An Analysis of the ^{3}He Proportional Counter Data from the Sudbury Neutrino Observatory Using Pulse Shape Discrimination

by

Ryan Martin

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

This thesis presents an independent analysis of the data from ^{3}He-filled proportional counters from the third phase of the Sudbury Neutrino Observatory (SNO) data. These counters were deployed in SNO's heavy water to independently detect neutrons produced by the neutral current interaction of ^{8}B solar neutrinos with deuterium. Previously published results from this phase were based on a spectral analysis of the energy deposited in the proportional counters. The work in this thesis introduces a new