectrum for this string is not shown in Figure 3.9

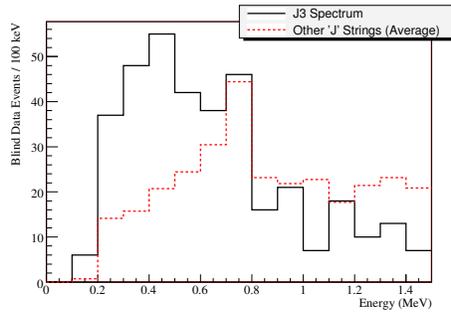
because the string was operated at different shaper thresholds which resulted in significant distortion in the low energy spectrum.

- **K5 (String 18):** This string was observed to produce events with pulse shapes similar to “J3-type” NNNA events. The energy spectrum of the string is not shown in Figure 3.9 because a leak in one of the connectors between two of the NCD segments resulted in a significant change in the gas gain in one counter. This string therefore has significant distortions in the neutron region energy spectrum even without the NNAs.
- **K2 (String 31) and M8 (String 1):** Both of these strings were observed (in hand scans of the NCD pulse shapes) to produce events with pulse shapes similar to “J3-type” NNAs.

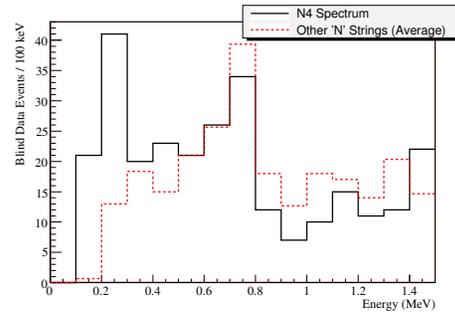
Searches for NNAs on strings besides those listed above were carried out by comparing the energy spectra (and the distributions of other pulse shape parameters) between the different strings [31]. These searches “detected” the previously identified NNNA strings, and no others. Additional searches for non-Poissonian event timing distributions [32], which correctly identified the known NNAs on K7, I7 and N4, also did not identify any other infected strings. For additional NNAs to have been present in the array and not detected by these searches, they would have to either have been very low rate or more-or-less uniformly distributed across the NCD array and fairly similar in spectrum to either the alpha or the neutron events. Again, however, while no additional NNAs were “discovered” by these searches, no limits on their possible number were produced either³. The negative results of these searches do, however, give some confidence that there was no gross effect from NNAs on strings other than those listed above.

Of the NNNA strings listed above, all but J3 and N4 were excluded from the analysis because of other observed defects (see Section 3.7). Although no causal relationships were determined between the presence of the NNAs and the other observed defects on these strings, it was deemed reasonable

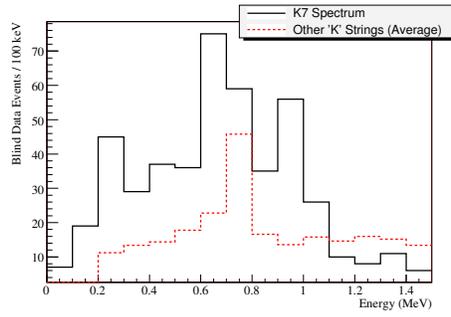
³ It is difficult in these types of searches to decide on the criteria for “detection” of an NNNA. In the burst search, for example, several strings (M8, K8, J6, N3, J3, and K2) showed non-Poissonian behaviour at 50% of the level of 0 and 8 [32], but below the criteria that had been set for a “detection.” It is equally difficult to define the necessary criteria on the “allowable” NCD behaviour that are necessary to set a limit on the possible number of NNAs. Limits *can* be placed on the NNNA contamination under specific assumptions about their uniformity both in the array and in time and their similarity to the alpha and neutron shapes [33, 34]. It was decided, however, that the assumption of any specific set of criteria was difficult to justify, and so these limits were not used in the published analysis.



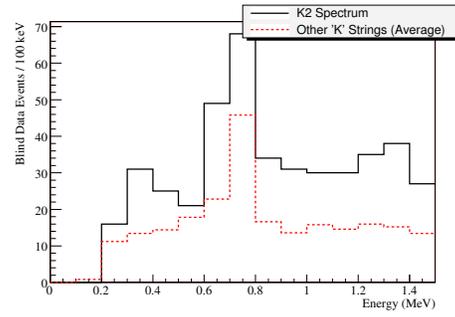
(a) J3 Spectrum.



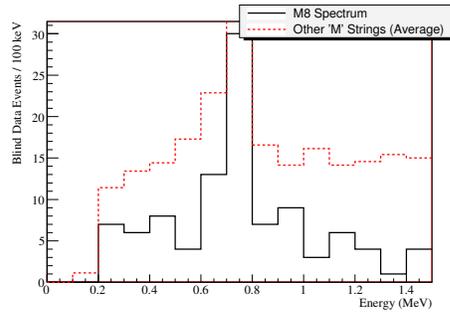
(b) N4 Spectrum.



(c) K7 Spectrum.



(d) K2 Spectrum.



(e) M8 Spectrum.

Figure 3.9: The energy spectra of some of the strings known to contain NNNAs compared to the average spectra of the other strings in the same NCD “ring” as the affected string (which were expected to have identical neutron signals). Note that K7, K2, and K5 were omitted from each other’s comparison spectra. Changes in the level of the alpha background above and beneath the neutron peak should be expected due to the different concentrations of alpha emitters in the different strings. The shapes of the alpha spectra, however, should be reasonably similar. The data shown are from the full NCD blinded data set (described in Section 4.3) with all of the data cleaning cuts applied, including some that were specifically designed to target NNNAs. In all cases, with the possible exception of M8, significant distortions appear to be present in the energy spectra of these strings.

to assume that the NNNAs were linked in some way to the observed defects and were therefore unlikely to appear on strings which did not display those defects. J3 and N4, on the other hand, were excluded from the analysis solely on the basis of their observed NNNAs. As the cause of these NNNAs was not known and they were only noticed based on their relatively high rates, there was the potential that similar behaviour may have occurred, undetected, on the strings that were included in the analysis. Therefore, in the previously published analysis [21] empirical PDF⁴ distributions representing the NNNA shapes observed on J3 and N4 were included in the fit to the data from the rest of the array.

The origin of the NNNA events remains unclear, although the different temporal behaviour and energy spectra of the different NNNAs suggest that they do not all have the same cause. The lack of evidence, in very careful searches, for similar events in the rest of the array is comforting, and means that any NNNAs present would have to have very specific distributions in time, space and energy. Still, the observation of these unexplained events on a significant portion of the array, coupled with the lack of a quantitative limit on the residual contamination in the rest of the array, lends some uncertainty to the prediction of the background shape in the rest of the array that was required in the signal extraction. This uncertainty was one of the primary motivations for the development of the robust signal extraction techniques which make up the bulk of the work in the SNO portion of this thesis, and which are described in the next chapter.

3.6 NCD Calibration

The response of the NCD array to physics events, especially to neutron captures, was measured and monitored through an extensive calibration program. In fact, the neutron energy PDF that was used in the ultimate signal extractions was drawn directly from calibration data.

⁴ “PDFs” or “Probability Distribution Functions” contain the assumed distributions of each signal and background population in each of the observed variables. In performing a fit to the data, these PDFs form the set of distributions that are compared to the data in each variable to determine the number of each class of signal and background events consistent with the observations.

3.6.1 Electronics Calibration

The response of the NCD electronics (the pre-amplifiers and higher) could be measured by injecting known waveforms into pulser inputs on each pre-amplifier. These “electronics calibrations” or “ECAs” were carried out weekly, with more extensive ECA scans monthly. The ECAs allowed electronics gain and linearity and the trigger thresholds to be monitored, and provided a means of measuring the amplification and offset parameters of the logarithmic amplifiers. The electronics parameters from a given ECA calibration were applied to the data between that calibration and the next; interpolation and back-calibration were not performed.

3.6.2 Physics Calibration

In addition to the ECAs, extensive neutron source calibrations were performed in order to study the response of the array to neutron events. Encapsulated sources ($^{241}\text{AmBe}$ and ^{252}Cf) could be moved in two two-dimensional planes (which contained the primary axes of the SNO detector) using the SNO calibration source manipulator system (described in [17]). Extensive scans designed to illuminate all of the NCD counters as uniformly as possible were carried out four times during the third phase of SNO, with shorter scans (to check array stability) every few months.

In addition, on two occasions, small amounts of ^{24}Na in the form of neutron-activated NaCl brine were mixed into the D_2O active volume. ^{24}Na beta decays to ^{24}Mg with a half-life of about 15 hours, releasing a 2.75 MeV γ -ray. These γ -rays had sufficient energy to photo-dissociate deuterons, and thus produced a very nearly uniform source of neutrons in the D_2O volume.

The physics calibrations allowed the absolute neutron capture efficiency and the neutron detection efficiency of the NCD array to be determined, both directly and through comparison with Monte Carlo. In addition, the physics calibrations allowed the shaper-ADC values to be converted to energies (the ADC values were scaled to place the neutron peak in neutron scans at 764 keV), and provided the neutron energy spectrum that was used as the neutron PDF in the NCD signal extraction.

3.7 Abnormal NCD Behaviour

Of the 36 ^3He strings deployed in the array, six strings were identified as behaving differently from the rest of the array in ways important to the analysis. As a result, the data from these strings, although not necessarily unrecoverable, were excluded from analysis to date. The affected NCDs were:

- **J3 and N4:** As discussed earlier, these strings exhibited significant numbers of “NNNA” events which passed the data cleaning cuts. These strings were omitted from the analysis on this basis.
- **K5:** This string developed a leak between one of its counter sections and the adjacent connecting volume, allowing ^3He to flow from the counter into this dead volume. This resulted in a gradual reduction of the gas gain of that counter, and hence a time-dependent distortion of the K5 neutron energy spectrum.
- **K2 and M8:** At various times during running, these strings showed behaviours consistent with the read-out electronics becoming electrically disconnected from the bulk of the NCD. When these strings were removed from the detector, investigation showed that the “resistive couplers” (springy connectors which were used to both electrically connect and provide impedance matching between the top of the NCD string and the bottom of the readout cable - see Figure 3.1) were only loosely connected to the NCD feed-throughs.
- **K7:** The energy calibration of this string was observed to be unstable during neutron calibrations.

In the rest of this thesis, data from “the NCD array” should be assumed to exclude these strings (and the ^4He strings) unless otherwise noted.

3.8 Understanding the Alpha Background

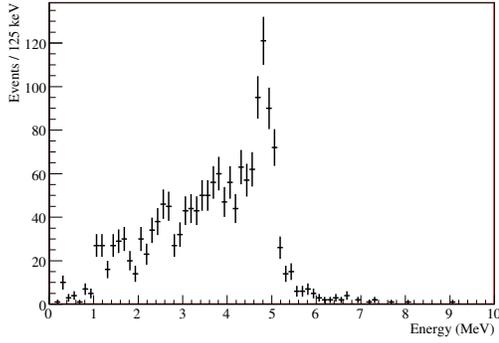
As mentioned earlier, the separation between the neutron signal and the (alpha dominated) background was primarily carried out by fits to the NCD energy spectrum. This meant that understanding the shape of the alpha energy spectrum was of critical importance. It was also one of the most challenging aspects of the analysis.

The in-situ data available from the ^4He strings to study the alpha background in the neutron energy region, while valuable, were quite limited. As can be seen in Figure 3.10, of the four ^4He detectors, only three produced useful data⁵. The energy calibration of the ^4He strings was also challenging, as they had no neutron peak to use for energy calibration. Instead, the energy scale and offset of each ^4He string was set by hand such that the ^{210}Po peak and the low energy threshold cut-off occurred at approximately the same energies as in the ^3He strings. It should also be noted that the three counters in I3 were not as well gain matched as in most strings, meaning that the alpha spectral features were broadened on this string. The main limitation of the ^4He data, however, was statistical; even with I7 included, there were only about 2600 events on the ^4He strings in the neutron energy window. More than two thirds of these events came from I6, which had an inordinately high level of ^{210}Po surface activity.

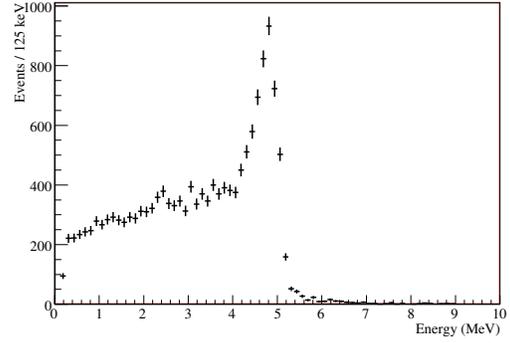
Attempting to determine the alpha energy spectrum in the neutron energy region from the ^4He strings is thus quite challenging. Even considering only surface and bulk contributions (and neglecting the other families of alphas described in Section 3.4.2), the statistical uncertainties on the extracted alpha PDFs were too large to be particularly useful.

Attempts were made to supplement the alpha data set using additional ^4He NCD strings that were run (in air) in the underground control room, and by taking data with a test NCD (which was made of stainless steel and had a feed-through to allow the insertion of sources) at Los Alamos National Laboratory. Both of these supplementary sources, however, ran in very different background environments, and with data acquisition and read-out systems which differed significantly from the

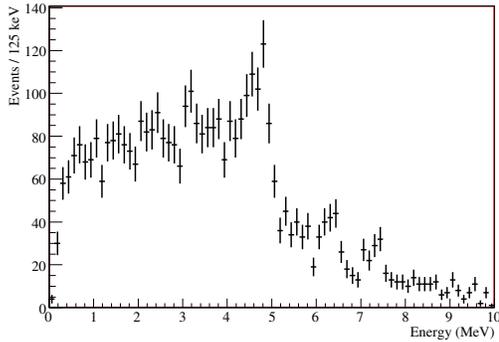
⁵ The shaper threshold of I7 was significantly raised to suppress continuing shaper bursts. This meant that events were no longer recorded in the neutron energy region, so the data from this string were excluded from the analysis.



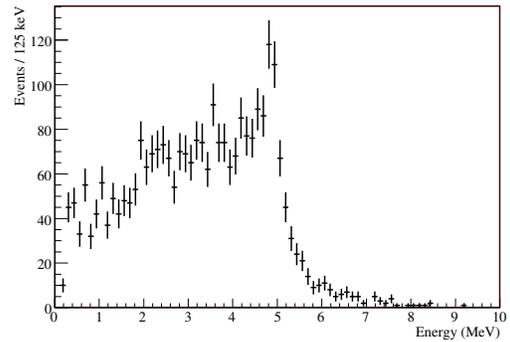
(a) String I7



(b) String I6



(c) String I3



(d) String I2

Figure 3.10: The ^4He alpha spectra. The energy scale of each string has been hand corrected so that the ^{210}Po peak and the low energy threshold cut-off occur at approximately the same energies as in the ^3He strings. The spectrum of I7 is distorted as a result of the string being operated with significantly raised shaper threshold for the majority of its operating period. Multiple ^{210}Po peaks can be seen in I3 as a result of gain mismatch between the three NCD counters which made up the string.

deployed array. These electronics changes, especially the differences in impedance matching in the electronics chain and differences in the reflection timing, affected the pulse shapes significantly, and hence made application of the data cleaning algorithms to this supplementary data difficult. The distribution of the alpha activity in the NCDs used to take this supplementary data was also quite different from the assumed distribution of alpha activity in the deployed array. Therefore, while this additional data were useful in studying the alpha behaviour, it could not be used to directly provide a background PDF for the signal extraction.

3.8.1 Alpha Monte Carlo

In order to produce background PDFs for the signal extraction, then, a good deal of effort was put towards modelling pulse production by the NCDs. The NCD pulse shape Monte Carlo, which was implemented as part of the SNOMAN [17] Monte Carlo package, is described in detail in [29] and [35].

The goal of the Monte Carlo was to model not only the energy spectra of the different classes of alphas and neutrons, but their pulse shapes as well. As a result, the Monte Carlo tracked not just the energy deposition of the primary particles in the NCD gas, but the full pulse production process including the propagation of the primary electrons to the anode wire, the avalanche amplification and charge collection, and the propagation of the resulting pulse through the electronics chain.

In the absence of the space charge effect, the energy spectrum of a given class of physics events would be relatively easy to model, as the charge recorded by the shapers would simply be proportional to the energy lost by the primary particle in the NCD gas. In the presence of a significant space charge effect, however, the charge recorded by the shapers, and hence the reconstructed energy of the event, depends on the details of the propagation of the primary particle, the drifting of the charge through the NCD gas, and the charge multiplication.

As described in Section 3.2, the “space charge effect” refers to the reduction in gas gain experienced by an electron which impinges on a portion of the anode wire which is surrounded by a cloud of positive ions produced by the multiplication of an earlier electron. In attempting to

model this process, the spatial distribution of the ion clouds generated by the electrons from a given segment of the primary track is critically important. It is also difficult to predict. In the Monte Carlo simulation, the lateral diffusion of the primary electrons as they drifted to the anode wire was modelled using known diffusion constants and drift speeds, while the lateral and radial spread of the ion density of an avalanche were parameters which were tuned using high energy alpha and neutron calibration data. For each segment of the anode wire, the space charge was assumed to form a complete cylinder around the wire, which slowly drifted outward according to the ion mobility. Within the cylinder of space charge, the change in the electric field could easily be calculated, and from this the change in the gas gain was determined using the Diethorn model [36].

This parameterization of the space charge effect was approximate. In reality, for example, the space charge cloud from a given electron avalanche was expected to have an RMS azimuthal spread of $\sim 20^\circ$ [37], and hence not wrap completely around the anode wire⁶. There may also have been electron capture effects as the primary electrons drift through the ions produced by the other primary ionizations and through any space charge ion charges present. On average, however, the Monte Carlo treatment of the space charge effect was thought to be sufficient for spectral modelling under standard running conditions [35].

Having developed a model of charge collection in a single event, the other required Monte Carlo input in simulating the alpha spectrum is the physical distribution of alphas within the NCDs. As little information is available about the distribution of alpha activity along the length of the NCDs, both the surface and bulk alpha distributions were assumed in the Monte Carlo to be uniform along the length of the NCDs⁷. Wire alphas and endcap alphas were included in the simulations, with their numbers estimated through fits to the alpha spectrum above the neutron energy region. The ratio of surface to bulk alpha activity was also determined via string-by-string fits to the high energy alpha spectrum. In these same high-energy alpha fits, it was found that the depth distribution of

⁶ This would result in a stronger but more local space charge effect, which might, for example, cause a significant change in tracks headed straight towards the wire.

⁷ The possibility of spectral distortions in the neutron energy region due to potentially enhanced activity in the endcap regions was considered unlikely, while spectral distortions on single strings due to local depositions of activity at fixed depth would likely have been detected by the NNNA searches described earlier.

the alpha activity within the nickel had a significant effect on the alpha spectrum. In particular, it was found that a better description of the data could be obtained by introducing both surface and bulk activity distributions which decayed exponentially with depth [29]. This distribution is difficult to motivate physically, but was accepted based on the improvement in the high energy fits.

As described above, the challenges in simulating the alpha energy spectrum are many. Nevertheless, based on comparisons with the ^4He data, with higher energy alphas on the ^3He strings, with neutron calibration data and with the supplementary alpha data described earlier, the Monte Carlo was deemed to provide a sufficiently accurate description of the data to be used in the signal extraction above 0.4 MeV⁸. Systematic uncertainties in the Monte Carlo alpha energy spectrum were set by varying the tuned parameters (the electron avalanche parameters described earlier, the ion mobility, which was measured using calibration neutron data, and the surface to bulk alpha activity ratios) and parameterizing their effect on the neutron region energy spectrum as polynomials⁹. Due to their questionable physical origin, the differences between the exponentially distributed surface and bulk activities and the more expected uniform bulk and purely surface distributions were also taken as systematic uncertainties. Finally, due to the much larger uncertainties in the modelling of individual pulse shapes (which will not be discussed here), the entire action of the scope-based data cleaning cuts on the Monte Carlo data was taken as a systematic uncertainty. The final Monte Carlo alpha spectrum, along with the range of variation encompassed by the systematic uncertainties described above, is shown in Figure 3.11.

3.9 Published NCD Results

The first, and to date only, SNO publication to make use of the NCD data was the “NCD PRL” [21]. This analysis consisted of a joint fit of the NCD-phase NCD and PMT data, with the neutral current neutron flux constrained to be equal in the two detection channels.

⁸ Below this point the systematic uncertainties in the Monte Carlo became difficult to reliably estimate. It is also in this low energy region where other backgrounds like electrons and other data cleaning leakage events were most likely to occur.

⁹ The alpha systematic uncertainty polynomials ranged from zeroth order for the avalanche gradient and ion mobility systematic uncertainties to as high as third order for the Po depth distribution and data cleaning systematic uncertainties.

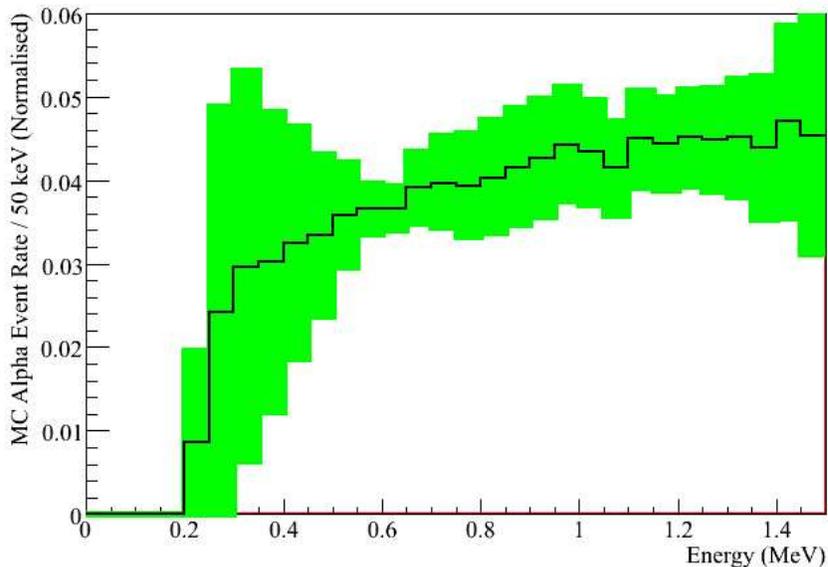


Figure 3.11: The Monte Carlo alpha spectrum near the neutron energy region. The line shows the central value Monte Carlo, while the band shows the range consistent with 1σ variations in the Monte Carlo systematic uncertainties.

The PMT-side signal extraction was carried out in a manner similar to the signal extractions in previous phases [20, 19],¹⁰ with the fitted observables being event energy, event radius, and the cosine of the angle between the reconstructed event direction and the position of the Sun. The PMT PDFs were taken from the Monte Carlo, and the energy spectrum of the charged current events (and hence the energy spectrum of the ^8B neutrinos) was not constrained in the fit, but was extracted from the data.

As discussed earlier, the only observable used in the NCD fits was event energy. The spectrum of events from the 2005 ^{24}Na spike calibration was taken to be the neutron PDF, while the Monte Carlo alpha PDF described in the previous section was used to describe the alpha backgrounds. The previously discussed NNNAs on J3 and N4 were represented in the fits by skew Gaussian curves. As was also described earlier, the Monte Carlo systematic uncertainty estimates were considered reliable above 0.4 MeV, so the NCD fits were carried out over the 0.4 - 1.4 MeV range.

The fit minimization was carried out using a Markov-Chain Monte Carlo type optimization

¹⁰ Note, however, that the optical characterization of the detector and the event position, direction, and energy estimators had to be updated to take into account the light shadowing of the NCDs.

procedure, which allowed the most important systematic uncertainties to vary within the fits. For the NCDs, these systematic uncertainties included the energy scale and resolution, the systematic uncertainties in the Monte Carlo alpha spectrum, and the widths and amplitudes of the two NNNA PDFs. These “floated” systematic uncertainties were constrained, via penalty terms in the likelihood, where external information was available (the energy scale, for example, was constrained to be within 1% of the nominal value, while the alpha spectrum systematic uncertainties were constrained to $\pm 1\sigma$). A complete description of the NCD PRL Markov-Chain Monte Carlo signal extraction procedure and results can be found in [38, 39].

The inclusion of the Monte Carlo systematic uncertainties and the NNNA PDFs created a good deal of freedom in the low energy part of the NCD background energy spectrum; indeed, the best fit included a large number of NNAs (571_{-175}^{+162}), while distorting the alpha spectrum significantly downward at low energy, as shown in Figure 3.12. Nevertheless, the neutron energy spectrum was sufficiently unique that the fitted number of NCD neutrons was relatively stable, with 983_{-76}^{+77} neutral current neutrons (in addition to the 185_{-22}^{+25} background neutrons expected based on external constraints (see Section 5.4)) being fitted. This corresponds to a total solar neutrino flux of $5.54_{-0.31}^{+0.33}(\text{stat})_{-0.34}^{+0.36}(\text{syst})\times 10^6\text{cm}^{-2}\text{s}^{-1}$ (see Section 5.3 for information about the conversion of the detected number of neutral current neutrons into neutrino flux).

On the PMT side, 267_{-22}^{+24} neutral current neutron captures, 1867_{-101}^{+91} charged current events and 171_{-22}^{+24} elastic scattering events were fitted above the 6.0 MeV (kinetic) energy threshold. Under the assumption that the solar neutrino flux is entirely ν_e , these correspond to fluxes of $1.67_{-0.04}^{+0.05}(\text{stat})_{-0.08}^{+0.07}(\text{syst})\times 10^6\text{cm}^{-2}\text{s}^{-1}$ and $1.77_{-0.21}^{+0.24}(\text{stat})_{-0.10}^{+0.09}(\text{syst})\times 10^6\text{cm}^{-2}\text{s}^{-1}$, respectively. The PMT-side data and signal extraction is described more fully in [21].

As discussed earlier, the most useful value in studying neutrino oscillations is the ratio of the charged current flux, ϕ_{CC} , and the neutral current flux, ϕ_{NC} . For the NCD phase results, this ratio is

$$\frac{\phi_{CC}}{\phi_{NC}} = 0.301 \pm 0.033. \quad (3.2)$$

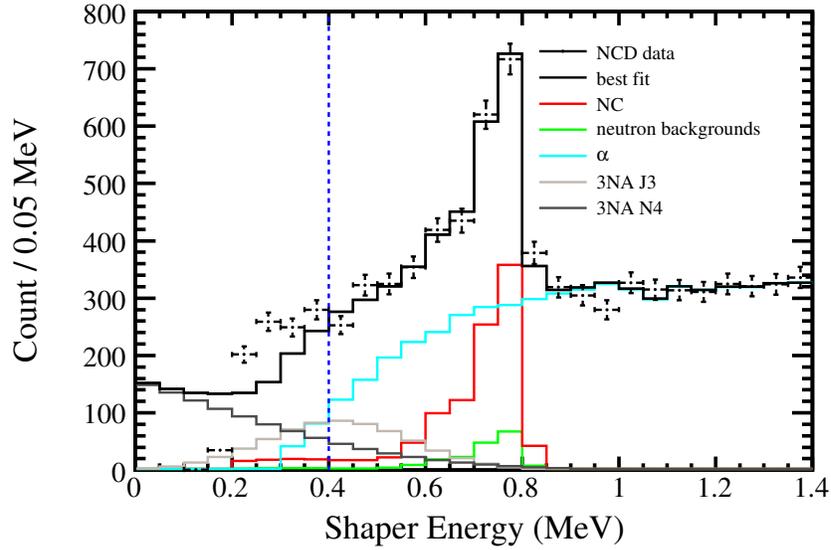


Figure 3.12: The fitted NCD energy spectrum from the NCD PRL fits. The freedom in the background PDFs can be seen to have resulted in a significant, and unphysical, downward distortion in the alpha energy spectrum which is taken up by an increase in the number of fitted NNNAs. This freedom in the background spectrum does not, however, result in any particular instability in the fitted number of neutrons.

The constraints on the solar neutrino oscillation parameters that resulted from including the NCD results in the global neutrino oscillation analysis were shown in Figure 2.9, earlier, while comparisons of the NCD phase flux results to the results from the previous phases of SNO were shown in Figure 2.7.

Chapter 4. Robust Determination of the NCD Neutron Number

As described in the previous Chapter, one of the most significant challenges in extracting the NCD neutron number was constraining the shape of the background energy spectrum. The systematic uncertainties in the Monte Carlo alpha spectrum were relatively large and difficult to reliably evaluate and the possibility of unexpected backgrounds, both physics and non-physics, could not be completely ruled out. In this Chapter, two alternative signal extraction procedures, the “polynomial method” and the “overlap method,” which were developed with the aim of successfully extracting the number of neutrons from the NCD energy spectrum while making minimal assumptions about the background energy spectrum, are introduced. With these robust techniques, the NCD energy fit can be extended to include the full neutron energy region, which both slightly increases the statistics in the extracted neutron number and includes in the fit additional information which may help to constrain the background behaviour under the neutron peak. The price of reducing the dependence of the fit on prior knowledge of the background shape is expected to be an increase in the uncertainty in the extracted number of neutrons; the benefit is a fit that is less susceptible to biases in the event that unexpected background shape did occur in the data. In addition to what is presented here, the history and development of the overlap and polynomial methods are also described in [40] and [41].

4.1 The “Polynomial Method”

A standard approach to the extraction of a well-defined and quickly varying signal from a relatively smoothly varying background is the use of a polynomial to model the background in the fit. In the case of the NCD energy spectrum, we must modify the polynomial to account for the low energy threshold cutoff (this is accomplished by multiplying the polynomial by an inverted Fermi Function,

the position and steepness of which are floated in the fit). The use of a polynomial to represent the background in the fit requires only three inputs from the experimenter: the shape of the signal (in this instance the neutron energy PDF, which is obtained from calibration data), the order of the polynomial to be used, and the parameter range over which the fit should occur.

4.1.1 Polynomial Order

Due to the uncertainty in the background spectrum, the order of the polynomial needed to accurately represent the background in the NCD fit is not known *a priori*; it is, however, of fundamental importance. Using a polynomial of insufficient order will in general result in fit biases, while using a polynomial with more degrees of freedom than needed unnecessarily inflates the fit uncertainties. The solution adopted here is to perform fits to the data with polynomials of different orders and to use information from the fits themselves to infer the best polynomial to use.

The obvious criteria to use in this type of *a posteriori* model selection is the “goodness of fit” between the models and the data. It is easily demonstrated, however, that adding an additional degree of freedom to a polynomial fit to data cannot result in a worse fit and almost always leads to a better fit. It is necessary, therefore, to determine whether the improvement in the fit resulting from an additional degree of freedom is large enough to justify that degree of freedom or not.

A standard method of making this determination involves the “likelihood ratio,”

$$\lambda = \frac{L_i}{L_j}, \tag{4.1}$$

where L_i and L_j are the likelihoods of the best fits of model i and model j , respectively, to the data. In the event that one of the models is a nested subset of the other (for example a 4th order polynomial can be considered a nested subset of a 5th order polynomial) and the model with fewer degrees of freedom is already a good description of the data, the logarithm of the likelihood ratio is well approximated by a χ^2 distribution [42]:

$$P(-2 \ln(\lambda)) \approx \frac{2^{-\nu/2}}{\Gamma(\nu/2)} (\lambda)^{\nu/2-1} e^{-\lambda/2}, \tag{4.2}$$

where ν is the number of additional parameters in model j compared to model i . Further, in the

regime of Gaussian statistics (within which the NCD energy spectrum lies for a broad range of possible bin widths), the log-likelihood is equivalent to the χ^2 with an additive constant [43]. It is therefore possible to replace the likelihood ratio in Equation 4.2 with a χ^2 difference:

$$P(\chi_i^2 - \chi_j^2) \approx \frac{2^{-\nu/2}}{\Gamma(\nu/2)} (\lambda)^{\nu/2-1} e^{-\lambda/2}, \quad (4.3)$$

where χ_i^2 and χ_j^2 are the χ^2 's of the best fits of models i and j to the data (here model j is assumed to have the larger number of free parameters).

Equation 4.3 thus provides a means by which the statistical significance of the decrease in χ^2 resulting from the use of a higher order polynomial can be evaluated. For example, the requirement that the addition of a single additional degree of freedom reduces the fit χ^2 by 1.0 or more has approximately a 68% chance of rejecting a superfluous parameter. Indeed, this is the origin of the standard “rule of thumb” in parameter significance testing that a parameter must produce a $\Delta\chi^2$ of at least 1.0 to be justified. In selecting the best polynomial to use, however, we cannot simply add orders one at a time until the $\Delta\chi^2$ of one addition fails the comparison criteria; if testing superfluous parameters this would work fine, but when testing polynomials of insufficient order to correctly describe the data it is possible that adding one order does not improve the fit much (or even at all) while adding the next order can make a significant improvement in fit¹. If the “true” polynomial order required to properly represent the data is high, this would lead to a significant probability that one of the χ^2 comparisons would fail the $\Delta\chi^2$ test, and hence that subsequent testing would be stopped, before an appropriate description of the background is reached.

To avoid this problem, it is necessary to scan all interesting polynomial orders and compare their fit χ^2 's. In the beginning, the lowest (0^{th}) order polynomial is assumed to correctly describe the background. The χ^2 of the 0^{th} order fit is sequentially compared with the χ^2 's of the higher order fits, and for each comparison if the $\Delta\chi^2$ between that fit and the reference (0^{th} order) fit is larger than the target $\Delta\chi^2$ for that number of additional degrees of freedom, the higher polynomial is taken

¹ As an example, consider the 3-point data set $y(1)=0$, $y(2)=1$, $y(3)=0$ with unit uncertainties (in fact, the demonstration works for any non-zero uncertainty, as long as the uncertainties on all three bins are equal). The best-fit χ^2 's for fits to 0^{th} and 1^{st} order polynomials can easily be determined analytically to both be equal to $\frac{2}{3}$. Thus, the first order term acts as a superfluous parameter, and tests only comparing adjacent χ^2 's would have a good chance of truncating the series at 0^{th} order, while a second order polynomial (of course) fits the data perfectly.

to be the correct description of the background. That higher order polynomial fit is then used as the reference fit in subsequent comparisons. The “reference fit” after all of the comparisons have been completed is taken to be the correct description of the background. The $\Delta\chi^2$ targets are set such that in a single comparison between models differing by n degrees of freedom, the chance of superfluous parameters being rejected is independent of n (i.e. based on Equation 4.3, the $\Delta\chi^2$ targets are set by locating a certain confidence level in the relevant χ^2 distribution). Changing the $\Delta\chi^2$ target with n is necessary because although the average $\Delta\chi^2$ for n added superfluous parameters is n , the cumulative probability distribution changes. For example, one superfluous parameter produces a $\Delta\chi^2$ less than 1.0 68.3% of the time, while for two extra parameters the 68.3% upper confidence limit on $\Delta\chi^2$ is 2.3, not 2.0. Thus, setting $\Delta\chi^2$ targets based only on a certain change in χ^2 per degree of freedom becomes significantly biased towards falsely accepting large n comparisons.

Unfortunately, the multiple comparisons between different polynomial orders necessary to select the “best” order in the manner described above incurs a significant trials factor in the probability that superfluous parameters are accepted in at least one of these comparisons. As can be seen in Table 4.1, when several superfluous polynomial orders are tested sequentially, the chance that none of them are accepted is significantly less than the chance that any one of them is rejected. This is because the average improvement in χ^2 when a superfluous degree of freedom is added is 1.0, even though 68% of the time the improvement is less than 1.0. Thus, relatively large jumps in χ^2 occur with some frequency, and with several subsequent parameter additions the chance of hitting one of these jumps and hence falsely accepting too high a polynomial order can become large. Figure 4.1 shows the same effect in tests of the polynomial order choice procedure using toy Monte Carlo energy spectra.

To mitigate the trials factor, and hence avoid an excessive probability of falsely accepting superfluous parameters, two complimentary approaches can be taken. First, the number of tested polynomial orders can be reduced, which lowers the number of χ^2 comparisons made, and hence the probability of a false acceptance. This is a reasonable thing to do, as there is no reason to

| | | | | | |
|-----------|------|------|------|------|------|
| n | 1 | 2 | 5 | 7 | 10 |
| P_{rej} | 0.68 | 0.57 | 0.43 | 0.39 | 0.35 |

Table 4.1: The probability, P_{rej} , that after sequentially testing nested models with up to n superfluous degrees of freedom, none of the superfluous models have been accepted by the χ^2 comparison tests. The acceptance criteria on χ^2 were set such that any single comparison between models differing by n superfluous degrees of freedom would have a 68% chance of rejecting the superfluous parameters. The calculation was carried out using repeated Monte Carlo sampling of the single-order χ^2 distribution, in order to properly calculate the conditional probability that a comparison with n superfluous parameters fails the χ^2 improvement test, given that the comparisons with all lower numbers of unnecessary parameters had also failed.

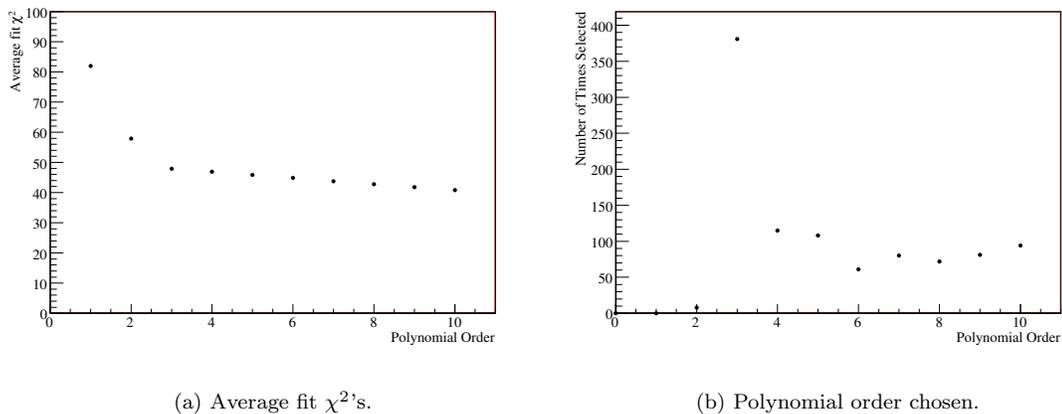


Figure 4.1: The results of testing the polynomial order selection algorithm against Monte Carlo data. Multiple high-statistics (~ 100 times the number of events in the NCD energy spectrum below 1.5 MeV) data sets were generated with the underlying third-order distribution $90 + 35x - 20x^2 + 10x^3$. The average χ^2 of the fit of each polynomial order to the fake data sets is shown in (a). As can be seen, as the polynomial order is increased, the fits improve dramatically with each order below the “true” order (the average χ^2 for order zero, ~ 5000 , is off-scale); above the true order, the χ^2 improves by approximately the expected 1.0 per superfluous parameter. The frequency with which a given polynomial order was selected (again based on χ^2 selection criteria that would reject a given set of n superfluous parameters 68% of the time in a single test) is shown in (b). As can be seen, the correct order is chosen 381 times, within error of the 390 times expected from Table 4.1; the remaining choices are approximately uniformly distributed over the higher orders, which is also as expected.

expect that a tenth order polynomial should be required to fit the background (the Monte Carlo and the data from the ${}^4\text{He}$ strings suggest that at most a third order polynomial should be needed to describe the alpha spectrum); we therefore truncate the polynomial order tests at sixth order, which reduces the chance of a false selection while still giving the fit significant ability to adapt to unexpected background shapes.

The second approach to reducing the trials factor is to increase the change in χ^2 required for additional parameters to be accepted. Table 4.2 shows the probability that the “correct” model is still accepted after consecutively adding and testing six superfluous parameters under different χ^2 acceptance criteria. As expected, the probability of correctly rejecting all six superfluous parameter tests increases as the probability of rejecting any single superfluous parameter is increased. It is therefore possible to find, through iteration, a single test superfluous parameter rejection probability, and hence a set of $\Delta\chi^2$ targets, that will produce any desired overall probability that multiple superfluous parameters tests are all rejected. We thus set the $\Delta\chi^2$ targets by choosing the total, rather than the single-trial, probability of accepting superfluous parameters².

| | | | | |
|---------------------------------------|------|------|------|------|
| Single test rejection probability. | 0.68 | 0.90 | 0.95 | 0.99 |
| Required $\Delta\chi^2$ for $n = 1$. | 1.0 | 2.7 | 3.8 | 6.6 |
| Six test rejection probability. | 0.42 | 0.77 | 0.88 | 0.97 |

Table 4.2: The change in the probability that χ^2 tests of six consecutive superfluous parameters are all rejected as the χ^2 target is changed. The $\Delta\chi^2$ targets are set based on the probability, shown in the first row, of rejecting any n superfluous parameters in a single comparison. The resulting $\Delta\chi^2$ target for $n = 1$ is shown in the second row to illustrate the dependence between the selected superfluous acceptance probability and the $\Delta\chi^2$ target (Note that, as described in the text, the $\Delta\chi^2$ target does not scale linearly with n). The final row shows the probability that consecutive tests of 6 nested superfluous parameters are all rejected. As can be seen, the probability of correctly rejecting multiple tests increases quickly as the single test rejection probability increases.

² There is an additional subtlety in implementing this algorithm: although we test up to polynomial order six, not all of these polynomials represent superfluous parameters; only those above the “correct” order are superfluous. Therefore, to give the correct probability of accepting superfluous parameters, the $\Delta\chi^2$ targets should be set based on the “correct” polynomial order (e.g. if a 4th order polynomial is chosen as the best description of the background, the $\Delta\chi^2$ targets should have been computed based on two, rather than six, sequential tests of superfluous parameters). Therefore the $\Delta\chi^2$ targets used to choose the correct polynomial order depend on which order is chosen. To resolve this co-dependence, the best polynomial order is chosen iteratively, first using $\Delta\chi^2$ targets appropriate for six superfluous parameter tests, and subsequently with $\Delta\chi^2$ targets set based on the polynomial order chosen in the last iteration. This is continued until a self consistent best polynomial and $\Delta\chi^2$ target are found (in practice, it was relatively rare for the selected best order to change beyond the first iteration).

The remaining question, then, is the total probability with which superfluous parameters should be rejected. Obviously, from the point of view of minimizing the uncertainty on the final neutron number, superfluous parameters should be avoided as completely as possible. As the $\Delta\chi^2$ requirements are made more stringent, however, the selection criteria will accept as sufficient models which more poorly describe the data; as an example, Table 4.3 shows the number of times insufficient, correct, and excessive polynomial orders were chosen in Monte Carlo experiments with different $\Delta\chi^2$ targets³. Thus in choosing the $\Delta\chi^2$ target, it is necessary to balance the desire to have the lowest possible polynomial order with the risk of bias associated with selecting a model which more poorly describes the background. The bias risks are likely relatively low, as those “incomplete” models must provide a reasonably good description of the data in order to be selected. Nevertheless, the conservative thing to do is to accept a relatively high probability of accepting superfluous parameters (and hence having unnecessarily high uncertainties). In keeping with the spirit of the “rule of thumb” that in a single parameter test a model is accepted if there is a “1-sigma” chance or less that the attendant improvement in χ^2 is due to statistical fluctuations alone, we choose here to assign the $\Delta\chi^2$ targets so that the total probability that all superfluous parameters be rejected is 0.683.

| | | | | |
|---|------|------|------|------|
| Total probability to reject superfluous models. | 0.68 | 0.90 | 0.95 | 0.99 |
| Frequency of selection of insufficient polynomial order | 39 | 148 | 242 | 411 |
| Frequency of selection of correct polynomial order | 624 | 743 | 707 | 565 |
| Frequency of selection of excessive polynomial order | 337 | 109 | 51 | 24 |

Table 4.3: The results of polynomial selection in Monte Carlo experiments similar to those described in Figure 4.1, under different $\Delta\chi^2$ targets. These results are only illustrative, as the probability of selecting polynomials of insufficient order will depend on the precise shape of the background. Still, it can be seen that while increasing the $\Delta\chi^2$ targets does indeed reduce the number of times superfluous parameters are accepted, it also increases the probability of choosing an insufficient description of the background.

³ Note that the excessive number of selections of excessive polynomial orders in the 0.99 column of Table 4.3 is due to numerical precision in the determination of the $\Delta\chi^2$ determination. In order to make the multi-trial selection probability of a superfluous model 0.99, the single trial probability had to be closer to 1.0 than the precision of the Monte Carlo numerical integration technique used.

Uncertainty Correction

A final, unexpected, effect that must be taken into account in the polynomial method is a slight broadening of the pulls distribution⁴ by the polynomial selection procedure. In Monte Carlo ensemble tests, it was observed that if many Monte Carlo data sets were fitted with a fixed polynomial order, the pulls distribution was correct (independent of polynomial order), while after performing the polynomial selection procedure inherent in the polynomial method, the distribution of pulls was slightly wider than expected. Investigation showed that this broadening was the result of a correlation, shown in Figure 4.2, between large changes in χ^2 and large changes in pull between fits of different polynomial orders. Thus, in selecting for large $\Delta\chi^2$ we tend to select for large changes in the neutron number (which can be either towards or away from the true number). This additional scatter is responsible for the observed widening of the pulls distribution.

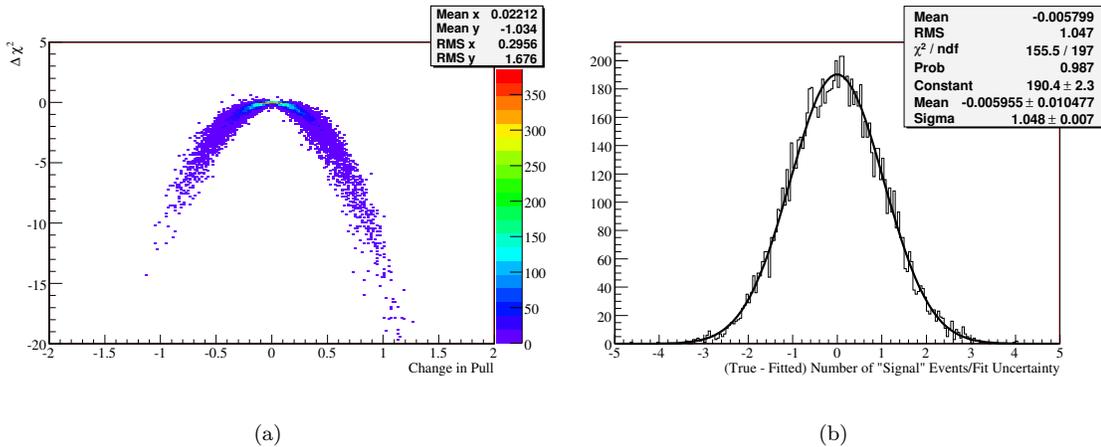


Figure 4.2: The relationship between the $\Delta\chi^2$ and the change in pull in the number of extracted “signal” events between fits with subsequent polynomial orders is shown in (a). A correlation is clearly present, and in selecting for large $\Delta\chi^2$'s, large changes in the pull are also selected, resulting in the slightly widened pulls distribution shown in (b). The Monte Carlo data sets used to produce this plot contained on average 7500 alpha background events below 1.5MeV (distributed according to the ^4He alpha spectrum) and 1000 signal events distributed according to “NNNA shape 1” in Figure 4.6.

Further investigation showed that the broadening effect was independent of neutron and back-

⁴ The “pull” of a fit is defined as the difference between the true and fitted value of a parameter divided by the fit uncertainty. If the uncertainties are properly estimated, the distribution of pulls for many fits should be described by a Gaussian curve with a standard deviation of 1.0.

ground statistics within at least a factor of ten of the expected statistics in the data, and did not depend on the signal shape. As a result, the width of the pulls distribution in Figure 4.2 was taken to be representative of the effect, and the uncertainties in subsequent polynomial fits were increased by 5% to compensate for it.

Polynomial Choice Summary

To summarize the polynomial fit procedure: the data are repeatedly fitted to the (calibration-data derived) neutron energy PDF, with different order polynomials (between 0th and 6th) used to represent the background. For each polynomial, the χ^2 between the data and the best fit model is recorded. The results of the different fits are then sequentially compared based on their χ^2 's, with a fit with more degrees of freedom being preferred only if the associated $\Delta\chi^2$ is larger than the target $\Delta\chi^2$ for a comparison over the appropriate number of degrees of freedom. The polynomial which passes the χ^2 comparison test relative to all preceding orders, and relative to which all subsequent orders fail, is taken to be the best description of the background. The $\Delta\chi^2$ targets are set such that the total probability that the selected polynomial order is higher than necessary is 0.683.

4.1.2 Binning

The binning of the data has very little effect on the polynomial fits; in Monte Carlo studies, no change was seen between fits with bin widths as small as 6.25 keV and as large as 50 keV. As bin widths get very large, the features of the neutron spectrum become less distinct and the ability to distinguish some background shapes may be degraded. As bin widths become narrower, the increase in the number of bins means that the fits take longer to perform. In what follows, the data are binned in 25 keV bins; this gives several bins across the neutron peak without unnecessarily slowing the fits.

4.1.3 Fit Range

The fit range of the polynomial fits is also not critical. Extending the fit range, especially to higher energies, has the potential to help constrain the background in the neutron energy region. This

constraint is only helpful, however, if the background is smoothly varying; if the background varies quickly, extending the fit range will only mean that higher order polynomials are needed to correctly describe the background. Based on the Monte Carlo and data from the ^4He strings, the alpha spectrum is expected to be relatively smooth below about 1.5 - 2.0 MeV. A “step” feature due to surface alpha events which are nearly perpendicular to the wire and hence experienced very large space charge effects occurs at energies above this. In order to avoid this known effect, the upper limit of the polynomial fit range is set at 1.5 MeV. The low end of the fit range is set by the requirement that the number of events in each bin must be high enough for Gaussian statistics, and hence the χ^2 tests that are used to select the polynomial order, to apply. Typically, Gaussian statistics are assumed to be a good approximation provided that all bins have more than 5-10 events [43]. In NCD data binned at 25 keV, this criteria is satisfied for bins above 0.15 - 0.2 MeV. Therefore, the lower limit of the fit range is chosen to be 0.2 MeV.

4.2 The “Overlap Method”

The Overlap Method for extracting the number of neutrons from the NCD signal was designed to capitalize on the unique energy spectrum of the NCD neutron signal. The idea is that by defining a comparison metric which describes the similarity between two distributions, one can determine the maximum and minimum number of neutrons consistent with the data under the condition that the residual background should not be too similar to the neutron shape. In this way, it proves possible to extract a neutron number while supplying no information at all about the background, other than that it is different from the neutrons (although as we shall see, adding information about the background does reduce the uncertainty in the extracted number of neutrons).

4.2.1 Method Description

The comparison metric used in the overlap method is a simple one called the “overlap,” O , which defines the similarity between two distributions L and K :

$$O_{(K,L)} = \frac{1}{\text{BW}} \frac{\sum_i K_i L_i}{\sum_i K_i \sum_i L_i}, \quad (4.4)$$

where X_i is the entry of distribution X in bin i , the sum spans the energy range over which the comparison is to be made, and BW is the bin width of the histograms (dividing by the bin width makes the overlap value independent of binning). In what follows, the overlap is always defined relative to the neutron energy PDF, P^N , so the second subscript on the overlap is dropped and the second normalization sum is superfluous as the PDF is pre-normalized. The overlap then becomes

$$O_L = \frac{1}{BW} \frac{\sum_i L_i P_i^N}{\sum_i L_i}. \quad (4.5)$$

Now, consider a data set comprised of neutron events (“ N ”) and background events (“ B ”). This later class will be dominated by alphas, but could also include instrumental events, etc. The overlap of the data (“ D ”) will be then be given by

$$\begin{aligned} O_D &= \frac{1}{BW} \frac{\sum_i D_i P_i^N}{\sum_i D_i} \\ &= \frac{1}{BW} \frac{\sum_i (N_i + B_i) P_i^N}{\sum_i (N_i + B_i)} \\ &= \frac{1}{N_N + N_B} (N_N O_N + N_B O_B). \end{aligned} \quad (4.6)$$

We see, therefore, that the linear definition of the overlap is useful in that it allows the overlap of a distribution to be determined from the sum of the overlaps of the constituent distributions.

In using the overlap as a signal extraction tool, we subtract a certain number of neutrons, n , from the data (using the assumed neutron PDF shape) and calculate the overlap value for the residual, $O_R(n)$. Using Equation 4.6, this is simply

$$O_R(n) = \frac{1}{N_N + N_B - n} ((N_N - n)O_N + N_B O_B). \quad (4.7)$$

We see that $O_R(n = N_N) = O_B$, so that if we knew O_B we could determine the number of neutrons in the data by determining the n for which $O_R(n) = O_B$ (this is actually algebraically calculable from O_D , O_N , O_B and $(N_N + N_B)$). Unfortunately, the same poor understanding of the background shape that necessitated consideration of the overlap method in the first place denies us precise knowledge of O_B . This, however, is the situation for which the overlap method was designed. In lieu of exact knowledge of the background shape and hence of O_B , the method provides a convenient way to map a prior probability distribution in O_B into a probability distribution in n .

Without a complete physical understanding of the background, the most straightforward probability distribution that can be assigned to the background is a flat distribution which is felt to span the physically reasonable set of possible background spectra. The upper and lower overlap limits, O_{lim}^{upper} and O_{lim}^{lower} , define the points at which it is felt, based on the set of backgrounds that have been observed, experience in modelling the response of the NCDs, and experience in working with similar types of detectors in other contexts, that the background looks too much like the neutron shape to be physically reasonable⁵.

Having established upper and lower limits on the overlap⁶ (and noting that higher allowed overlaps will be associated with lower estimates of the number of neutrons, and vice versa), Equation 4.7 can be used to determine the maximum and minimum number of neutrons consistent with the data such that the residual overlap does not exceed those limits:

$$n_{lower} = N_N + N_B \frac{O_B - O_{lim}^{upper}}{O_N - O_{lim}^{upper}} \quad (4.8)$$

and

$$n_{upper} = N_N + N_B \frac{O_B - O_{lim}^{lower}}{O_N - O_{lim}^{lower}}. \quad (4.9)$$

We see that if $N_B = 0$, the confidence region collapses to N_N , the true neutron number⁷, that for $O_B = O_{lim}^{upper}$ or $O_B = O_{lim}^{lower}$ we have $n_{lower} = N_N$ or $n_{upper} = N_N$, respectively, and that $n_{lower} \leq N_N \leq n_{upper}$ provided $O_{lim}^{lower} \leq O_B \leq O_{lim}^{upper}$, as desired. The size of the confidence interval is given by

$$n_{upper} - n_{lower} = N_B \frac{(O_{lim}^{upper} - O_{lim}^{lower})(O_N - O_B)}{(O_N - O_{lim}^{upper})(O_N - O_{lim}^{lower})}. \quad (4.10)$$

Overall, then, a signal extraction based on overlap limits contains all of the features one might expect from a signal extraction:

⁵ The upper limit, which can be thought of in terms of a neutron-shaped ‘‘bump’’ appearing on an otherwise innocuous background is more intuitive than the lower limit. However, a neutron-shaped depression in an otherwise smooth background (which leads to the lower limit on the overlap) is just as physically unreasonable as a neutron-shaped bump.

⁶ The actual determination of these limits is discussed in Section 4.2.5

⁷ Note, however, that this situation is not practically achievable due to statistical fluctuations; these will be discussed later, but it should be noted at this point that statistical fluctuations about the assumed neutron shape are counted as part of the background.

- The extracted neutron confidence interval contains the true number of neutrons provided that the overlap of the background falls within the overlap limits.
- A more precise knowledge of the background (reflected in a smaller spread in the overlap limits) leads to a more precise determination of the number of neutrons.
- Both a larger amount of background and a greater similarity between the background and neutron shapes (reflected in a smaller difference between the overlap limits and the neutron overlap) lead to greater uncertainty in the number of neutrons.

Finally, we note that simple confidence limits on the residual overlap yield simple confidence limits on the number of neutrons. Thus, while we expect that the true number of neutrons lies within the region we determine, we have no reason to expect that the true value is more likely to occur near the middle of the region than it is to occur near either of the ends.

As can be seen in Equation 4.10, the size of the confidence region for the extracted number of neutrons depends linearly on the number of background events. One might expect, therefore, that the precision of the neutron number determination could be significantly improved by subtracting from the data the large component of the alpha background that is understood based on modelling and the ^4He data. Separating the background into N_A alpha events (by definition these are the alpha events that follow the predicted energy spectrum; any alphas which do not fall into the next category) and N_o “other” events we find that

$$O_R(n) = \frac{1}{N_N + N_o + N_A - n} ((N_N - n)O_N + N_o O_o + N_A O_A) \quad (4.11)$$

from whence comes

$$n_{upper} = N_N + N_A \frac{O_A - O_{lim}^{lower}}{O_N - O_{lim}^{lower}} + N_o \frac{O_o - O_{lim}^{lower}}{O_N - O_{lim}^{lower}} \quad (4.12)$$

(with n_{lower} similar) and

$$n_{upper} - n_{lower} = \frac{(O_{lim}^{upper} - O_{lim}^{lower})}{(O_N - O_{lim}^{upper})(O_N - O_{lim}^{lower})} (N_A(O_N - O_A) + N_o(O_N - O_o)). \quad (4.13)$$

Subtracting the alpha background (using the alpha shape determined from the Monte Carlo and/or the ^4He strings and normalizing it by fitting above the neutron energy region, for example) is equivalent to setting $N_A = 0$ in these equations. We therefore recover Equations 4.7-4.10 with $N_B \rightarrow N_o$ and $O_B \rightarrow O_o$. Because $N_o < N_B$, for the same overlap limits, and under the assumption that O_o is not significantly nearer O_N than O_B was, subtracting the alpha component does indeed improve the precision of the neutron extraction very significantly⁸.

Unfortunately, subtraction of the alpha background brings to the fore an inherent difficulty with the normalization scheme used in the above derivation; there is nothing to prevent the denominator in Equation 4.11 from being small or negative. When a large number of alphas are present this does not pose a problem, but once the alphas are subtracted ($N_A \rightarrow 0$), the denominator is just $(N_N - n + N_o)$. As the “other events” category now includes, by definition, the component of the alpha background that we do not understand, any NNNAs that may be present, and the statistical fluctuations in the original data, N_o can be small or negative. In this case, as n is scanned over the region near N_N , a discontinuous divergence is encountered across which the sign of the overlap changes. To circumvent this problem, a new normalization convention must be introduced in the definition of the overlap:

$$O_L^{new} = \frac{1}{\text{BW}} \frac{\sum_i L_i P_i^N}{\sum_i |L_i|}, \quad (4.14)$$

so the overlap of the alpha-subtracted data is now

$$O(n) = \frac{((N_N - n)O_N + N_o O_o)}{\sum_i |(N_N - n)P_N^i + N_o P_o^i|}, \quad (4.15)$$

where the “other background PDF” P_o^i (normalized such that $\sum_i P_o^i = 1$) has been introduced, and the subscript R has been dropped to differentiate overlaps calculated using the old and new normalization conventions⁹. Unfortunately, the new normalization convention has destroyed the linearity which led to the analytic relations derived earlier between n and the overlap limit. However,

⁸ The assumption that the overlap limit would not need to be changed after the subtraction of the alpha background is, of course, a poor one. The following discussion, however, renders this point moot.

⁹ We note that it is still possible to have a divergence with this new normalization if $N_o = 0$. However, this would require that the data agree exactly with the assumed alpha and neutron shapes, to the extent of having no statistical fluctuations. This possibility is statistically negligible.

the feature that allowed the overlap to be used as a signal extraction technique, that is the monotonic decrease of $O(n)$ as n increases, remains:

$$\begin{aligned}
\frac{dO(n)}{dn} &= -\frac{O_N}{\sum_i |D_i|} - \frac{((N_N - n)O_N + N_o O_o)}{(\sum_i |D_i|)^2} \frac{d}{dn} \left(\sum_i |D_i| \right) \\
&\leq -\frac{O_N}{\sum_i |D_i|} + \frac{((N_N - n)O_N + N_o O_o)}{(\sum_i |D_i|)^2} \\
&= -\frac{O_N}{\sum_i |D_i|} + \frac{O(n)}{\sum_i |D_i|} \\
&\leq 0.
\end{aligned} \tag{4.16}$$

The second line in this sequence follows from the observation¹⁰ that $-1 \leq \frac{d}{dn} \sum_i |D_i| \leq 1$, and the last line follows from the assumption¹¹ that $O(n) \leq O_N$. Given that $O(n = N_N) = O_o$, this monotonic decrease shows that the true number of neutrons will still lie between the upper and lower neutron limits defined by the overlap method provided that the true overlap of the background lies between the overlap limits.

To summarize, in order to be able to subtract the alpha component from the data it is necessary to introduce a normalization convention to the definition of the overlap which sacrifices the linearity of the equation and hence the analytic relationships between the overlap limits and the limits of the neutron confidence region. However, the confidence region can still be defined numerically (by scanning n to find the limiting positions of $O(n)$), and it can still be demonstrated that provided the background overlap lies between the overlap limits, the true neutron number will fall within the extracted confidence interval. The “new” overlap thus still represents a viable signal extraction tool and, based on Monte Carlo studies, the subtraction of the alpha background thus enabled does indeed result in a substantial reduction in the width of the extracted confidence interval, even after the overlap range is expanded to take into account the significant reduction in our knowledge of the background after alpha subtraction. As a result, it is the alpha subtraction method with the new normalization convention that is used in the following analysis.

¹⁰ If all the D_i are positive, $\frac{d}{dn} (\sum_i |D_i|) = 1$; if all the D_i are negative, $\frac{d}{dn} (\sum_i |D_i|) = -1$.

¹¹ It is technically possible to have $O(n) > O_N$ when $O_o > O_N$; this could occur, for example, for a “single-bin” background immediately beneath the neutron peak. We neglect this possibility here as it will certainly lie outside of our assumed overlap interval.

Statistical Uncertainty

As was noted after Equations 4.9 and 4.8, if the number of background events in the data is zero, the neutron confidence interval returned by the overlap method collapses to the number of events in the data. That is to say, the overlap method as described to this point provides a confidence region describing the number of neutrons present in a specific data set; the statistical uncertainty associated with the conversion between the number of neutrons present in a specific data set and the “true” number expected based on the underlying distribution is not included. It is therefore necessary to widen the confidence region returned by the overlap method by an amount equivalent to adding in quadrature the width of the original confidence region and the statistical uncertainty in the number of extracted neutrons in order to include this effect.

4.2.2 Fit Range

As the overlap method (by design) does not impose any information about the background other than its dissimilarity to the neutron shape, there is no benefit in including the data outside the neutron energy region in the fit (any data outside this region will have a “zero” weighting when multiplied by the neutron PDF value, and so will not contribute to the overlap results). Therefore, the overlap fits are performed over 0.2 - 1.0 MeV.

4.2.3 Binning

The bin width used in the overlap fits is important. In the ideal situation, where the data are well described by the assumed background and neutron shapes, in the vicinity of the true neutron number the residual being used to calculate the overlap will be dominated by statistical fluctuations. These fluctuations, which on average will themselves have an overlap of zero, will thus determine the normalization used in Equation 4.14. In determining the upper and lower limits on the neutron number, neutrons can be thought of as being added to, and subtracted from, the residual at the true neutron number until the residual overlaps meet the overlap limits. The normalization increases or decreases the contribution of each added neutron to Equation 4.14, and hence changes the number

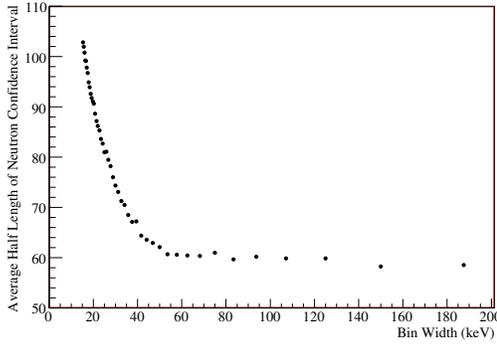
of neutrons that must be added and subtracted to reach the overlap limits. Thus, as the bin widths narrow and the sum of the absolute values of the statistical fluctuations in the bins increases, the width of the neutron confidence region will also increase. From this point of view, then, the bin widths should be as large as possible.

On the other hand, as the bin widths get very large, the distinct neutron shape is “washed out” by the binning. This means that the overlap of a given physics distribution with the neutron shape increases as bin widths increase (in the limit of single bin histograms, any background shape has the same overlap as the neutron shape). Thus, in order to encompass the same set of physics backgrounds, the overlap limits must be increased as the bin width increases. This effect increases the neutron confidence interval as bin width increases, and hence makes smaller bin widths preferable.

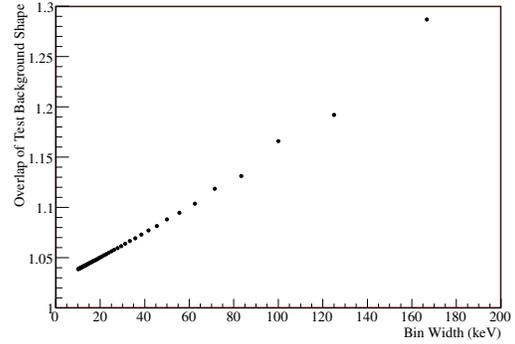
Figure 4.3 gives an illustrative example of the opposing effects described above. In this Figure, we see that when both of these effects are considered, the width of the neutron confidence interval (and hence the size of the neutron number uncertainty) has a shallow minimum for bin widths in the 50 - 90 keV range. This result is slightly dependant on the shape of the background, as the rate with which the overlap of the background increases with bin width depends on the background shape being considered. The change in the background overlap with bin width is quite slow, however, compared to the effect of the changing sum of the statistical uncertainties, so the position of the minimum is unlikely to change significantly. We therefore choose 50 keV (approximately at the beginning of the steep increase due to the statistical effect) as the bin width for the overlap fits.

4.2.4 Assumed Alpha Shape

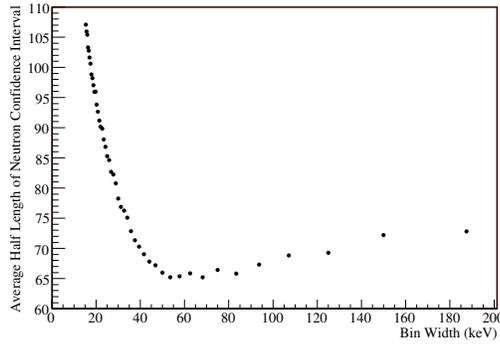
The background shape that was assumed in carrying out the overlap fits was obtained from data from the ^4He strings. For ease of manipulation, the ^4He data were fitted using the polynomial method to provide an analytic form for the assumed background shape. The polynomial method selected a first order polynomial to describe the ^4He data; the data and fit are shown in Figure 4.4. For reference, the (non-normalized) best fit form of the background, B , as a function of the energy,



(a) Constant Overlap Limit.



(b) Change in Overlap with Bin Width.



(c) Constant Physics Limit.

Figure 4.3: The effect of bin size on the overlap fits. In (a), the half-width of the neutron confidence interval is shown for overlap fits to Monte Carlo data sets (each data set consisted of about 7500 alpha background events between 0 and 1.5 MeV drawn from the ^4He distribution and 1000 neutron events drawn from the energy spectrum of the ^{24}Na spike data) with different bin widths and constant overlap limits. Each point on the plot represents the average of 1000 Monte Carlo tests. The increase in the neutron confidence interval at small bin widths as a result of the increase in the total statistical fluctuations in the data, as described in the text, can clearly be seen. In (b), the increase in the overlap of a constant physics data shape (in this case a broad Gaussian centred at just below the neutron peak; this is “NNNA shape 1” as described in Figure 4.6) as the bin width increases is shown. Finally, (c) shows the half-widths of the neutron confidence regions resulting from Monte Carlo data sets similar to those in (a), except that this time the overlap limits for each binning are scaled by the results shown in (b) (i.e. (c) shows the change in confidence region with binning when a constant background shape limit, as opposed to a constant overlap limit, is applied). A shallow minimum is visible at around a bin width of 50 keV.

E , (in MeV) is

$$B(E) = \frac{16.87 + 6.18E}{e^{(56.847(0.199272-E))} + 1}. \quad (4.17)$$

In the overlap fits, the normalization of this assumed background shape was determined by fitting to the data above the overlap window (i.e. from 1.0 - 1.5 MeV).

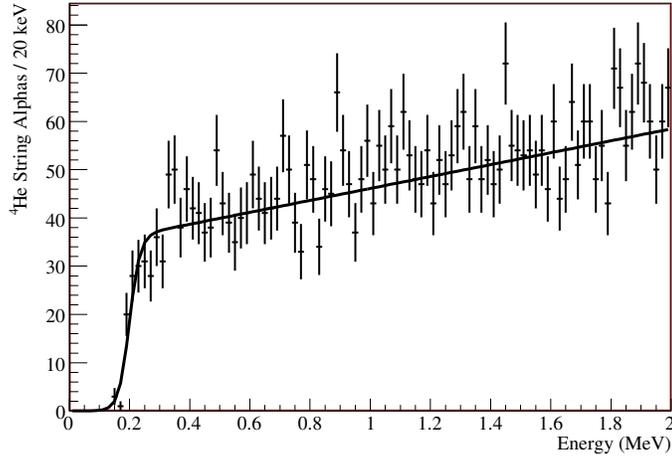


Figure 4.4: The ${}^4\text{He}$ string based alpha spectrum, and the fit to the first order polynomial (convolved with an inverted Fermi Function to simulate the threshold cutoff) that was used as the assumed background shape in the overlap fits.

It should be noted that initial attempts were made to use the Monte Carlo-derived alpha PDF that was used in the NCD PRL as the assumed alpha shape. However, as can be seen in Figure 4.5, the ${}^4\text{He}$ shape provides a better description of the data, and hence a narrower neutron confidence interval, when tested on a blinded NCD data set¹².

4.2.5 Defining the Overlap Limits

In defining the overlap limits, one can take two approaches; the limits can be based either on the range of possible background shapes which seem *a priori* reasonable, or they can be based on a rough expectation of what the background shape actually is. In the first case, the overlap limits are set by quantifying the point at which the background seems “too similar” to the neutron shape to

¹² The blind data used here were produced using the “Wilkerson-Elliott blindness procedure” which was used in all SNO analyses. In this procedure, an unknown fraction (<30%) of events are excluded at random from the data, and the muon-follower data cleaning cut is intentionally partially spoiled to allow an unknown number of muon follower events (which are primarily neutrons) into the analyzed data.

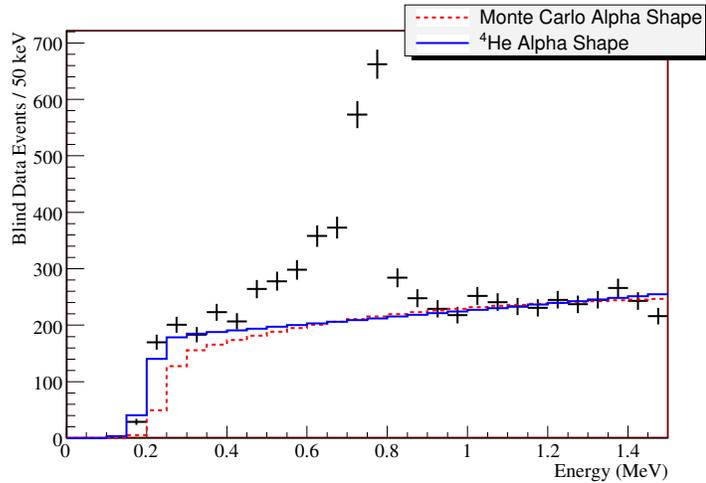


Figure 4.5: A comparison of the extrapolated ^4He and Monte Carlo alpha spectra compared to the (blind) ^3He data. The assumed spectra are fitted to the data between 1.0 and 1.5 MeV and extrapolated to lower energy, as is done in the overlap method. The ^4He shape can be seen to give better agreement with the blind data at low energies (where the neutron spectrum gives very little contribution).

be physically acceptable. In this case, the overlap method simply serves as a mechanism to map our prior expectations for the types of background that seem reasonable into a range of allowed neutron numbers. The range of reasonable backgrounds upon which the overlap limits are based is, of course, a somewhat subjective notion which will vary from person to person. Still, there is some consistency in the background behaviours that people consider reasonable, and in the event that no other information is available, the overlap method allows a signal extraction to be carried out on this basis alone.

In the second case, the overlap limits are based on a firmer understanding of the expected background shape. If, for example, one was able to define even roughly the PDFs that might be used in a more conventional signal extraction, the overlap limits could be based on the range of background overlaps allowed by the background PDFs and their systematics. Used in this way, the overlap method will yield the “correct” answer in any situation where the conventional signal extraction would, and thus simply provides an additional method of evaluating systematic uncertainties (in a very conservative way). More than that, however, the overlap method “automatically includes” in that systematic uncertainty any background shape within that overlap range, not just those that

would be considered in traditional systematics determination. The overlap method thus has the potential to be correct even where a traditional analysis might not be. This generality, of course, comes at the cost of sacrificing some specific information about the expected background shape, presumably with an attendant increase in neutron uncertainty. One would not, therefore, opt for this approach when the systematics are small and well defined. When the systematics are not well defined, however, the overlap approach could provide a reasonable alternative.

Before proceeding further, it is useful to develop an intuition for the behaviour of the overlap variable. In general, of course, the overlap is larger for those backgrounds which have a larger fraction of their area under the neutron peak. By definition, the overlap of a flat line with the neutron shape is 1.0, while the overlap of the neutron shape with itself is 3.06 and the overlap of the assumed alpha shape (with the “assumed background subtraction” turned off) is 1.04. To provide further examples, a set of hypothetical “NNNA shapes” was prepared and their overlaps calculated; the results are shown in Figure 4.6. These shapes will also be used in the next section to test the response of the different signal extraction routines to unexpected shapes in the backgrounds. The maximum achievable overlap (for a single bin background directly beneath the neutron peak) is 4.99.

If the overlap limits are to be set using simply our intuition about reasonable behaviours of the overlap, Figure 4.6 is all that is required. In the opinion of the author, for example, shapes 2, 6, and 7 are too similar to the neutron shape to be physically reasonable, while shapes 1 and 5 (and their inverses) seem reasonable. Therefore, the overlap limits might be chosen as 1.2 and -1.2.

To proceed along the second path, it is necessary to determine the overlaps of the background PDFs that were used in the published NCD PRL. These PDFs consisted of a Monte Carlo alpha spectrum, with eight associated polynomial systematic uncertainties, and two “NNNA PDFs” produced by fitting Gaussian curves to the spectra of the known NNNA strings after subtracting the average spectrum from “clean” strings. As discussed earlier, the PDFs and their associated systematics were only vetted in the 0.4 - 1.4 MeV energy range. However, the produced PDFs do extend down to 0.2 MeV, and that full range was used in determining their overlaps. The overlaps of the

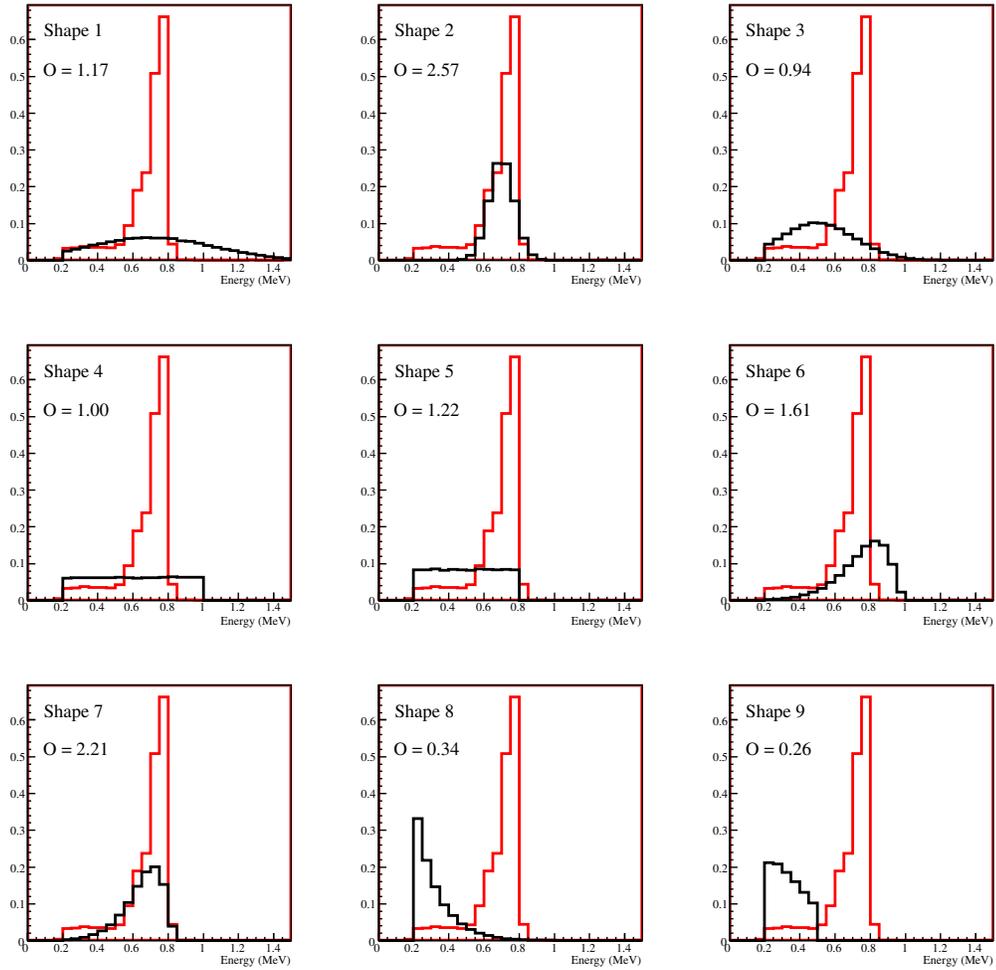


Figure 4.6: The “NNNA” shapes used in calculating the sensitivity of the signal extraction methods to unexpected backgrounds. The 2005 ^{24}Na energy spectrum is overlaid with each background shape for comparison, and the overlap, O , of each shape with the neutron spectrum is also given. The overlap is seen to have the expected behaviour, with higher overlaps for those shapes which “look more” like the neutron shape.

different PDFs are shown in Table 4.4.

| PDF | Overlap |
|--|---------|
| Central Alpha Monte Carlo | -0.14 |
| Alpha MC $+1\sigma$ Po alpha depth systematic | -0.24 |
| Alpha MC -1σ Po alpha depth systematic | 0.31 |
| Alpha MC $+1\sigma$ Bulk alpha depth systematic | -0.48 |
| Alpha MC -1σ Bulk alpha depth systematic | 0.44 |
| Alpha MC $+1\sigma$ Drift time systematic | -0.24 |
| Alpha MC -1σ Drift time systematic | 0.03 |
| Alpha MC $+1\sigma$ Avalanche width offset systematic | -0.32 |
| Alpha MC -1σ Avalanche width offset systematic | 0.08 |
| Alpha MC $+1\sigma$ Avalanche gradient offset systematic | -0.14 |
| Alpha MC -1σ Avalanche gradient offset systematic | -0.14 |
| Alpha MC $+1\sigma$ Po/bulk fraction systematic | -0.32 |
| Alpha MC -1σ Po/bulk fraction systematic | 0.11 |
| Alpha MC $+1\sigma$ Ion mobility systematic | -0.14 |
| Alpha MC -1σ Ion mobility systematic | -0.14 |
| Alpha MC $+1\sigma$ Data cleaning systematic | 0.38 |
| Alpha MC -1σ Data cleaning systematic | -0.27 |
| String 26 NNNA | 0.48 |
| String 0 NNNA | 0.46 |

Table 4.4: The overlaps of the PDFs used in the signal extraction published in the NCD PRL. The systematic PDFs themselves are shown in [35].

The systematic uncertainties in the Monte Carlo alpha energy spectrum were parameterized as multiplicative, rather than additive, changes to the spectrum [35]. This means that the overlap of the alpha shape produced by combinations of these systematics will not necessarily lie within the overlap range spanned by the individual corrections. To check for this possibility, combinations of the different systematic uncertainties, with each of them randomly varied according to the $\pm 1\sigma$ constraints, were applied to the Monte Carlo alpha spectrum and the overlap recorded. The result is shown in Figure 4.7. In Figure 4.7, we see that overlap limits of -0.5 and 0.6 spans the bulk of the possible overlaps. To be conservative, however, (and to have symmetric overlap limits), we choose the overlap limit of -1.0 and 1.0 as spanning practically all of the backgrounds allowed by the Monte Carlo.

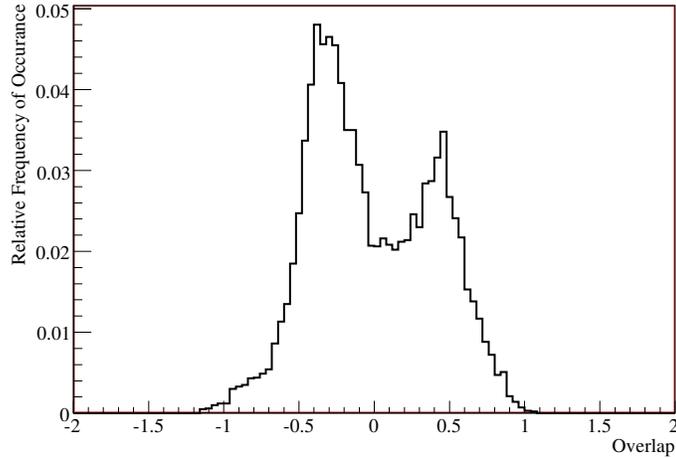


Figure 4.7: The spread in the overlaps of the different Monte Carlo alpha spectra admitted by the Monte Carlo systematic uncertainties.

In summary, then, we have produced two possible sets of overlap limits using two different techniques. The first technique, in which the overlap limits are chosen based on the author’s intuition for the range of background shapes that are physically reasonable, yields overlap limits of -1.2 and 1.2. The second technique, which bases the overlap limits on the assumption that the true background should not be more like the neutron shape than what might be expected from the PDFs used in the NCD PRL signal extraction (including their systematic uncertainties) yields overlap limits of -1.0 and 1.0.

4.3 Testing the Methods

To ensure that the polynomial and overlap methods function as desired and to gain some intuition for how the fits behave when faced with unexpected backgrounds, it is helpful to carry out studies using Monte Carlo data. Beyond that, the methods can be tested using a blinded set of actual NCD data.

In the simplest tests, the polynomial and overlap methods were used to fit fake data sets composed only of neutrons and alphas drawn from the assumed distributions. Again, each data set contained on average 1000 neutrons and 7500 alphas below 1.5 MeV. As can be seen in Figure 4.8, both

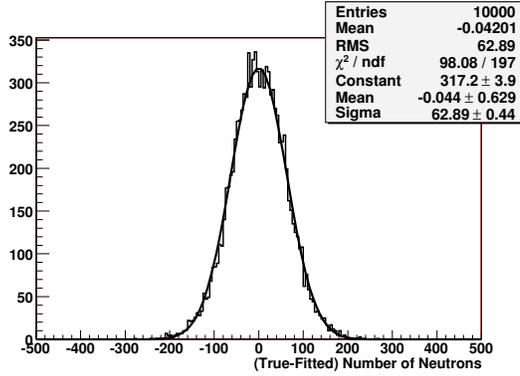
the polynomial and overlap methods have mean biases less than one event, and consistent with zero¹³. Recalling that if the statistical spread in the bias distribution is correctly described by the fit uncertainties, the pulls distribution should have a width (“sigma”) of 1.0. As can be seen, the pulls distribution of the polynomial method is correct. The distribution of pulls for the overlap method has a width of less than one; this is not unexpected, however, as it simply shows that the spread in the chosen overlap limits was larger than the spread required to describe the background present in the simulated data sets. Repeating the overlap tests with overlap limits of -0.65 and 0.65, for example, increases the width of the overlap pulls distribution to 1.03, while increasing the overlap range to -1.2 to 1.2 narrows the width of the pulls distribution to 0.59.

Having established that both signal extraction methods work as anticipated on data drawn from well understood signal and background distributions, it is now interesting to study their behaviour when unexpected background shapes are present. To this end, Monte Carlo ensemble tests were carried out with “fake NNNAs” added to the data. The NNNAs were drawn from the spectra shown in Figure 4.6, and an average of 200 NNNAs was added to each fake data set (in addition to 1000 neutrons and 7300 alphas drawn from the expected distributions). The choice of 200 NNNAs is meant to be illustrative; the effect of the NNNAs should scale roughly as the number present. Each fake data set was then fitted using the polynomial method, the overlap method¹⁴ and a “standard method,” which was a simple binned likelihood fit to the expected neutron and alpha PDFs. Comparisons between the results of the polynomial, overlap and standard fits in both bias and pull are shown in Figures 4.9-4.14.

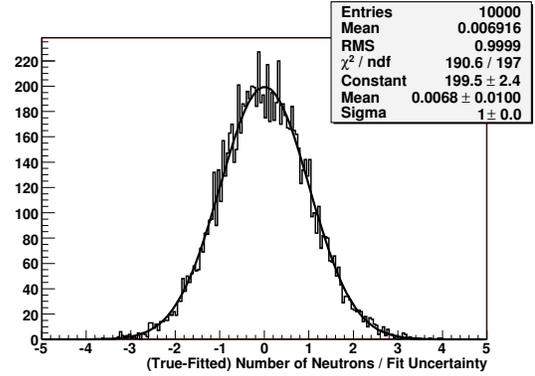
In these tests, it appears that both the overlap and polynomial tests react to the “unexpected” backgrounds in the expected manner. In the overlap fits, the presence of the unexpected background shape shifts the position of the true neutron number with respect to the overlap-defined neutron limits (increasing the size of the “biases”), while at the same time increasing the neutron limit spread

¹³ In order to easily compare the overlap method with the polynomial method, a “fitted number of neutrons” for the overlap method was defined as the average of the upper and lower neutron number limits, with a “fit uncertainty” equal to half of the spread between the limits. Note that there is no reason to expect the resulting distributions to have small biases; all that should really be expected is that if the true background overlap lies between the overlap limits, the overlap pulls should lie between -1 and 1.

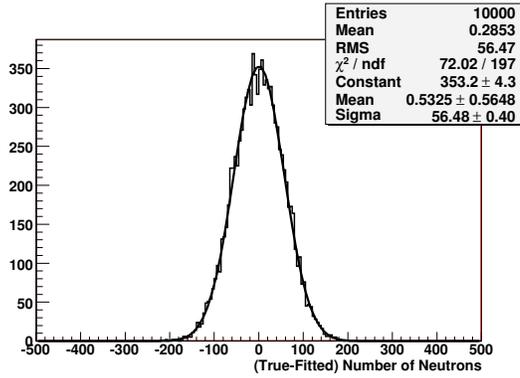
¹⁴ The overlap limits used in these comparisons were -1.0 and 1.0.



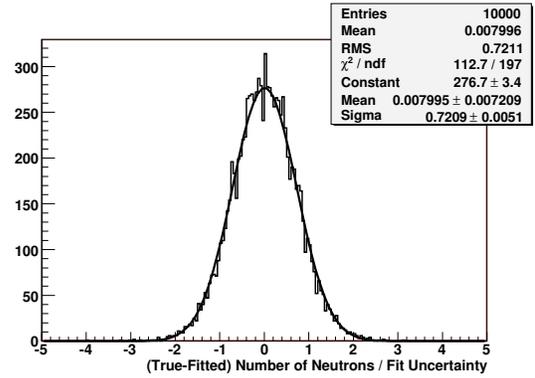
(a) Polynomial Bias.



(b) Polynomial Pull.



(c) Overlap Bias.



(d) Overlap Pull.

Figure 4.8: The results of simple Monte Carlo ensemble tests of the polynomial and overlap methods (the overlap limits were -1.0 and 1.0). Each fake data set consisted of, on average, 1000 neutrons and 7500 alphas below 1.5 MeV drawn from the assumed alpha and neutron spectra. The fitted width of the polynomial pulls distribution is 1.00005 ± 0.007 .

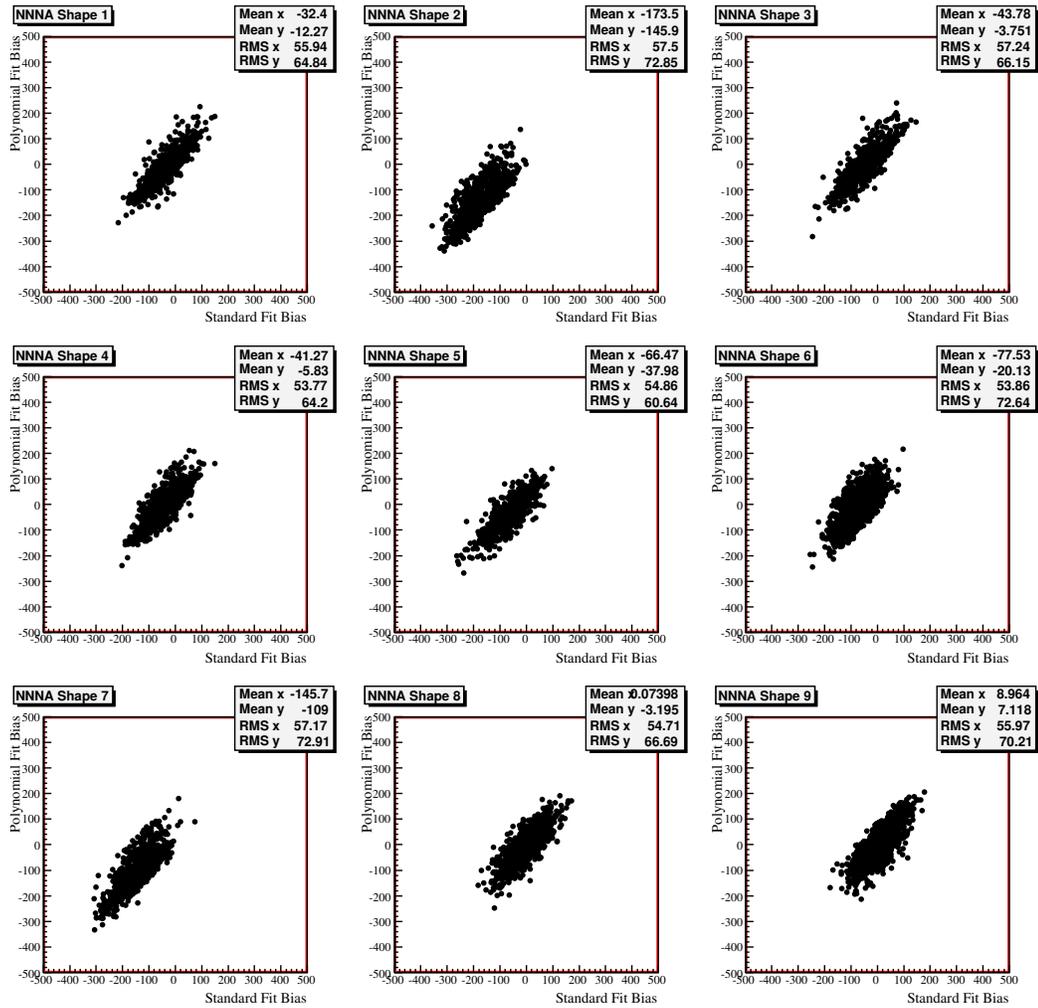


Figure 4.9: Comparison of the effect of potential NNNAs on polynomial and standard fit biases. Note that the average biases of the polynomial method are significantly smaller than in the standard signal extraction, while the RMS spread in the bias (which should be approximately equal to the fit uncertainty) is larger for the polynomial method.

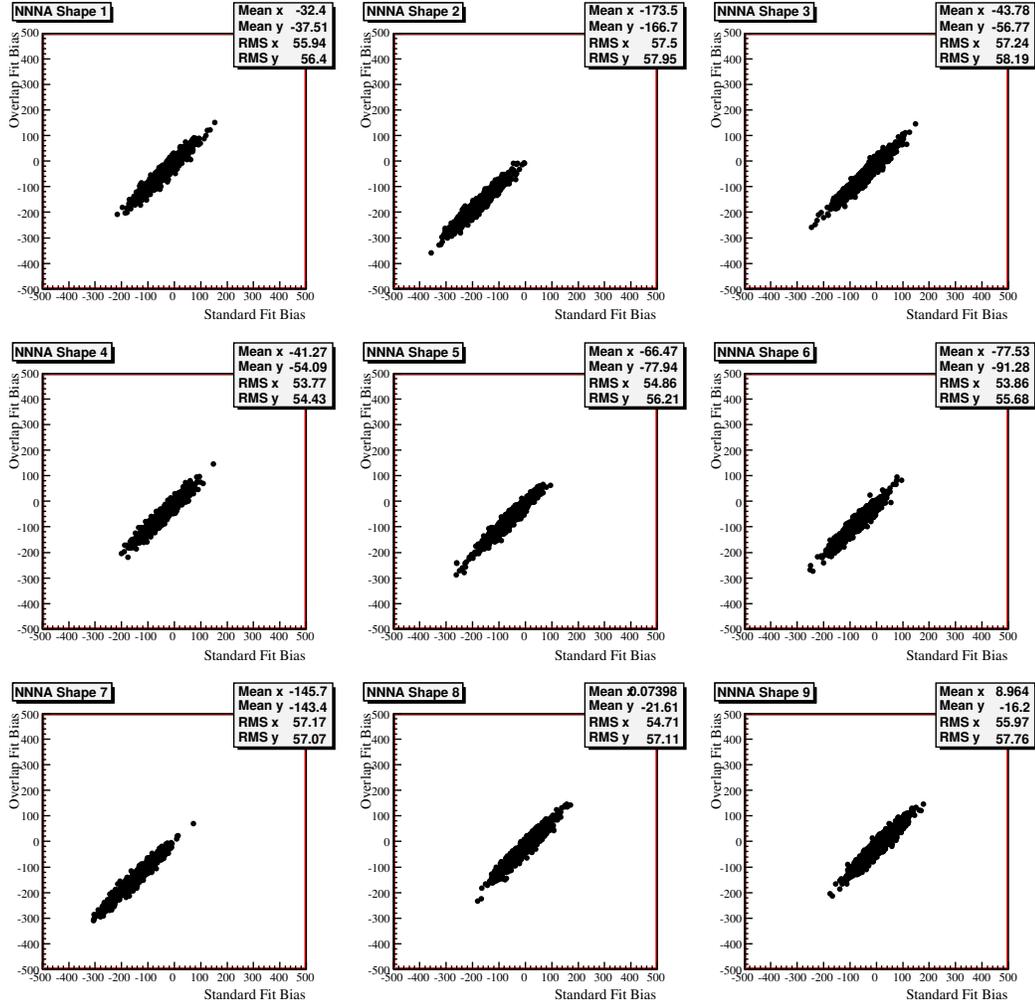


Figure 4.10: Comparison of the effect of potential NNNA's on overlap and standard fit biases. The behaviour of the standard and overlap fits can be seen to be quite correlated. The overlap biases can be seen in several cases to be larger than the biases in the standard fits. It is difficult, however, to make meaningful comparisons with the means of the overlap fit ranges, as the chief way in which the overlap method responds to unexpected backgrounds is by increasing the width of the neutron confidence interval. This should be more evident in the pulls distribution (Figure 4.13), and in the fraction of test fits in which the true number of neutrons is contained in the neutron confidence interval (Table 4.5).

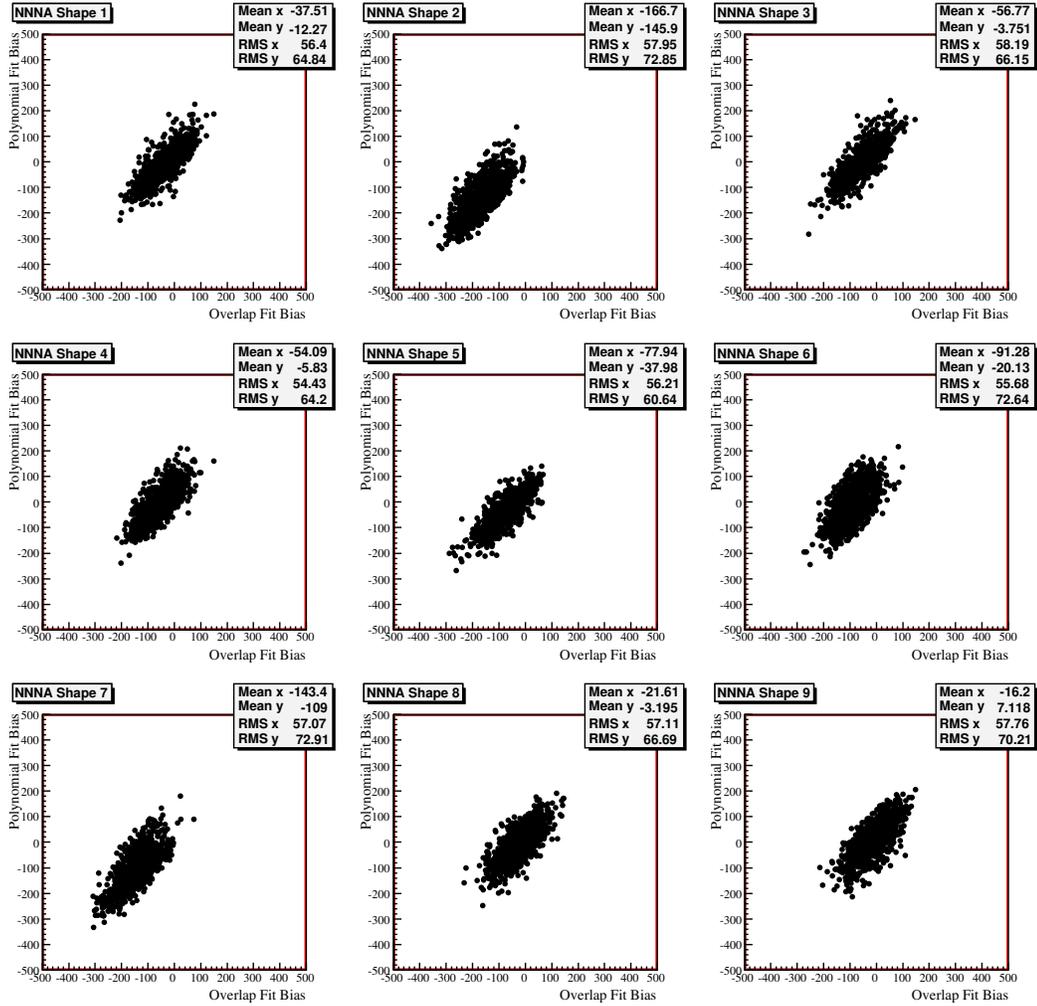


Figure 4.11: Comparison of the effect of potential NNNA's on overlap and polynomial fit biases. As expected from the previous comparisons, the polynomial fits have significantly smaller average biases.

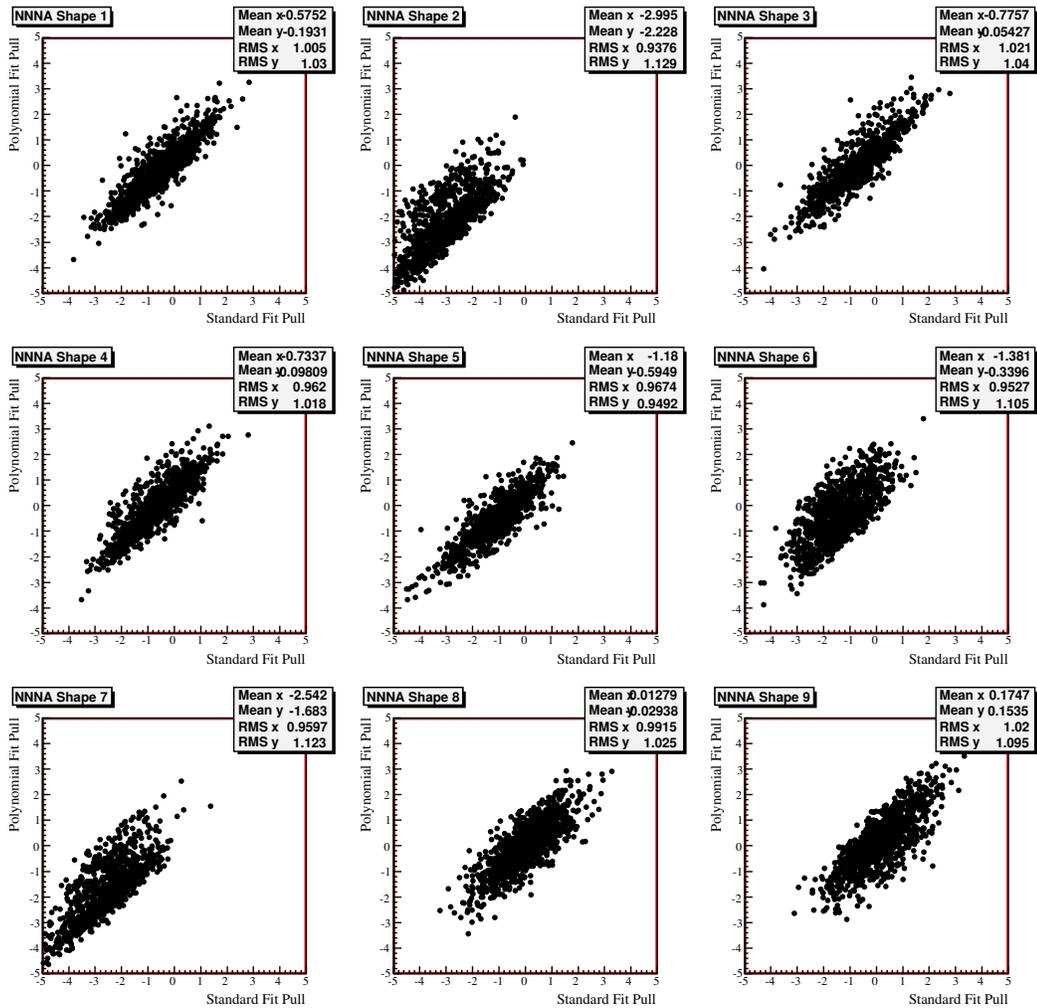


Figure 4.12: Comparison of the effect of potential NNAs on polynomial and standard fit pulls. Some deviation from an RMS of 1.0 can be seen in both the standard and polynomial fits when NNAs are present; however, the deviations are not particularly large indicating that the fit uncertainties do a reasonably good job of predicting the statistical spread in the fit values even when NNAs are present. The polynomial fits have significantly smaller average pulls.

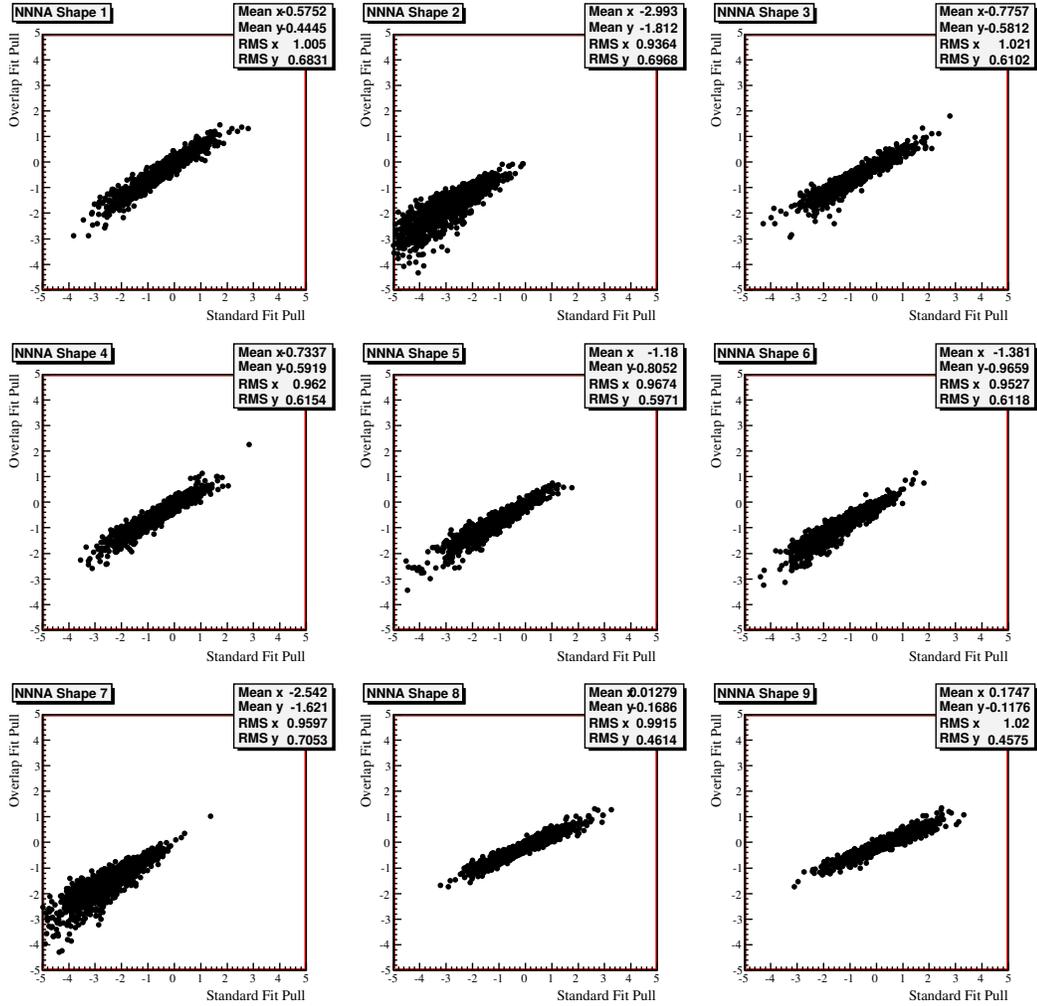


Figure 4.13: Comparison of the effect of potential NNNA on overlap and standard fit pulls. Here the “pull” of the overlap is defined as the mean of the confidence interval divided by half the width of the interval. Based on the definition of the overlap neutron limits, all that should really be expected of such a “pull” is that it should be between -1 and 1 when the background overlap falls between the overlap limits. Still, the mean pull of the overlap fits is seen to be smaller than the mean pull of the standard fit (in spite of the larger average overlap biases in Figure 4.10), indicating that the widths of the overlap confidence intervals are properly increasing in the presence of the NNNA to help to take their effect into account.

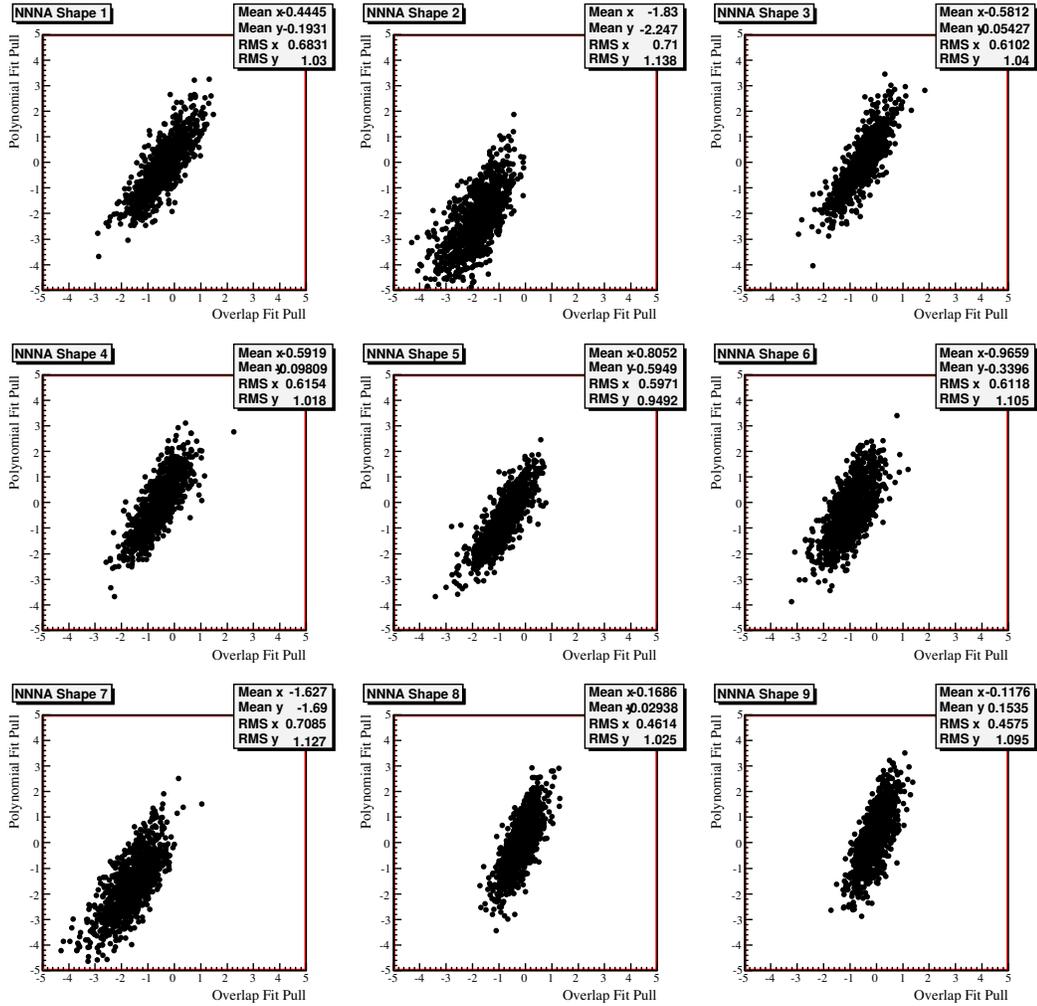


Figure 4.14: Comparison of the effect of potential NNNA's on overlap and polynomial fit pulls. In some cases the polynomial method has a smaller mean bias and in others the overlap fit has a smaller mean bias. While comparisons to the “pulls” of the overlap fits should again not be taken too seriously, this does indicate that the preferable method can depend on the shape of the background.

(which reduces the size of the “pulls”). The overlap limits used in these fits (-1.0, 1.0) are small compared to the overlaps of most of the NNNA test shapes; as a result, the increase in the width of the neutron confidence region is not necessarily sufficient to encompass the change produced by the NNNA. Table 4.5 shows the fraction of test data sets for which the true number of neutrons falls within the overlap neutron confidence intervals for overlap fits using different overlap limits. As can be seen in the Table, the neutron limits tend to encompass the true neutron number provided that the overlap limits encompass (or almost encompass) the background overlap, as expected.

| | | Overlap Limits | | |
|---------------------------|----------|----------------|-----------|-----------|
| | | -1.0, 1.0 | -1.2, 1.2 | -2.0, 2.0 |
| NNNA Shape (NNNA Overlap) | none (-) | 0.85 | 0.92 | 1.00 |
| | 1 (1.17) | 0.79 | 0.87 | 1.00 |
| | 2 (2.57) | 0.11 | 0.19 | 0.89 |
| | 3 (0.94) | 0.76 | 0.87 | 1.00 |
| | 4 (1.00) | 0.76 | 0.86 | 1.00 |
| | 5 (1.22) | 0.65 | 0.77 | 1.00 |
| | 6 (1.61) | 0.52 | 0.67 | 0.99 |
| | 7 (2.21) | 0.19 | 0.31 | 0.94 |
| | 8 (0.34) | 0.96 | 0.99 | 1.00 |
| | 9 (0.26) | 0.96 | 0.99 | 1.00 |
| No NNNA uncertainty | | 69 | 97 | 210 |

Table 4.5: The fraction of Monte Carlo ensemble tests in which the true neutron number was contained in the neutron confidence region defined by the overlap method. It can be seen that the overlap fit limits have a good chance of encompassing the true neutron number provided that the overlap limits encompass, or nearly encompass, the background overlap. The effect is not perfect because the overlap of a specific NNNA background can differ from the overlap of the underlying distribution due to statistical fluctuations. The average “uncertainty” (the half width of the neutron confidence interval) for fits to Monte Carlo data sets with no NNNA added are shown to give an idea of how they change with overlap limits.

The polynomial tests also react as expected: as seen in Table 4.6, in the presence of the hypothetical NNNA, higher order polynomials are chosen, on average, than when the NNNA are not present. The fits also appear to be correctly adapting to the unexpected backgrounds, as seen in the reduced average biases and pulls compared to the standard fits.

Finally, the inclusion of the “standard” fitting technique in the ensemble tests allows comparison

| NNNA | None | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
|------------|------|------|------|------|------|------|------|------|------|------|
| Avg. Order | 1.76 | 1.99 | 2.45 | 2.10 | 1.98 | 1.95 | 3.29 | 2.58 | 2.97 | 3.04 |

Table 4.6: The average polynomial order selected in the ensemble tests which included hypothetical NNNA's (the number in the first row refers to the number of the NNNA shape as shown in Figure 4.6). It can be seen that the polynomial method chooses higher order polynomials, on average, when the NNNA's are present.

of the performance of the new methods with standard techniques. Both the polynomial and overlap fits have smaller pulls in the presence of the NNNA's than does the “standard” fit, at the cost of larger fit uncertainties. In the overlap method, the fit biases tend to be as large or larger than the standard fits, but with the fit “uncertainties” in the overlap fits increasing to reduce the pull. The polynomial method tends to have smaller biases and slightly larger uncertainties than the standard fits. In general, then, both the polynomial and overlap methods appear to be fulfilling their intended purpose of increasing the robustness of the fit results against unexpected background shapes. As might be expected, this increased robustness, which is achieved by relying less strongly on prior knowledge of the background energy spectrum, comes at the cost of increasing the fit uncertainties.

Tests Using Blind Data

Finally, the new fit procedures were tested using a blinded set of NCD data. As described earlier, this blind data set consisted of a subset of about 93% of the NCD data from which an unknown fraction (<30%) of individual events was randomly excluded, and for which the data cleaning cuts designed to reject secondary events resulting from the passage of a muon through the detector were intentionally spoiled to allow an unknown number of single muon follower events (typically neutrons) into the data set. Only data from the ^3He strings that were included in the NCD PRL were included in this blind data set.

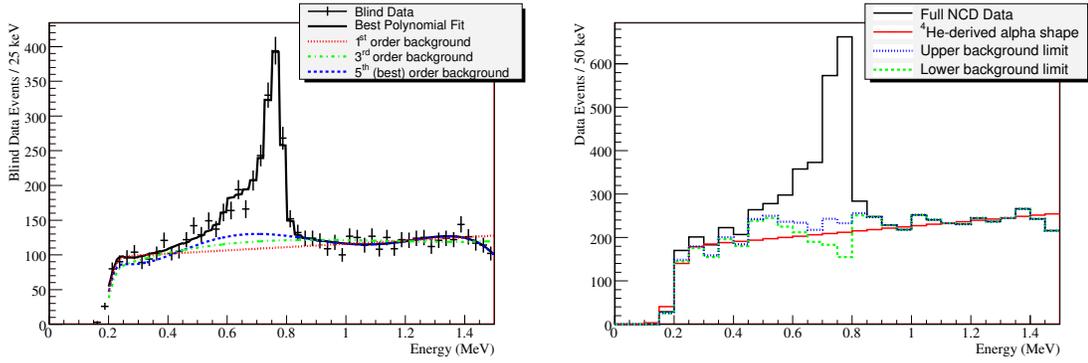
The best fits to the blind data set using the polynomial method and the overlap method with overlap limits -1.0 and 1.0 are shown in Figure 4.15. The polynomial fit chose order five as the best description of the background, with 1188 ± 73 neutrons fitted and a χ^2 of 32.0/45, while the overlap method returned (1290, 1537) as the neutron confidence interval. The standard fit with

the ^4He background PDF fitted 1399 ± 57 neutrons with a χ^2 of 64.5/50. Based on Figure 4.15, it appears that the discrepancy between the polynomial fit and the other fit methods is due to an apparent “step” in the background spectrum just below the neutron peak. The standard fit, lacking the freedom in the PDFs included in the NCD PRL, was unable to account for this background and as a result “undershot” the data in this area, fitting a larger number of neutrons as a result. As can be seen in Figure 4.15 and Table 4.7, the lower order polynomial fits also encountered this problem; as the polynomial order increases, the background shape begins to form a “bump” at energies just below the neutron peak, and the number of fitted neutrons decreases. A bump can also clearly be seen in the upper and lower residual backgrounds determined by the overlap fit. In looking at Figure 4.15 (b), it is clear that the overlap limits used were not sufficient to encompass the background present; the lower residual background clearly exhibits the expected inverted neutron peak, but the upper residual background does not show a significant (positive) peak. We thus expect the upper limit on the neutron confidence region to be safe while the lower limit should be too high. And, in comparison with the polynomial fit, at least, this is what is observed.

| Polynomial Order | Fitted Neutrons | Fit χ^2 /NDF |
|------------------|--------------------|-------------------|
| 0 | 1287^{+59}_{-58} | 100.4/50 |
| 1 | 1370^{+62}_{-61} | 63.0/49 |
| 2 | 1308^{+64}_{-63} | 47.0/48 |
| 3 | 1278^{+67}_{-66} | 44.8/47 |
| 4 | 1286^{+68}_{-67} | 44.5/46 |
| 5 | 1188^{+73}_{-73} | 32.0/45 |
| 6 | 1190^{+73}_{-73} | 31.9/44 |

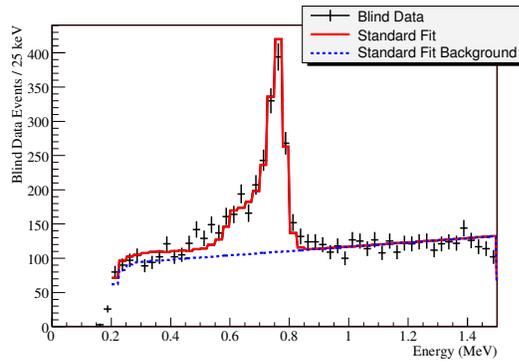
Table 4.7: The results of fits to the blind data set using different polynomial orders to represent the background. Order 5 was selected by the polynomial fitting procedure as the best description of the background in the data.

That the chosen overlap limits do not appear to adequately describe the background present in the blind NCD data is disappointing, and indicates that the true background in the NCD data was more “neutron-like” than expected based on the alpha background model. Indeed, as shown in Figure 4.16 (a), even using -1.2 and 1.2 as the overlap limits (as suggested earlier by subjectively



(a) Polynomial Fit.

(b) Overlap Fit.



(c) Standard Fit.

Figure 4.15: The results of fits to the blind data set. The polynomial fit results including the best fit 5th order background are shown in (a). The 1st and 3rd order polynomial backgrounds are also shown for comparison. The upper and lower limit backgrounds determined by the overlap limit are shown in (b), along with the ⁴He-based alpha shape that was subtracted in calculating the overlaps. Finally, the results of the “standard” PDF-based fits (including only the ²⁴Na neutron spectrum and ⁴He alpha spectrum) is shown in (c) for comparison.

examining the test NNNA shapes) appears to give too low an upper background limit. This seems to be the result of a step or bump in the background at around 0.5 MeV, which increases the fraction of the residual background that falls beneath the neutron peak and hence the background overlap¹⁵. To continue to use the overlap method in the NCD signal extraction, then, a new set of overlap limits must be produced. This can be accomplished by changing the limits and examining the upper and lower limit backgrounds until evidence of a neutron peak is shown in both¹⁶. Carrying out this procedure yields overlap limits of -0.3 and 1.7, the upper and lower limit backgrounds of which are shown in Figure 4.16 (b)¹⁷. The increase in the width of the neutron confidence interval associated with this change in overlap limits is shown in Table 4.8.

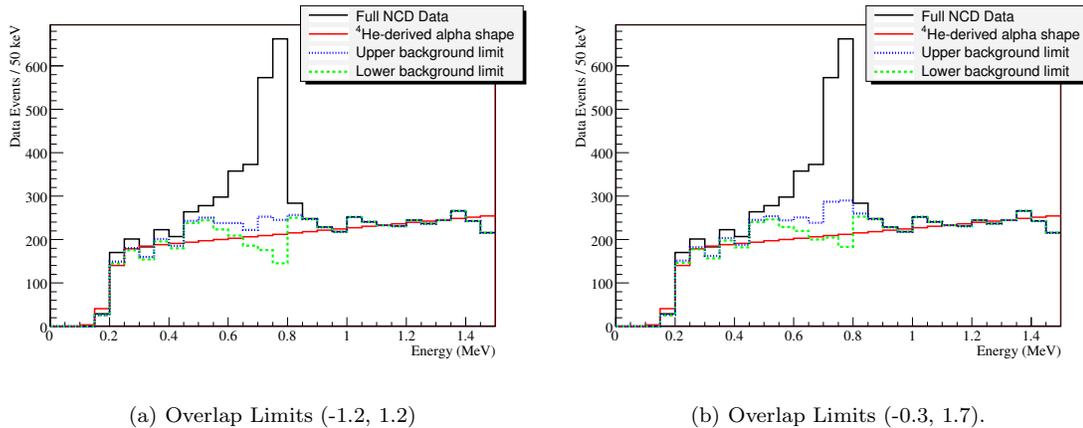


Figure 4.16: The upper and lower limit backgrounds of fits to the blind data with different overlap limits. In (a), the upper limit overlap does not show definite evidence of the neutron shape, while it does in (b). Also, as the neutron “dip” in (a) appears to be larger than required, the lower overlap limit has also been raised between (a) and (b). Therefore, -0.3 and 1.7 seem to be a reasonable set of overlap limits to describe the background present in the blind data.

¹⁵ It should be noted that the observed effect could also be due to the higher energy “shoulder” in the neutron spectrum being much larger in the data than in the ²⁴Na PDF. In this case, the adaptive fit techniques would be attempting to accommodate the inaccuracy in the neutron PDF as background, and hence biasing the fits. There is, however, no physical reason that the NCD neutron energy spectrum should be systematically different during the ²⁴Na running than during neutrino running, while there are plenty of plausible mechanisms to account for the observed distortion as a background effect. It is therefore reasonable to proceed under the assumption that what is observed is an effect in the background and not in the signal.

¹⁶ That this procedure is possible (and necessary) points out the fact that the overlap metric is not identical to what appears to be “neutron like” by eye. The latter seems to weight the few bins around the neutron peak much more heavily. It is possible, therefore, that an improved comparison metric could be devised that corresponds more closely to how the judgement of “neutron-ness” is made by eye.

¹⁷ Note that the neutron number was not used (or known) in setting the overlap limits “by eye,” so there was little risk of consciously (or unconsciously) biasing the results.

| Overlap Limits | Neutron Confidence Interval |
|----------------|-----------------------------|
| (-1.0, 1.0) | (1290, 1537) |
| (-1.2, 1.2) | (1253, 1565) |
| (-0.3, 1.7) | (1119, 1450) |

Table 4.8: The neutron confidence intervals extracted from the blind data set using different overlap limits. It can be seen that as the upper overlap limit is raised (allowing more neutron shape to appear in the upper limit background, as seen in Figure 4.16), the lower limit of the neutron confidence region decreases.

Having to define the overlap limits using the blind data in this way is frustrating and highlights the difficulty in setting these limits. This is not, however, necessarily a weakness of the method itself as much as it is a symptom of the same poor understanding of the background shape that motivated consideration of the overlap method in the first place. Thus, while it would certainly have been preferable to use overlap limits that were set independently of the data, it turns out that this was not possible because of the same unexpected behaviour in the background that was feared when the overlap method was developed.

With this third choice of overlap limits, the polynomial and overlap methods agree well. One further check that can be done is to examine the stability of the fit results when known NNAs and alpha events are added to the data sets. Given that rate of capture of neutral current neutrons on the different ^3He strings is expected to be essentially identical¹⁸, including the data from J3 and N4 (the “known” NNA strings) in the blind data set is expected to increase the number of detected neutrons by approximately $2/30 = 6.7\%$. Adding the ^4He data should not change the neutron number but will change the statistical fluctuations of the background. Comparing fits to the blind data before and after adding these additional events is not straightforward, as the fit results are expected to be partially statistically correlated. As a result, Monte Carlo ensemble tests were used to predict the expected change and statistical fluctuation in the fitted neutron number in these tests, under the assumption that the added events were either all alphas (in the case of the ^4He events) or the expected number of neutrons with the balance of the added events alphas (in the case of

¹⁸ Monte Carlo studies of uniform neutron capture indicate that the number of neutrons captured by each ^3He string is almost exactly identical, in spite of the different string lengths. Non-uniform neutron backgrounds will make this estimate slightly incorrect, but it is still more than good enough for the present comparison.

the NNNA strings). In the case of the ^4He events we fully expect the observed change in the fit parameters to agree with the Monte Carlo predictions¹⁹, while for the NNNA data the changes in the fit results will agree with the Monte Carlo predictions if the added NNNAs do not systematically affect the fits. The predicted and observed changes in the fit parameters under the addition of the data from the NNNA stings and the data from the ^4He strings is shown in Table 4.9.

| | Variable | Observed Change | Expected Change |
|--------|-----------------------------|-----------------|-----------------|
| NNNAs | Polynomial Neutron Number | 64 | 81 ± 18 |
| | Overlap Upper Neutron Limit | 142 | 81 ± 17 |
| | Overlap Lower Neutron Limit | -44 | 74 ± 25 |
| Alphas | Polynomial Neutron Number | -48 | 10 ± 50 |
| | Overlap Upper Neutron Limit | -10 | 3.2 ± 32 |
| | Overlap Lower Neutron Limit | -33 | -24 ± 50 |

Table 4.9: The changes observed after adding either the data from the ^4He strings (“Alphas”) or the data from the known NNNA strings (“NNNAs”) to the blind data set compared to the changes expected based on Monte Carlo studies of adding similar numbers of known events (neutrons and alphas) to the (real) blind data set. When the NNNA string data were added, the number of neutrons returned by the polynomial method increased by roughly the amount expected from the Monte Carlo, indicating that the NNNAs did not pull the polynomial fits any more than alpha events (although the chosen polynomial order did increase by one). On the other hand, adding the NNNAs had a much more significant effect on the overlap method than would be expected from the addition of alpha and neutron events; that is to say, the addition of background events with an unexpected distribution caused the overlap method to return a wider neutron confidence interval, as it was designed to do. On the other hand, when the ^4He events were added to the blind data, both the polynomial and overlap fits behaved as expected, with the change in the polynomial fit being 1.2σ from expectation. The non-zero central value in the expected change in the polynomial neutron number with the addition of the ^4He data was not expected, and is likely due to a statistical fluctuation in the (real) blind data that is “covered up” by the addition of more background events; repeating the tests with random initial data sets in place of the true blind data yielded expected changes in the polynomial neutron number consistent with zero. The asymmetric expected change in the overlap neutron limits is the result of the asymmetric overlap limits (-0.3, 1.7) being used.

4.4 Robust Signal Extraction Summary

Two new signal extraction techniques, the polynomial method and the overlap method, have been developed for use in extracting the number of neutron events from the NCD energy spectrum. The

¹⁹ In fact, this is not really a valid test of the overlap method, as the ^4He spectrum was used to produce the alpha shape that is assumed in calculating the overlaps. In the case of the overlap method, then, the addition of the ^4He events to the blind data can test the implementation of the overlap algorithm, but not the algorithm itself.

methods were explicitly developed to capitalize on the unique and well-defined neutron energy spectrum, while minimizing their dependence on specific prior knowledge of the background spectrum. This has the effect of making the polynomial and overlap methods more robust than standard PDF fits against unexpected shapes in the background. In the case of the polynomial fits, the background fit is able to adapt to an unexpected background shape by choosing higher order polynomials if necessary. In the case of the overlap method, the presence of such an unexpected shape results in the inflation of the neutron confidence region such that the true number of neutrons remains within the confidence region (provided that the overlap of the actual background lies within the overlap limits). In both methods, the price for the increased fit robustness is an increase in the fit uncertainty compared to a standard PDF fit.

The two methods have very different conditions for yielding the correct (within uncertainty) number of neutrons. The polynomial method will be correct provided that the background is smoothly varying enough to be well described by a polynomial, while the overlap method is correct if the overlap of the background in the data falls between the overlap limits. This means that the polynomial method can be expected to work well for slowly varying backgrounds, while the overlap technique might be better for most rapidly varying backgrounds.

Three different methods of determining the overlap limits were explored. Examining the background PDF used in the signal extraction published in the NCD PRL yielded overlap limits of -1.0 and 1.0, while examining a set of example NNNAs using the author’s experience in modelling the NCDs and their backgrounds suggested that -1.2 and 1.2 might be useful limits. Examination of the blind NCD data, however, suggested that the presence of a “bump” in the background at slightly lower energy than the neutron peak rendered both of these ranges too narrow²⁰. Examining the upper and lower background limits in the blind data set suggested that overlap limits of -0.3 and 1.7 should span the actual background present in the blind data set, and hence (likely) in the full NCD data set.

²⁰ Note that this observation does not necessarily imply that the published NCD results were significantly biased; it appears that the narrower energy range used in fitting the published data, combined with the freedom in the published PDFs allows the fit to adapt to any background “bump” in much the same way as the polynomial method does.

Both the polynomial and overlap methods have been tested using Monte Carlo data, including data containing example NNNA's, and seem to function as designed. Tests on a blind set of NCD data also yielded reasonable results, once the overlap limits were adjusted as described above. The fits also reacted as expected both to the addition of data from the known NNNA strings and to the inclusion of data from the ^4He strings with the blind data set. In all, both the polynomial and overlap methods appear to function as designed, and as a result these methods appear to represent valid new techniques for extracting the number of neutrons from the poorly defined NCD background.

Chapter 5. Robust NCD Measurement of the Total Solar Neutrino Flux

Having developed two methods with some ability to extract the number of neutrons from a poorly defined background, it is now possible to use these methods to extract the number of neutrons from the full NCD signal, and hence to make an NCD-only deduction of the neutral current flux. In doing so, additional systematic uncertainties associated with the use of the ^{24}Na PDF as the neutron shape must be evaluated and a number of conversion factors dealing with detector live time, efficiency, and neutron backgrounds must be derived.

5.1 The NCD Neutron Number

Applying the polynomial to the full NCD data set yields 1156_{-77}^{+78} neutrons (from a 5^{th} order polynomial with a fit χ^2 of 33.9 in 45 degrees of freedom), while applying the overlap method (with overlap limits of -0.3 and 1.7) to the full NCD data yield a neutron confidence interval of (1058, 1394). The fitted polynomial and overlap limit backgrounds are shown in Figure 5.1, while the results of the polynomial fits of different order are shown in Table 5.1.

From these results it seems that the full NCD data are substantially similar to the blind data, as expected. The spectral feature which appeared in the overlap backgrounds at around 0.5 MeV in the blind data appears to still be present in the full data set. The polynomial method still chooses a higher order background, presumably to accommodate this feature, and the overlap limits chosen based on the blind data result in overlap limit backgrounds which appear by eye to nicely bracket the range of reasonable backgrounds. The polynomial and overlap results are in good agreement, with the polynomial result being more precise.

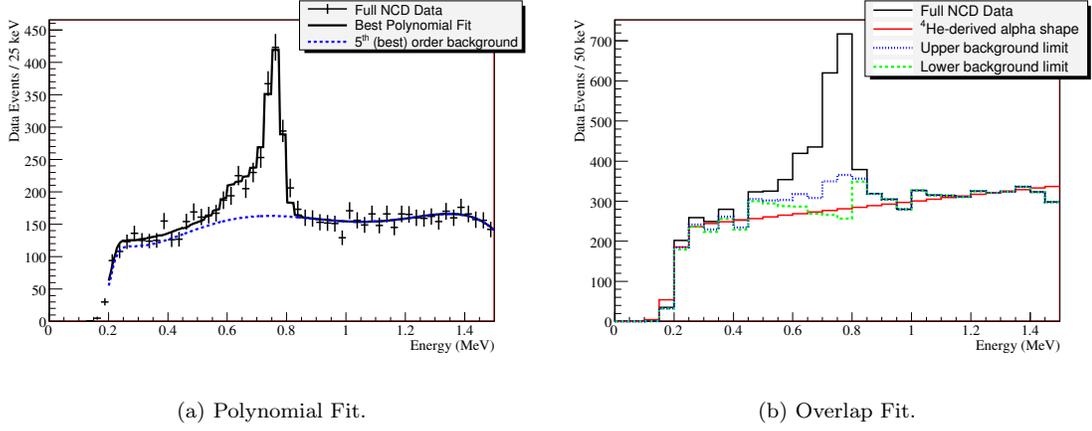


Figure 5.1: The results of fits to the full NCD data set. The polynomial fit results including the best fit 5th order background are shown in (a). The upper and lower limit backgrounds determined by the overlap limit are shown in (b), along with the ⁴He-based alpha shape that was subtracted in calculating the overlaps.

| Polynomial Order | Fitted Neutrons | Fit χ^2 /NDF |
|------------------|------------------------------------|-------------------|
| 0 | 1213 ⁺⁶¹ ₋₆₁ | 115.7/50 |
| 1 | 1322 ⁺⁶⁵ ₋₆₄ | 56.7/49 |
| 2 | 1258 ⁺⁶⁸ ₋₆₆ | 43.8/48 |
| 3 | 1246 ⁺⁷¹ ₋₇₀ | 43.3/47 |
| 4 | 1245 ⁺⁷² ₋₇₂ | 43.4/46 |
| 5 | 1156 ⁺⁷⁸ ₋₇₇ | 33.9/45 |
| 6 | 1156 ⁺⁷⁸ ₋₇₇ | 33.9/44 |

Table 5.1: The results of fits to the NCD data set using different polynomial orders to represent the background. Order 5 was selected by the polynomial selection procedure as the best description of the background in the data.

5.2 Neutron Number Systematic Uncertainties

A significant advantage to the use of the overlap and polynomial methods in determining the number of neutrons detected by the NCDs is that many possible fit systematic uncertainties, especially those involving the shape of the background spectrum, are automatically treated by the fits and are therefore included in the fit “statistical” uncertainties. In fact, the only uncertainties in the extracted number of neutrons which must be treated independently of the fit itself are the systematic uncertainties associated with the use of the ^{24}Na spectrum to represent the total neutron spectrum in the data. The potential contributions to that systematic are considered individually below.

5.2.1 Energy Resolution

As the 2005 ^{24}Na spike was taken over a relatively short period of time (~ 4 days), it is possible that the energy calibration of the NCDs was different during the ^{24}Na data taking period than it was, on average, during neutrino data taking. In particular, it is possible that slow drifts in detector or electronics gain, which are periodically corrected by energy calibrations, could have resulted in a broadening of the energy resolution in the data compared to that of the ^{24}Na spike data.

As described earlier, the NCD energy scale was set by two calibrations. First, calibration pulses were injected into the pre-amplifiers as part of the “NCD ECA” calibrations in order to measure the gain and offset of each electronics channel, thus providing a conversion between Q , the charge recorded by the shaper, and q , the charge collected by the anode wire. Second, neutron source runs were used to determine the relationship between q , the collected charge, and E , the energy deposited in the NCD gas, by fitting the position of the neutron peak in “q,” and setting that peak position to be 764 keV (note that some average space charge effect was included in this conversion). The energy of an NCD event with shaper charge Q was thus taken to be [44]

$$E = F_{q2E} \left(\frac{Q - a_0}{a_1} \right), \quad (5.1)$$

where a_0 and a_1 are the ECA offset and slope parameters, and F_{q2E} is the charge-to-energy conversion taken from the neutron calibration data. The calibration constants were determined individually

for each string, and were applied to data taken from the time of that calibration until the next calibration. Note that it was not possible to combine a_1 and F_{q2E} because the ECA and neutron calibrations were not done at the same time, and because there is a strong correlation between a_0 and a_1 .

If any array-wide drift occurred in either the electronics gain or the NCD gain over time, one would expect to see a corresponding trend in the calibration constants described above. However, as can be seen in Figure 5.2, no such trend exists. Instead, it appears that the calibrations changed in a random manner within a relatively narrow spread. This suggests that no significant array-wide drifts occurred, and that the changes between the measurements simply reflect fluctuations in the energy scale. Depending on the time scale of these fluctuations, then, the average energy resolution during the ^{24}Na data taking period may have been the same as, or narrower than, the average resolution during neutrino running.

To estimate the systematic uncertainty due to the ^{24}Na possibly having a narrower energy resolution than the data, the ^{24}Na distribution was widened by convolving it with the distributions shown in Figure 5.2 (this was done by shifting the energy of each event in the PDF by an amount determined by randomly sampling the Figure 5.2 distributions). This is equivalent to assuming that the ^{24}Na energy resolution was perfect and that Figure 5.2 gives a good measure of the relative width of the energy resolution in the data. The fits were repeated using this broadened neutron PDF, and the changes in the neutron fit parameters were taken as the associated systematic uncertainties¹. As can be seen in Figure 5.3, the changes in the F_{q2E} parameter had a greater effect on the neutron PDF than did the changes in the a_0 and a_1 parameters, and this was reflected in a significantly larger effect of F_{q2E} on the neutron fits. In the case of the polynomial fit, for example, convolving the PDF with the F_{q2E} parameter changes increased the number of fitted neutrons by 1.9%, while the a_0 and a_1 changes increased the fit value by only 0.4%. Because the F_{q2E} PDF broadening

¹ Here, and in the evaluation of other systematic uncertainties where the systematic effect was applied by a Monte Carlo procedure, there was not a unique systematically varied ^{24}Na PDF due to fluctuations in the Monte Carlo systematic uncertainty application. In these cases, the “apply systematic and refit” procedure was repeated 100 times, and a distribution of systematically varied fit parameters created. The systematic uncertainty was then set by defining the range about each central value parameter which contained 68% of the systematically varied fit parameters, and the half-width of that range was taken to be the systematic uncertainty.

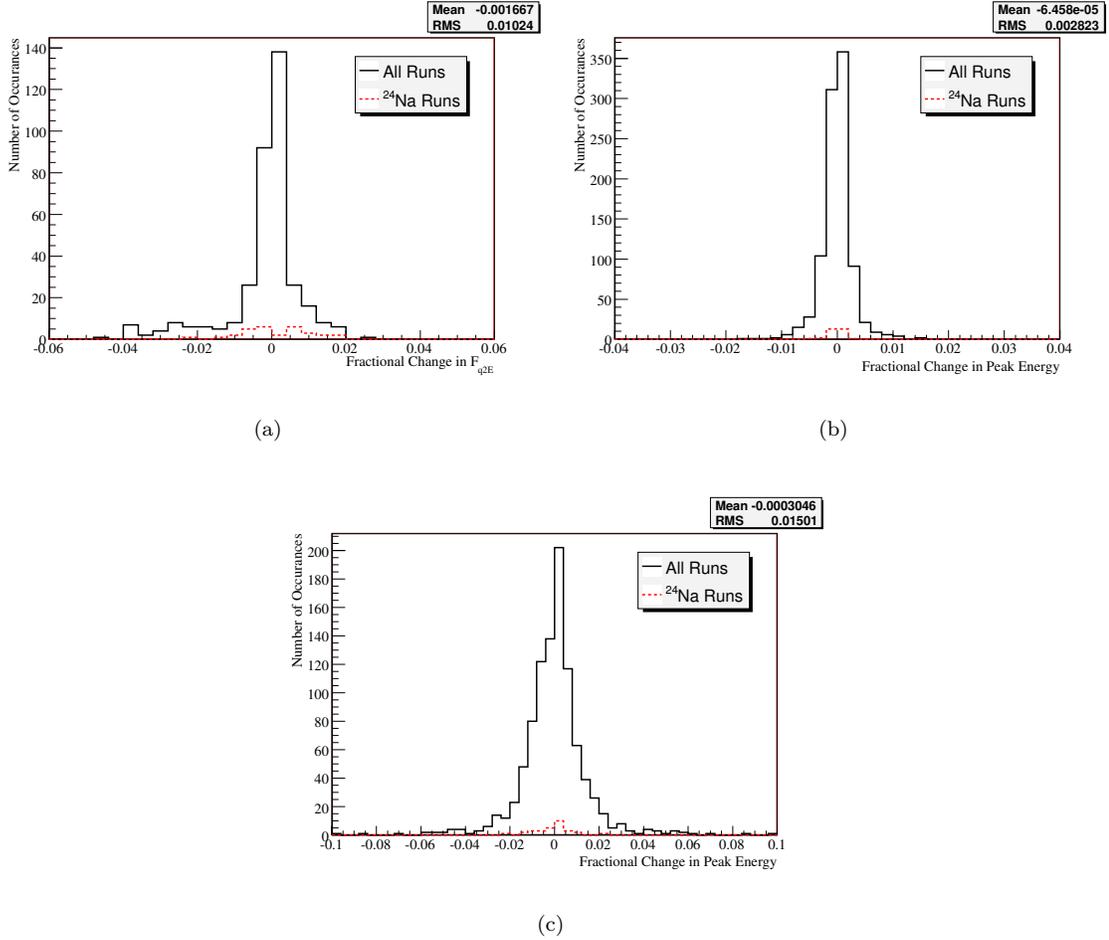


Figure 5.2: The fractional changes in the NCD energy calibration constants between subsequent measurements. The changes shown are between subsequent measurements of that constant on the same string; however, the changes for all of the analyzed strings are shown on the same plot. The changes in F_{q2E} are shown in (a). Because the changes in the a_0 and a_1 parameters were correlated, (b) and (c) show the fractional changes in the reported energy for an event in shaper ADC bin 130 (near the neutron peak) and ADC bin 20 (near the low end of the neutron tail), respectively, that result from the correlated changes of these parameters with each measurement. It can be seen that the changes in the a_0 and a_1 parameters have a larger fractional effect on lower energy events, primarily due to the proportionately larger influence of the a_0 offset on these events. The changes in the calibration constants between measurements before and after the 2005 ^{24}Na spike are also shown for comparison. The means and RMS's shown are for the full set of neutrino runs.

is independent of energy while the a_0 and a_1 broadening is energy dependent, these tests are not evaluating exactly the same systematic, and hence both effects are included in Table 5.2 and the overall uncertainty.

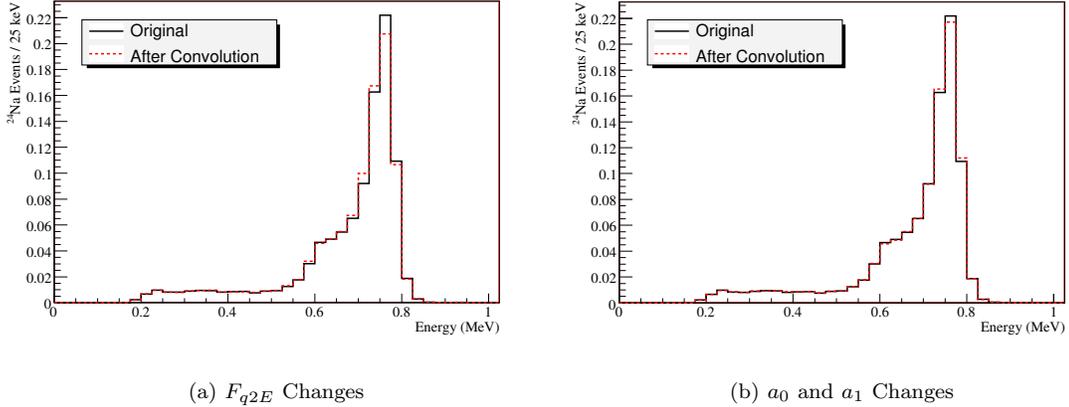


Figure 5.3: The change in the ^{24}Na energy PDF after smearing with (a) the observed changes in the F_{q2E} calibration constant and (b) the observed changes in the a_0 and a_1 constants. The smearing was accomplished by randomly choosing, from the distributions used to create Figure 5.2, a change in the relevant energy calibration constant for each event in the ^{24}Na . In the case of the a_0 and a_1 constants, applying the change to the ^{24}Na event required converting the event energy back to q using a representative F_{q2E} . After this, the “old” a_0 and a_1 constants were used to determine Q , and the “new” a_0 and a_1 constants (and the same representative F_{q2E}) were then used to convert that Q back to the (jittered) E . In both of these smearings, no attempt was made to correlate the string from which the event was taken with the string from which the change in calibration constant occurred.

5.2.2 Energy Scale

It is also possible that the calibration of the energy scale was slightly different during ^{24}Na running than it was, on average, during neutrino running. The possible change in the energy scale between the ^{24}Na data and the neutrino data can be estimated by comparing the mean change in the calibration parameters before and after the ^{24}Na spike and during neutrino running (i.e. by comparing the means of the distributions shown in Figure 5.2). The mean change in the F_{q2E} parameter between measurements before and after the ^{24}Na spike is 0.1%, while the corresponding change due to a_0 and a_1 is 0.1% in ADC bin 20 and 0.03% in ADC bin 130 (see the Figure 5.2 caption for approximate

calibrations of these ADC bins). In the total data, the mean change in F_{q2E} is -0.2%, and the mean changes due to a_0 and a_1 are less than 0.1%. The difference in the mean energy scale between the ^{24}Na spike and the full NCD data (dominated by the F_{q2E}) can therefore reasonably be assumed to be less than 0.3%. To estimate the effect of this potential change in energy scale on the number of fitted neutrons, the ^{24}Na PDF was shifted up and down (event-by-event) in energy by 0.3% and the data re-fitted. The largest change in each neutron fit parameter as a result of these shifts was taken to be the energy scale related systematic uncertainty associated with that parameter; these are tabulated in Table 5.2.

5.2.3 MUX Threshold Effects

During the time of the 2005 ^{24}Na spike, the MUX thresholds of all of the analyzed strings were set to their nominal values. During the rest of the neutrino running period, however, the MUX thresholds on various strings were occasionally raised by different amounts to suppress occasional periods of increased noise. The effect on the data is expected to be small, as only 2.8% of neutron events are expected to have been captured by one of these strings during a time when the threshold was raised [45]. Nevertheless, as can be seen in Figure 5.4, the neutron spectrum will be different on strings with higher MUX thresholds than in the ^{24}Na PDF.

To investigate the possible effect of this MUX threshold distortion on the signal extraction, the ^{24}Na PDF was rebuilt, with events randomly rejected (based on the known fraction of array live time spent in each raised MUX threshold setting and the energy dependant raised MUX threshold efficiencies inferred from the spectra shown in Figure 5.4) to simulate ^{24}Na data taking with MUX threshold live time fractions equivalent to the neutrino data taking period. The effect of the simulated threshold increases on the neutron PDF are so small as to be indistinguishable by eye, and the smallness of the change is reflected in the effect on the neutron fits. In fact, as seen in Table 5.2, the spectral effect of the raised MUX thresholds changed the neutron fit parameters by 0.2% or less.

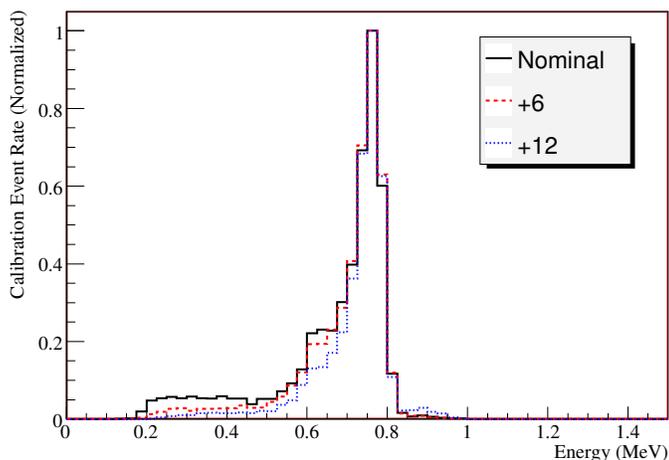


Figure 5.4: The summed spectrum of AmBe source data at several positions within the SNO detector. For each source position, data were taken with all of the MUX thresholds set to their nominal values, raised by 6 units and raised by 12 units (a “MUX threshold unit” is the software setting used to determine the MUX threshold). As the MUX threshold is increased, the neutron spectrum is distorted as more and more lower energy events fail to trigger the MUX (the curves in the Figure are normalized such that the neutron peak heights are equal; it may actually be that even at the peak energies the MUX efficiency is reduced.)

5.2.4 ^{24}Na Statistics

The 2005 ^{24}Na spike data contain approximately 17,000 events in the neutron energy region. As this is a large but not overwhelming number of events, the statistical uncertainty in the ^{24}Na PDF could have an important effect on the signal extraction. To test this, the bin entries of the ^{24}Na PDF were repeatedly jittered (according to Poisson Statistics) and the data re-fitted with these new PDFs. As shown in Table 5.2, this leads to a systematic uncertainty of just less than 1.0%.

5.2.5 ^{24}Na Non-Uniformity

A significant concern for the NCD neutron efficiency measurement was the possibility that the ^{24}Na brine may not have mixed completely uniformly within the D_2O volume. In particular, there was concern that a “dead layer” free from brine might have existed along the inside of the acrylic vessel. This would have increased the brine concentration in the interior regions of the detector and hence increase the apparent neutron capture efficiency. From the point of view of the neutron PDF, non-

uniformity of the brine in the detector could change the string-to-string distribution of neutron events (which, along with the z-distribution of the neutrons within a string, is known from tests with point source data to have significant effects on the neutron spectrum²). Thankfully, however, the neutron PDF turns out to be relatively insensitive to such non-uniformities; it has been shown [46] that re-weighting the contributions of the different strings according to a PMT-derived estimate of the ^{24}Na non-uniformity changes the results of polynomial-type fits by less than 0.1%. The effect on the overlap limits is expected to be similar, and the smallness of the uncertainty makes it unnecessary to re-evaluate the systematic in these cases - 0.1% will simply be adopted for the overlap limit uncertainties as well.

5.2.6 Backgrounds in the ^{24}Na

A final consequence of using a data-derived PDF for the neutron shape is the possibility that events besides ^{24}Na neutrons will be present in the PDF. These could include alpha events, neutron events from sources besides the ^{24}Na , and non-physics backgrounds.

Alpha Contamination

The rate of alpha decays in the NCD array should not have changed during the ^{24}Na calibration, and we know that alphas dominated the neutron region during neutrino data taking. We can therefore use the neutrino data rate to estimate the alpha contamination of the ^{24}Na PDF. Treating every event below 1 MeV in the neutrino data as an alpha event, we find an alpha rate of 15.8/day, which translates to 64 alpha events below 1 MeV in the ^{24}Na neutron PDF, or about 0.3% of the total PDF. This represents a direct uncertainty on the extracted number of neutrons, as it means that some small number of alphas in the data will be “recruited” as neutrons to make up this portion of the PDF (in fact, this should be a correction rather than an uncertainty, but it is small enough to simply be taken as an uncertainty).

² Presumably these effects are dominated by calibration differences between strings and by gain differences between the counters within a string.

Other Neutrons

Non-uniform background neutron sources could influence the ^{24}Na PDF in a manner similar to non-uniform mixing of the ^{24}Na brine. The total neutron background rate, however, is much less than the alpha rate, meaning that the “other neutron” contribution to the PDF will be less than 0.3%, from above. This is a much smaller non-uniformity than the non-uniformity of the ^{24}Na brine itself, which was already found to be negligible.

Other Backgrounds

Finally it is possible that the ^{24}Na PDF could contain non-physics events which pass the data cleaning cuts. If the rate of these events was the same during ^{24}Na data taking as during neutrino running the total number of such events would be negligible (and would already be included in the alpha contamination systematic discussed above). It has been shown, however, that the rate of events failing the data cleaning cuts decayed during the ^{24}Na data taking period with a half-life very similar to that of ^{24}Na [46]. It is therefore possible, if there is significant leakage through the data cleaning cuts, that the leakage events occurred in the ^{24}Na PDF at a rate proportionally comparable to their rate in data. This would be a most unfortunate situation, as the leakage events in the data would then be fitted out as neutrons. The possibility of such a “just so” scenario seems remote, however, and will be neglected here as it has been in other NCD analyses. Nevertheless, the lack of quantitative estimates of the number of spurious events passing the data cleaning cuts means that no quantitative limit can be placed on their effect.

5.2.7 Neutron PDF Systematic Uncertainties Summary

A summary of the different systematic uncertainties associated with the use of the ^{24}Na neutron PDF are shown in Table 5.2. As can be seen, the systematic uncertainties are dominated by the possible differences in the average energy scale and resolution between the ^{24}Na neutron PDF and the data, and by the statistical uncertainty in the neutron PDF. The magnitude of these systematic uncertainties as evaluated here is very similar to (but slightly smaller than) their magnitudes as

evaluated in [21] and [46].

| | Neutron Fit Parameter | | |
|---------------------------------------|-----------------------|---------------|---------------|
| | Polynomial Fit | Overlap Lower | Overlap Upper |
| Energy Scale | 0.5% | 1.7% | 1.0% |
| Energy Resolution (F_{q2E}) | 1.9% | 1.6% | 1.0% |
| Energy Resolution (a_0 and a_1) | 0.4% | 0.2% | 0.1% |
| MUX Threshold | 0.2% | 0.1% | 0.1% |
| ^{24}Na Non-Uniformity | 0.1% | 0.1% | 0.1% |
| ^{24}Na PDF Statistics | 0.9% | 0.8% | 0.5% |
| Alpha Contamination | 0.3% | 0.3% | 0.3% |
| Total | 2.2% | 2.5% | 1.5% |

Table 5.2: The final contributions of the different ^{24}Na PDF systematic uncertainties to the different neutron extraction parameters.

In applying these systematic uncertainties to the polynomial and overlap fits, they are assumed to represent the standard deviations of Gaussian curves centred about the central value fit parameter. In the case of the polynomial method, where the fit statistical uncertainties are interpreted in a similar way, the statistical and systematic errors can then be combined in quadrature in the usual way. For the overlap method, however, recall that the confidence region should not be expected to represent the width of a Gaussian distribution with the most probable value of the true neutron number at the centre; to be consistent with the flat prior that was assumed on the background overlap, all that can be expected is that the true number of neutrons should lie between the upper and lower neutron limits. In considering the systematic uncertainties on the overlap neutron limits, then, the upper neutron limit should be shifted further up to include the systematic uncertainty; taking the systematic uncertainty as describing the width of a Gaussian centred about the original upper overlap neutron limit, the new upper neutron limit should be chosen such that 68.3% of the systematic distribution lies below that value. As 68.3% of a normal Gaussian distribution lies below 0.48σ , for a systematic uncertainty of 1.5%, the upper overlap limit should be raised by $1.5 \times 0.48 = 0.7\%$. The lower overlap neutron limit should be lowered in a similar way. Note, however, that additional systematic uncertainties which will follow should be added in quadrature with these

before the limits are shifted. For now, therefore, we will keep track of them simply as an uncertainty on each limit. Thus, after including the systematic uncertainties due to the use of the ^{24}Na neutron PDF, the number of neutrons fitted using the polynomial technique is $1156_{-77}^{+78}(\text{stat}) \pm 25(\text{syst})$, while the neutron confidence interval defined by the overlap method is $(1058 \pm 26, 1394 \pm 21)$.

5.2.8 Verification with 2006 ^{24}Na Data

As a check that the systematic uncertainty estimates described above are reasonable, the data were refitted using the 2006 ^{24}Na spike data as the neutron PDF. This spike has fewer events than the 2005 spike, and so should be expected to have even larger uncertainties due to statistical fluctuations in the PDF. It was also recorded using a different set of energy calibration constants, so any offsets in energy scale might be expected to be different. As the 2006 spike also took place during a limited period of time, however, the energy resolution uncertainty might be expected to be similar (under the model that the resolution is “perfect” during the ^{24}Na data taking period relative to the changes observed during the rest of the data taking).

After switching to the 2006 ^{24}Na PDF, the number of neutrons returned by the polynomial fit decreased by 0.8%, the upper overlap neutron limit decreased by 2.1% and the lower overlap neutron limit decreased by 2.2%. These changes are reasonably consistent with the expectations from Table 5.2. Even if the extreme assumption that the energy resolution is perfect during the ^{24}Na data taking periods is accepted and the energy resolution contributions are removed from Table 5.2, the changes expected based on the from the systematic uncertainties are 1.1%, 1.2%, and 1.9%, respectively, which is not so different from that observed (the largest deviation is 1.8σ in the the upper overlap limit). This suggests that the evaluation of the systematic uncertainties was indeed reasonable.

5.3 Live Fraction and Detection Efficiency

Not all neutrons that were captured by the NCDs during a period of time when the detector was operating produced events which were included in the analysis. Some neutron captures failed to trigger either the MUX or shaper systems and hence were not recorded, while others occurred while

some portion of the system was reading out a previous event and was hence unable to accommodate the new event. Still other events were successfully recorded but then erroneously rejected by the data cleaning algorithms. The number of neutrons extracted from the NCD data must therefore be corrected to account for the detector live fractions and the trigger and data cleaning efficiencies.

5.3.1 Shaper Live Fraction

The shaper dead time after an event on any one string was quite short ($236 \pm 5 \mu\text{s}$), and applied to all of the strings [28]. The 1.4×10^6 shaper triggers during the NCD phase thus translated to $330 \pm 7\text{s}$ of dead time (or $(9.9 \pm 0.2) \times 10^{-4}\%$ of the total run time) during neutrino running. This dead time is negligible.

5.3.2 MUX Live Fraction

The MUX dead time per event was ~ 1 ms [28], but was somewhat variable. In order to determine the MUX live time, then, a system of live time clocks was used. One of these “live time scalers” (all of which were based on a local 10 MHz oscillator) was never inhibited, and thus gave a measure of the full live time of the run. A second scaler was inhibited while the MUX system was locked out by the MUX controller, and hence recorded the MUX live time during the run. Comparing the output of these clocks for each run and performing a time weighted average gives a total average MUX live fraction of 0.9980 ± 0.0001 for the NCD data taking period [47]. During the period of time the random pulser (described in the next section) was running, the random pulser could be used to verify the live time scaler result; the two methods agree to within $7 \times 10^{-3}\%$.

5.3.3 Scope Live Fraction

The readout time for a scope trace, and hence the dead time associated with a scope event, was quite long ($\sim 0.6\text{s}$). To mitigate this dead time, two scopes were used, with a flip-flop bit used to alternate MUX events between the scopes. Thus, in principle, any event which occurred more than 0.6s from the second most recent scope event was expected to find a live scope. In practice, however, the scope read-out time was quite variable and the flip-flop bit did not always function exactly as

intended. Therefore, an empirical scope live fraction determination was required.

The scope live fraction is defined as the fraction of time that an event triggering the MUX would find a live scope. Thus, partial MUX events can be used to define the times during which the scope system was dead, while scope events define the times during which the system was live³. To provide the random sampling necessary to obtain an accurate measurement of the scope live fraction, a “random pulser” system was installed [28], which randomly pulsed a spare NCD channel (with a height-adjustable square wave pulse) at an average rate of 0.01Hz. For the data taking period during which the random pulser operated, then, the scope live fraction could be determined by dividing the number of scope events on the random pulser channel by the sum of the number of scope and partial MUX events.

Unfortunately, the random pulser was not installed during the first part of the data taking period, meaning that empirical scope live fraction measurements were not available during this time. In addition, bursts of events were occasionally observed on the random pulser channel [48]. In order to confirm the random pulser results and to extrapolate them to the complete data taking period, a semi-empirical scope live fraction determination was developed. In this semi empirical method, the scope and partial MUX events were used to determine the scope live fraction as a function of the time since the most recent event on each of the scopes, as shown in Figure 5.5. The average scope live fraction for the data taking period was then integrated by sampling this distribution using a Monte Carlo technique in which the real timeline of the data taking period was repeatedly randomly sampled. For each randomly sampled time, the probability that the scope was live was determined (from the live fraction distribution shown in Figure 5.5(c)) based on the difference between the sampled time and the most recent (real) event on each scope. In this way, the average scope live fraction as experienced by a uniform signal could be calculated. During the period of time when the random pulser was operating, the scope live fraction determined by the random pulser (0.960 ± 0.001) was in good agreement with the results from the semi-empirical calculation (0.959 ± 0.001). Applying

³ Note that the scope live fraction is not simply given by the number of scope events divided by the sum of the number of scope events and the number of partial MUX events because the partial MUX events tended to occur in bursts, while neutrino events are expected to be uniform in time. Using this simple procedure would therefore overweight the bursty data periods and underestimate the live fraction relevant to the neutrino signal.

the semi-empirical technique to the entire data taking period yielded an average scope live fraction of 0.957 ± 0.004 [49].

5.3.4 Shaper Threshold Efficiency

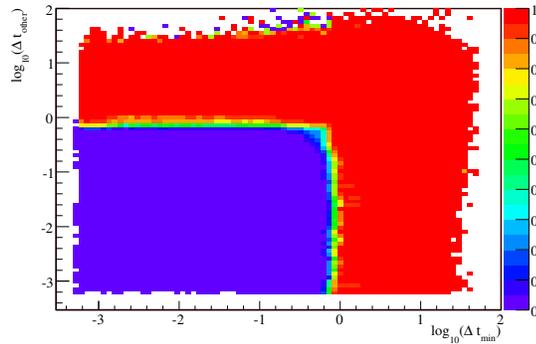
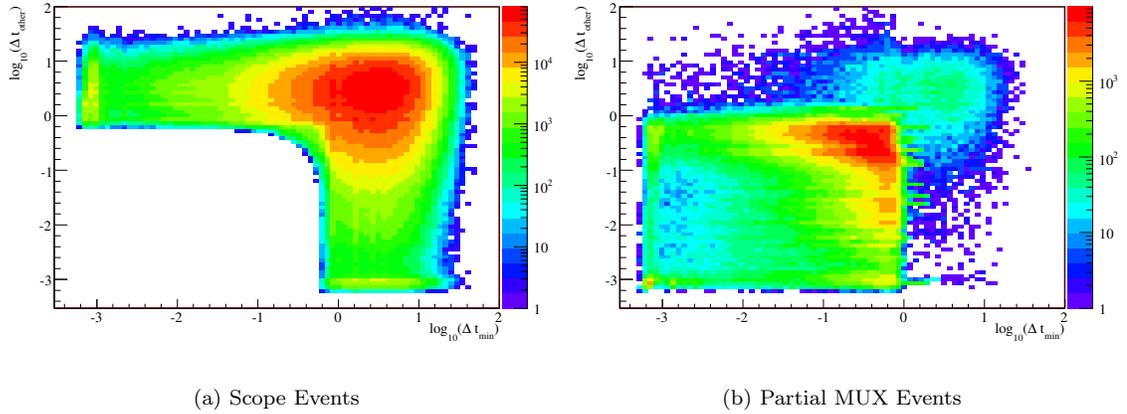
The shaper thresholds were intentionally set significantly below the lower edge of the neutron spectrum, and on the NCDs used in the analysis the shaper thresholds were never raised. Thus, when no energy cut is applied to the data, the shaper threshold efficiency was assumed to be 1.0.

5.3.5 MUX Threshold Efficiency

The probability that a neutron event had sufficient amplitude to pass the MUX threshold was measured by comparing, in dead time corrected AmBe point source calibration data, the number of events detected by the MUX and shaper systems. Given that the shaper threshold efficiency was assumed to be perfect, the MUX efficiency could be inferred by the difference in the number of observed events. This was done individually for each of the strings at each of the MUX threshold settings (the MUX thresholds on some of the strings were raised from time to time to mitigate episodes of bursty behaviour), and a live time weighted average was calculated. Over the full neutron energy range, the average MUX threshold efficiency was 0.98284 ± 0.00037 [45].

5.3.6 Data Cleaning Efficiency

The data cleaning cuts rejected a small but finite fraction of real neutron events, and these “false rejections” must be corrected for. The efficiency with which neutron events passed the data cleaning cuts was estimated from the fraction of events in neutron calibration runs (AmBe and Cf) which passed the data cleaning cuts. The use of high rate source data limited the fraction of non-neutron events in the data, and hence the possibility that the sacrifice was overestimated because of correct rejections of non-physics events. This effect was further reduced through the use of timing cuts which selected neutron-like events (i.e. events which appeared to be part of a fission burst for Cf source data and those preceded by a PMT event consistent with a 4.4 MeV γ -ray for the AmBe source). The average efficiency for neutron events to pass the data cleaning cuts was thus determined to be



(c) Scope Live Fraction

Figure 5.5: The empirical distributions used in calculating the semi-empirical scope live fraction. Δt_{min} is the time (in seconds) between the given event and the most recent event recorded by either scope, while Δt_{other} is the time in seconds between the most recent event recorded by either scope and the most recent preceding event on the other scope. (a) and (b) show the timing distributions of the scope events and the partial MUX events, respectively. Overall we see that the scopes tend to be live if an event is more than ~ 0.6 s after the previous event on at least one of the scopes, as expected. However, there are also unexpected events, including some events for which the scope system was dead more than 10s after the most recent event on either scope. Dividing the data shown in (a) by the sum of the data in (a) and (b) yields the average scope live fraction as a function of the time from the most recent events on both scopes. This distribution, shown in (c), was then used in the semi-empirical scope live fraction determination described in the text.

0.98338 ± 0.00021 [50]; this number was assumed to be stable in time.

5.3.7 Efficiency and Live Fraction Summary

A summary of the live fraction and threshold efficiency corrections is given in Table 5.3. The total correction can be multiplied directly with polynomial fit central values and the overlap neutron limits, with the uncertainties combined in the standard manner. The polynomial fit result for the number of neutron captures by the NCDs is then $1252_{-83}^{+85}(\text{stat}) \pm 28(\text{syst})$, and the overlap neutron capture confidence interval is $(1146 \pm 29, 1510 \pm 24)$.

| Source | Correction Factor |
|-----------------------------------|-----------------------|
| LF_{shaper} | – |
| LF_{MUX} | 0.9980 ± 0.0001 |
| LF_{scope} | 0.957 ± 0.004 |
| ϵ_{shaper} | – |
| ϵ_{MUX} | 0.98284 ± 0.00037 |
| $\epsilon_{\text{data cleaning}}$ | 0.98338 ± 0.00021 |
| Total Correction | 0.923 ± 0.004 |

Table 5.3: A summary of the NCD live fractions (LF) and efficiencies (ϵ) discussed in the text.

5.4 Neutron Backgrounds in the NCDs

When a neutron was captured by the NCDs, no distinction could be made between those neutrons which resulted from neutral current interactions and those originating from other sources. In order to extract the neutral current flux, therefore, the number of neutron captures in the NCDs produced by sources other than the ${}^8\text{B}$ solar neutrino neutral current interaction must be estimated and subtracted from the observed number. The additional sources of neutrons which have been identified are briefly described below. The rate of neutron production by each of these background sources was evaluated, and the capture efficiencies of those neutrons by the NCD array were estimated using point- and distributed-source verified Monte Carlo (with the exception of the uniform sources, where the ${}^{24}\text{Na}$ -derived neutron capture efficiency, described in the next Section, was used directly). The resulting expected numbers of neutron captures in the NCDs are given in Table 5.4.

- **D₂O Radioactivity:** Gamma rays with energies greater than 2.2 MeV have the potential to photo-dissociate a deuteron, producing a free neutron. Thus, radioactive impurities in the heavy water volume, in particular ²⁰⁸Tl and ²¹⁴Bi, have the potential to provide a neutron background that is indistinguishable from the neutral current neutrons. The levels of these contaminants were monitored through PMT-based *in-situ* analyses and through water assays [51], and their neutron contributions calculated by Monte Carlo.
- **NCD Bulk and Cables:** Gammas from radioactive impurities in the NCDs and NCD read-out cables could propagate to the heavy water volume and produce background neutrons. The ²³⁸U and ²³²Th chain impurities in the NCD components were assayed [28], and the resulting background contributions calculated.
- **NCD “Hot-spots”:** Three positions on two NCD strings were identified by the PMT-based *in-situ* low energy background analysis as having relatively high-level local contaminations of ²³⁸U and ²³²Th chain activity, which were likely the result of handling during deployment. Upon the removal of the NCDs, these hot-spot sections were counted in an external alpha counting system and assayed by elution of the contamination in a modified version of the water assay technique in order to accurately determine the background contribution of this activity [51].
- **H₂O and Acrylic Vessel Radioactivity:** Gamma rays from decays in the light water outside of the acrylic vessel and from the acrylic vessel itself could also propagate into the heavy water and produce neutrons through photo-dissociation. The contributions of these reactions are constrained through *in-situ* measurements of the light water and acrylic vessel activity, and through radio-assays of the light water contamination [52].
- **Acrylic Vessel Neutrons:** During the construction of the SNO experiment, the acrylic vessel was exposed for a significant period of time to (HEPA filtered) mine air, which has high levels of ²²²Rn. As a result, radon daughter isotopes plated onto the inner and outer surfaces of

the acrylic vessel, where some of them were slightly embedded through recoil from subsequent alpha decays. This resulted in long-lived ^{210}Po activity (supported by ^{210}Pb with a 22a half-life) on the acrylic vessel surfaces, which produces a 5.3 MeV alpha particle. These alphas are energetic enough to produce neutrons via (α, n) interactions with ^{13}C in the acrylic and ^{17}O and ^{18}O in the water. Because the ^{210}Pb daughters do not produce decay products energetic enough to be directly detected by the PMT system, it was not possible to directly monitor this background. However, during Phase II of SNO, the neutron capture distance in the salty D_2O was short enough that the external-source neutron contribution could be extracted from the radial profile of the detected events [20]; from this the acrylic vessel neutron production rate was determined. In the NCD phase, the neutron capture length was too long for the external neutron component to be identified. As a result, the salt phase measurement was used. To confirm that this was reasonable, the alpha activity of the inner surface on the neck of the acrylic vessel was measured with silicon detectors before and after the NCD phase [53]. This activity was found to be consistent (within a large uncertainty which was applied to the extrapolated salt acrylic vessel neutron rate), supporting the hypothesis that the production rate of neutrons by acrylic vessel surface contamination was constant between Phases II and III.

- **NCD (α, n)** Radon daughter activity on the outer surfaces of the NCDs had the potential to produce neutrons through (α, n) interactions with ^{17}O and ^{18}O in the water. To estimate this contribution, the surface activities of several sections of the NCDs were counted with a custom-built gas proportional chamber [54].
- **Atmospheric Neutrinos** Cosmic-ray induced atmospheric neutrinos could also interact with deuterium to produce neutrons. The flux of these neutrinos is much lower than the solar neutrinos, but their energies and hence their interaction cross-sections are higher. This background contribution is evaluated using Monte Carlo calculations based on published estimates of the atmospheric neutrino flux [55].

- ***hep* Neutrinos** The *hep* neutrino flux is predicted to be 718 ± 162 times less than the ^8B flux [3], with an average neutral current interaction cross-section that is 2.8 ± 0.4 times larger [25]. Scaling the expected neutron capture rate from ^8B (described below) thus gives $4.1 \pm 1.1(\text{syst})$ as the expected *hep* contribution to neutron capture in the NCDs.
- **Other Neutron Backgrounds** Finally, a variety of other sub-dominant neutron production mechanisms, each of which is expected to produce less than two neutron captures in the NCD array over the course of the NCD phase, have been evaluated [55]. These include neutron production by cosmogenically created isotopes, the neutral current interactions of reactor anti-neutrinos, geo-neutrinos and CNO solar neutrinos, deuteron photo-dissociation by cosmogenically produced ^{16}N , and the spontaneous fission of and (α, n) interactions produced by background radioactivity in the heavy water.

| Source | NCD Captures |
|--------------------------------|------------------|
| D ₂ O | 33.3 ± 5.5 |
| NCD bulk and Cables | 39.3 ± 15.0 |
| NCD “Hot-spots” | 74.7 ± 7.5 |
| H ₂ O and A.V. Bulk | 23.8 ± 17.1 |
| A.V. Surface | 23.7 ± 16.8 |
| NCD Surface | 2.1 ± 0.47 |
| Atmospheric ν | 15.8 ± 3.1 |
| <i>hep</i> Solar ν | 4.1 ± 2.3 |
| Other | 2.7 ± 0.4 |
| Total | 219.4 ± 33.4 |

Table 5.4: A summary of the neutron events expected in the NCDs from sources other than the neutral current interactions. The uncertainties on the individual contributions are systematic only; both statistical and systematic uncertainties are included in the total.

The expected number of background neutron captures must be subtracted from the number of NCD neutron captures developed earlier. The uncertainty in the number of background neutron captures becomes part of the systematic uncertainty in number of neutrino-induced neutron captures. The number of neutron captures in the NCDs resulting from ^8B solar neutrino neutral current

interactions is thus determined to be $1033_{-83}^{+85}(\text{stat}) \pm 44(\text{syst})$ based on the results of the polynomial fit, and in the range $(927 \pm 44, 1290 \pm 41)$ based on the overlap method.

5.5 Live Time and Capture Efficiency

Two final numbers are required to convert the number of observed ^8B neutral current neutron captures into a ^8B neutron production rate: the length of time for which the detector was operated and the efficiency with which uniformly produced neutrons were captured by the NCDs.

5.5.1 Live Time

The total live time included in the NCD phase data set was measured using both 10 MHz and 50 MHz scaler clocks (similar to the 50 MHz clock described earlier in the determination of the MUX live fraction). This live time was then corrected for the periods of time that were removed offline by muon- and burst-related data cleaning cuts (about 2% in total). The total live time thus determined was 385.17 ± 0.14 days [56]. This number was verified using the data from the NCD random pulser (during the period of time for which the random pulser was operational) and the data from a 5 Hz random trigger that was applied to the PMTs.

5.5.2 Capture Efficiency

The efficiency with which the NCD array captured neutrons produced uniformly throughout the D_2O volume was determined in two ways. The primary determination used the ^{24}Na spikes; given that the injected activity of ^{24}Na was known⁴, the number of neutrons produced by the ^{24}Na could be calculated. The neutron capture efficiency was then the ratio of the number of neutrons produced to the number detected (corrected for threshold efficiencies and dead time to avoid double counting these). The neutron capture efficiency thus determined was 0.211 ± 0.007 , with the uncertainty dominated by the measurement of the injected brine activity and by uncertainty in the uniformity of the distribution of the brine in the heavy water volume [58].

⁴ Samples of the injected brine were counted in a sealed container at the centre of the SNO detector (where their activities could be compared to sources of known strength), and on a low background germanium detector in the SNOLAB facility to determine the injected activity [57].

To confirm this distributed source measurement, the SNOMAN Monte Carlo (which had been extensively verified against neutron point source calibration data) was also used to calculate the capture efficiency for uniform neutrons. The Monte Carlo capture efficiency was 0.210 ± 0.003 [57], in good agreement with the ^{24}Na value.

5.5.3 Live Fraction and Detection Efficiency Summary

The rate of ^8B neutral current interactions in the heavy water volume can now be calculated by dividing the number of ^8B neutral current neutron captures in the NCDs by the live time and the capture efficiency. Uncertainties in the correction factors become part of the systematic uncertainty in the neutral current rate. The polynomial fit results give a ^8B solar neutrino neutral current interaction rate of $12.71_{-1.02}^{+1.05}(\text{stat}) \pm 0.69(\text{syst}) \text{ d}^{-1}$, while the overlap method gives a ^8B solar neutrino neutral current interaction rate in the range $(11.41 \pm 0.66, 15.87 \pm 0.73)$.

5.6 NCD ^8B Flux Result

Having determined the neutral current interaction rate, it is now possible to determine the total solar neutrino flux.

5.6.1 Neutron Rate to Flux Conversion

The conversion from neutron production rate to solar neutrino flux is carried out by calculating, using the SNOMAN Monte Carlo, the *expected* neutron production rate in SNO and then scaling to the *observed* neutron production rate. For the NCD phase, the Monte Carlo used a ^8B neutrino flux of $5.145 \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$, and the Winter-Friedman ^8B neutrino energy spectrum [59]. The neutrino-deuteron neutral current interaction cross-section was taken from [60], with radiative corrections as calculated in [61]⁵. The SNO heavy water target is taken to be a sphere of radius 600.54 cm with an average deuteron density (after correcting for the NCD displacement) of $6.6352 \times 10^{28} \text{ m}^{-3}$. The Monte Carlo includes the effects of the Earth's orbital eccentricity, properly weighted by live time.

⁵ In order to do this, the partial radiative corrections in [60] had to be removed by applying a correction factor of $1./1.024$ [62].

Taking all of this into account, the average expected neutron production rate in the SNO heavy water during the NCD data taking period is 12.96 ± 0.13 neutrons per day [62].

5.6.2 The NCD ^8B Flux Measurement

The total ^8B flux can now be obtained by scaling the Monte Carlo neutrino flux by the ratio of the predicted and observed neutron production rates. The ^8B flux as measured by the polynomial method is thus determined to be $5.04_{-0.40}^{+0.42}(\text{stat}) \pm 0.28(\text{syst}) \times 10^6 \text{cm}^{-2}\text{s}^{-1}$, and the range of fluxes allowed by the results of the overlap method is $(4.53 \pm 0.27, 6.30 \pm 0.28) \times 10^6 \text{cm}^{-2}\text{s}^{-1}$. Now that all of the systematic uncertainties are included, it is possible to combine the systematic uncertainties with the confidence range (as described in Section 5.2.7) to determine a final overlap based flux confidence region of $(4.40, 6.43) \times 10^6 \text{cm}^{-2}\text{s}^{-1}$. A summary of the contributions to the systematic uncertainty of each of the flux parameters is given in Table 5.5.

| | Flux Parameter | | |
|--------------------------------------|----------------|---------------|---------------|
| | Polynomial Fit | Overlap Lower | Overlap Upper |
| ^{24}Na PDF Systematics | 2.67% | 3.09% | 1.17% |
| Detection Efficiency & Live Fraction | 0.52% | 0.53% | 0.50% |
| Neutron Background | 3.23% | 3.60% | 2.59% |
| Live Time & Capture Efficiency | 3.32% | | |
| Target & Cross Section | 1.0% | | |
| Total Systematic | 5.46% | 5.89% | 4.51% |

Table 5.5: The contributions of the different sources of systematic uncertainty to the ^8B flux parameters. Note that the fractional contribution of a given systematic to the total flux can be different than its fractional value as evaluated in earlier sections due to the subtraction of the neutron backgrounds. Note also that the systematic uncertainties are not dominated by any single source, but have approximately equal contributions from several sources.

As seen in Figure 5.6, the polynomial and overlap results agree well with the previously published SNO results and with the Standard Solar Model prediction. The total uncertainty for the polynomial method is 9.9%. This is slightly larger than the 8.8% total uncertainty in the NCD PRL [21]. However, this is a relatively modest increase in uncertainty given the relative simplicity and robustness of the technique compared to the published signal extraction. It is interesting to note

that the central value fit results of the polynomial method and the previously published NCD phase analysis differ at approximately the 1σ level. Although not particularly alarming, this is a larger change than one would expect in fitting identical data using similar fitting procedures. This suggests either that there was tension between the PMT and NCD data in the published fit, resulting in a central fit value that was pulled relative to the NCD-only best fit, or that the NCD background spectrum determined by the polynomial fit was slightly different than the backgrounds that the PDFs used in the published analysis could accommodate. In either case, however, the relatively small difference suggests that any effect in the published result is unlikely to be more than 1σ .

As predicted, the overlap method had a significantly larger uncertainty (the half range of the total confidence interval was 18.6% of the central value) compared to either the polynomial fit or the published result. Still, the confidence region agrees well with the other results, and 20% sensitivity seems reasonable given how little information about the backgrounds had to be supplied to the method.

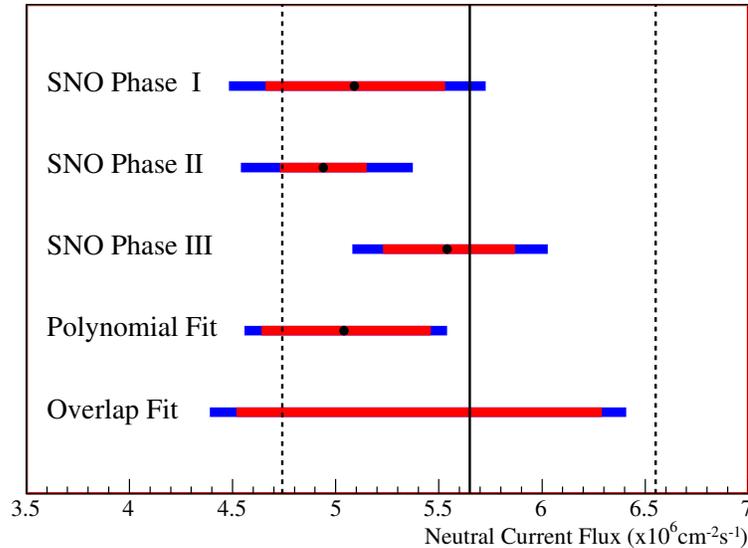


Figure 5.6: A comparison of the results of the polynomial and overlap methods with the previously published SNO neutral current flux results [19, 20, 21] and the BS05 SSM predictions [3]. The current results can be seen to agree well with both the previously published results and the SSM prediction.

5.7 SNO Conclusion

Two new signal extraction techniques have been developed which allow the number of detected neutrons to be extracted from NCD data using minimal information about the poorly understood background spectra in the NCDs. This has the effect of making these signal extraction techniques more robust against unexpected behaviour in the background at the cost of increasing the uncertainty in the extracted number of neutrons. These new robust signal extraction techniques were applied to the NCD data, yielding a total ^8B solar neutrino flux confidence interval of $(4.40, 6.43)\times 10^6\text{cm}^{-2}\text{s}^{-1}$. From the overlap method and a ^8B solar neutrino flux of $5.04_{-0.40}^{+0.42}(\text{stat}) \pm 0.28(\text{syst})\times 10^6\text{cm}^{-2}\text{s}^{-1}$ from the polynomial method. The published NCD phase flux result $(5.54_{-0.31}^{+0.33}(\text{stat})_{-0.34}^{+0.36}(\text{syst})\times 10^6\text{cm}^{-2}\text{s}^{-1})$ agrees well with the overlap confidence interval and is higher than the polynomial result by approximately 1σ . This deviation appears to be the result of an unexpected step or bump in the background of the NCD data at around 0.5 MeV. Overall, though, the fairly good agreement between the published analysis and the results of the robust signal extraction methods confirms that the published result was not unduly affected by unexpected background behaviour.

Chapter 6. The SNO+ Experiment

The SNO+ experiment is a new liquid scintillator based neutrino experiment that is being constructed in SNOLAB using the SNO detector. A liquid scintillator is an organic liquid that produces light when excited by the passage of a charged particle. In fact, liquid scintillators produce significantly more light than the water Čerenkov process, which will make SNO+ sensitive to events of much lower energy than those that were detectable by SNO. This low energy sensitivity will allow SNO+ to investigate a new regime of interesting physics.

Building SNO+ from the remaining SNO apparatus will mean that SNO+ can be constructed much more quickly, and at much lower expense, than a comparable experiment in another location. SNO+ will also inherit many of the features that made SNO successful, including the relatively high photocathode (PMT) coverage, the excellent shielding provided by the large amount of external light water, the low rate of cosmogenic backgrounds afforded by the very deep location, and the very clean environment that was maintained within the SNO laboratory. Therefore, the completion of the SNO experiment presented a compelling opportunity to create a new large-volume liquid scintillator experiment, and as a result the SNO+ project was begun.

6.1 SNO+ Physics

The sensitivity of a scintillator experiment to low energy events combines with some important aspects of the SNOLAB site to give SNO+ the ability to make some extremely interesting physics measurements. The different physics signals of interest to SNO+ are briefly introduced below.

6.1.1 Solar Neutrinos

Although the Solar Neutrino Problem is now resolved, there is still interesting physics to be done by observing solar neutrinos. With an energy threshold below 0.5 MeV, SNO+ will be sensitive to a number of different solar neutrino fluxes. Of particular interest will be the *pep* and CNO neutrinos.

pep Neutrinos

As discussed in Section 2.1, the major advantage of the SNO experiment compared to previous solar neutrino experiments in studying neutrino oscillations was SNO's ability to normalize its charged current flux measurement with its own neutral current flux measurement. This made SNO "model independent," whereas the other solar neutrino experiments had to rely on the Standard Solar Model (SSM) for their normalization. As SNO+ will detect solar neutrinos only through the elastic scattering interaction (Equation 2.3), it too will have to be normalized against the SSM. This, in turn, means that the ability of SNO+ to study solar neutrino oscillations could be limited by the precision with which the neutrino flux is predicted by the SSM. From the point of view of neutrino oscillations, then, the two fluxes which are most precisely predicted by the SSM, namely the *pp* and the *pep* fluxes (see Figure 2.2), are the most interesting. The *pp* neutrinos are not measurable in SNO+ due to the background from ^{14}C intrinsic to the scintillator, but a SNO+ measurement of the *pep* neutrinos should be possible.

Using the *pep* neutrinos to probe the solar neutrino survival probability is particularly interesting because, at 1.44 MeV, the *pep* flux falls on the upturn of the solar neutrino survival probability (Figure 2.8). From the point of view of MSW oscillations this region is interesting because the energy at which the survival probability upturn occurs depends on the neutrino mass splitting. Therefore, a measurement of the *pep* neutrino flux, and hence the electron neutrino survival probability at 1.44 MeV, could help to improve the precision of the solar neutrino measurement of Δm_{12}^2 .

Probing the upturn in the electron neutrino survival probability is also interesting because it is in this region that the detection of new physics is most likely. This is because at low neutrino energies

the solar matter density is insufficient to produce significant changes in the neutrino propagation Hamiltonian, and what are essentially vacuum oscillations occur. At high energies, on the other hand, the resonant MSW transition is fully “turned on” and the survival probability is essentially independent of the neutrino-matter interaction. That is, as long as the neutrinos are interacting with the solar matter strongly enough for the resonant transition to occur, the actual mechanism through which the neutrinos are interacting is relatively unimportant. Therefore, it is in the transition region between the low energy vacuum oscillations and the “full matter effect” at higher energy that the details of the neutrino-matter interaction have the greatest effect on the survival probability.

A number of non-standard models have been proposed, including non-standard neutrino-matter interactions [63], and variations in the neutrino mass depending on the local neutrino density[64], which can greatly affect the electron neutrino survival probability in the transition region while leaving the high and low energy regions within current experimental bounds (see Figure 6.1). Measurement of the *pep* electron neutrino survival probability would help to constrain these models.

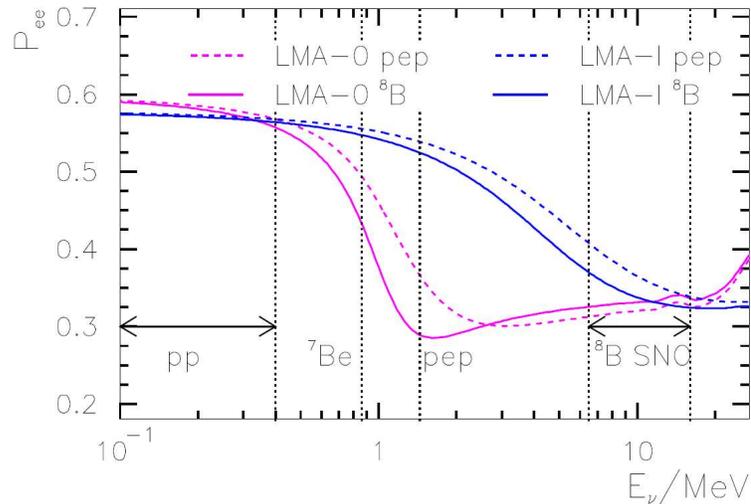


Figure 6.1: The solar electron neutrino survival probability curve with non-standard neutrino-matter interactions. The difference between the non-standard model solar electron neutrino survival probability (“LMA-0”) and the standard MSW curve (“LMA-1”) can be seen to be greatest in the transition region, while the models agree well in the high and low energy regions. The solid and dotted lines represent survival probabilities averaged over different neutrino production regions in the Sun. Figure from [63].

For further discussion of the implications of low energy solar neutrino measurements on our

understanding of the Standard Solar Model and on neutrino oscillations, the reader is referred to [65].

Monte Carlo simulations suggest [66] that SNO+ should be able to make an 8.5% (statistical) measurement of the solar neutrino flux with 1 kT·a of data¹, which would thus provide a measurement of the solar neutrino survival probability at 1.44 MeV of similar precision. This should begin to provide constraints on the non-standard neutrino interaction models.

It should be noted that the depth of the SNOLAB site has an important impact on the SNO+ *pep* measurement. For experiments in shallower locations, cosmogenic backgrounds, in particular ¹¹C, provide a difficult background to the *pep* measurement [68]. At the depth of SNO+, however, this cosmogenic background is almost negligible.

It should also be noted that the BOREXINO experiment (which has recently published a low energy solar neutrino survival probability based on a measurement of the ⁷Be neutrino flux [22] and presented a low energy ⁸B neutrino flux measurement [69]) will also attempt to measure the *pep* flux. Although the ¹¹C is a significant background in BOREXINO, the collaboration has developed a neutron-tagging technique [70], which may allow the *pep* measurement to be made. Ultimately, however, the larger size and lack of ¹¹C background in SNO+ should permit SNO+ to make a more precise measurement².

CNO Neutrinos

Where a measurement of the *pep* neutrinos would provide interesting information about neutrino physics, measurement of the CNO neutrinos would provide interesting information about solar physics. As mentioned in Section 2.1, there are two possible chains of nuclear reactions that account for the energy production of the Sun: the *pp* chain, which produces the *hep*, ⁸B, *pep*, ⁷Be, and *pp* neutrinos, and the CNO cycle. *pp* chain reactions are thought to be responsible for generating

¹ These simulations were carried out assuming that SNO+ would have no ability to discriminate between alpha and beta events based on their timing distributions, while recent studies [67] suggested that the alpha - beta discrimination might be higher than 90%. Also, these simulations were carried out at background levels higher than those which have subsequently been demonstrated by BOREXINO. It is likely, therefore, that a repeated analysis would demonstrate an improved SNO+ *pep* sensitivity.

² Provided that SNO+ achieves radio-purity comparable to BOREXINO.

about 98.5% of the Sun’s energy, with only 1.5% being supplied by the CNO chain [71]. The CNO contribution is still quite uncertain, however, and depends strongly on the density of the “metals” (i.e. N and O, and especially C) in the solar interior. Thus, a measurement of the CNO neutrino flux would constrain not only the contribution of the CNO cycle to the solar energy generation, but also provide a measurement of the solar metallicity. That solar metallicity measurement, in turn, would be helpful in resolving the discrepancy which has recently developed between spectroscopic determinations of the solar metallicity and the metallicities required in solar models to correctly reproduce helioseismological observations [72]³.

Monte Carlo simulations predict that SNO+ should be able to make approximately a 20% measurement of the CNO flux after 1 kT·a of solar neutrino data. This would be good enough to confirm the CNO contribution to solar energy generation and to provide some information to the metallicity controversy. It is again possible that BOREXINO will make a measurement of the CNO flux before SNO+, but again the larger mass and lower cosmogenic background in SNO+ should ultimately allow SNO+ to make a more precise measurement.

6.1.2 Reactor Anti-neutrinos

Nuclear reactors used for power generation emit large fluxes of electron anti-neutrinos as a result of the beta decay of the fission products. Because the number of decays is related to the thermal power produced by the reactor and this thermal power is carefully measured, the anti-neutrino flux leaving each reactor is well known. The energy spectrum of the anti-neutrinos emitted by a reactor is also known quite accurately, although it does depend slightly on the composition of the nuclear fuel used.

It should be noted at this juncture that anti-neutrinos are detected in a liquid scintillator through the inverse beta decay reaction

$$p + \bar{\nu}_e \rightarrow n + e^+. \tag{6.1}$$

The positron is produced with a kinetic energy 1.804 MeV less than the energy of the anti-neutrino,

³ When the metallicities supported by the two measurements are used in the SSM, the predicted CNO neutrino flux changes by 30-40% [72].

E_ν , and is annihilated by a nearby electron almost immediately, so that a total of $(E_\nu - 0.782)$ MeV is deposited in the detector. The neutron diffuses through the detector for $\sim 200 \mu\text{s}$ before being captured by another nucleus, usually another proton, which releases another burst of energy (a 2.2 MeV γ -ray in the case of a proton capture). This “delayed coincidence signal” of two events separated by a fairly well defined time interval provides a powerful method for extracting anti-neutrino signals from single event backgrounds.

As they propagate through the Earth, reactor anti-neutrinos undergo what are essentially vacuum oscillations. The oscillation lengths for reactor anti-neutrinos are on the order of 100km and depend on the energy of the neutrino, as shown in Equation 2.8. This relatively accessible oscillation length, combined with the well understood source flux and spectrum, makes detecting reactor anti-neutrinos with a detector a few hundred kilometres from the reactor an extremely powerful method for studying neutrino oscillations.

Indeed, this was the primary physics motivation for the construction of the KamLAND detector, a kilo-tonne scale liquid scintillator experiment in Japan, which has succeeded in measuring the reactor anti-neutrino oscillation spectrum [23] and extracting from it the most precise measurement to date of Δm_{12}^2 . KamLAND sits at an average distance of 180km from several Japanese and Korean nuclear reactors, with 9 reactor sites within 300km. This gives KamLAND a high reactor anti-neutrino flux, but “averages out” the neutrino oscillation features.

SNO+, on the other hand, will be 240 km from the Bruce nuclear generating station and about 330 km from reactors in Pickering and Darlington. This greater distance will combine with the lower power produced by these reactors to give SNO+ a much lower anti-neutrino interaction rate than KamLAND⁴. However, it happens that the SNO+ reactor-detector baselines are such that the second oscillation minimum in the Bruce reactor energy spectrum coincides with the third oscillation minimum in the Pickering and Darlington oscillation spectra, producing a clear dip in the expected SNO+ reactor anti-neutrino spectrum (see Figure 6.2). This Figure also shows that the position

⁴ In fact, the SNO+ reactor anti-neutrino interaction rate is expected to be approximately 147 events/kT·a after neutrino oscillations [73], which is approximately 30% of the KamLAND rate.

of this combined oscillation feature is quite sensitive to Δm_{12}^2 . The presence of a fairly obvious oscillation feature which has good sensitivity to the mass splitting gives SNO+ better sensitivity to Δm_{12}^2 than might be predicted given the low statistics expected in the measurement. It has even been suggested [74] that the SNO+ measurement of Δm_{12}^2 may rival the KamLAND measurement in precision, even given the reduced statistics of the SNO+ measurement.

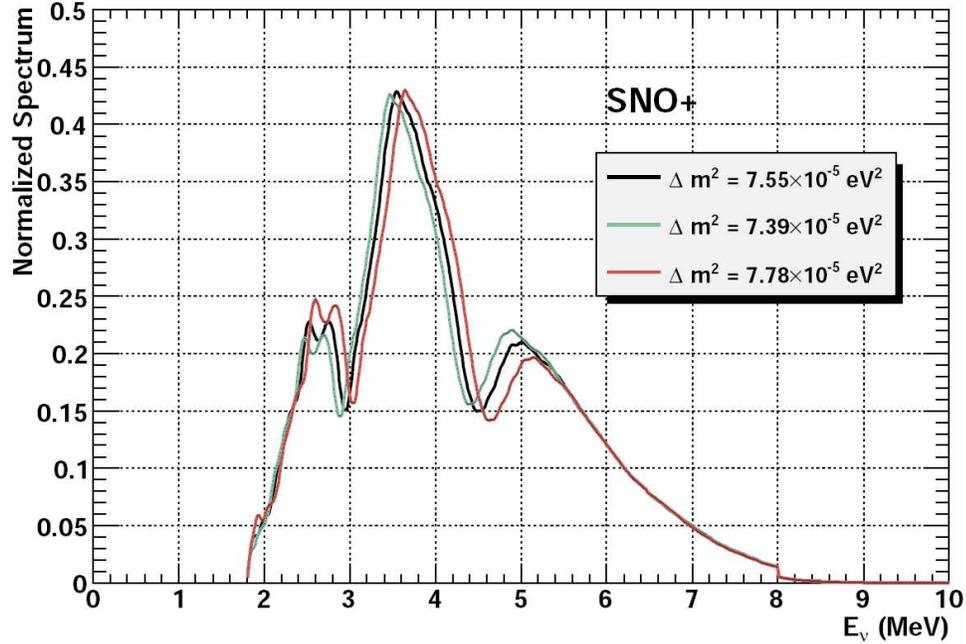


Figure 6.2: The expected spectrum of reactor anti-neutrinos at SNO+ for three different values of Δm_{12}^2 . The spectrum can be seen to contain clear oscillation features. The chosen values of Δm_{12}^2 are approximately the current global (total solar + KamLAND) best fit value and the 1σ variations about this value. It can be seen that the oscillation features in the expected SNO+ spectrum are sufficiently sensitive to the oscillation parameters to produce noticeable differences in the SNO+ spectrum even within the uncertainty bounds of the KamLAND-dominated global result. Figure from [74].

6.1.3 Geo-Neutrinos

The energy released from nuclear decays of long lived isotopes⁵ within the Earth’s mantle and crust is thought to be responsible for replenishing a large portion (50-100%) [75] of the heat emitted by the Earth. As can be seen, however, this contribution is not well known, and constraining it would provide useful information for those attempting to understand the Earth’s thermal history. SNO+

⁵ Especially ^{40}K and isotopes in the ^{238}U and ^{232}Th decay chains.

may be able to contribute to answering this question, as each of these decay chains contains at least one beta decay, and each beta decay produces an electron anti-neutrino. These geologically produced anti-neutrinos, which are generally referred to as “geo-neutrinos” should be detectable in SNO+.

The geo-neutrino energies are quite low (see Figure 6.3), but are well within the reach of scintillator experiments. Also, as discussed earlier, the coincidence signal used to detect anti-neutrinos in a scintillator experiment makes the measurement robust against most backgrounds.

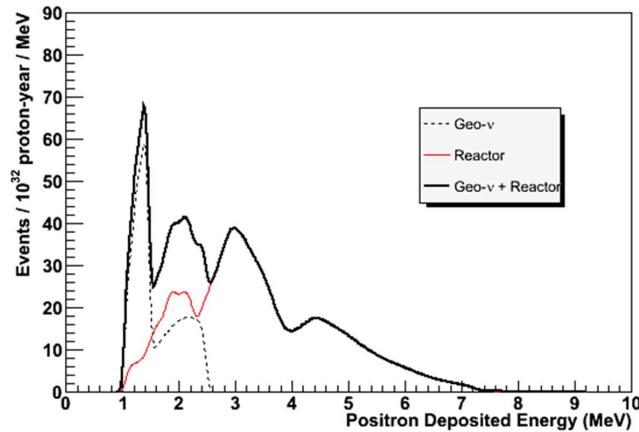


Figure 6.3: The expected spectrum of geo-neutrinos in SNO+. The reactor anti-neutrino background is also shown for comparison.

The KamLAND experiment has already published initial studies of the geo-neutrino flux at Kamioka [76, 23]. These measurements are important first steps in our understanding of these terrestrial anti-neutrinos but, because of the low rate of the geo-neutrino signal, background from the high KamLAND reactor anti-neutrino rate discussed above, and (α, n) backgrounds produced by alpha decays of radon daughters dissolved directly into the KamLAND scintillator⁶, it will be some time before the KamLAND measurement is able to usefully constrain the radiogenic heat production.

⁶ (α, n) reactions can produce a coincidence background through the scattering of the high energy neutron when it is first produced (this produces a signal through the recoil of the scattered proton) followed by the neutron being captured by another proton. KamLAND is currently working to remove these radon daughters from their scintillator via re-distillation; once this is complete the only background to the geo-neutrino measurement should be the reactor anti-neutrinos.

Accurately measuring the geo-neutrino flux should be easier in SNO+ than in KamLAND. For one thing, the rate of geo-neutrino interactions in SNO+ is predicted to be significantly higher than at KamLAND (42 /kT·a for SNO+ vs. 31 /kT·a for KamLAND [77, 73]), due to differences in the regional geology of the two sites. SNO+ has also taken care in the design of the scintillator handling and purification systems to prevent, as much as possible, the ingress of radon, which should prevent (α, n) reactions from significantly affecting the SNO+ geo-neutrino measurement⁷. Also, as discussed earlier, the reactor anti-neutrino flux at SNOLAB is significantly lower than at Kamioka (the location of KamLAND). The combination of these factors should make it possible for SNO+ to make a fairly precise measurement of the geo-neutrino flux at SNOLAB.

Another important point about a SNO+ geo-neutrino measurement is that the interpretation of a geo-neutrino measurement at SNOLAB would be relatively straightforward. This is because the SNOLAB local geography consists of uniform crust which, because of the Sudbury area's unique geological history and the extensive geological studies that have been carried out as a result, is among the best understood in the world. This makes the contribution of the local crust to the SNO+ geo-neutrino flux fairly well constrained, and hence the de-convolution of the crust and mantle contributions to the geo-neutrino signal is less uncertain at SNO+ than it is at other locations. Having the ability to compare geo-neutrino flux measurements from Sudbury, Kamioka⁸ and Gran Sasso⁹ would also be useful in de-convolving the geo-neutrino contributions from the different regions of the Earth's interior.

6.1.4 Supernova Neutrinos

When a star undergoes a Type IIA supernova explosion, more than 99% of its gravitational binding energy is carried away by neutrinos [79] in a neutrino burst which emits more neutrinos in the span of a few seconds than are released in total during the rest of the star's life. The neutrinos are

⁷ The design of the SNO+ scintillator purification and handling systems closely follows the design used in the BOREXINO experiment [78], which successfully produced scintillator with radon daughters at the level of a few thousand decays/kT·a [22].

⁸ Kamioka sits at the boundary between continental crust, which is believed to have relatively high levels of the ²³⁸U, ²³²Th and ⁴⁰K geo-neutrino emitters, and oceanic crust, which is expected to be almost devoid of these chemicals.

⁹ Gran Sasso is the location of the BOREXINO experiment, which should make a relatively low background, but fairly low statistics, measurement of the geo-neutrino flux [77].

predicted to be evenly divided between the flavours and between particles and anti-particles, and to have an average energy of about 15 MeV [80]. The neutrino flux is so large, in fact, that even a supernova at Galactic distances should produce a large number of events in SNO+. Indeed, a total of 24 neutrinos from supernova 1987a, which occurred about 50 kpc away from Earth in the Large Magellanic Cloud, were detected by three neutrino detectors which were active at the time (see [81] and references therein). A supernova effectively provides a “point source” of neutrinos both in space and in time which propagate out through the supernova remnants and then traverse a long but well defined path through the interstellar medium to Earth. Supernova neutrinos can thus carry with them information about the supernova itself, about the interstellar medium through which they have propagated, and about neutrino propagation over very large distances. As a result, they are interesting to many different areas of physics. As an example, the information from the 24 neutrinos from supernova 1987a has been used for everything from testing neutrino oscillation models [81] to setting limits on neutrino mass [82] to constraining the size of possible compact extra dimensions in the universe [83]. In fact, [82] is entitled “Yet Another Paper on Sn1987a....”. If neutrinos from a supernova within our Galaxy were detected, the much greater number of observed events would provide even more information to these diverse areas of physics.

As a large volume scintillator experiment, SNO+ would be a very good supernova detector. The reactions that could be used by SNO+ to detect supernova neutrinos are listed in Table 6.1 below, along with the number of interactions of each type expected from a “standard” type IIA 3×10^{53} erg supernova 10 kpc from the Earth. As can be seen, a SNO+ detection of a Galactic supernova would have good statistics and would provide some ability to distinguish $\nu_e, \bar{\nu}_e, \nu_{\mu,\tau}$ and $\bar{\nu}_{\mu,\tau}$. This flavour separation is expected to provide even more information about supernova physics and neutrino oscillations.

6.1.5 Neutrinoless Double Beta Decay

Many people consider the most important open question in neutrino physics to be the question of whether neutrinos are Majorana or Dirac particles. Dirac particles, which include most familiar

| Interaction | # Events / kT | Reference |
|--|---------------|-----------|
| $\nu_e + e^- \rightarrow \nu_e + e^-$ | 8 | [80] |
| $\bar{\nu}_e + e^- \rightarrow \bar{\nu}_e + e^-$ | 3 | [80] |
| $\nu_{\mu,\tau} + e^- \rightarrow \nu_{\mu,\tau} + e^-$ | 4 | [80] |
| $\bar{\nu}_{\mu,\tau} + e^- \rightarrow \bar{\nu}_{\mu,\tau} + e^-$ | 2 | [84] |
| $\bar{\nu}_e + p \rightarrow n + e^+$ | 263 | [84] |
| $\nu_e + {}^{12}\text{C} \rightarrow {}^{12}\text{N} + e^-$ | 27 | [85] |
| $\bar{\nu}_e + {}^{12}\text{C} \rightarrow {}^{12}\text{B} + e^+$ | 7 | [85] |
| $\nu_x + {}^{12}\text{C} \rightarrow {}^{12}\text{C}^*(15.11 \text{ MeV}) + \nu_x$ | 58 | [85] |
| $\nu_x + p \rightarrow \nu_x + p$ | 273** | [86] |

Table 6.1: The expected number of neutrinos detected through different interaction channels in SNO+ for a 10 kpc type IIA supernova, assuming MSW oscillations. Note that there is one high statistics charged current channel and one high statistics neutral current channel. (**assuming a 0.2 MeV threshold)

particles, are distinct from their corresponding anti-particles. Majorana particles, on the other hand, are identical to their corresponding anti-particles except for their helicities. Many theorists believe that neutrinos should be Majorana particles, as the theory of Majorana masses admits a mechanism that naturally explains the smallness of the neutrino mass [87].

Attempts to detect the (possible) Majorana nature of neutrinos focus around the double beta decay process (“ $2\nu\beta\beta$ ”),

$$(A, Z) \rightarrow (A, Z + 2) + 2e^- + 2\bar{\nu}_e. \quad (6.2)$$

If neutrinos are Majorana particles, the anti-neutrino emitted by one of the neutrons can be absorbed as a neutrino by the other. The resulting process, in which no neutrinos are emitted, is neutrinoless double beta decay (“ $0\nu\beta\beta$ ”),

$$(A, Z) \rightarrow (A, Z + 2) + 2e^-. \quad (6.3)$$

One experiment [88] claims to have detected neutrinoless double beta decay in ${}^{76}\text{Ge}$ with a neutrino mass in the range 0.27-0.37 eV (90%C.L.). The initial claims by this experiment [89], however, were extremely controversial [90], and experiments to confirm or refute the claim are considered necessary.

SNO+ will be able to make an important contribution to this effort by loading ${}^{150}\text{Nd}$, a double

beta decay isotope, into the liquid scintillator. The large mass of isotope thus obtained will be sufficient to give the SNO+ double beta decay search competitive, and potentially leading, sensitivity to neutrinoless double beta decay compared to other proposed and existing experiments. The SNO+ double beta decay search is described in detail in Chapter 7.

6.1.6 Physics Summary

The SNO+ experiment has the potential to carry out interesting studies of a number of different physics signals. These include the low energy solar neutrinos, geo-neutrinos, reactor anti-neutrinos and supernova neutrinos, as well as neutrinoless double beta decay.

Of these topics, all but the potential for the neutrinoless double beta decay search had been fully appreciated for some time. Indeed, even the 2001 NSERC Subatomic Physics Five Year Plan includes a scintillator experiment built from the SNO detector (called 'SNO+' purely by coincidence) to “study low energy neutrinos” [91]. As described in the next chapter, however, the potential sensitivity of the experiment to neutrinoless double beta decay only began to be appreciated in the late 1990's. The addition of the neutrinoless double beta decay search to this already very interesting physics program provided the final impetus to stimulate the formation of the SNO+ collaboration and the development of the SNO+ experiment.

6.2 Experimental Design and Development

Having established that the SNO+ experiment was scientifically interesting, it remained to be demonstrated that the experiment was practically realizable. The two main prerequisites to the successful staging of the experiment were the identification of a liquid scintillator suitable for long term use in the acrylic vessel and the design of a hold-down system to hold the now buoyant acrylic vessel (A.V.) in position.

6.2.1 Scintillator Development

The liquid scintillators used in large scintillator experiments typically have two or three components: a solvent, which is an aromatic hydrocarbon which is efficiently excited by the passage of charged

particles, and which forms the bulk of the scintillator; a fluor, typically 2,5-diphenyloxazole (PPO) at the level of a few grams per litre, which “collects” (typically via dipole interactions) the excitations from the solvent and emits them as photons with high efficiency; and sometimes a wavelength shifter which, at the level of a few mg/L, captures the light emitted by the fluor and re-emits it at longer wavelengths where there is less chance of it being re-absorbed by the other scintillator components. In addition, scintillators are sometime diluted in inert solvents like mineral oil to reduce their cost or improve their chemical compatibility.

Unfortunately, most aromatic liquids, like the pseudocumene liquid scintillator that has been used in previous large liquid scintillator experiments, are quite aggressive towards acrylic. Pseudocumene diluted to 20% in mineral oil had been used in contact with acrylic for extended periods in the Palo Verde experiment [92], but even there anecdotal accounts suggest visible changes occurred in both the acrylic and the scintillator immediately adjacent over time. In the higher-stress application of the SNO+ acrylic vessel, such degradation may not be acceptable.

Therefore, a campaign was begun to test different scintillation solvents for suitability for use in SNO+. An ideal scintillator would be acrylic compatible, have good light output, provide little attenuation to its own emitted light and have as high a density as possible (to reduce the A.V. buoyancy problem). Initial testing involved exposing small samples of acrylic to the different solvents and monitoring for changes in the acrylic and in the optical behaviour of the scintillator. The initial tests also compared the amount of light produced by the different scintillators (with 2 g/L PPO fluor) under irradiation from a γ -ray source. In these initial tests two scintillators stood out: diisopropylnaphthalene (DIN) and linear alkyl benzene (LAB). Of these, DIN had a slightly higher light output and a higher density (0.99 g/cm³ vs 0.86 g/cm³ for LAB), but a slight increase in the amount of 90° scattering observed in the acrylic-exposed DIN raised fears that chemical attack might be causing small particles of acrylic to become suspended in the scintillator. In addition, measurements of the optical attenuation lengths of the candidate scintillators showed that LAB exhibited significantly less attenuation than DIN in the important wavelength region, and was thus

more suited for use in large detectors like SNO+. As a result, LAB was selected for further, more careful, testing. A complete description of the initial scintillator tests can be found in [93]

Acrylic Compatibility

In order to determine more conclusively whether LAB was compatible with acrylic, two more series of compatibility studies were begun. In one test, carried out at Brookhaven National Laboratory, bars of acrylic were stressed by bending them across a fulcrum to a surface stress of 3000 PSI, in a test technique described in [94]. Absorbent pads were used to hold solvent in contact with the stressed surface. A high level of tensile stress (for comparison, the maximum design stress in the SNO A.V. (in tension) was 600 PSI) increases the susceptibility of the acrylic to chemical attack to such a degree that this test is considered to be passed if no change is observed in the acrylic after a solvent exposure of 4 hrs. To confirm the sensitivity of the test, acrylic was first exposed to pseudocumene, which, as mentioned earlier, is known to be incompatible with acrylic. The acrylic bar exposed to pseudocumene failed (broke) after one hour. Next, pseudocumene diluted to 20% in mineral oil (which, as mentioned above, was known to be “acrylic compatible” from the Palo Verde experiment) was tried. The diluted pseudocumene passed the test, but the exposed acrylic was observed to have softened after 3 days of exposure. Finally, a test with LAB showed no change in the acrylic after more than 30 days exposure. In this series of tests, then, LAB was found to have less interaction with acrylic than diluted pseudocumene, a known “acrylic compatible” solvent.

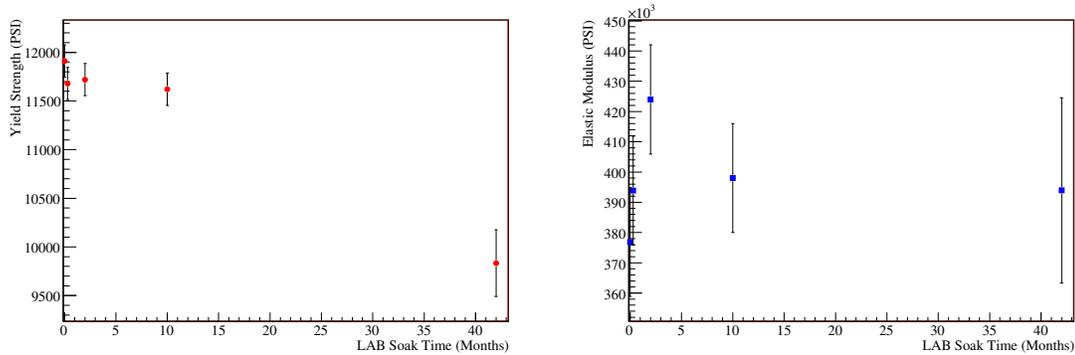
The second series of compatibility tests were based on the ASTM D 543 chemical compatibility test [95]. In these tests, acrylic “dogbones” (which were designed to be destructively tested in order to determine the acrylic elastic modulus and failure stress as described in [96]), were exposed to LAB for long periods of time. There were three series of dogbones:

- **Series I:** A series of of dogbones cut from a commercial (cast, UV absorbing) acrylic sheet. These dogbones were soaked, unstressed, in pure LAB. Five sets of four dogbones were made.
- **Series II:** A series of dogbones cut from a sheet of SNO acrylic, with all of the machining

done at Queen’s University. These dogbones were mounted on custom built strain jigs which stressed the outside surface of the dogbones to 1200 PSI (twice the SNO design maximum) and soaked in pure LAB. Seven sets of four dogbones were made, of which five were mounted on strain jigs and soaked (the others were broken as a control).

- **Series III:** A larger series of 600 dogbones constructed from SNO acrylic. Half of these dogbones were produced¹⁰ with a SNO-type acrylic-to-acrylic bond in the centre of the “neck” of the dogbone. These dogbones were soaked, on strain jigs, in three groups; one group in water (as a control), one group in SNO+ scintillator (LAB + 2 g/L PPO), and one group in neodymium-loaded SNO+ scintillator¹¹.

The first series of dogbones, which have now all been broken, soaked in LAB for up to 42 months. As seen in Figure 6.4, the acrylic properties initially appeared to be stable, but in the final test a decrease was observed in the failure stress; the elastic modulus appears to have remained unchanged.



(a) Yield Stress

(b) Elastic Modulus

Figure 6.4: The results of the Series I LAB - acrylic compatibility tests. The elastic modulus of the acrylic appears to be unchanged with exposure, while the failure stress in the final group of dogbones tested exhibited a decrease. Four dogbones were broken in each test; the mean results of each test are shown with error bars representing the RMS spread between the individual dogbones measured within that test.

This pattern is repeated in the current results of the Series II tests, as seen in Figure 6.5. Two

¹⁰ By Reynold’s Polymer Technology, the same company which bonded the SNO A.V.

¹¹ This group is currently soaking in scintillator only; the Nd-carboxylate will be added at a later date.

sets of Series II dogbones remain to be tested.

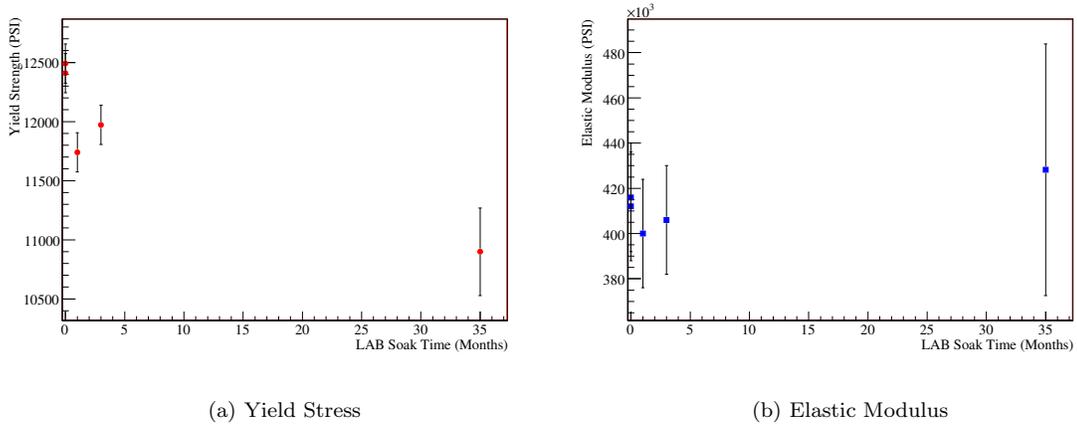


Figure 6.5: The current results of the Series II LAB - acrylic compatibility tests. As in the Series I tests, the elastic modulus of the acrylic appears to be unchanged with exposure, while the failure stress appears to decrease.

Finally, the Series III tests have just been started, with the intention of maintaining them as a long-term exposure tests for the duration of the SNO+ experiment. A single set of dogbones have so far been tested, after 10 month's exposure, with no change seen in the LAB-exposed samples (both bonded and non-bonded). There was, however, an unexplained decrease of about 15% in the failure stress of the non-bonded water exposed acrylic.

The observed change in the acrylic failure stress after prolonged exposure to LAB is unsettling, but, at least at the observed levels, it is likely tolerable. The most important quantity identified in the finite element calculations that have been used to study the acrylic vessel stability is the elastic modulus, and that still appears to be stable after LAB exposure. In addition, the majority of the acrylic vessel is, by design, under compression rather than under tension, so the failure stress under tension is less important [97]. At some point, of course, continued weakening of the acrylic through LAB exposure would become troublesome. Efforts are being made to understand the cause of the observed change in the acrylic failure stress¹², and the remaining dogbones will continue to be tested periodically to monitor for continuing changes in the acrylic. The observed, and likely extrapolated,

¹² It could, for example, be caused by the formation of surface cracks as the result of chemical attack by the LAB, or it could simply be the result of stresses induced in the acrylic surfaces as they swell due to LAB absorption. Clearly these different underlying mechanisms have different implications for the SNO+ acrylic vessel.

changes in the failure stress will also be included in the finite element analysis of the acrylic vessel to confirm that they will not compromise the integrity of the vessel.

Scintillator Purification

Concurrently with the development of the acrylic compatibility tests, investigations were carried out into purification methods for the LAB. These purification trials, which were principally based on techniques that proved highly successful in the purification of the scintillator in the BOREXINO experiment [78], showed that the standard vacuum distillation and Al_2O_3 column purification techniques were effective both in improving the optical transmission of the LAB and in reducing the level of radioactive contamination in the scintillator [66]. Based on this information, a full scale vacuum distillation based scintillator purification plant has been designed¹³, and construction activities for the installation of the plant in SNOLAB have begun.

Optical Characterization

The amount of light detected by a large liquid scintillator experiment like SNO+, after the passage of a charged particle, is more difficult to predict than it would be in a smaller detector. Very pure scintillators can have absorption and scattering lengths of many meters at the important wavelengths, so in smaller detectors the amount of light detected basically depends only on the amount of light initially produced and on the detector geometry. In a larger experiment like SNO+, on the other hand, the path lengths for photons through the scintillator are of order of the attenuation lengths, so a typical detected photon has been absorbed and re-emitted by the scintillator more than once (1.4 times on average for photons produced at the centre of the detector in the SNOMAN simulations (described below)). The wavelength and direction of the photon change with each re-emission. Thus, precise measurements of long attenuation and scattering lengths in the scintillator are necessary to achieve a reliable prediction of the scintillator light output, and Monte Carlo simulations must be used.

¹³ The SNO+ scintillator purification plant was designed by the same company which designed and built the BOREXINO scintillator purification plant.

In order to provide such simulations for SNO+, the SNOMAN particle energy deposition and photon tracking code, which was developed for the SNO experiment and which, as a result, already contained excellent models of the SNO geometry and the optical response of the PMT array, was modified to include scintillator optical effects. As scintillator is much more optically active than the water for which the code was originally designed, this required fairly significant modifications to the photon production and propagation routines. The scintillator optical properties which are important to the light output of a large scintillation experiment, and which were therefore accommodated in the simulation include:

1. **Primary Light Output:** This is the average number of photons initially produced when a given amount of energy is deposited in the scintillator (this number includes both the fraction of the energy deposited in the scintillator that creates electronic excitations capable of producing photons and the transfer and emission efficiencies described below). The light output depends on the composition of the scintillator cocktail and the type of particle responsible for the excitation.
2. **Absorption Length:** The (wavelength dependant) average distance travelled by a photon before it is absorbed by one of the scintillator components. The absorption of each scintillator component must be treated separately because the re-emission probability is different for each component.
3. **Scattering Length:** The (wavelength dependant) average distance travelled by a photon before it scatters (typically by Rayleigh scattering in clean scintillator), changing its direction but not its wavelength. The scattering length is common to the scintillator cocktail, and can vary with the component concentrations.
4. **Quantum Efficiency:** The efficiency with which a scintillator component converts electronic excitations into emitted photons. The fluor especially is chosen to have a high quantum efficiency. The quantum efficiency also gives the probability that an absorbed photon is re-

emitted.

5. **Non-Radiative Transfer Efficiency:** The component-specific efficiency with which a molecular electronic excitation in a molecule of one scintillator component is transferred (typically by dipole interactions) to another scintillator component. This is especially important in the transfer of primary particle excitations from the solvent to the fluor, but it also plays a role in the re-emission of captured photons. The non-radiative transfer efficiency varies with the concentrations of the different scintillator components.
6. **Emission Spectrum:** The distribution which describes the relative probabilities of a photon emitted by a given scintillator component having different wavelengths.
7. **Emission Time:** The average time between the absorption of a photon by a scintillator component and its subsequent re-emission. This is not important in light output predictions, but it is necessary to simulate event reconstruction within the detector.

In a fairly extensive campaign, the properties listed have been measured or estimated for the LAB + 2 g/L PPO SNO+ scintillator cocktail (and the same cocktail with 0.1% Nd-carboxylate) [98, 93]; work is ongoing to improve some of these measurements, especially the long absorption and scattering lengths in LAB and Nd-LAB, and the LAB quantum efficiency. With the current best measurements incorporated, the SNOMAN-based simulation suggests that the light output for the SNO+ solar neutrino phase is adequate (~ 520 PMT hits for a 1 MeV electron at the centre of the detector) and fairly robust against the uncertain measurements described above.

6.2.2 Acrylic Vessel Hold-Down

Apart from its apparent acrylic compatibility, the other major contributing factor in selecting LAB from the scintillator candidates studied in the original scintillator selection trials was the emerging understanding that the $\sim 15\%$ density mismatch between LAB and the surrounding light water could likely be tolerated by the SNO acrylic vessel. This began with an engineering opinion [99] that because the A.V. had been designed to support, with a wide safety margin, an internal over-

density of $\sim 10\%$ and because the junction at the neck of the vessel was re-enforced such that the vessel should behave structurally as a complete sphere, there was no fundamental reason that the vessel should not be expected to tolerate a $\sim 15\%$ internal under-density with a carefully designed hold-down system.

Continued work resulted in the creation of several candidate hold-down net designs, an example of which is shown in Figure 6.6. Finite-element calculations have been carried out based on these designs [100] and indicate that the A.V., thus anchored, would withstand even 100% internal under-density compared to the surrounding light water, provided that it were perfectly spherical. Surveying of the real A.V. in order to determine its actual deviation from sphericity, and studies of second order effects like potential increases in pressure on the A.V. beneath rope junctions are ongoing. Nevertheless, it appears that the $\sim 15\%$ internal under-density associated with the use of LAB should be safely tolerated by the SNO+ acrylic vessel, and this opinion was supported by an independent engineering review of the A.V. hold-down design [101].

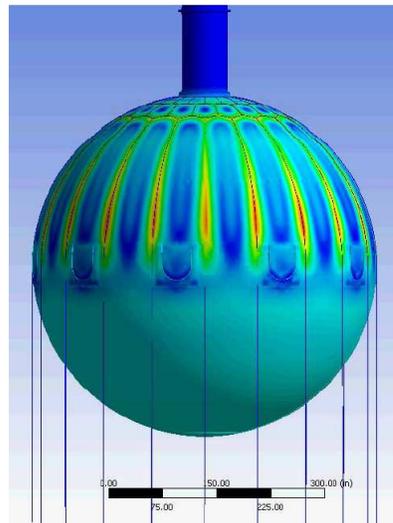


Figure 6.6: The results of a finite-element analysis of one of the candidate SNO+ acrylic vessel rope net hold-down configurations. The colour indicates the level of stress in the acrylic. Figure from [100].

6.2.3 SNO+ Current Status and Outlook

The SNO+ project is moving forward. The collaboration currently has about 65 members at 14 institutions in Canada, the United States, Great Britain, Portugal, and Germany. The project has received annual development funding from the Canadian Natural Sciences and Engineering Research Council since 2005, and support from the National Science Foundation for SNO+ research activities in the U.S. since 2007. The Canadian FedNor northern development agency provided additional seed funding in 2008. In June of 2009, capital funding for the SNO+ project was received from the Canada Foundation for Innovation to cover the final design, purchase and installation of the acrylic vessel hold-down net, the purchase and installation of the scintillator purification system, and the purchase of the liquid scintillator. Preparation for the installation of the scintillator purification plant has already begun, and work to prepare the SNO+ cavity for the installation of the acrylic vessel hold-down is also underway.

On the current schedule, the SNO+ detector is predicted to be filled and operational by 2011. After a few months of “shakedown” operation, (natural) neodymium carboxylate will be added to the scintillator and the neutrinoless double beta decay experiment will be carried out. Later, perhaps in 2015, the neodymium will be removed and the SNO+ solar neutrino phase will begin. If at any point during this time enriched neodymium were to become available it would be installed at once, as would nanoparticle neodymium loading if it should be perfected. The geo-neutrino and reactor anti-neutrino signals should be detectable during both phases, and if a Galactic supernova were to occur it would be detectable in either phase. On the current schedule, then, the exciting scientific journey of the SNO+ experiment should begin in a few years’ time.

Chapter 7. The SNO+ Neutrinoless Double Beta Decay Search

7.1 Introduction to Neutrinoless Double Beta Decay

“Two neutrino” double beta decay ($2\nu\beta\beta$) is a rare second order weak process,

$$(Z, A) \rightarrow (Z + 2, A) + e_1^- + e_2^- + \bar{\nu}_{e1} + \bar{\nu}_{e2} \quad (7.1)$$

which is observed when single beta decay from the initial (Z, A) nucleus to the intermediate $(Z+1, A)$ nucleus is energetically disallowed, but the double beta decay is not. This situation occurs naturally in some even-even nuclei, and two-neutrino double beta decay has in fact been observed for some isotopes, including ^{48}Ca , ^{76}Ge , ^{96}Zr , and ^{150}Nd [102].

Another potential decay mode for these double beta decay isotopes is neutrinoless double beta decay ($0\nu\beta\beta$),

$$(Z, A) \rightarrow (Z + 2, A) + e_1^- + e_2^-. \quad (7.2)$$

In this process, rather than emitting two anti-neutrinos, the decaying nucleons instead exchange a virtual neutrino¹. In order for this process, which explicitly violates lepton number conservation, to occur, the neutrino must be massive and invariant under charge conjugation. Such “Majorana” neutrinos differ from the corresponding anti-neutrinos only in their helicity. As discussed in the earlier sections, neutrino oscillation experiments have demonstrated that neutrinos are massive; the search for neutrinoless double beta decay thus seeks to determine whether or not neutrinos are invariant under charge conjugation (i.e. whether they are “Majorana Particles” or “Dirac Particles”).

¹ It is also possible that neutrinoless double beta decay could involve the emission of a neutral boson, $(Z, A) \rightarrow (Z + 2, A) + e_1^- + e_2^- + \chi$, or arise via right-handed weak currents or the exchange of exotic heavy particles [103, 104]. Such exotic processes are, however, beyond the scope of this thesis and will not be discussed further.

7.1.1 The Neutrinoless Double Beta Decay Rate

A detailed description of the theoretical mechanism behind neutrinoless double beta decay is beyond the scope of this work. The overview given below highlights the experimentally important points, and is largely drawn from the excellent reviews given in [103], [105] and [104]. For more details, the reader is referred to those works and the references contained therein.

The half-life for neutrinoless double beta decay, $T_{1/2}^{0\nu}$, is given to good approximation (for $0^+ \rightarrow 0^+$ transitions, which form the experimentally interesting cases) by

$$T_{1/2}^{0\nu} = \left[G^{0\nu}(E_0, Z) |M^{0\nu}|^2 \langle m_\nu \rangle^2 \right]^{-1}, \quad (7.3)$$

where the matrix element $M^{0\nu}$ and the effective neutrino mass $\langle m_\nu \rangle$ will be discussed further below.

The phase space factor is given by

$$G^{0\nu} \propto \int F(Z, \epsilon_1) F(Z, \epsilon_2) p_1 p_2 \epsilon_1 \epsilon_2 \delta(E_0 - \epsilon_1 - \epsilon_2) d\epsilon_1 d\epsilon_2, \quad (7.4)$$

where $F(Z, \epsilon)$ is the Fermi Function, E_0 is the energy released during the double beta decay, and the p_i 's and ϵ_i 's are the momenta and total energies, respectively, of the two emitted electrons. $G^{0\nu}$ is in principle exactly calculable.

In the case of two neutrino double beta decay, the matrix elements are given by the standard overlap integrals between the initial, intermediate, and final nuclear states ($|i\rangle$, $|m\rangle$, and $|f\rangle$, respectively):

$$M^{2\nu} = \sum_m \frac{\langle f | \sigma \tau_+ | m \rangle \langle m | \sigma \tau_+ | i \rangle}{E_m - (M_i + M_f)/2}, \quad (7.5)$$

where the sum is over the (1^+) states of the intermediate nucleus with energy E_m , M_i and M_f are the nuclear masses of the initial and final states, τ_+ is the isospin raising operator (describing the conversion of neutrons to protons) and σ is a Pauli matrix (which describes the spin changes inherent to this Gamow-Teller type transition). Therefore, the matrix element of the two-neutrino double beta decay can essentially be calculated as the product of the standard matrix elements for two independent beta decays. In neutrinoless double beta decay, on the other hand, this simplification is not possible because the requisite neutrino exchange couples the two decays; the emission of a

Majorana neutrino at position x , which is subsequently re-absorbed at position y , adds the following to the matrix element:

$$-i \int \frac{d^4q}{(2\pi)^4} e^{-iq(x-y)} \bar{e}(x) \gamma_p P_L \frac{q^\mu \gamma_\mu + m}{q^2 - m^2} P_L \gamma_\sigma e^c(y), \quad (7.6)$$

where q is the momentum 4-vector and m is the mass of the neutrino, $\bar{e}(x)$ and $e^c(y)$ are electron creation operators, the γ 's are the standard Dirac matrices, and $P_L = (1 - \gamma^5)/2$ is the chirality projector which ‘‘picks out’’ the proper chirality component (which for Majorana neutrinos is similar to ‘‘picking out’’ the neutrino or anti-neutrino component, as discussed above) at the two vertices. Note the presence of the neutrino mass term in the numerator. Integrating this expression over the neutrino energy and momentum yields the ‘‘neutrino potential’’ terms

$$H_\pm(r) = \frac{2m}{\pi r} \int_0^\infty dq \frac{q \sin(qr)}{w(w + E_m - (M_i + M_f)/2 \pm (\epsilon_1 - \epsilon_2)/2)}, \quad (7.7)$$

where $r = x - y$ and $w = \sqrt{\vec{q}^2 + m^2}$. These potentials are then averaged to give $H(r) = (H_+(r) + H_-(r))/2$, which appears in the matrix element:

$$|M^{0\nu}| = \langle f | \sum_{lk} H(r_{lk}) \tau_{+l} \tau_{+k} \left(\vec{\sigma}_l \cdot \vec{\sigma}_k - \frac{g_V^2}{g_A^2} \right) | i \rangle, \quad (7.8)$$

where the sum is over nucleons which could participate in the decay, g_V and g_A are the vector and axial coupling constants, respectively, and the multiplicative neutrino mass factor m has been factored out². The $0\nu\beta\beta$ half-life is thus

$$T_{1/2}^{0\nu} = \left[G^{0\nu}(E_0, Z) |mM^{0\nu}|^2 \right]^{-1}. \quad (7.9)$$

The neutrino mass m which multiplies Equation 7.9 now needs to be considered in more detail, in light of the misalignment between the neutrino mass and flavour eigenstates discussed in Section 2.4. The neutrino coupling at each of the electron creation vertices in the neutrinoless double beta decay occurs through the weak interaction, which in this case couples to the electron neutrino flavour eigenstate, and from Equation 2.4, we see that $|\nu_e\rangle = U_{e1} |\nu_1\rangle + U_{e2} |\nu_2\rangle + U_{e3} |\nu_3\rangle$. Thus, the neutrinoless double beta decay matrix element actually contains three ‘‘copies’’ of Equation 7.8,

² This formula is actually an approximation considering only s-type electron wave-functions and non-relativistic nucleons, but is considered quite accurate for $0^+ \rightarrow 0^+$ transitions.

one for each mass eigenstate, with each of them weighted by the mass of that eigenstate times the square³ of the appropriate entry in the PMNS matrix. The actual mass term in Equation 7.9 is thus a sum,

$$m = \sum_j m_j U_{ej}^2, \quad (7.10)$$

where m_j is the mass of the j^{th} neutrino mass eigenstate. The neutrino mass and nuclear matrix element terms in Equation 7.9 are usually then separated to yield Equation 7.3 with the “effective Majorana neutrino mass” defined as

$$\langle m_\nu \rangle = \left| \sum_j m_j U_{ej}^2 \right|. \quad (7.11)$$

A final important point is that for Majorana neutrinos, the PMNS matrix is more general than the Dirac case introduced in Section 2.4. In particular, the Majorana PMNS matrix contains two additional complex phases (provided that there are three light neutrino mass eigenstates) such that

$$\hat{U}_{\text{Majorana}} = \hat{U}_{\text{Dirac}} \times \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\alpha} & 0 \\ 0 & 0 & e^{i\beta} \end{pmatrix}. \quad (7.12)$$

These complex terms thus appear (squared) in the sum in Equation 7.11 and lead to the possibility of destructive interference between the mass eigenstates in the effective neutrino mass. The neutrino mass splittings that have been measured by the oscillation experiments can, however, still be used to set upper and lower limits on $\langle m_\nu \rangle$ as a function of the minimum neutrino mass. These are shown in Figure 7.1.

7.1.2 Implications of Detecting Neutrinoless Double Beta Decay

The detection of neutrinoless double beta decay would have important implications for particle physics and cosmology. First of all, the detection of neutrinoless double beta decay would demonstrate that neutrinos are Majorana particles. This would confirm the existence of the two additional complex phases in the PMNS matrix, making it more natural that large CP violation in the lepton

³ The squared term is necessary because there are two neutrino coupling vertices.

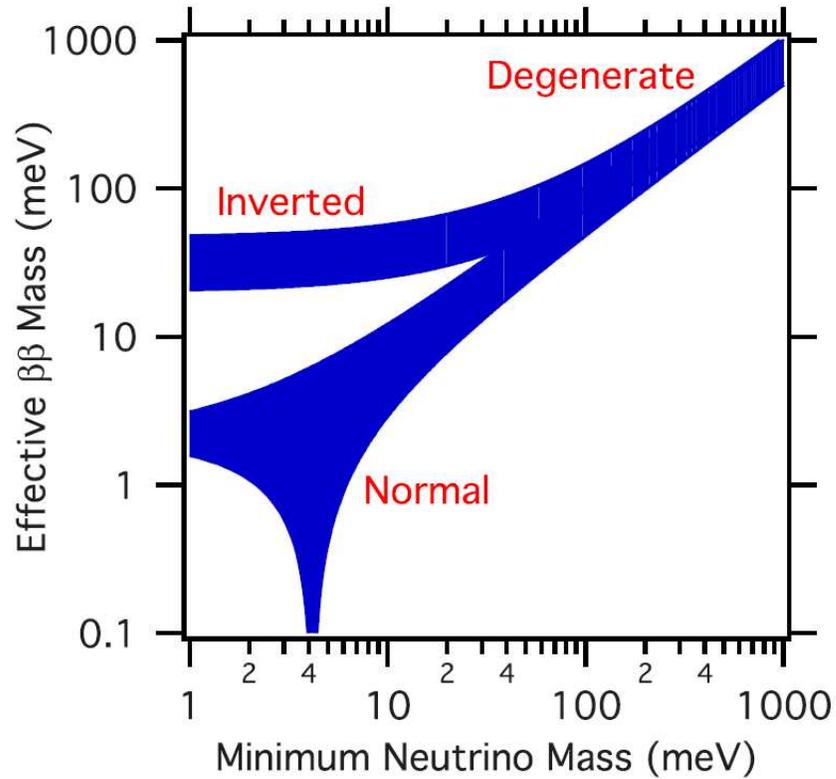


Figure 7.1: The limits on the effective Majorana neutrino mass as a function of the minimum neutrino mass. As the splittings between the mass eigenstates are known but their ordering is not, there are two potential scenarios, the “normal” hierarchy, in which the largest mass splitting occurs between the two heaviest mass eigenstates and the “inverted” hierarchy, in which the largest mass splitting occurs between the two lightest mass eigenstates. These scenarios are essentially indistinguishable if the minimum neutrino mass is large compared to both mass splittings, which is referred to as the “degenerate” scenario. Figure from [105].

sector could be responsible for the observed matter-antimatter asymmetry in the Universe. Majorana neutrinos would also admit the “see-saw mechanism” [87], which naturally explains a small but finite neutrino mass (which is otherwise difficult to explain theoretically).

The second important implication of the detection of neutrinoless double beta decay would be a measurement of the neutrino mass. Although the effective Majorana neutrino mass that would be measured does not necessarily correspond exactly to the physical neutrino mass (because of the possible cancellations due to the complex phases discussed above), given a measurement of $\langle m_\nu \rangle$, the limiting relationships shown in Figure 7.1 can be used to determine an allowed range for the physical neutrino mass.

7.1.3 Calculation of Double Beta Decay Matrix Elements

In order to convert between observed neutrinoless double beta decay rates (or rate limits) and bounds (or limits) on the effective Majorana neutrino mass, it is necessary to evaluate the phase space and nuclear matrix element terms in Equation 7.3. As mentioned earlier, the phase space term can be calculated exactly. The matrix element term, on the other hand, is much more challenging to calculate, especially for spin-deformed nuclei like ^{150}Nd . Two main techniques are used in numerically evaluating $M^{0\nu}$: Nuclear Shell Model (NSM) calculations, and calculations based on the Quasiparticle Random Phase Approximation (QRPA). Again, only very brief outlines of these approaches are given here, with the reader referred to [103, 105, 104] and the references therein for more information.

In NSM calculations, the nucleons are individually modelled and their interactions described by an effective Hamiltonian which is tuned to properly reproduce the observed energy levels splittings and transition probabilities of that nucleus. NSM calculations can in principle account for nuclear deformation. However, computers are not yet powerful enough to treat more than a few nucleons at a time in NSM calculations; the effects of the remaining nucleons must then be treated in some approximate way.

In QRPA, on the other hand, more significant approximations are made which means that the

calculation requires more empirical tuning but which allows more nucleons to be included in the calculation. At its most basic level, QRPA involves replacing interacting pairs of nucleons with boson “quasi-particles” and modelling the excited intermediate nuclear states as phonon-like states in a vacuum. These approximations mean that a larger number of nucleons can be treated in a given calculation (although there is still usually assumed to be an “inert core” of inner nucleons), but they also mean that the particle-particle correlations are no longer treated properly. This has an effect on the calculated nuclear transition rates, and so an empirical correction factor g_{pp} , which must be tuned against data from other types of nuclear transitions, is introduced to compensate. It turns out that the predicted rate of $2\nu\beta\beta$ decays depends strongly on g_{pp} , and much of the work on QRPA involves attempting to lessen its sensitivity to this parameter and on developing new ways to fix g_{pp} . One popular recent approach is renormalized QRPA or RQRPA. Nuclear deformations can in principle be included in QRPA and RQRPA calculations “by hand,” but this has not yet been attempted for $0\nu\beta\beta$ matrix elements.

The uncertainties in both NSM and QRPA calculations of $M^{0\nu}$ are both relatively large and poorly defined. For lack of a better method of estimating the uncertainties, the spread in the $M^{0\nu}$ values calculated by different methods and different evaluations using the same method is often used. In the smaller, undeformed (and presumably simpler to calculate) nuclei, like ^{76}Ge , the calculated values of $M^{0\nu}$ differ by factors of more than three [103, 104]; in the larger spin-deformed nuclei, the uncertainties in the matrix elements are likely larger, but have not yet been even approximately defined.

^{150}Nd , the isotope of interest to SNO+, falls into this category of large, spin-deformed nuclei. Several attempts at evaluating its matrix element using QRPA/RQRPA have been made [106, 107, 108, 109], all of which have neglected the deformity of the nucleus. In addition, a single NSM calculation has been attempted [110], with the basis of nucleons severely truncated in order to make the calculation tractable. The calculation did, however, attempt to include nuclear deformation effects. The results of these different matrix element calculations (in terms of the $0\nu\beta\beta$ half-life that

they predict for $\langle m_\nu \rangle = 50$ meV) are shown in Table 7.1.

| Calculation Method | $T_{1/2}^{0\nu}$ for $\langle m_\nu \rangle = 50$ meV (a) | Reference |
|--------------------|---|-----------|
| QRPA | 1.0×10^{25} | [106] |
| RQRPA | 2.0×10^{25} | [107] |
| RQRPA | 2.0×10^{25} | [108] |
| RQRPA | $(2.0 - 4.4) \times 10^{25}$ | [109] |
| QRPA | $(1.6 - 3.5) \times 10^{25}$ | [109] |
| NSM | 4.2×10^{26} | [110] |

Table 7.1: The $T_{1/2}^{0\nu}$ values predicted for ^{150}Nd at $\langle m_\nu \rangle = 50$ meV for several different calculations of $M^{0\nu}$. In some cases, a range of allowed values (set by the level of variance in the numerical calculation, uncertainty in the experimental measurements of $T_{1/2}^{2\nu}$ that were used to tune the calculations, and by uncertainty in the axial coupling constant used) is shown. In all of these calculations the nucleus is assumed to be spherical.

7.2 Neutrinoless Double Beta Decay in SNO+

The search for neutrinoless double beta decay involves attempting to detect a small, mono-energetic neutrinoless double beta decay peak at the endpoint of a much larger two-neutrino double beta decay continuum. Double beta decay experiments thus require good energy resolution (so that the $2\nu\beta\beta$ continuum is not “smeared out” to such a degree that the $0\nu\beta\beta$ signal region is not statistically accessible at an interesting level), low background in the $0\nu\beta\beta$ energy region, and large mass (so that a reasonable number of the extremely rare $0\nu\beta\beta$ events might be expected).

In SNO+, the double beta decay search will be carried out by suspending neodymium in the LAB scintillator. Natural neodymium contains 5.9% ^{150}Nd , which is a double beta decay isotope decaying to ^{150}Sm with an endpoint energy of 3.37 MeV. The ^{150}Nd $2\nu\beta\beta$ half-life has been measured to be $(9.2 \pm 0.8) \times 10^{18}$ a⁴, and the current experimental 90% C.L. upper limit on the $0\nu\beta\beta$ half-life is 1.8×10^{22} a [102]. At the 0.1% (w/w) loading that optical studies suggest will be tolerable in SNO+, the 860 tonnes of SNO+ liquid scintillator would hold 48 kg of ^{150}Nd isotope. This is a competitive

⁴ There may in fact be an additional uncertainty on the $2\nu\beta\beta$ half-life due to the fact that the NEMO3 experiment [102], which is responsible for the best current $T_{1/2}^{2\nu}$ measurement, explicitly rejected events with co-incident γ -rays [111]. To the extent of the NEMO3 γ -ray detection efficiency, then, the reported $T_{1/2}^{2\nu}$ includes only decays to the ground state of ^{150}Sm . Therefore, for the purposes of SNO+, it may be necessary to add (some fraction of) the $1.33_{0.26}^{0.45} \times 10^{20}$ a branch to the first ^{150}Sm 0^+ excited state [112] to the $2\nu\beta\beta$ rate inferred from the NEMO3 result.

amount of isotope compared to other near-term double beta decay experiments (see Section 7.5.1). The SNO+ collaboration is also investigating the possibility of obtaining neodymium enriched in ^{150}Nd ⁵; neodymium enriched to 50% ^{150}Nd would mean 430kg of isotope in SNO+, making the experiment one of the largest double beta decay searches currently envisioned.

7.3 Development of Neodymium Loaded Liquid Scintillator

In early 2004, the idea of dissolving neutrinoless double beta decay isotope into a SNO-based scintillator experiment began to receive significant attention within the SNO+ working group. A similar idea, involving the dissolution of Xe gas into a large liquid scintillation detector had been proposed earlier [114, 115]. The principle advantage to this type of experiment would be scale; even a few percent loading in a kilo-tonne scale scintillator experiment gives tens of tonnes of material, well beyond the near-term reach of any other double beta decay search technique. Xe was an interesting element for this type of experiment because of its relatively high ($\sim 2\%$) solubility in scintillator and the relative ease with which it can be enriched in the ^{136}Xe double beta decay isotope. Xe also has the advantage that it does not chemically interact with the scintillator and does not significantly absorb light in the wavelength regime important to liquid scintillator function. On the other hand, the predicted $0\nu\beta\beta$ rate for ^{136}Xe is quite low, even at large loading, and the predicted ratio of $0\nu\beta\beta / 2\nu\beta\beta$ rates is also low⁶. In addition, an endpoint energy of ~ 2.5 MeV makes the Xe measurement particularly sensitive to the 2.6 MeV γ -ray background from external ^{208}Tl , and to internal ^{214}Bi background.

Other double beta decay isotopes exist which have higher endpoint energies, better $0\nu\beta\beta / 2\nu\beta\beta$ ratios, and higher rates than ^{136}Xe [103]. However, the loading of these isotopes into liquid scintillator is more challenging. One potential solution which was studied extensively by the SNO+ group during this time was the possibility of loading the double beta decay isotope into the liquid

⁵ Neodymium enrichment is currently only possible through the atomic vapour laser isotope separation (AVLIS) process [113]. A large AVLIS facility called MENPHIS for uranium enrichment has been developed in France; discussions about the potential to use this facility for neodymium enrichment are ongoing, but at the present time are not promising. Possible alternative AVLIS facilities in Russia are also being investigated.

⁶ As will be seen, the $2\nu\beta\beta$ signal is one of the dominant backgrounds to the $0\nu\beta\beta$ signal in scintillator experiments, so the $0\nu\beta\beta / 2\nu\beta\beta$ ratio is a very important consideration in choosing an isotope for a scintillator based neutrinoless double beta decay search.

scintillator as a nanocrystalline suspension. The impetus here was the observed effect that for very small nanoparticles, the size of the particles can have a significant effect on the optical activity of the constituent atoms. In particular, as the nanoparticle radius decreases, in general so does the wavelength of the absorbed and emitted light [116]. The thought, then, was that for sufficiently small nanoparticles the optical absorption might be pushed beyond the important region for optical transmission in a liquid scintillator. The scattering length is also expected to decrease as the particles are made smaller [117], so for sufficiently small nanoparticles one might expect to regain the optical benefits of Xe loading, but with the decay benefits of another isotope⁷.

To investigate this possibility, suspensions of various types of nanoparticles in water and pseudocumene⁸ were obtained from a commercial source⁹. These nanoparticles were all nominally smaller than 3nm diameter and were loaded into the liquids at approximately 0.5% by mass using proprietary surfactant techniques. By this time, ¹⁵⁰Nd had emerged as a leading candidate for the double beta decay experiment based on its high endpoint energy, high predicted $0\nu\beta\beta / 2\nu\beta\beta$ ratio, reasonable natural abundance, and high predicted $0\nu\beta\beta$ decay rate, so neodymium was paid particular attention. As was hoped, the nanoparticle neodymium suspensions exhibited significant reductions in neodymium atomic absorption, as can be seen in Figure 7.2.

The stability of the nanoparticle suspensions was variable, however, with many of them (especially those in water) forming visible crystalline precipitates after periods of days to weeks. The Nd₂O_{3(np)} in pseudocumene has not formed a visible precipitate after more than 5 years. Changes were observed in its optical attenuation over time, however, along with the potential plate-out of material onto the sides of the storage vessel. Also, the lack of optical activity in the neodymium meant that the loading level claimed by the company could not be confirmed. A brief summary of the nanoparticle investigations can be found in [118].

Although the nanoparticle loading technique is promising, and is in fact still being pursued by

⁷ Nanoparticle loading is also attractive as it would likely be relatively material non-specific, meaning that a number of different double beta decay isotopes could be run in sequence.

⁸ This is the liquid scintillator used in the BOREXINO experiment and, diluted with mineral oil, in the KamLAND experiment.

⁹ Applied NanoWorks, Watervliet NY. Now called Auterra Inc.

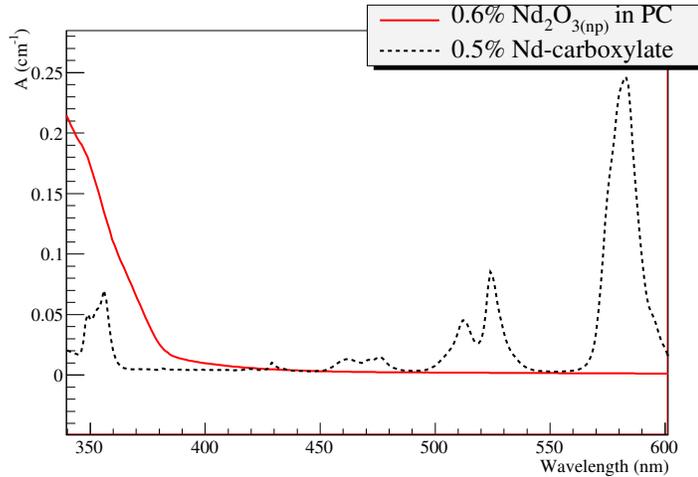


Figure 7.2: The optical attenuation of 0.6% $\text{Nd}_2\text{O}_3(\text{np})$ suspension in pseudocumene compared to the attenuation measured for neodymium loaded into the LAB scintillator (discussed below) in organometallic form (as Nd-carboxylate); the measurement was made at 0.1% neodymium loading and scaled to 0.5%. The increase in attenuation in the nanoparticle sample at low wavelengths is the absorption of the pseudocumene; the corresponding LAB absorption has been subtracted out of the Nd-carboxylate reference. The strong neodymium absorption lines visible in the carboxylate sample appear to be absent in the nanoparticle suspension.

some members of the SNO+ collaboration, the instability seen in the nanoparticle suspensions, coupled with the necessity of having a viable liquid scintillator experiment demonstrated and ready to begin when the SNO experiment was completed, meant that other avenues of neodymium loading were pursued. In particular, work following previously developed (although still experimental) techniques of loading neodymium [119] and gadolinium [120] into organic solution via organometallic compounds quickly led to the successful creation of a stable neodymium-carboxylate loaded liquid scintillator. As shown in Figure 7.2, the neodymium atomic absorption is fully present in these organometallic compounds. Parallel efforts to add scintillator and neodymium optical properties to the SNOMAN Monte Carlo simulation code (described in Section 6.2.1) showed that neodymium loading of 0.1% w/w might still be possible in spite of this absorption (see Figure 7.3).

Work is ongoing to confirm the light output of the actual neodymium-carboxylate loaded LAB scintillator that will be used in SNO+ will be sufficient at 0.1% loading. The primary light output of neodymium-loaded scintillator has been shown to be equal to unloaded scintillator to within

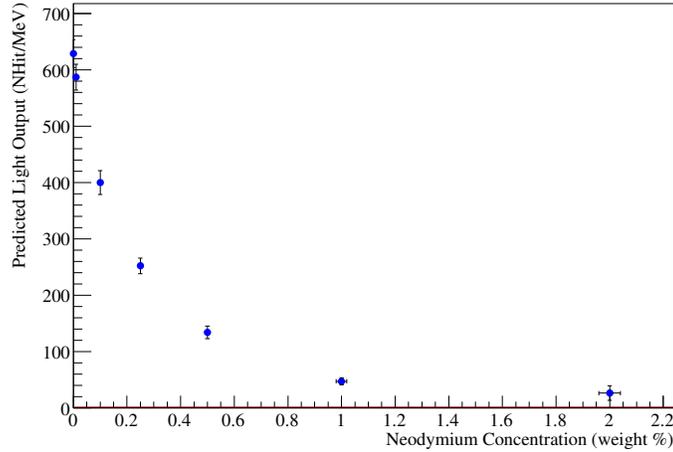


Figure 7.3: The simulated light output of central electrons SNO+ as a function of neodymium loading. This plot was made using the scintillator optical properties of pseudocumene with 1.5 g/L PPO wavelength shifter as measured by the BOREXINO collaboration [121] and the neodymium absorption spectrum as measured in a $\text{NdCl}_3(\text{aq})$ solution.

$\sim 10\%$ through direct measurements in the SNO detector¹⁰. The overall simulated light output of the neodymium-loaded scintillator, however, is currently lower (~ 90 PMT hits for a 1 MeV central electron) than would be ideal, and quite dependant on some of the scintillator optical properties that are the most challenging to measure accurately (it is especially sensitive to the long scintillator absorption and scattering lengths and the re-emitted photon spectra). It may well be that with improved measurements of these properties, the simulated light output in the SNO+ double beta decay phase will increase to match the early predictions shown in Figure 7.3; if not, the neodymium loading may have to be reduced to increase the amount of detected light. For now, however, it is assumed that the target of 400 NHit/MeV at 0.1% loading target can be met.

¹⁰ A clear acrylic cylinder “source” was constructed, filled with scintillator, and deployed into the light water filled SNO detector. By placing an AmBe neutron source near the cylinder, the light output of the scintillator could be observed using the SNO photomultiplier array. Repeating this measurement with both neodymium-loaded and unloaded scintillator allowed the primary light outputs to be compared (the cylinder was too small for the other scintillator optical properties to be important). See [122] for more information.

7.4 Simulations to Estimate the SNO+ Sensitivity to Neutrinoless Double Beta Decay

In order to estimate the sensitivity of the SNO+ experiment to neutrinoless double beta decay (and the other physics signals of interest) a standalone Monte Carlo simulation was developed which allowed the sensitivity of the experiment to be studied under different detector conditions.

7.4.1 Simulation Technique

The physics sensitivity simulations discussed here were separate from the detailed SNOMAN-based simulations of the SNO+ light output described earlier. Work on producing a detailed Monte Carlo (including event-by-event simulation of the energy deposition and light propagation within the detector, and employing the same event reconstruction algorithms that will be used on the data) is ongoing. These studies will give a more detailed understanding of the factors influencing the SNO+ physics sensitivity. However, these detailed simulations are not yet sufficiently developed to be used in physics sensitivity predictions, and (especially in the early stages of an experiment) a great deal can be learned from significantly simpler and faster simulations. The simulations described here are just such simpler simulations, and they have been used to demonstrate the SNO+ physics potential and to guide the development of the experiment.

In these simulations, high statistics spectra of the expected energy deposition in the scintillator by the radioactive backgrounds of interest were first simulated using the EGS4 (a standard particle interaction code) portion of SNOMAN, and the expected energy spectra of solar neutrino interactions (including the effects of MSW oscillations) were calculated using the QPhysics neutrino propagation and elastic scattering code developed by the SNO collaboration. These spectra were generated only once and stored as binned histograms. The $2\nu\beta\beta$ spectrum was taken as having the approximate analytic form,

$$\frac{dN}{dK} \sim K(T_0 - K)^5 \left(1 + 2K + \frac{4K^2}{3} + \frac{K^3}{3} + \frac{K^4}{30} \right), \quad (7.13)$$

where K is the sum of the electron energies and T_0 is the endpoint energy of the decay [123]¹¹. Together, these distributions provided a set of “known spectra” upon which the simulations were based.

When performing a simulation, the “known spectra” described above were used to create two additional distributions: the “data distributions” from which the fake data sets were drawn, and the PDFs that were used to fit the fake data sets. These new distributions were created by applying a simulated detector energy resolution to the “known spectra.” The standard energy resolution applied was a Gaussian smearing (applied as a numerical convolution between the binned energy deposition spectra and a Gaussian curve), the width of which was determined by counting statistics on the number of PMTs expected to detect photons for an event of a given energy¹². It was possible in the simulations to apply different energy resolutions to the data distributions and the PDFs (in order to study the effects of an improper understanding of the detector energy resolution) and to add exponential “tails” to the data distributions (in order to test the experiment’s sensitivity to possible systematic effects in energy reconstruction).

To generate fake data sets, each of the data distributions described earlier were scaled to contain the average number of events of that type desired in the fake data set. For each bin in the fake data set, the average contribution of a given class of events was then simply the entry of the same bin in the relevant (scaled) data distribution¹³. The true contribution of that class of event to that bin in the fake data set was then determined by randomly drawing a number from a Poisson distribution with that mean. Especially for large signals, this method is much quicker computationally than filling the data histogram by repeatedly randomly sampling the data distribution histograms to

¹¹ This formula is approximate in that non-relativistic Coulomb effects are assumed in the Fermi Functions; including the relativistic effects numerically has been shown to shift the peak in the $2\nu\beta\beta$ spectrum lower in energy and steepen the spectrum near the endpoint [104].

¹² Energy reconstruction in SNO+, as in SNO, will consist of a mapping between the energy deposited in the scintillator by an event and the number of PMTs which detect photons as a result (“NHits”). The limit on the obtainable energy resolution is then set by the statistical fluctuations in the number of PMTs hit. For example, at a light output of 400 NHit/MeV, 2 MeV electrons are expected to produce an average of 800 PMT hits with a 1σ statistical fluctuation of $\sqrt{800}$, or about 28. So the 1σ fractional energy resolution is expected to be approximately $1/\sqrt{\text{NHit}}$. The actual resolution will be broadened somewhat by effects such as the position-dependant energy response of the detector. These detailed effects will be one of the main subjects of study of the detailed simulation code once it is completed.

¹³ Note that for this procedure to work properly, the fake data histogram and the data distribution histograms had to be identically binned.

simulate individual events.

7.4.2 Backgrounds Simulated

Only internal backgrounds (i.e. backgrounds originating from the LAB scintillator itself) were included in this simulation. External backgrounds, principally neutrons and γ -rays produced in the light water and acrylic vessel which then propagate into the scintillator volume will also be present in the experiment, but their numbers decay quickly with distance from the edge of the detector. Independent simulations [124] suggest that the rate of external background events, dominated by ^{208}Tl from the acrylic vessel, will equal the rate of $0\nu\beta\beta$ events approximately 1.5 m in from the acrylic vessel. For this reason, SNO+ will likely operate with a fiducial radius of 4.5 m (i.e. only those events which reconstruct within 450 cm of the centre of the detector will be included in the analysis), and neglecting the external backgrounds within this fiducial radius is a reasonable first approximation.

As shown in Figure 7.4, many potential internal backgrounds, including ^{40}K , ^{39}Ar , ^{85}Kr , and all of the isotopes in the ^{238}U chain and the ^{232}Th chains were included in the simulations. These isotopes (along with ^{14}C which has a lower energy endpoint) are believed to account for the radiogenic backgrounds observed in the KamLAND [125] and BOREXINO [22] experiments. Most of these backgrounds have decay energies somewhat lower than the $0\nu\beta\beta$ peak and were included in the simulation to make possible low energy solar neutrino sensitivity studies [66]. As will be seen, the SNO+ neutrinoless double beta decay search is essentially insensitive to the energy spectrum below 2.5-3.0 MeV, so the only internal backgrounds important to the double beta decay sensitivity are ^{208}Tl in the ^{232}Th chain and ^{214}Bi in the ^{238}U chain.

Additional internal backgrounds from *in situ* cosmogenic activation are also possible, but studies of the potential cosmogenic activation of both scintillator [68] and neodymium [126, 127] have so far failed to identify any isotopes with half-lives long enough and decay energies high enough to pose a significant concern to the double beta decay measurement through *in-situ* cosmogenic activation¹⁴.

¹⁴ Although with a 285 day half-life, ^{144}Ce (which supports ^{144}Pr which has a 3 MeV beta decay) produced in the neodymium before it is brought underground may be a concern. This is currently under investigation.

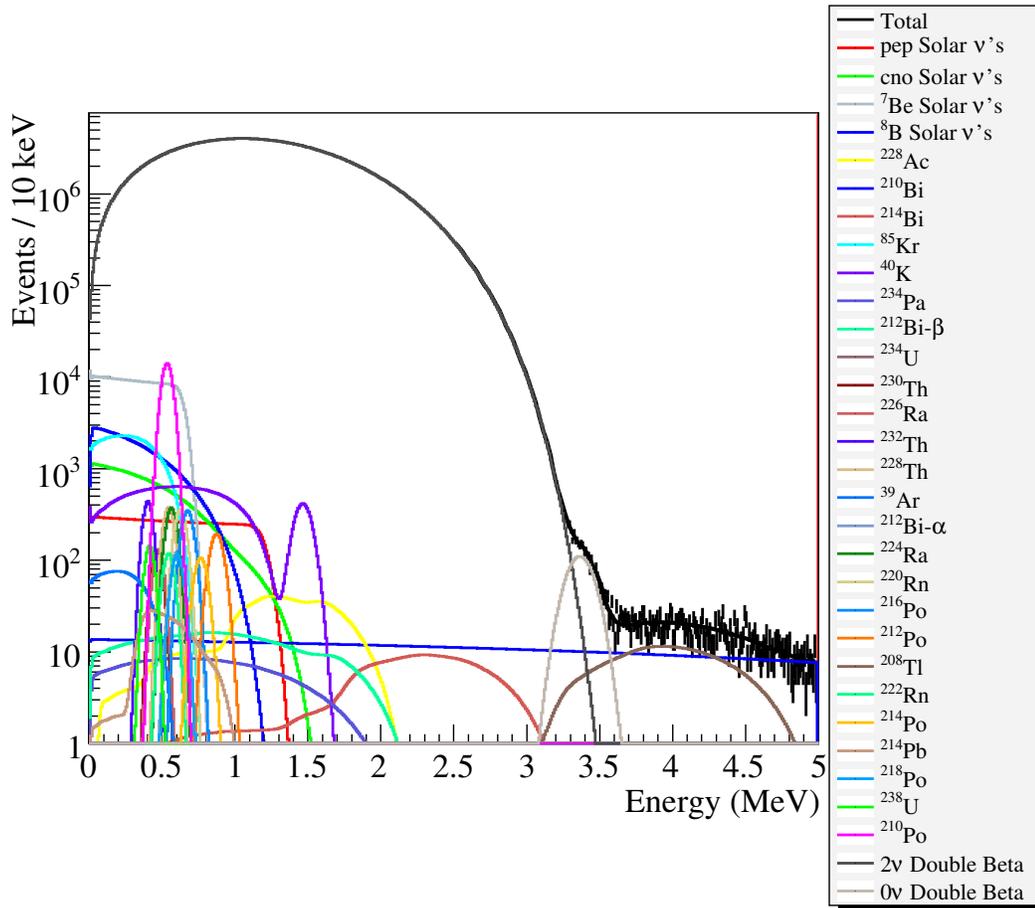


Figure 7.4: A spectrum generated by the SNO+ simulation including low energy signals and backgrounds. The backgrounds below about 3.0 MeV are unimportant to the determination of the SNO+ $0\nu\beta\beta$ sensitivity (due to the overwhelming $2\nu\beta\beta$ statistics in this region), but have an important impact on the sensitivity to low energy solar neutrinos. Due to space restrictions, not all simulated decays are listed in the legend. The event rates in this simulation differ slightly from the levels described in the text. This plot includes the detector energy resolution expected at a light output of 400 NHit/MeV, and assumes $\langle m_\nu \rangle \sim 170$ meV. The spectra of any low energy backgrounds which may be added to the detector along with the neodymium are not included in the simulation.

A final (and rather amusing, from the point of view of one who was also involved with SNO) important background to the $0\nu\beta\beta$ signal in SNO+ will be the elastic scattering interaction of ^8B solar neutrinos.

7.4.3 Fit Range

The energy range over which the fits to the energy spectrum are performed might be expected to have some effect on the predicted sensitivity of the experiment. In particular, as the lower limit of the fit range is extended to lower energies, more and more of the $2\nu\beta\beta$ signal is included in the fit, and hence the constraint on the $2\nu\beta\beta$ background in the $0\nu\beta\beta$ signal region might be expected to improve. Similarly, raising the upper limit of the fit range might be expected to give a statistically better constraint on the ^8B elastic scattering background. As can be seen in Figure 7.5, including the signal below about 3.0 MeV in the fit gives limited statistical gain (because the number of $2\nu\beta\beta$ events increases quickly with decreasing energy). The ^8B elastic scattering, on the other hand, has a low enough rate that there is benefit in extending the fit range up to about 5.0 MeV. In the actual SNO+ experiment, the fits will likely be performed over as wide a background range as possible in order to ensure the best result possible statistically. Here, however, in the interests of computational speed the fit range was restricted to 3.0 - 5.0 MeV.

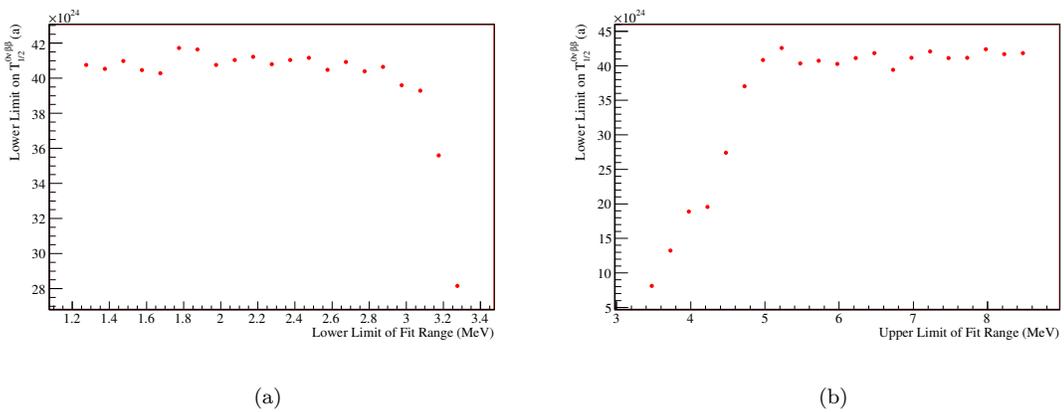


Figure 7.5: The effect of fit range on the SNO+ $0\nu\beta\beta$ fits. In (a), the upper limit of the fit range was held at 5.0 MeV and the lower limit was varied. In (b) the lower limit of the fit range was fixed at 2.8 MeV and the upper limit was varied. As can be seen, there is little statistical benefit in extending the fit range below 3.0 MeV or above 5.0 MeV.

7.4.4 Quantifying the $0\nu\beta\beta$ Sensitivity

The sensitivities of neutrinoless double beta decay experiments are typically expressed in terms of the limit that the experiment would set on the $0\nu\beta\beta$ signal in the absence of a detection. Thus, to determine the SNO+ sensitivity using the simulations described above, no $0\nu\beta\beta$ signal was included in the fake data sets, but the $0\nu\beta\beta$ PDF was included in the fits. The fitted number of $0\nu\beta\beta$ events and its uncertainty were then used to define a neutrinoless double beta decay limit at the required confidence level.

Thus, defining the $0\nu\beta\beta$ limit for a single fake data set was simple. There is, however, an additional subtlety associated with defining the neutrinoless double beta decay sensitivity of an experiment. This is because statistical fluctuations will cause variation in the $0\nu\beta\beta$ limits set by different fake data sets (or equivalently by different runs of the same real experiment). In order to define an overall sensitivity for an experiment, then, it is necessary to define not only the confidence limit being used, but also the frequency with which that confidence limit, as set by a single iteration of the experiment, will satisfy the stated criteria¹⁵. There appears to be a consensus in the community that $0\nu\beta\beta$ sensitivities be given at the 90% C.L. and a frequency of 50%.

There is also some uncertainty about whether $0\nu\beta\beta$ limits should be given in terms of limit on the half-lives (or equivalently decay rates) of the isotope in question, or whether they should be converted into effective Majorana neutrino mass using Equation 7.3. The latter is preferable in that it allows experiments using different isotopes to be compared according to the most interesting physics parameter; it suffers the drawback, however, that the limits thus inherit the uncertainty in the calculation of the matrix elements.

In Section 7.5, the SNO+ $0\nu\beta\beta$ limits given are the levels at which there is a 50% chance (based on multiple simulations) of a single experiment achieving a 90% C.L. on $T_{1/2}^{0\nu}$ that high or higher.

Tables of the average $0\nu\beta\beta$ fit central values and uncertainties will also be given, so that other

¹⁵ That is, it is necessary to state limits of the form “the experiment is expected to set an X%C.L. upper limit on the $0\nu\beta\beta$ signal that is less than Y Z% of the time.”

confidence limits or frequencies can easily be computed if desired¹⁶. The conversion to limits on the effective Majorana neutrino mass will only be given in the section involving comparisons to other experiments; for a given matrix element calculation, however, the conversion is easily made, using the half-life values for $\langle m_\nu \rangle = 50$ meV in Table 7.1. In Section 7.6, the $0\nu\beta\beta$ levels given are the 95%C.L. 50% frequency upper limits on $T_{1/2}^{0\nu}$ ¹⁷

7.5 Simulated SNO+ $0\nu\beta\beta$ Sensitivity

The “baseline” sensitivity of the SNO+ experiment to neutrinoless double beta decay is defined to be the sensitivity when all of the detector parameters are at their nominal values. The nominal detector conditions assumed for SNO+ are:

- A light output of 400 NHit/MeV, taken from the optical simulations described earlier.
- A neodymium loading of 0.1% (w/w), with a ¹⁵⁰Nd isotopic abundance of 5.6% in natural neodymium and 50% in enriched neodymium.
- A ²⁰⁸Tl rate of 383 events/kT·a (equivalent to 8.3×10^{-18} g/g ²³²Th in the scintillator)¹⁸.
- A ²¹⁴Bi rate of 6682 events/kT·a (equivalent to 1.7×10^{-17} g/g ²³⁸U in the scintillator).
- A ⁸B elastic scattering rate of 2099 events/kT·a¹⁹. The ⁸B elastic scattering energy spectrum was assumed to have MSW distortions with $\Delta m_{12}^2 = 8 \times 10^{-5}$ eV² and $\tan^2 \theta_{12} = 0.45$.
- An assumed $T_{1/2}^{2\nu}$ of 9.2×10^{25} a (the central value of the best current measurement [102]), yielding 1.7×10^7 $2\nu\beta\beta$ events/kT·a for natural neodymium loading and 1.5×10^8 $2\nu\beta\beta$ events/kT·a for enriched neodymium.

¹⁶ It is possible to do this using just these averages because the uncertainties on the fitted numbers of $0\nu\beta\beta$ events are remarkably tightly grouped and independent of the statistical fluctuations in the central values of the fit. Thus, the distribution of upper limits at a given confidence level is essentially identical to the distribution of fit central values shifted upwards by a given amount. For example, if the distribution of fitted central values is a Gaussian of width σ centred about zero, the distribution of 90% upper confidence limits will be a Gaussian of width σ centred about 1.3σ ; thus 50% of the 90% C.L. upper limits will lie within 1.3σ of zero, while 90% of them will lie within $2 \times 1.3\sigma$ of zero.

¹⁷ The change in confidence level relative to Section 7.5 was due to a coding error which resulted in two-sided, rather than one-sided, 90%C.L.'s being applied in Section 7.6.

¹⁸ The levels of ²³²Th and ²³⁸U are taken to be the levels observed in the BOREXINO experiment. BOREXINO observes $(1.6 \pm 0.1) \times 10^{-17}$ g/g ²³⁸U equivalent in their scintillator using the ²¹⁴Bi - ²¹⁴Po delayed coincidence and $(6.8 \pm 1.5) \times 10^{-18}$ g/g ²³²Th equivalent using the ²¹²Bi - ²¹²Po delayed coincidence [22].

¹⁹ This is the Standard Solar Model prediction after MSW oscillations with $\Delta m_{12}^2 = 8 \times 10^{-5}$ eV² and $\tan^2 \theta_{12} = 0.45$.

- A live time of 1 kT·a.
- Identical energy resolution in the data and PDF distributions (i.e. perfect knowledge of the energy spectra that underlie the data).
- Where conversion to effective neutrino mass is necessary, a matrix element corresponding to $T_{1/2}^{0\nu} = 4.0 \times 10^{25}$ a at $\langle m_\nu \rangle = 50$ meV (the upper edge of the QRPA/RQRPA calculations, although the current range of calculations extends from a factor of four lower to a factor of ten higher) is used.

These conditions can be assumed to apply to all simulation results shown in this thesis unless stated otherwise.

Under these default conditions, the 50% frequency 90%C.L. upper limits that SNO+ could set on the $0\nu\beta\beta$ signal are shown in Figure 7.6(a) for different live times and for both natural and enriched neodymium. These results are also given in Tables 7.2 and 7.3 along with the corresponding average numbers of fitted $0\nu\beta\beta$ events and their uncertainties. Figure 7.6 (b) shows the same limits in terms of $\langle m_\nu \rangle$.

The first goal for the next generation of neutrinoless double beta decay experiments is to test the controversial result from [88], which claims detection of $0\nu\beta\beta$ with $\langle m_\nu \rangle \gtrsim 270$ meV at 90%C.L.. This should be well within the reach of the SNO+ double beta decay search using either natural or enriched neodymium. For interest, Figure 7.7 shows simulated SNO+ spectra near the endpoint for both natural and enriched neodymium at $\langle m_\nu \rangle = 270$ meV.

Looking beyond the 270 meV target, the next interesting $\langle m_\nu \rangle$ target for neutrinoless double beta decay searches is the 10-40 meV range, which (as shown in Figure 7.1), tests the possibility of an inverted neutrino mass hierarchy. Depending on the actual ^{150}Nd matrix element, SNO+ could probe the upper edges of this interesting region with natural neodymium, or explore a reasonable fraction of it using enriched neodymium.

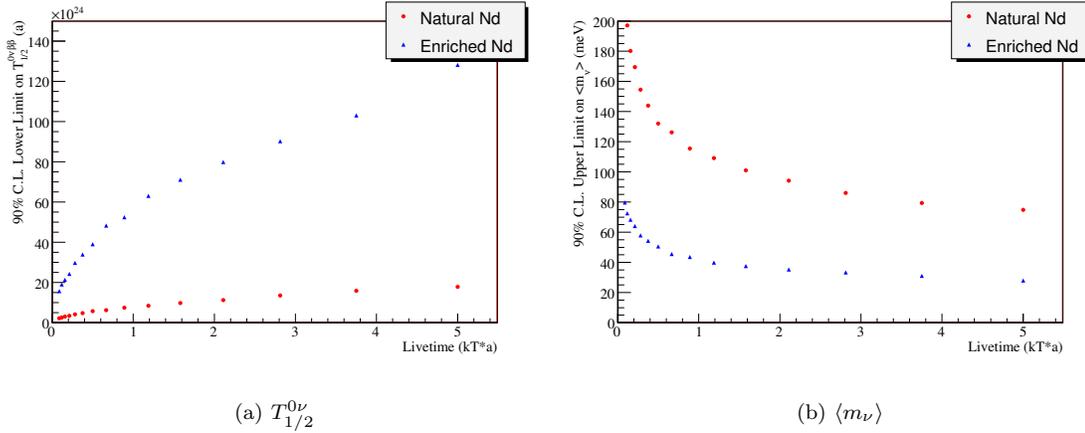


Figure 7.6: The baseline SNO+ $0\nu\beta\beta$ sensitivity as a function of live time. As SNO+ will have a real mass of 860kg, a fiducial volume cut that rejects events in $\sim 50\%$ of the active volume, and will likely spend about 25% of its time performing calibrations and other maintenance tasks, it will take approximately three real years of running to accumulate one kT·a of data. Note that the range of variation in the ^{150}Nd matrix element calculations to date admit variation in the $\langle m_\nu \rangle$ curves shown in (b) within a factor of 2 lower and a factor of 3 higher. 1000 fake data sets were thrown and fitted in generating each point on the plots.

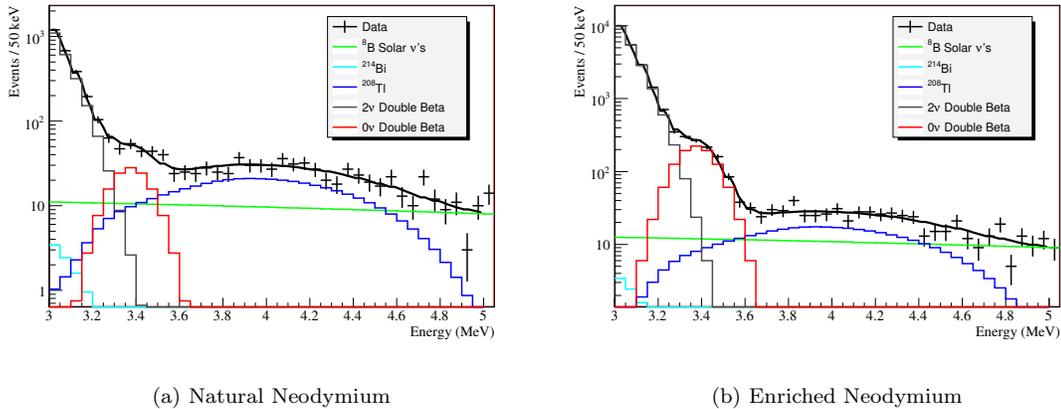


Figure 7.7: The simulated SNO+ spectrum near the endpoint for $\langle m_\nu \rangle = 270$ meV after one kT·a of data taking. The $0\nu\beta\beta$ peak is easily detected with both natural and enriched neodymium, with the fitted number of $0\nu\beta\beta$ events 7σ from zero after one kT·a of data in the case of natural neodymium, and 25σ from zero in the case of enriched neodymium.

| Live time (kT·a) | $T_{1/2}^{0\nu}$ Limit (x10 ²⁴ a) | $\langle m_\nu \rangle$ Limit (meV) | $\langle \text{NFit}_{0\nu\beta\beta} \rangle$ | $\langle \Delta \text{NFit}_{0\nu\beta\beta} \rangle$ |
|---------------------|---|--|--|---|
| 5.00 | 17.8 | 75 | 0.0 ± 1.1 | 34.10 ± 0.02 |
| 3.75 | 15.9 | 79 | -1.1 ± 0.9 | 29.57 ± 0.02 |
| 2.81 | 13.5 | 86 | -0.4 ± 0.8 | 25.64 ± 0.02 |
| 2.11 | 11.3 | 94 | 0.6 ± 0.7 | 22.29 ± 0.02 |
| 1.58 | 9.8 | 101 | 0.3 ± 0.6 | 19.36 ± 0.02 |
| 1.19 | 8.4 | 109 | 0.5 ± 0.5 | 16.83 ± 0.02 |
| 0.89 | 7.5 | 115 | -0.2 ± 0.4 | 14.60 ± 0.02 |
| 0.67 | 6.3 | 126 | 0.3 ± 0.4 | 12.71 ± 0.02 |
| 0.50 | 5.7 | 132 | -0.5 ± 0.3 | 11.00 ± 0.02 |
| 0.38 | 4.8 | 144 | -0.2 ± 0.3 | 9.58 ± 0.02 |
| 0.28 | 4.2 | 155 | -0.2 ± 0.3 | 8.33 ± 0.02 |
| 0.21 | 3.5 | 169 | 0.1 ± 0.2 | 7.31 ± 0.02 |
| 0.16 | 3.1 | 180 | -0.1 ± 0.2 | 6.35 ± 0.02 |
| 0.12 | 2.6 | 197 | 0.1 ± 0.2 | 5.52 ± 0.02 |
| 0.09 | 2.2 | 214 | 0.2 ± 0.1 | 4.84 ± 0.03 |

Table 7.2: The baseline SNO+ sensitivity to $0\nu\beta\beta$ with natural neodymium for different run times. The second (third) column gives the 90% C.L. lower (upper) limit that could be placed on $T_{1/2}^{0\nu}$ ($\langle m_\nu \rangle$) 50% of the time after that amount of live time. The fourth and fifth columns give the average fit value ($\langle \text{NFit}_{0\nu\beta\beta} \rangle$) and fit uncertainty ($\langle \Delta \text{NFit}_{0\nu\beta\beta} \rangle$) for the number of $0\nu\beta\beta$ events in the simulations at each live time; these values can be used to calculate other sensitivity limits if desired.

7.5.1 Comparison with Other $0\nu\beta\beta$ Experiments

A review of the $0\nu\beta\beta$ experiments currently being built or considered is given in [105]. Only a few of these experiments with the greatest potential to produce comparative $0\nu\beta\beta$ limits in the near- and medium-term are discussed here.

- **GERDA:** The GERDA experiment [128] will search for double beta decay in ^{76}Ge using enriched germanium detectors. The experiment is similar to, and uses the enriched germanium detectors from, the Heidelberg-Moscow experiment [129] from which the current $0\nu\beta\beta$ detection claim originated. GERDA will also use the enriched germanium detectors from the IGEX experiment [130]. GERDA will ultimately operate using approximately 50 kg of enriched germanium, and is expected to reach a sensitivity of $\langle m_\nu \rangle \lesssim 125$ meV ($T_{1/2}^{0\nu}(^{76}\text{Ge}) \gtrsim 1.4 \times 10^{26}$ a)

| Live time (kT·a) | $T_{1/2}^{0\nu}$ Limit ($\times 10^{24}$ a) | $\langle m_\nu \rangle$ Limit (meV) | $\langle \text{NFit}_{0\nu\beta\beta} \rangle$ | $\langle \Delta \text{NFit}_{0\nu\beta\beta} \rangle$ |
|---------------------|---|--|--|---|
| 5.00 | 128.1 | 28 | -0.6 ± 1.4 | 42.85 ± 0.03 |
| 3.75 | 103.1 | 31 | 2.9 ± 1.1 | 37.22 ± 0.03 |
| 2.81 | 90.2 | 33 | 2.0 ± 1.0 | 32.30 ± 0.03 |
| 2.11 | 79.9 | 35 | 0.8 ± 0.9 | 28.02 ± 0.03 |
| 1.58 | 71.1 | 38 | -0.1 ± 0.8 | 24.29 ± 0.03 |
| 1.19 | 63.1 | 40 | -0.8 ± 0.7 | 21.05 ± 0.03 |
| 0.89 | 52.5 | 44 | 0.1 ± 0.6 | 18.32 ± 0.03 |
| 0.67 | 48.3 | 46 | -1.1 ± 0.5 | 15.89 ± 0.03 |
| 0.50 | 39.1 | 51 | 0.1 ± 0.4 | 13.85 ± 0.03 |
| 0.38 | 33.9 | 54 | -0.0 ± 0.4 | 12.04 ± 0.03 |
| 0.28 | 29.8 | 58 | -0.2 ± 0.3 | 10.44 ± 0.03 |
| 0.21 | 24.3 | 64 | 0.3 ± 0.3 | 9.17 ± 0.03 |
| 0.16 | 21.4 | 68 | 0.1 ± 0.2 | 7.96 ± 0.03 |
| 0.12 | 19.0 | 72 | -0.2 ± 0.2 | 6.93 ± 0.03 |
| 0.09 | 15.8 | 80 | 0.1 ± 0.2 | 6.07 ± 0.03 |

Table 7.3: The baseline SNO+ sensitivity to $0\nu\beta\beta$ with enriched neodymium for different run times. The second (third) column gives the 90%C.L. lower (upper) limit that could be placed on $T_{1/2}^{0\nu}$ ($\langle m_\nu \rangle$) with 50% frequency. The fourth and fifth columns give the average fit value and uncertainty for the number of $0\nu\beta\beta$ events in the simulations at each live time; these values can be used to calculate other sensitivity limits if desired.

with 100 kg·a of data. The experiment is currently under construction and is expected to be fully operational within the next few years.

- **MAJORANA:** The Majorana experiment [131] will be a larger (~ 500 kg) enriched germanium experiment which should reach a sensitivity of $\langle m_\nu \rangle \lesssim 60$ meV ($T_{1/2}^{0\nu}({}^{130}\text{Te}) \gtrsim 5.5 \times 10^{26}$ a) after about one live year of data taking. Majorana is currently in the development/proposal stage and will likely not be operational for several years.
- **CUORE:** The CUORE experiment [132] will search for $0\nu\beta\beta$ in ${}^{130}\text{Te}$ using an array of cryogenic TeO_2 crystal bolometers. The CUORE experiment will contain 740 kg of TeO_2 , which should allow the experiment to reach a sensitivity of $\langle m_\nu \rangle \lesssim 45$ meV ($T_{1/2}^{0\nu}({}^{130}\text{Te}) \gtrsim 2.5 \times 10^{26}$ a) in 5 years' exposure. The use of enriched tellurium, which is currently being explored, could improve the sensitivity to $\langle m_\nu \rangle \lesssim 30$ meV. CUORE has received Italian funding and is currently

under construction, with data taking scheduled to begin in 2011. The prototype CUORICINO experiment [133] has been operating for several years and has reached a sensitivity of $\langle m_\nu \rangle \lesssim 400$ meV ($T_{1/2}^{0\nu}(^{130}\text{Te}) \geq 3.0 \times 10^{24}$ a).

- **EXO200:** The EXO200 experiment [134] will search for $0\nu\beta\beta$ in ^{136}Xe using a liquid xenon time projection chamber which will measure both scintillation light and charge production. The experiment, which will contain 200 kg of enriched xenon (80% ^{136}Xe), is currently being deployed and should become operational in the near future. The experiment is expected to reach a sensitivity of $\langle m_\nu \rangle \lesssim 300$ meV.
- **MOON:** MOON [135] is a proposed experiment which will search for neutrinoless double beta decay using thin foil sources sandwiched between thin position sensitive detectors (possibly multi-wire proportional chambers) and plastic scintillator calorimeters. MOON will study double beta decay in ^{82}Se , ^{100}Mo , and/or ^{150}Nd , and has a target design sensitivity of $\langle m_\nu \rangle \lesssim 50$ meV. The MOON experiment is likely several years away from operation.
- **SuperNEMO:** SuperNEMO [136] will build on the successful operation of the NEMO3 detector [102], to search for $0\nu\beta\beta$ in ^{82}Se and/or ^{150}Nd . SuperNEMO will consist of thin foil sources within gas tracking chambers to allow reconstruction of the $\beta\beta$ tracks. SuperNEMO will contain 100-200 kg of isotope and should reach a sensitivity of $\langle m_\nu \rangle \lesssim 50$ -100 meV ($T_{1/2}^{0\nu} \gtrsim 10^{26}$ a). SuperNEMO is currently in the proposal stage and will likely not be operational for several years.

As can be seen from the list above, there is a set of “near-term” experiments which should soon become operational and which are expected to reach $0\nu\beta\beta$ sensitivities corresponding to $\langle m_\nu \rangle \lesssim 150$ -300 meV in the next three to five years. Another generation of experiments with sensitivities in the $\langle m_\nu \rangle \lesssim 50$ meV region are currently at the proposal stage or beginning construction.

With natural neodymium, SNO+ offers significantly better $0\nu\beta\beta$ sensitivity than the near-term generation of experiments (for any of the currently computed matrix elements) on an approximately

similar time scale. However, the next generation of experiments, likely led by CUORE, would soon eclipse this sensitivity. If enriched neodymium could be obtained, however, SNO+ would have competitive and (depending on the matrix elements) perhaps leading sensitivity even compared to the medium term generation of experiments.

7.6 Impact of Detector Parameters on the SNO+ $0\nu\beta\beta$ Sensitivity

In designing different aspects of the SNO+ experiment, from the target scintillator light output and radio-purity levels to the calibration systems, it is necessary to understand how the sensitivity of the experiment changes with different detector parameters. One of the most important uses of these simple simulations has been to provide feedback to guide the development of various aspects of the SNO+ project. Some of these studies are described below.

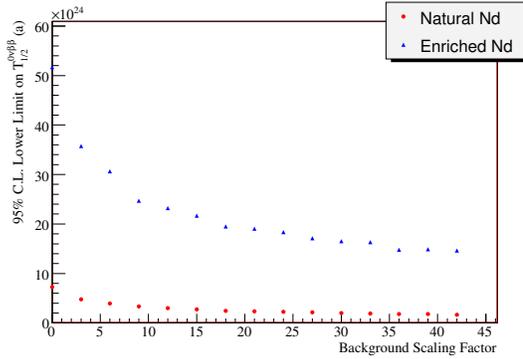
7.6.1 Background Levels

As previously mentioned, three backgrounds (apart from the irreducible $2\nu\beta\beta$ background) are relevant to the SNO+ double beta decay experiment: ^{208}Tl , ^{214}Bi , and the ^8B solar neutrinos. From Figure 7.7, it might reasonably be expected that at the baseline SNO+ levels, ^{208}Tl and the ^8B solar neutrinos are the dominant backgrounds, with ^{214}Bi being significantly less important. Figure 7.8 shows the effect of varying the level of each of these backgrounds in turn²⁰.

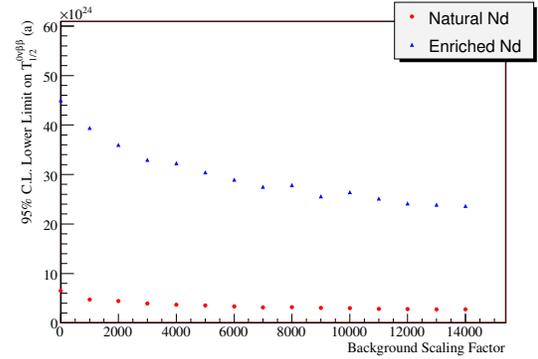
If the Tl and Bi backgrounds were both completely eliminated (but their PDFs left in the fit) the SNO+ $T_{1/2}^{0\nu}$ lower limits would improve to 7.7×10^{24} a for natural neodymium and 52×10^{24} a for enriched; if the ^8B solar neutrino background could simultaneously be reduced to 50% (10%) of its nominal value, the $T_{1/2}^{0\nu}$ limits would increase further to 10.2×10^{24} a (17.6×10^{24} a) for natural neodymium and 64.2×10^{24} a (94.8×10^{24} a) for enriched.

It is also interesting to note that varying the ^{150}Nd $2\nu\beta\beta$ half life within its $\sim 9\%$ measurement uncertainty has only a very small effect on the $0\nu\beta\beta$ limits.

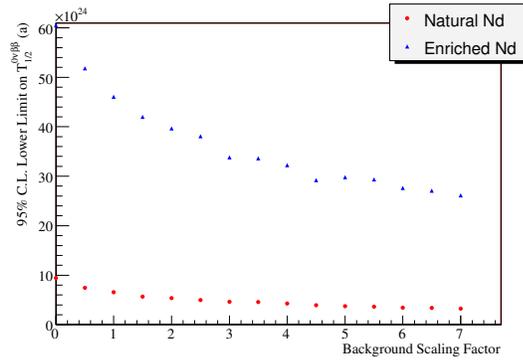
²⁰ The magnitude of the ^8B solar neutrino signal could in principle be reduced by confining the double beta decay isotope within a BOREXINO style “balloon” in the centre of SNO+.



(a) Varying ^{208}Tl Level



(b) Varying ^{214}Bi Level



(c) Varying ^8B Solar ν Level

Figure 7.8: The effect of varying the background levels on the SNO+ $0\nu\beta\beta$ limits. A background scaling of 1.0 corresponds to the baseline level of that background. Varying the important backgrounds one at a time confirms that ^{208}Tl and the ^8B solar neutrinos, along with the $2\nu\beta\beta$ signal, are the dominant backgrounds to the $0\nu\beta\beta$ signal. ^{214}Bi has such a weak effect that its level must be increased by orders of magnitude before it begins to noticeably impact the $0\nu\beta\beta$ limit.

7.6.2 Light Output/Energy Resolution

As discussed earlier, one of the most important parameters of a $0\nu\beta\beta$ experiment is the energy resolution, as this determines the extent to which the $2\nu\beta\beta$ signal is “smeared into” the $0\nu\beta\beta$ energy region. In SNO+, the energy resolution will primarily be determined by the amount of light collected by the detector. Figure 7.9 shows the lower limits on $T_{1/2}^{0\nu}$ that could be achieved in SNO+ after 1 kT·a of data with different light output levels. Note that the light output is changed in both the data and the PDFs so that while the energy resolution is varied it is still assumed to be perfectly understood.

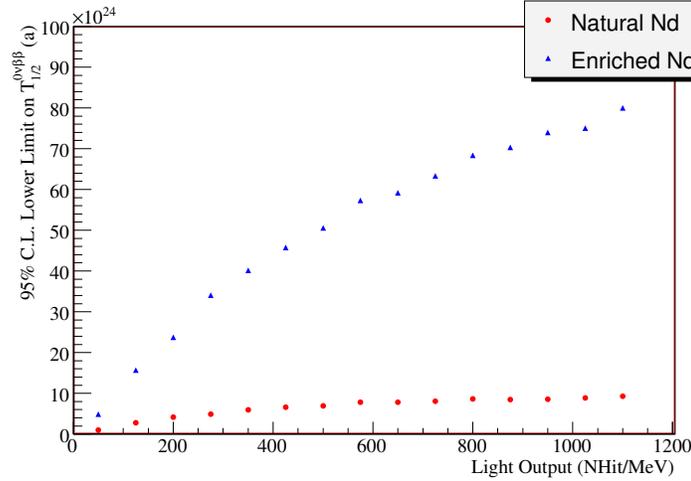


Figure 7.9: The effect of energy resolution on the SNO+ $0\nu\beta\beta$ limits (for an otherwise “baseline” detector and 1 kT·a of data). Increasing the light output (which improves the fractional energy resolution according to $1/\sqrt{N\text{Hit}}$) improves the $0\nu\beta\beta$ limits by reducing the amount of $2\nu\beta\beta$ signal in the $0\nu\beta\beta$ energy region. Note that it is assumed in all cases that the energy response of the detector is perfectly understood.

7.6.3 Neodymium Loading Levels

The effect of changing the amount of ^{150}Nd in the detector can be seen by comparing the limits set using enriched and natural neodymium. Nevertheless, it may be useful to have specific plots of the $0\nu\beta\beta$ limits as a function of neodymium loading, as once the light output of the neodymium-loaded scintillator is better predicted, an optimization in the neodymium loading level should be performed (optimization is necessary because increasing the loading increases the amount of isotope

but decreases the light output). These plots are shown in Figure 7.10.

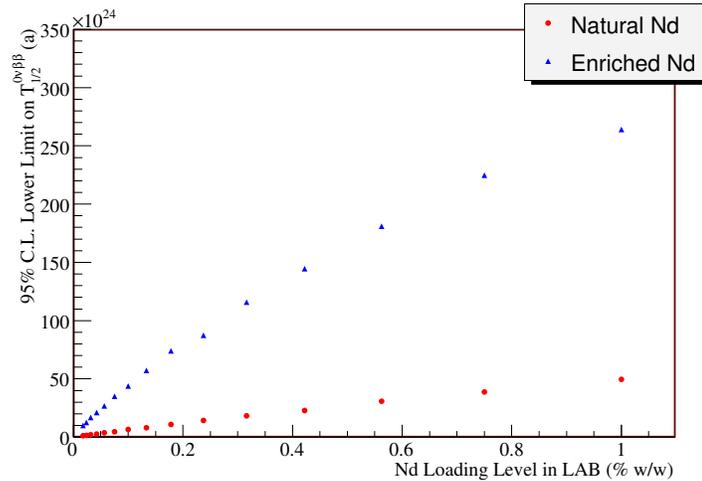


Figure 7.10: The effect of the neodymium loading level on the SNO+ $0\nu\beta\beta$ limits. Note that natural neodymium is assumed to be 5.6% (w/w) ^{150}Nd , while enriched neodymium is taken to be 50% ^{150}Nd . Although obvious, it is worth noting again that if the loading could be increased by a factor of 10 without reducing the light output (using nanoparticles, perhaps), 1% loading with natural neodymium offers approximately the same “performance” as 0.1% loading with enriched neodymium.

7.6.4 Energy Resolution Systematics

In setting a limit on the $0\nu\beta\beta$ signal in a data set, it is very important that the spectral shapes of the backgrounds be well understood. In particular, any small distortions in the $2\nu\beta\beta$ spectrum, if not included in the $2\nu\beta\beta$ PDF, can have a significant effect on the fitted number of $0\nu\beta\beta$ events and hence the $0\nu\beta\beta$ limit. One obvious potential cause of an imperfect understanding of the background spectra would be an incorrect characterization of the energy response of the detector. It is therefore useful to attempt to get an idea of what precision will be necessary in the energy calibration of SNO+ in order not to negatively affect the $0\nu\beta\beta$ limits.

The first check to be made is to determine how well the average energy resolution of the detector must be understood. This can be done by simply changing the width of the energy resolution function that was used to broaden the PDF distributions in the simulations relative to the resolution that was applied to the data (note however, that both energy resolutions were still assumed to be Gaussian). As can be seen in Figure 7.11 (b), overestimating the width of the detector energy response in making

the PDFs results in a positive bias in the number of fitted $0\nu\beta\beta$ events, as the fitter attempts to fill the empty space beneath the (broader) $2\nu\beta\beta$ PDF; underestimating the width of the energy resolution has the opposite effect. As the 95% C.L. upper limit on the average number of fitted $0\nu\beta\beta$ events (which is used to create the $T_{1/2}^{0\nu}$ limit) is taken to be 1.6σ above the average fitted $0\nu\beta\beta$ value, these biases in the fitted $0\nu\beta\beta$ values are reflected in the $T_{1/2}^{0\nu}$ limits, as shown in Figure 7.11 (a). In order to keep this mean bias in the fitted number of $0\nu\beta\beta$ events less than about 5 (or about 0.3σ in any single experiment) and hence prevent a large effect on the $0\nu\beta\beta$ limits, the energy resolution must be known to within about 0.1% (absolute) at the endpoint. As the overall fractional energy resolution at the endpoint is 2.7%, this is equivalent to the requirement that the average energy resolution be known to better than $\sim 4\%$ of its value. As SNO, with its lower light output, was able to model the average energy resolution at approximately the 1.5% level [20], this seems to be an achievable requirement for SNO+.

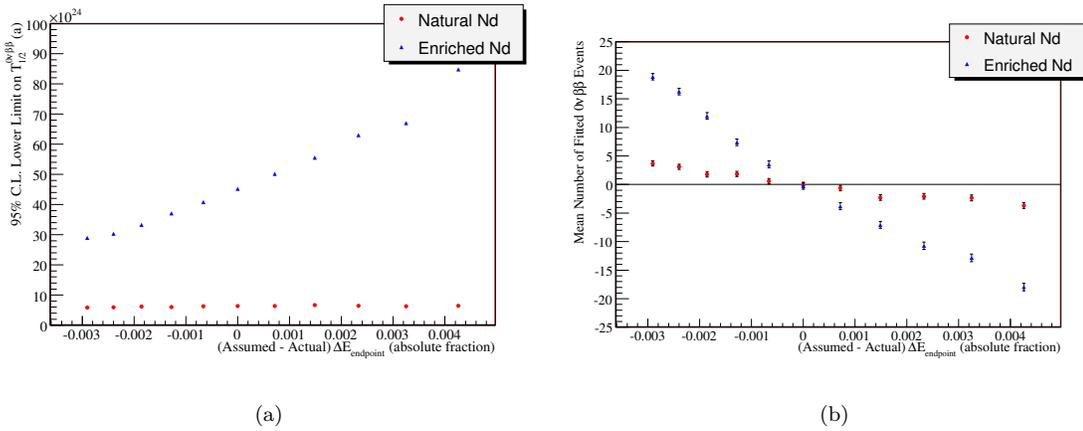


Figure 7.11: The effect of energy resolution on the SNO+ $0\nu\beta\beta$ limits. The change in the 95% C.L. upper limits on $T_{1/2}^{0\nu}$ that could be set by SNO+ as the difference between the true and assumed energy resolutions is varied is shown in (a). The changes in the $T_{1/2}^{0\nu}$ limits are primarily driven by changes in the mean number of $0\nu\beta\beta$ events fitted from the fake data sets (which actually contained no $0\nu\beta\beta$ signal), shown in (b). The effect is larger in the case of enriched neodymium because the (relatively flat and hence less sensitive to energy resolution) ^{208}Tl and ^8B solar neutrino backgrounds have a fractionally larger effect in the fits with natural neodymium. Note that the error bars in (b) represent the uncertainty in the mean number of events fitted in 1000 fake data sets, not the uncertainty in the fitted number of $0\nu\beta\beta$ events in any one trial; the single trial uncertainty on the number of fitted $0\nu\beta\beta$ events is ~ 16 for natural neodymium and ~ 19 for enriched.

A much more dangerous energy systematic for SNO+ will likely be a small fraction of events which, because of detector asymmetries, do not follow the average Gaussian energy resolution distribution but instead fall as outliers²¹. To study this potential effect, exponential tails were added to the energy response function which was applied to the data distribution (but not to the response function applied to the PDFs), so that the probability, P , of reconstructing an event in the data with true energy E at energy E' became

$$P(E', E) = (1.0 - A) \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(E-E')^2}{\sigma^2}} + \frac{A \cdot B}{2} e^{-|B(E-E')|}, \quad (7.14)$$

where $\sigma = \sqrt{E/N_h}$ and N_h is the expected number of NHits/MeV (400 by default). The A and B parameters describe the fraction of events which fall into the exponential tails and the (inverse) decay length of the tails, respectively. To illustrate the effect, the response of the detector to a mono-energetic signal with and without exponential tails added to the response function is shown in Figure 7.12.

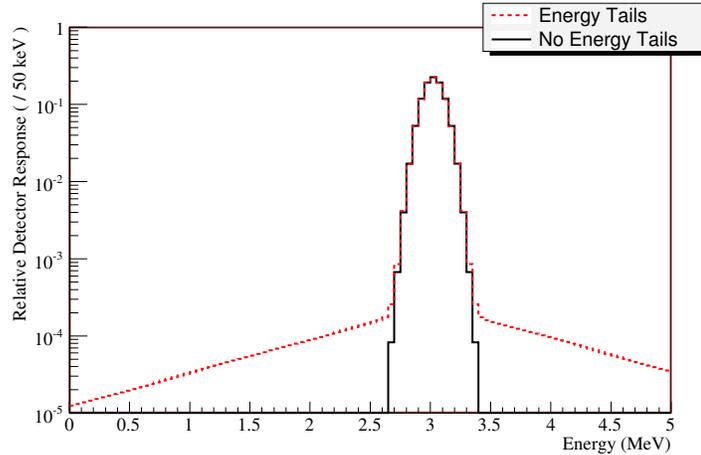


Figure 7.12: The simulated energy response of SNO+ to a mono-energetic signal at 3.0 MeV with and without the addition of exponential tails (with $A = 0.01$, $B = 1.0 \text{ MeV}^{-1}$).

From Figure 7.12, one might reasonably expect the effects of the A and B parameters to be correlated; that is, it seems reasonable to expect that SNO+ could tolerate a much larger fraction

²¹ As an example, consider an event having a reconstructed position directly adjacent to an AV hold-down rope, but being in reality 15cm (about the position resolution that can be expected in SNO+) away from it. If the energy estimation code attempts to correct for the assumed shadowing of the rope, it will significantly overestimate the energy of that event. Such “unusual” events can result in the addition of non-Gaussian “tails” in the energy resolution.

of backgrounds in these exponential tails (i.e. a larger A) when the tails decay quickly (i.e. when B is larger) compared to when they decay slowly. As shown in Figure 7.13, that is exactly the case; for natural neodymium, more than 1% of $2\nu\beta\beta$ events can be accommodated in exponential tails with a decay length of 0.15 MeV without negatively affecting the $0\nu\beta\beta$ limit, while less than $1 \times 10^{-4}\%$ of events can be allowed to fall in exponential tails with 1 MeV decays (like those shown in Figure 7.12).

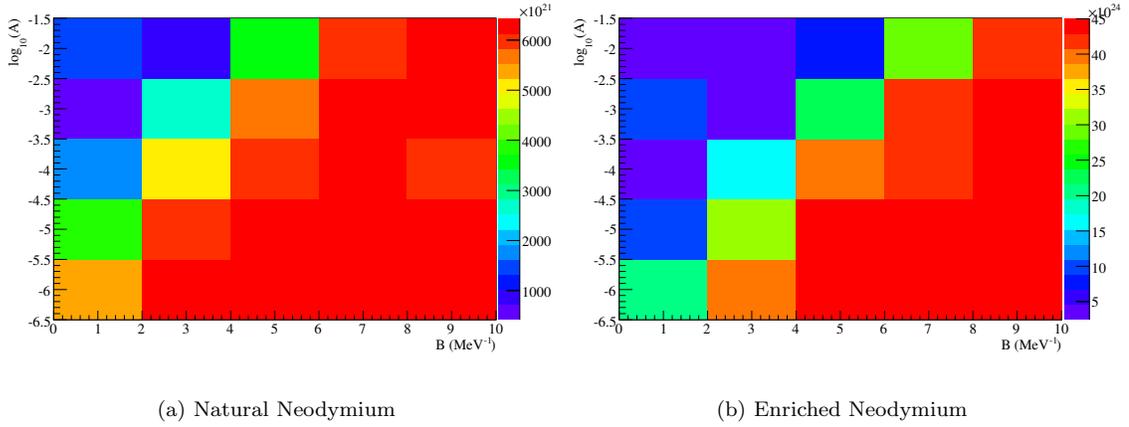


Figure 7.13: The effect of energy resolution tails on the SNO+ $0\nu\beta\beta$ limits. The colour axis shows the $T_{1/2}^{0\nu}$ sensitivity limit as a function of the A and B exponential tail parameters (defined in the text and Equation 7.14). The red areas in the lower right portion of both plots correspond to the baseline SNO+ sensitivities; the lower $T_{1/2}^{0\nu}$ limits in the top left corners show the reduction in $0\nu\beta\beta$ sensitivity. As expected, quickly decaying tails in the resolution (larger B) have a smaller effect on the $0\nu\beta\beta$ fits, and so more events (larger A) can be allowed to fall into such tails without affecting the $0\nu\beta\beta$ limits. As in Figure 7.12, natural neodymium is less sensitive to energy resolution effects, presumably because of the smaller effect of the $2\nu\beta\beta$ background in natural neodymium relative to the other sources of background.

As shown in Figure 7.14, the correlated behaviour of the A and B parameters can be simplified by projecting the effects of these parameters onto the “spillover” axis, where spillover is defined to be the number of $2\nu\beta\beta$ events which are smeared into the 3.0-3.7 MeV window (which contains the $0\nu\beta\beta$ signal) by the data energy resolution function in excess of the number expected by the PDF energy resolution function²². As can be seen in the Figure, the correlated statements about

²² That is, spillover is the difference between the true and assumed fraction of $2\nu\beta\beta$ events which fall into the 3.0-3.7 MeV window, multiplied by the total number of $2\nu\beta\beta$ events in the data set.

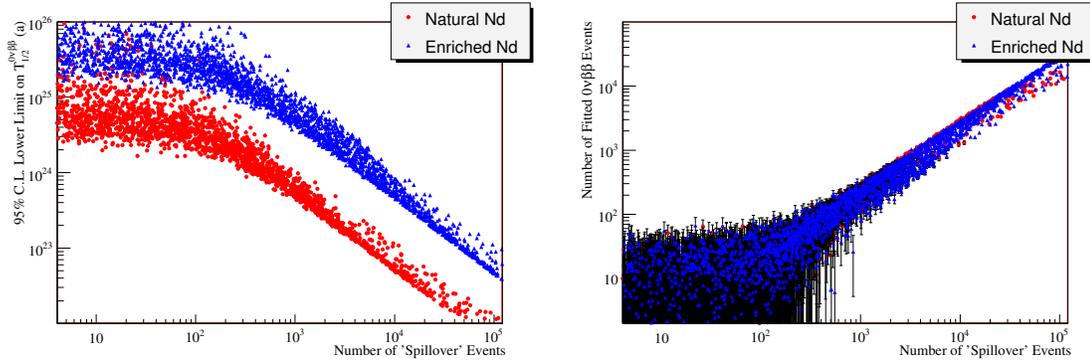
the allowable limits in the A and B parameters can be replaced by the much simpler statement that fewer than about 100 unexpected $2\nu\beta\beta$ events can be mis-reconstructed into the 3.0-3.7 MeV window.

Achieving this low number of mis-reconstructed events will likely be the greatest challenge to the successful operation of the SNO+ experiment. Thankfully, however, it is really only the upward tails²³ which are of concern to the experiment²⁴, and upward tails are much easier to suppress in event reconstruction than downward tails. This is because while there are a number of detector geometry effects which can artificially reduce the amount of light collected for (and hence the apparent energy of) a given event compared to the detector average, there are very few physical processes which will significantly increase the amount of collected light. Upward tails in reconstructed event energy, then, tend to be the result of overcorrection for the light loss mechanisms in estimating the energy of an event (as in the example given earlier of an event which has its position falsely reconstructed as being adjacent to an A.V. hold-down rope), rather than real non-statistical upward fluctuations in the amount of light collected. To mitigate this effect, then, in creating the SNO+ energy estimators every effort should be taken to ensure that the estimators return the lowest energy consistent with the observed event. This essentially sacrifices those events which actually do have additional light loss and accepts the attendant downward tails in the reconstructed energy as the price for minimizing the number of events whose energies are over estimated.

Another concern about spillover events is the potential, shown in Figure 7.14(b), that they can lead to false “detections” of $0\nu\beta\beta$ at very high confidence levels. As discussed earlier, when unexpected events are present in the $0\nu\beta\beta$ energy region, one of the only PDFs that the fit can use to attempt to accommodate these events is the $0\nu\beta\beta$ PDF; thus, in the presence of large numbers of unexpected backgrounds, the fitted number of $0\nu\beta\beta$ events can be far from zero. This effect is illustrated in Figure 7.15, which shows endpoint fits for simulated SNO+ experiments with different numbers of spillover events. As shown in Figure 7.14(c), the imperfect description of the unexpected

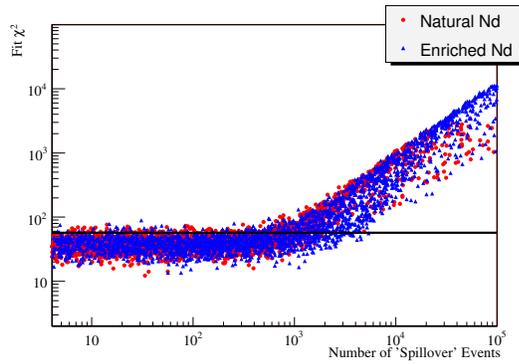
²³ Meaning those tails which extend the main distribution towards higher energy.

²⁴ This is simply due to the much higher statistics in the $2\nu\beta\beta$ signal below the $0\nu\beta\beta$ peak compared to the background statistics above the peak. Smearing a tiny fraction of the $2\nu\beta\beta$ signal upwards has a much larger effect in the $0\nu\beta\beta$ energy window than smearing a large fraction of the higher energy backgrounds downward.



(a)

(b)



(c)

Figure 7.14: The effect of energy resolution “tails” as projected on the spillover axis. Each point in these plots represents a single simulated SNO+ experiment. Different A and B parameters were used for each simulated experiment, with $\log_{10}(A)$ drawn randomly from $(-5, -1)$, and B drawn randomly from $(0.5, 8.5)$. It can be seen in (a) that the range of $T_{1/2}^{0\nu}$ limits in Figure 7.13 projects relatively neatly onto the spillover axis. The fitted $0\nu\beta\beta$ values, shown in (b), demonstrate that, as in the case of the average energy resolution systematic (Figure 7.11), the non-Gaussian tails principally affect the $T_{1/2}^{0\nu}$ limits by increasing the number of $0\nu\beta\beta$ events fitted out. The error bars plotted in (b) are the single fit uncertainties, so at high spillover the simulated SNO+ experiments would “falsely detect” $0\nu\beta\beta$ decay. As shown in (c), however, the attempt by the fits to use $0\nu\beta\beta$ events to compensate for the presence of the additional spillover events is not perfect, resulting in an increase in the fit χ^2 's. The horizontal line in (c) indicates the χ^2 level above which the fits (which all have 35 degrees of freedom) have a χ^2 probability of less than 1%.

$2\nu\beta\beta$ smearing provided by the $0\nu\beta\beta$ shape leads to an increase in the fit χ^2 , such that the most egregious false detections could be rejected by a cut on the fit χ^2 probability. There is a worrisome region, however, between a few hundred and a few thousand spillover events, where false detections are possible without unduly poor fit χ^2 's. This means that it would not be impossible for SNO+ to unwittingly claim a false detection of $0\nu\beta\beta$ as the result of an insufficient understanding of the detector energy response. This possibility does not, of course, affect the ability of the SNO+ experiment to interpret its results in terms of upper limits on the $0\nu\beta\beta$ process, but it is something that the collaboration will have to be extremely conscious of in the event that a $0\nu\beta\beta$ detection claim is ever considered.

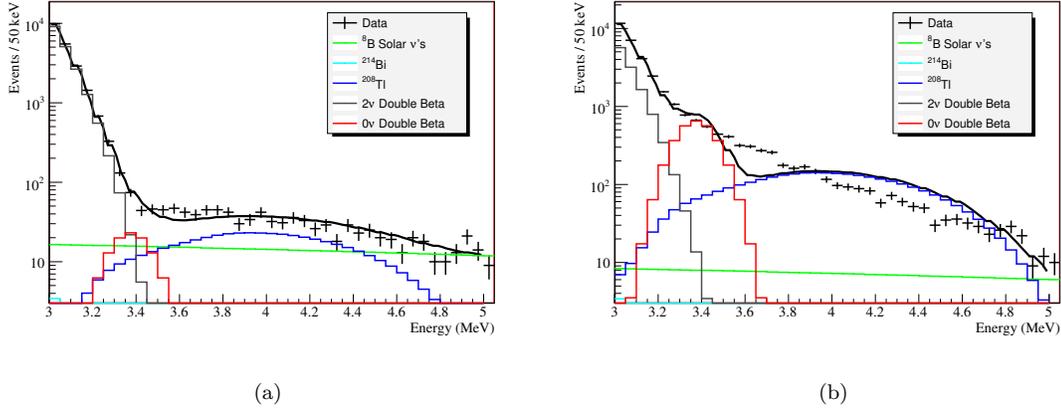


Figure 7.15: The $0\nu\beta\beta$ fit region for two simulated SNO+ experiments with natural neodymium which show false “detections” of $0\nu\beta\beta$ as a result of spillover events. In (a) a simulated data set with $A = 0.0005$ and $B = 3.0$ is shown. The fit to this data set, which contains 542 spillover events but no real $0\nu\beta\beta$ events, returns 107 ± 25 $0\nu\beta\beta$ events (for a 4.3σ “detection”) with a χ^2 of 50 in 35 degrees of freedom. The fit slightly under-predicts the number of events in the 3.5-3.8 MeV region, but on the whole there is little in this plot to raise concern about the validity of the fit. (b) shows a more extreme case, with $A=0.01$ and $B=3$ giving 10832 spillover events. In this case, 3027 ± 73 $0\nu\beta\beta$ events are fitted out, but the fit would certainly be rejected as a bad description of the data, as it has a χ^2 of 953 in 35 degrees of freedom.

Energy Resolution Systematics Conclusion

It remains to be seen whether the SNO+ energy response characterization can be carried out sufficiently precisely to provide interesting $0\nu\beta\beta$ limits or demonstrated robustly enough to make

credible $0\nu\beta\beta$ detection claims. The detailed simulations under development may help to provide some information about the potential for success, and will certainly provide a means by which event reconstruction algorithms can be developed, but it is unlikely that success can be gauged in a meaningful way until SNO+ is built and operated. At the current time, no fundamental factor has been identified which would prevent SNO+ from achieving the required understanding of its energy resolution. Nevertheless, the development of event reconstruction algorithms and plans for the energy calibration of the detector will have to be a major focus of the collaboration as the project moves forward.

7.6.5 Pileup Backgrounds and Excited State Decays

Apart from precisely understanding the energy response of the detector, in order to produce interesting and accurate limits on the $0\nu\beta\beta$ process it will also be necessary for SNO+ to precisely predict the underlying shape of the $2\nu\beta\beta$ spectrum. Two effects which could have an important impact on this underlying shape are $2\nu\beta\beta$ decays to excited states of ^{150}Sm and pileup between $2\nu\beta\beta$ events and low energy backgrounds.

Excited State Decays

The $2\nu\beta\beta$ decay of ^{150}Nd to the first 0^+ excited state of ^{150}Sm (0.740 MeV above the ground state) has been observed via detection of the de-excitation γ -rays [112], and determined to have a half-life of $(1.33_{-0.26}^{+0.45})\times 10^{20}$ a. Other excited state $2\nu\beta\beta$ decays of ^{150}Nd are possible but have not been observed; limits on several of these processes are also set in [112]. In SNO+, which will have virtually 100% γ -ray detection efficiency, the energy observed in such an excited state $2\nu\beta\beta$ decay will be the sum of the $2\nu\beta\beta$ decay energy and the energy of the de-excitation γ -rays²⁵. Thus, SNO+ would observe the energy spectrum of the lower endpoint excited state $2\nu\beta\beta$ spectrum shifted upwards by the γ energy such that the endpoint of the excited state decay coincides almost exactly with the endpoint

²⁵ Decays to the first 0^+ excited state typically release one 0.33 and one 0.41 MeV γ , with an $\sim 2\%$ branch to a single 0.74 MeV photon [112].

of the ground state decay²⁶. As lower endpoint $2\nu\beta\beta$ spectra approach their endpoints more steeply, excited state decays have the potential to affect the shape of the $2\nu\beta\beta$ spectrum at the endpoint, as shown in Figure 7.16. Based on the quick calculations used to produce Figure 7.16 (and after applying energy resolution smearing in the standard way), it is found that if all of the $2\nu\beta\beta$ decays were ground state decays there would be 2.0×10^4 (2.3×10^3) events/kT·a in the 3.0-3.7 MeV window for enriched (natural) neodymium, while the corresponding numbers would be 6.4×10^4 (7.3×10^3) events/kT·a if all of the $2\nu\beta\beta$ decays went to the 0.74 MeV excited state. The spillover studies described in Section 7.6.4 suggest that there must be fewer than ~ 100 unexpected events in the 3.0-3.7 MeV window. In order to meet this constraint, the branching ratio to the 0.74 MeV excited state decay must be known to about 0.2% absolute (or 3% of its value) for enriched neodymium, while for natural neodymium it need only be measured to 2% absolute or 30% of its value (similar to the precision of the existing measurement). Finally, shifting the excited state endpoint down by 1 keV in the calculation (to simulate very much exaggerated electron quenching) changes the number of $2\nu\beta\beta$ events in the spillover window to 6.5×10^4 events/kT·a for enriched neodymium. This is a relatively small change, and suggests that the additional energy “lost” through electron quenching for each additional electron produced need only be measured to ~ 100 eV (even less precision is required in the case of natural neodymium). These targets should be confirmed by more careful simulation of the actual effect of the change in the $2\nu\beta\beta$ endpoint shape on the $0\nu\beta\beta$ limits, but these “spillover” based estimates give an idea of the level of understanding of the excited state decays that will be required in SNO+.

Decays to other ^{150}Sm excited states are also expected, but have so far not been observed; in [112], lower limits between $2\times 10^{20}\text{a}$ and $8\times 10^{20}\text{a}$ are placed on the half-lives of decays to other 0^+ and 2^+ excited states. Calculations of the $2\nu\beta\beta$ spectra resulting from decays to these other excited states were not performed. However, the calculations described above for decays to the first 0^+ excited state suggest that if lower limits of about $4\times 10^{20}\text{a}$ ($4\times 10^{21}\text{a}$) could be set on the half-lives

²⁶ The endpoints do not coincide exactly because of electron quenching in the scintillator. This is the observed effect that two electrons of a given energy create slightly less total light in a scintillator than a single electron with twice the energy.

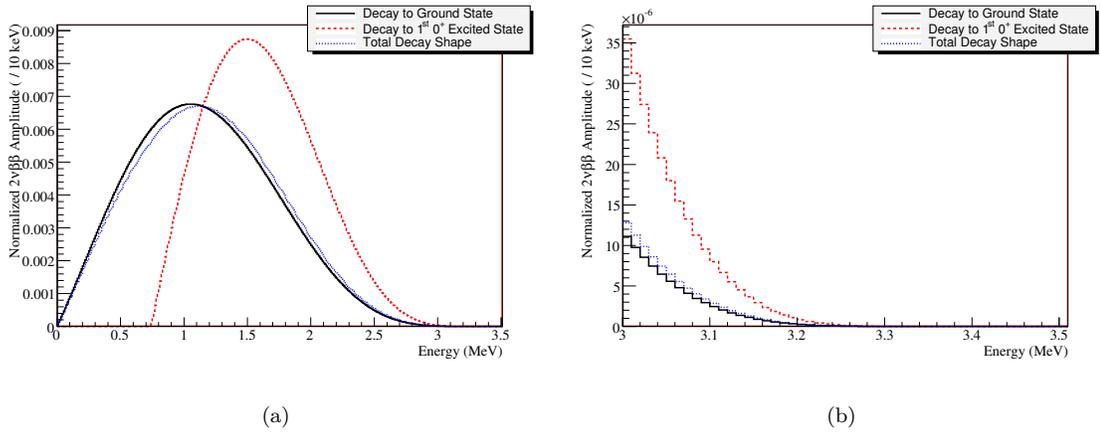


Figure 7.16: The $2\nu\beta\beta$ spectrum over the full energy range (a) and near the endpoint (b), showing the difference between the spectrum generated by decays to the ground state of ^{150}Sm and the spectrum generated by decays to the first 0^+ excited state (neglecting electron quenching). The “total” spectrum is a weighted sum of the ground state and excited state spectra, assuming a 7% branching ratio to the excited state.

of the decays to these states, their effect on the endpoint of the $2\nu\beta\beta$ spectrum would be negligible for a SNO+ double beta decay search with natural (enriched) neodymium. Again, these constraints are similar to the current experimental limits for natural neodymium, but more stringent branching ratio measurements would be required if SNO+ were to use enriched neodymium.

Pileup Backgrounds

Another effect which might change the shape of the $2\nu\beta\beta$ spectrum in the endpoint region is pileup. In SNO+, an event window will be approximately 400 ns wide, and all PMT hits which occur during this window are recorded as being from a single event. Thus, if two events should occur within about 400 ns of one another, they will appear in the data stream as a single event. Thus, if two events occur 200 ns apart and the first event is sufficiently energetic to trigger the detector, 400 ns of the first event and about the first 200 ns of the second event will be recorded²⁷. If the first event does not trigger the detector but the second one does, 400 ns of the second event will be recorded along

²⁷ The shape of the roll-off at the end of the 400 ns event window (or even the length of the window) is not particularly well understood. For the purposes of discussion the window is assumed here to be square and exactly 400 ns; as pileup will ultimately be treated phenomenologically, the exact shape of the trigger window is unimportant.

with the “tail” of the first event²⁸.

Pileup is particularly worrisome in the context of pileup between a $2\nu\beta\beta$ event and another background event. This has the potential to artificially increase the energy of a (hopefully) small fraction of the $2\nu\beta\beta$ events, distorting the $2\nu\beta\beta$ spectrum and adding a “bump” in the $0\nu\beta\beta$ energy region, as illustrated in Figure 7.17 for two potential real pileup scenarios. In order to get an initial estimate of the severity of the pileup problem, quick simulations were performed in which it was assumed that the full energy of events that occurred within a 400 ns window (i.e. 200 ns before and 200 ns after) of a $2\nu\beta\beta$ event was summed with the energy of the $2\nu\beta\beta$ event. This rather extreme assumption can be expected to overestimate the average effect of each pileup event but underestimate the total number of pileup events. However, these assumptions simplify the calculations considerably, as the pileup spectrum is then simply the convolution of the $2\nu\beta\beta$ spectrum and the spectrum of the other background events. Also, under this assumption, the probability $P(n)$ of the pileup of n of the other background events with a $2\nu\beta\beta$ event is simply the Poisson probability of n other background events occurring in a given 400 ns window,

$$P(n) = \frac{e^{-m} m^n}{n!}, \quad (7.15)$$

where m is the mean number of other backgrounds in a 400 ns window.

Any relatively high rate background will form significant numbers of $2\nu\beta\beta$ pileup events. Table 7.4 shows the fraction of $2\nu\beta\beta$ events expected to have associated pileup backgrounds at different background rates. The fraction of these pileup events which end up in the $0\nu\beta\beta$ energy region, of course, depends on the energy of the secondary events; the higher the energy of the secondary event, the lower the energy of the $2\nu\beta\beta$ events that pileup would push into the $0\nu\beta\beta$ window. As the rate of $2\nu\beta\beta$ events increases steeply as energy decreases, the tolerable rate of potential “pileup partners” decreases quickly as their energy increases.

The seriousness of a pileup background can be estimated using the “spillover” metric described in Section 7.6.4 (i.e. by determining the number of unexpected events that appear in the 3.0-3.7

²⁸ Although the majority of the light from an event is expected to be collected in the first 200 ns after the event, scattering and re-emission processes will lead to some photons “trickling in” up to 500 ns after the event.

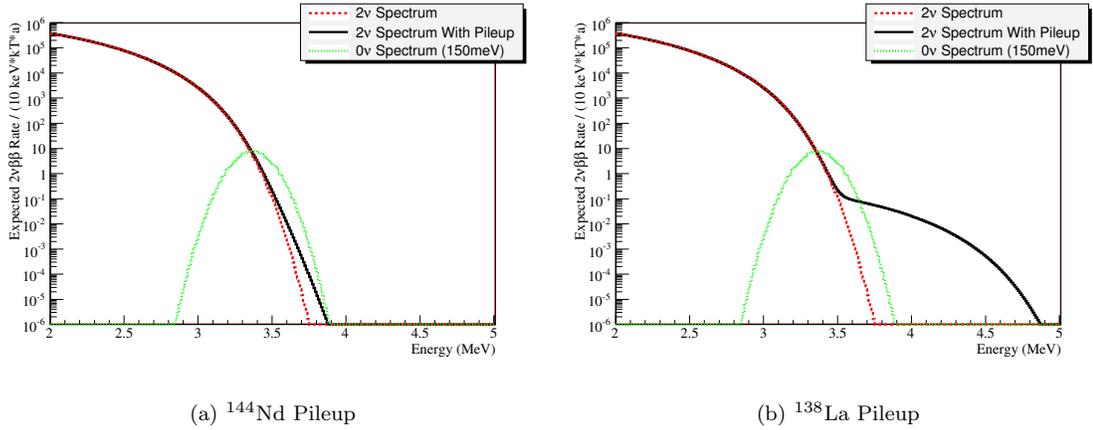


Figure 7.17: The distortion of the $2\nu\beta\beta$ spectrum for enriched neodymium from two potential pileup sources: the 1.9 MeV alpha decay of ^{144}Nd (with an assumed alpha quenching factor of 10) and the 1.4 MeV electron capture of ^{138}La . The ^{144}Nd rate is assumed to be 9500 Hz (about what would be expected if the ^{144}Nd abundance were not reduced in the enrichment process) and the ^{138}La rate is taken to be 0.9 Hz (which is the rate expected at the upper allowed limit in Table 7.5). The $0\nu\beta\beta$ spectrum at $\langle m_\nu \rangle = 150\text{meV}$ is shown for comparison.

MeV window as the result of pileup after applying energy resolution). The pileup of two $2\nu\beta\beta$ events with themselves, for example, will add about 45 events/kT·a to the spillover window for enriched neodymium, and less than one event/kT·a for natural neodymium. These pileup rates are tolerable. At the levels targeted for the SNO+ solar phase [66], the internal ^{238}U and ^{232}Th chain and ^{40}K backgrounds will be negligible from a pileup perspective. External backgrounds at their currently predicted rates [124] will contribute significant, but likely tolerable, pileup (the external backgrounds will likely cause more pileup events than the double $2\nu\beta\beta$ pileup, but less than the ^{144}Nd discussed below). The ^{14}C intrinsic to the scintillator is a potential problem, depending on its level. In atmospheric carbon, ^{14}C occurs at approximately the 1 ppt level; this would be disastrous for SNO+ as each $2\nu\beta\beta$ event would have on average more than 60 ^{14}C events (which are beta decays with a 0.156 MeV endpoint) in coincidence with it. Thankfully, however, in scintillator produced from fossil fuels (which have been underground for a very long time) the ^{14}C level is significantly reduced; the BOREXINO scintillator, for example, is observed to have a ^{14}C concentration of $1.94 \times 10^{-18}\text{g/g}$ [137]. At this level, the expected rate of pileup events appearing in the spillover window is should be

| Background Rate (Hz) | Pileup Fraction (%) |
|----------------------|----------------------|
| 1 | 4.0×10^{-5} |
| 10 | 4.0×10^{-4} |
| 100 | 4.0×10^{-3} |
| 1000 | 0.040 |
| 10000 | 0.40 |
| 100000 | 3.8 |

Table 7.4: The fraction of events in SNO+ which will have a background event with a ± 200 ns window. The pileup fraction depends approximately linearly on the event rate at low rate, but becomes non-linear quickly once the rate increases. Note that at higher rates, double and triple pileup events can also become significant.

approximately $2 / \text{kT} \cdot \text{a}$ for enriched neodymium and $< 1 / \text{kT} \cdot \text{a}$ for natural neodymium²⁹. Therefore, ^{14}C pileup events should not be a problem in SNO+ as long as the LAB ^{14}C content does not exceed about 1×10^{-17} g/g.

The addition of the neodymium to the scintillator has the potential to introduce additional low energy backgrounds. ^{144}Nd , for example, which has 21.8% natural abundance, decays with a relatively short 2.3×10^{15} a half-life to produce 1.9 MeV alpha particles (which, due to alpha quenching in the scintillator produce approximately the same amount of light as 0.19 MeV electrons). The short half-life makes the rate of ^{144}Nd pileup with $2\nu\beta\beta$ events relatively high, with 58 “spillover” pileup events/kT·a produced with natural neodymium and 514 events/kT·a produced with enriched neodymium (assuming that the ^{144}Nd concentration is not reduced by the enrichment, which it likely would be). Neodymium, like all rare earth elements, can also contain significant contamination from other rare earth elements which can be difficult to remove chemically. Some potential contaminants which have the potential to contribute significantly to the pileup background have been identified [138], and their pileup potential has been evaluated using the spillover metric [139]. Based on the spillover analysis³⁰, limits were set on the concentrations of these different elements allowable in the

²⁹ In calculating the ^{14}C pileup spectrum, as with the other β -decay pileups simulated, the spectrum is simulated by scaling the ^{210}Bi spectral shape to have the appropriate endpoint. This is again a fairly crude approximation, but it should be sufficient to give a first estimate of these pileup effects.

³⁰ The analysis in [139] has been modified slightly here to bring it into agreement with the ^{150}Nd half-lives and loading level assumed in this thesis and to incorporate the technique described earlier of scaling the ^{210}Bi spectrum to represent all beta spectra.

SNO+ neodymium in order that pileup problems are avoided. These concentrations are listed in Table 7.5. As chemically separating and removing these elements is challenging, assays of neodymium from different suppliers are being carried out [140, 141, 142] in the hope of finding neodymium with reduced levels of these contaminants. As seen in Table 7.6, different samples differ markedly in their levels of contamination by other rare earth elements. Although the current measurements are only precise enough to constrain their concentrations at 2 and 100 times the required levels, respectively, samples were identified in which no Lu or La activity was seen. Another sample was determined to have ^{235}U chain activity at less than 20 times the target. This suggests that it may yet be possible to identify a source of neodymium satisfying all of the purity targets given in Table 7.5.

| Background | Limit |
|----------------------------------|--|
| ^{14}C | $1 \times 10^{-17} \text{g/g}$ in LAB |
| $^{85}\text{Kr}/^{39}\text{Ar}$ | 0.05Bq/m^3 (total) |
| La | $5 \times 10^{-7} \text{g/g}$ in Nd |
| Lu | $1 \times 10^{-7} \text{g/g}$ in Nd |
| Sm | $5 \times 10^{-7} \text{g/g}$ in Nd |
| Gd | $2 \times 10^{-2} \text{g/g}$ in Nd |
| $^{235}\text{U}/^{227}\text{Ac}$ | $5 \times 10^{-11} / 1 \times 10^{-18} \text{g/g}$ in Nd |

Table 7.5: The contamination levels of different isotopes/elements at which pileup would be a negligible effect in SNO+ (at these levels, <10 pileup events are expected to be produced in the spillover window in one $\text{kT}\cdot\text{a}$ with enriched neodymium). Where the isotope is not specified, natural isotopic abundances are assumed. It should be noted that these are rough limits based on the approximate modelling described in the text and [139], and that more complete simulations of the pileup effect may shift these limits slightly.

Even if it proves impossible to achieve the neodymium purity targets defined in Table 7.5 exactly, if the contaminant concentrations can be reduced to within a factor of 100 or so of the targets (as certainly appears to be possible), SNO+ should still be a viable experiment. In this situation, the pileup background resulting from these backgrounds, along with the irreducible ^{144}Nd , ^{14}C , and double $2\nu\beta\beta$ pileup, would have to be dealt with in the analysis. Some pileup events can likely be rejected by data analysis cuts based on the abnormal distribution of arrival times of the detected photons. This approach is currently being investigated, and although simulations to evaluate its

| Sample | Background | | |
|--------|-------------------------------------|-------------------------------------|---|
| | La ($\times 10^{-6}$ g La/g Nd) | Lu ($\times 10^{-7}$ g Lu/g Nd) | ^{227}Ac ($\times 10^{-8}$ g ^{235}U /g Nd) |
| 1 | 195 ± 19 | 13 ± 1 | 0.86 ± 0.10 |
| 2 | 18.1 ± 2.2 | not measured | 94.8 ± 0.52 |
| 3 | 0 ± 19 | 0 ± 2.3 | 2.4 ± 0.18 |
| 4 | 0 ± 25 | 0 ± 4.6 | 2.9 ± 0.18 |
| 5 | 175 ± 42 | 0 ± 17 | 52 ± 2.9 |
| 6 | 240 ± 40 | 0 ± 15 | 52 ± 2.9 |

Table 7.6: The results from germanium detector-based assays of neodymium salts. Sample 1 was NdCl_3 , with the other samples being Nd_2O_3 of different nominal purities and from different suppliers. The samples and measurements are described in detail in [140] (Sample 1), [141] (Sample 2), and [142] (Samples 3-6). Where the isotope is not specified, natural isotopic abundances are assumed. Samarium and gadolinium were not included in this study as their decays do not result in the production of detectable γ -rays.

potential are in the very early stages, it is the opinion of the author that this approach could reduce the pileup background by up to, but not much more than, a factor of 10. The remaining irreducible pileup backgrounds will have to be incorporated into the analysis by including them in the PDFs. The actual pileup background in SNO+ can be sampled by randomly triggering the detector throughout data running (similar to the random “PulseGT” triggers which were used in SNO), and these observed pileup backgrounds can then be added to the Monte Carlo events used to create the SNO+ signal extraction PDFs. Treated in this way, pileup background would affect the $0\nu\beta\beta$ sensitivity statistically rather than systematically (i.e. in a manner more similar to a well understood degradation in detector energy resolution like that shown in Figure 7.9 than to the energy systematic uncertainties that were described by the spillover metric), and hence would have a much reduced impact on the experiment.

One final complication of pileup backgrounds in SNO+ will be their effect on calibration activities. In order to achieve the very precise understanding of the detector energy response necessary to set interesting $0\nu\beta\beta$ limits, energy calibrations with statistics larger than the 10^7 - 10^8 $2\nu\beta\beta$ events expected will be required. In order to achieve this number of events in a reasonable time (and also to prevent the 1-10 Hz of $2\nu\beta\beta$ decays from being a significant background to the calibrations),

calibration data rates approaching or exceeding 1000 Hz will be required. At this rate, $\sim 0.05\%$ of the calibration events (or between 5,000-50,000 events) will pile up with another calibration event, producing a pileup distortion in excess of what is expected in the data on the high energy side of the calibration spectra. Thus, interpreting the calibration data will rely on modelling, specifically on modelling the pileup during the high rate calibration periods. This in turn will require random triggering of the detector at a significant rate during calibration running.

In summary, the pileup of events in SNO+ has the potential to provide a significant challenge to the double beta decay portion of the experiment. Although it is reasonable to expect that, provided the radioactive impurities in the neodymium can be kept to a reasonable level, the pileup backgrounds can be accommodated in the analysis with minimal impact on the $0\nu\beta\beta$ limit, pileup backgrounds will certainly have to be carefully considered by the collaboration as the experiment moves forward.

7.7 Summary of the SNO+ Neutrinoless Double Beta Decay Search

The search for neutrinoless double beta decay through the loading of the ^{150}Nd double beta decay isotope into the SNO+ liquid scintillator presents an exciting opportunity for the double beta decay community. The SNO+ active volume is sufficiently large that, even at relatively modest 0.1% loading of the scintillator, large masses of isotope are achievable. Using natural neodymium, SNO+ should be able to confirm or refute the controversial claim of detection of $0\nu\beta\beta$ with $\langle m_\nu \rangle \gtrsim 270\text{meV}$, in approximately 1 kT·a of running (or in approximately three real years) for all currently calculated values of the ^{150}Nd nuclear matrix element. If neodymium enriched to 50% in ^{150}Nd could be obtained, SNO+ would conclusively test the existing claim in the matter of a few months, and, depending on the ^{150}Nd nuclear matrix element, proceed to probe the $\langle m_\nu \rangle = 10\text{-}40\text{ meV}$ region predicted by the observed neutrino mass splittings in the case of the inverted neutrino mass hierarchy. This makes SNO+ competitive, both in sensitivity and in time, with other current neutrinoless double beta decay experiments.

Ensuring the success of the SNO+ neutrinoless double beta decay experiment will not be simple, as an extremely precise understanding of the background energy spectra, especially the ^{150}Nd $2\nu\beta\beta$ energy spectrum, will be required in order to make interesting and credible $0\nu\beta\beta$ measurements. This will require the energy resolution of the SNO+ detector to be characterized and monitored very carefully, and the $2\nu\beta\beta$ spectral shape underlying the energy resolution to be well understood. It will also be necessary to identify sources of, or to produce, neodymium with very low levels of rare earth element contamination. These challenges are significant but by no means insurmountable, and with a good deal of careful and well-planned effort, the SNO+ collaboration should be able to make a very significant contribution to the global search for neutrinoless double beta decay.

7.8 SNO+ Conclusion

The SNO+ experiment represents an exciting opportunity for SNOLAB, for Canada, and for the particle astrophysics community in general. With relative ease and for relatively little expense it is possible to construct, from the remaining infrastructure of the SNO experiment, a new and extremely capable liquid scintillator neutrino detector. Using this detector, the SNO+ collaboration will be able to study the low energy solar neutrinos, reactor anti-neutrinos, geo-neutrinos, and, with luck, supernova neutrinos. These studies will provide important information to several different physics disciplines including neutrino physics, solar physics, and geo-physics. In addition, by loading the liquid scintillator with double beta decay isotope, the SNO+ collaboration will be able to carry out a very competitive, and possibly world-leading, search for neutrinoless double beta decay. Thus, the same experiment which helped to demonstrate that neutrinos have mass will, in a new form, continue to lead the global effort to better understand these particles. I wish the collaboration and the experiment all the best.

References

- [1] H. Bethe. Energy production in stars. *Phys. Rev.*, 55:103(L), 1939.
- [2] J.N. Bahcall. Solar Models: An Historical Overview. *Nucl. Phys. B - Proc. Suppl.*, 118:77–86, 2003.
- [3] J.N. Bahcall, A.M. Serenelli, and S. Basu. New Solar Opacities, Abundances, Helioseismology, and Neutrino Fluxes. *Astrophys. J.*, 621:L85–L88, 2005.
- [4] B. T. Cleveland et al. Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector. *Astrophys. J.*, 496:505–526, 1998.
- [5] J. N. Abdurashitov et al. Measurement of the Solar Neutrino Capture Rate by the Russian-American Gallium Solar Neutrino Experiment During One Half of the 22-year Cycle of Solar Activity. *J. Exp. Theor. Phys.*, 95:181–193, 2002.
- [6] M. Altmann et al. Complete Results for Five Years of GNO Solar Neutrino Observations. *Phys. Lett.*, B616:174–190, 2005.
- [7] W. Hampel et al. GALLEX Solar Neutrino Observations: Results for GALLEX IV. *Phys. Lett.*, B447:127–133, 1999.
- [8] Y. Fukuda et al. Solar Neutrino Data Covering Solar Cycle 22. *Phys. Rev. Lett.*, 77:1683–1686, 1996.
- [9] J. Hosaka et al. Solar Neutrino Measurements in Super-Kamiokande-I. *Phys. Rev.*, D73:112001, 2006.

- [10] J.N. Bahcall and R. Davis. An Account of the Development of the Solar Neutrino Problem. In D.D. Barnes, C.A. Clayton and D. Schramm, editors, *Essays in Nuclear Astrophysics*, pages 243–285. Cambridge University Press, 1982.
- [11] J. N. Bahcall, N. Cabibbo, and A. Yahil. Are Neutrinos Stable Particles? *Phys. Rev. Lett.*, 28:316, 1972.
- [12] A. Cisneros. Effect of Neutrino Magnetic Moment on Solar Neutrino Observations. *Astrophys. Space Sci.*, 10:87–92, 1971.
- [13] V. Gribov and B. Pontecorvo. Neutrino Astronomy and Lepton Charge. *Phys. Lett.*, 28B:493, 1969.
- [14] L. Wolfenstein. Neutrino Oscillations in Matter. *Phys. Rev.*, D17:2369–2374, 1978.
- [15] S.P. Mikheyev and A.Yu. Smirnov. Resonant Amplification of ν Oscillations in Matter and Solar-Neutrino Spectroscopy. *Nuovo Cimento*, 9C:17, 1986.
- [16] H. H. Chen. Direct Approach to Resolve the Solar-Neutrino Problem. *Phys. Rev. Lett.*, 55:1534–1536, 1985.
- [17] J. Boger et al. The Sudbury Neutrino Observatory. *Nucl. Inst. Meth.*, A449:172, 2000.
- [18] Q.R. Ahmed et al. Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301, 2002.
- [19] B. Aharmim et al. Measurement of the ν_e and Total B-8 Solar Neutrino Fluxes with the Sudbury Neutrino Observatory Phase I Data Set. *Phys. Rev.*, C75:045502, 2007.
- [20] B.R. Aharmin et al. Electron Energy Spectra, Fluxes, and Day-Night Asymmetries of ^8B Solar Neutrinos from the 391-Day Salt Phase SNO Data Set. *Phys. Rev. C.*, 72:055502, 2005.
- [21] B.H. Aharmin et al. Independent Measurement of the Total Active ^8B Solar Neutrino Flux using an Array of ^3He Proportional Counters at the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 101:111301, 2008.

- [22] C. Arpesella et al. Direct Measurement of the Be-7 Solar Neutrino Flux with 192 Days of Borexino Data. *Phys. Rev. Lett.*, 101:091302, 2008.
- [23] S. Abe et al. Precision Measurement of Neutrino Oscillation Parameters with KamLAND. *Phys. Rev. Lett.*, 100:221803, 2008.
- [24] T. Araki et al. Measurement of Neutrino Oscillation with KamLAND: Evidence of Spectral Distortion. *Phys. Rev. Lett.*, 94:081801, 2005.
- [25] J.N. Bahcall. *Neutrino Astrophysics*. Cambridge University Press, 1989.
- [26] R. Martin. Numerical Methods and Approximations for the Calculation of Solar Neutrino Fluxes in Three Flavours Focused on Results for the Sudbury Neutrino Observatory. Master's thesis, Queen's University, 2006. UMI-MR-15302.
- [27] P. Amsler et al. (Particle Data Group). Leptons. PL **B667**, 1. URL:<http://pdg.lbl.gov>, 2008.
- [28] J.F. Amsbaugh et al. An Array of Low-Background ^3He Proportional Counters for the Sudbury Neutrino Observatory. *Nucl. Inst. Meth.*, A554:1054–1080, 2005.
- [29] H.S. Wan Chan Tseung. Simulation of NCD Alpha Background Energy PDFs. SNO internal document MANN-78S64R, November 2007.
- [30] S. McGee. Skimmed NNNA Events. SNO internal document MANN-6ZKS8A, March 2007.
- [31] B. Monreal. Stringy NNNA Hunt on Many Pulse-Shape Parameters. SNO internal document MANN-767T5F, August 2007.
- [32] B. Monreal. Search for Bursty or Non-Poissonian NNAs. SNO internal document MANN-76LKR3, August 2007.
- [33] B. Monreal. The Prior Probability of Low-Badness NNAs. SNO internal document MANN-77Q4DM, October 2007.

- [34] A. Wright. Update to the String-by-String Rate Analysis. SNO internal document MANN-7DJQ2S, April 2008.
- [35] N. Obath et al. NCD MC Readiness Report for the First NCD Paper, v3.1. SNO internal document MANN-79662C, April 2008.
- [36] H.S. Wan Chan Tseung. A Simulation of Space Charge Effects in the NCDs. SNO internal document, April 2004.
- [37] H.S. Wan Chan Tseung. The Motion of Electrons in NCD Gas. SNO internal document MANN-6VYDCG, November 2007.
- [38] B. Jamieson. SNO NCD Phase Signal Extraction on Unblinded Data with Integration over Systematic Nuisance Parameters by Markov-Chain Monte-Carlo - Version 1.1. SNO internal document MANN-7NCU9S, May 2008.
- [39] J. Goon et al. Signal Extraction of the Solar Neutrino Neutral-Current Flux with the Sudbury Neutrino Observatory Neutral Current Detectors. *J. Phys.: Conf. Ser.*, 136(4):042010 (1pp), 2008.
- [40] A. Wright. Extracting the NCD Neutron Flux From an Unknown Background. SNO internal document MANN-73HQNM, May 2007.
- [41] A. Wright. Update on Polynomial and Overlap Fits Under Different PSA Cuts. SNO internal document MANN-7AWRDE, January 2008.
- [42] Likelihood Ratio Tests. *NIST/SEMATECH e-Handbook of statistical Methods*, Section 8.2.3.3. Available at www.itl.nist.gov/div989/handbook/index.html.
- [43] Y.M. Yao et al. Statistics. *J. Phys. G*, 33:1, 2006.
- [44] B. VanDevender. NCLI Parameters, Uncertainties, and Covariance. SNO internal document, April 2007.

- [45] E. Guillian. MUX Threshold Efficiency Determination for the NCD. SNO internal document MANN-7D7UJ3, March 2008.
- [46] A. Wright. Systematic Uncertainties Associated with the Use of the ^{24}Na Neutron PDF. SNO internal document MANN-7DHHKC, April 2008.
- [47] E. Guillian. The Average MUX Live Fraction for the Solar Neutrino Data. SNO internal document MANN-7DLJVS, April 2008.
- [48] A. Wright. A Collection of Scope Live Fraction Documents. SNO internal document MANN-7AQL4S, January 2008.
- [49] E. Guillian, A. Wright, and K. Boudjemline. The NCD Scope Live Fraction. SNO internal document MANN-7BW3Z7, May 2008.
- [50] E. Guillian. Some Correction Factors of Converting Extracted Neutron Numbers to the Solar Neutrino Flux. SNO internal document MANN-7DSE7Y, April 2008.
- [51] H.M. O’Keeffe. *Low Energy Background in the NCD Phase of the Sudbury Neutrino Observatory*. PhD thesis, University of Oxford, 2008.
- [52] B. Cai, C. Kraus, and A. Poon. External Neutrons for the NCD Phase. SNO internal document MANN-7E8TX9, May 2008.
- [53] C. Kraus. A.V. Alpha Counting - 2007. SNO internal document MANN-7DDNCP, April 2008.
- [54] S. McGee et al. Limit on Neutron Production from Activity on NCD Outer Surface. SNO internal document MANN-7A2MHW, January 2008.
- [55] J.A. Formaggio. Other Background Sources for the NCD Phase of the SNO Experiment. SNO internal document MANN-7AMQGM, February 2008.
- [56] J. Detwiler et al. Livetime for the Full NCD Data Set. SNO internal document MANN-7APU7D, January 2008.

- [57] J. Loach. *Measurement of the Flux of 8B Solar Neutrinos at the Sudbury Neutrino Observatory*. PhD thesis, University of Oxford, 2008.
- [58] E. Guillian. The Determination of the Neutral Current NCD Neutron Capture Efficiency with ^{24}Na Data. SNO internal document MANN-7E7LF8, April 2008.
- [59] W. T. Winter et al. The B-8 Neutrino Spectrum. *Phys. Rev.*, C73:025503, 2006.
- [60] S. Nakamura et al. Neutrino Deuteron Reactions at Solar Neutrino Energies. *Nucl. Phys.*, A707:561–576, 2002.
- [61] A. Kurylov, M.J. Ramsey-Musolf, and P. Vogel. Radiative Corrections in Neutrino Deuterium Disintegration. *Phys. Rev.*, C65:055501, 2002.
- [62] C. Kraus. Corrections for the FNP. SNO internal document, May 2008.
- [63] A. Friedland, C. Lunardini, and Carlos Peña-Garay. Solar Neutrinos as Probes of Neutrino - Matter Interactions. *Phys. Lett.*, B594:347, 2004.
- [64] V. Barger, P. Huber, and D. Marfatia. Solar Mass-Varying Neutrino Oscillations. *Phys. Rev. Lett.*, 95:211802, 2005.
- [65] J.N. Bahcall and C. Peña-Garay. A Road Map to Solar Neutrino Fluxes, Neutrino Oscillation Parameters, and Tests for New Physics. *JHEP11*, 004, 2003.
- [66] S. Quirk. Purification of Liquid Scintillator and Monte Carlo Simulations of Relevant Internal Backgrounds in SNO+. Master's thesis, Queen's University, 2007.
- [67] E. O'Sullivan et al. Scintillation Timing Results from Bucket Data and Bench Top Measurements. SNO+ internal presentation doc-308-v1, April 2009.
- [68] T. Hagner et al. Muon Induced Production of Radioactive Isotopes in Scintillation Detectors. *Astropart. Phys.*, 14:33–47, 2000.

- [69] G. Bellini et al. Measurement of the Solar ^8B Neutrino Flux with 246 Live Days of Borexino and Observation of the MSW Vacuum-Matter Transition. 2008.
- [70] M. Balata et al. CNO and *pep* Neutrino Spectroscopy in Borexino: Measurement of the Deep-Underground Production of Cosmogenic ^{11}C in an Organic Liquid Scintillator. *Phys. Rev.*, C74:045805, 2006.
- [71] J. Bahcall, M. Gonzalez-Garcia, and C. Peña-Garay. Does the Sun Shine by *pp* or CNO Fusion Reactions? *Phys. Rev. Lett.*, 90:131301, 2003.
- [72] C. Peña-Garay and A. Serenelli. Solar Neutrinos and the Solar Composition Problem. arXiv:0811.2424, 2008.
- [73] M. Chen. Geo-neutrinos in SNO+. *Earth, Moon, and Planets*, 99(1-4):221–228, 2006.
- [74] E. Guillian. The Sensitivity of SNO+ to Δm_{12}^2 Using Reactor Anti-neutrino Data. arxiv:0809.1649, 2008.
- [75] G. Fiorentini et al. Geo-Neutrinos: A Short Review. *Nucl. Phys. Proc. Suppl.*, 143:53–59, 2005.
- [76] T. Araki et al. Experimental Investigation of Geologically Produced Antineutrinos with KamLAND. *Nature*, 436:499–503, 2005.
- [77] C.G. Rothschild, M.C. Chen, and F.P. Calaprice. Antineutrino Geophysics with Liquid Scintillator Detectors. *Geophys. Res. Lett.*, 25:103, 1998.
- [78] J. Benziger et al. A Scintillator Purification System for the Borexino Solar Neutrino Detector. *Nucl.Inst. Meth. A*, 587:277, 2008.
- [79] A. Burrows, D. Klein, and R. Gandhi. The Future of Supernova Neutrino Detection. *Phys. Rev.*, D45:3361–3385, 1992.
- [80] C.J. Virtue. SNO and Supernovae. *Nucl. Phys. B. (Proc. Suppl.)*, 100:24–29, 2001.

- [81] C. Lunardini and A. Yu. Smirnov. Neutrinos from SN1987A: Flavor Conversion and Interpretation of Results. *Astropart. Phys.*, 21:703–720, 2004.
- [82] P.J. Kernan and L.M. Krauss. Yet Another Paper on Sn1987a: Large Angle Oscillations, and the Electron Neutrino Mass. *Nucl. Phys.*, B437:243–256, 1995.
- [83] C. Hanhart et al. Extra Dimensions, SN1987a, and Nucleon Nucleon Scattering Data. *Nucl. Phys.*, B595:335–359, 2001.
- [84] L. Cadonati, F. P. Calaprice, and M. C. Chen. Supernova Neutrino Detection in BOREXINO. *Astropart. Phys.*, 16:361–372, 2002.
- [85] J. Busenitz et al. Proposal for U.S. Participation in KamLAND. 2003. Available online http://bama.ua.edu/~andreas/ps_files/kamland_prop.pdf.
- [86] J.F. Beacom, W.M. Farr, and P. Vogel. Detection of Supernova Neutrinos by Neutrino Proton Elastic Scattering. *Phys. Rev.*, D66:033001, 2002.
- [87] M. C. Gonzalez-Garcia and M. Maltoni. Phenomenology with Massive Neutrinos. *Phys. Rept.*, 460:1–129, 2008.
- [88] H. Klapdor-Kleingrothaus and I.V. Krivosheina. The Evidence for the Observation of $0\nu\beta\beta$ Decay: The Identification of $0\nu\beta\beta$ Events from the Full Spectrum. *Mod. Phys. Lett. A*, 21(20):1547–1566, 2006.
- [89] H. V. Klapdor-Kleingrothaus et al. Evidence for Neutrinoless Double Beta Decay. *Mod. Phys. Lett.*, A16:2409–2420, 2001.
- [90] C.E. Aalseth et al. Comment on 'Evidence for Neutrinoless Double Beta Decay'. *Mod. Phys. Lett.*, A17:1475–1478, 2002.
- [91] P.S. Vincet et al. Report of the Canadian Subatomic Physics Five-Year Planning Committee. Available www.ap.smu.ca/~butler/lrp/, June 2001.

- [92] A. Piepke. Final Results from the Palo Verde Neutrino Oscillation Experiment. *Prog. Part. Nucl. Phys.*, 48:113–121, 2002.
- [93] V. Novikov. Scintillator Development for SNO+. SNO+ internal document doc-319-v1., September 2006.
- [94] J.D. Stachiw. *Handbook of Acrylics for Submersibles, Hyperbaric Chambers, and Aquaria*. Best Publishing Co., 2003.
- [95] Designation D 543. Standard Practices for Evaluating the Resistance of Plastics to Chemical Reagents. In *Annual Book of ASTM Standards 2005*, page 28. ASTM International, 2005.
- [96] Designation D 638. Standard Test Method for Tensile Properties of Plastics. In *Annual Book of ASTM Standards 2005*, page 50. ASTM International, 2005.
- [97] D. Rogowsky. Private Communication, April 2009.
- [98] A. Wright et al. The Scintillation and Optical Properties of Linear Alkyl Benzene. SNO+ internal document doc-219-v1., September 2008.
- [99] B. Brewer. SNO+ Engineering Feasibility Report. SNO+ report by SNC-Lavalin Nuclear Inc., October 2005.
- [100] A. Hallin. Nonlinear Buckling Analysis of Acrylic Vessel. Internal SNO+ report., March 2008.
- [101] C. Hearty et al. Report of the Engineering Review of the SNO+ Acrylic Vessel Hold-down Design. Report to SNO+ by hold-down review committee., September 2008.
- [102] A.S. Barabash. NEMO 3 Double Beta Decay Experiment: Latest Results. *J. Phys.: Conf. Ser.*, 173:012008 (10pp), 2009.
- [103] S. R. Elliott and P. Vogel. Double Beta Decay. *Ann. Rev. Nucl. Part. Sci.*, 52:115–151, 2002.
- [104] P. Vogel. Nuclear Physics Aspects of Double Beta Decay. arXiv:0807.2457, 2008.

- [105] F.T. Avignone, S.R. Elliott, and J. Engel. Double Beta Decay, Majorana Neutrinos, and Neutrino Mass. *Rev. Mod. Phys.*, 80:481–516, 2008.
- [106] A. Staudt, K. Muto, and H. V. Klapdor-Kleingrothaus. Calculation of 2ν and 0ν Double-Beta Decay Rates. *Europhys. Lett.*, 13:31–36, 1990.
- [107] A. Faessler and F. Šimkovic. Double Beta Decay. *J. Phys. G*, 24:2139–2178, 1998.
- [108] J. Toivanen and J. Suhonen. Renormalized Proton-Neutron Quasiparticle Random-Phase Approximation and Its Application to Double Beta Decay. *Phys. Rev. Lett.*, 75:410–413, 1995.
- [109] V.A. Rodin, A. Faessler, F. Simkovic, and P. Vogel. Erratum: Assessment of Uncertainties in QRPA $0\nu\beta\beta$ Nuclear Matrix Elements [Nucl. Phys. A 766, 107 (2006)]. 2007. arXiv:0706.4304v1, 2007.
- [110] J.G. Hirsch, O. Castaños, and P.O. Hess. Neutrinoless Double Beta Decay in Heavy Deformed Nuclei. *Nucl. Phys. A*, 582(1-2):124 – 140, 1995.
- [111] N. Fatemighomi. Private Communication, July 2009.
- [112] A.S. Barabash et al. Investigation of $\beta\beta$ Decay in ^{150}Nd and ^{148}Nd to the Excited States of Daughter Nuclei. *Phys. Rev.*, C79:045501, 2009.
- [113] A. P. Babichev et al. Development of the Laser Isotope Separation Method (AVLIS) for Obtaining Weight Amounts of Highly Enriched ^{150}Nd Isotope. *Quantum Electronics*, 35(10):879–890, 2005.
- [114] R. S. Raghavan. New Approach to the Search for Neutrinoless Double Beta Decay. *Phys. Rev. Lett.*, 72(10):1411–1414, Mar 1994.
- [115] B. Caccianiga and M. G. Giammarchi. Neutrinoless Double Beta Decay with Xe-136 in BOREXINO and the BOREXINO Counting Test Facility. *Astropart. Phys.*, 14(1):15 – 31, 2000.

- [116] S. Link and M.A. El-Sayed. Optical Properties and Ultrafast Dynamics of Metallic Nanocrystals. *Ann. Rev. Phys. Chem.*, 54(1):331 – 366, 2003.
- [117] J.D. Jackson. *Classical Electrodynamics*. John Wiley & Sons, Inc., 1999.
- [118] A. Wright. Nanoparticle Loading Report. SNO+ internal document doc-320-v1, March 2005.
- [119] C. Buck et al. Metal Loaded Liquid Scintillators. 2004. Available www.scientificcommons.org/20562176.
- [120] M. Yeh, A. Garnov, and R.L. Hahn. Gadolinium-Loaded Liquid Scintillator for High-Precision Measurements of Antineutrino Oscillations and the Mixing Angle, θ_{13} . *Nucl. Inst. Meth. A*, 578(1):329 – 339, 2007.
- [121] F. Masetti, F. Elisei, and U. Mazzucato. Optical Study of a Large-Scale Liquid-Scintillator Detector. *J. Lum.*, 68(1):15 – 25, 1996.
- [122] H.M. O’Keeffe et al. Analysis of the Scintillator Bucket Data. SNO+ internal presentation doc-306-v1., April 2009.
- [123] F. Boehm and P. Vogel. *Physics of Massive Neutrinos*. Cambridge University Press, 1987.
- [124] J. Maneira. External Background in SNO+. SNO+ internal document doc-43-v1., August 2007.
- [125] Y. Kishimoto. Liquid scintillator purification. In B. Cleveland, R. Ford, and M. Chen, editors, *Topical Workshop on Low Radioactivity Techniques LRT2004*, pages 139–198. AIP Conference Proceedings, 2005.
- [126] K. Zuber. Nd-related Issues. SNO+ internal presentation doc-310-v1., April 2009.
- [127] A. Wright. Nd Cosmogenics: A Preliminary Look. SNO+ internal presentation doc-39-v1., August 2007.

- [128] S. Schonert et al. The GERmanium Detector Array (GERDA) for the Search of Neutrinoless Beta Beta Decays of Ge-76 at LNGS. *Nucl. Phys. Proc. Suppl.*, 145:242–245, 2005.
- [129] A. Balysh et al. The Heidelberg-Moscow Double Beta Decay Experiment with Enriched ^{76}Ge . First Results. *Phys. Lett. B*, 283(1-2):32 – 36, 1992.
- [130] C. E. Aalseth et al. The Igex ^{76}Ge Neutrinoless Double-Beta Decay Experiment: Prospects for Next Generation Experiments. *Phys. Rev.*, D65:092007, 2002.
- [131] R. Gaitskell et al. White Paper on the Majorana Zero-Neutrino Double-Beta Decay Experiment. 2003.
- [132] R. Ardito et al. CUORE: A Cryogenic Underground Observatory for Rare Events. hep-ex/0501010, 2005.
- [133] C. Arnaboldi et al. A New Limit on the Neutrinoless DBD of ^{130}Te . *Phys. Rev. Lett.*, 95:142501, 2005.
- [134] R. Neilson. Exo-200. *J. Phys. Conf. Ser.*, 136(4):042052 (1pp), 2008.
- [135] T Shima et al. MOON for a Next-Generation Neutrino-less Double-Beta Decay Experiment; Present Status and Perspective. *J. Phys.: Conf. Ser.*, 120(5):052055 (3pp), 2008.
- [136] A. S. Barabash. NEMO-3 and SuperNEMO Double Beta Decay Experiments. *J. Phys. Conf. Ser.*, 39:347–349, 2006.
- [137] J.C. Maneira. *Calibration and Monitoring for the Borexino Solar Neutrino Experiment*. PhD thesis, University of Lisbon, 2001.
- [138] B. Cleveland. Rare-earth Backgrounds for Nd Phase. SNO+ internal document doc-117-v1., January 2008.
- [139] A. Wright. A Study of Event Pileup Effects in SNO+. SNO+ internal document doc-118-v2., January 2008.

- [140] B. Cleveland and I. Lawson. Measurement of the γ Activity of a NdCl_3 Sample with the SNOLAB Ge Detector. SNO+ internal document doc-120-v1., March 2008.
- [141] B. Cleveland and I. Lawson. Measurement of the γ Activity of a Nd_2O_3 Sample with the SNOLAB Ge Detector. SNO+ internal document doc-122-v1., August 2008.
- [142] B. Cleveland and I. Lawson. Measurement of the Activity of Four Nd_2O_3 Samples from Stanford Materials with the SNOLAB Ge Detector. SNO+ internal document doc-220-v2., November 2008.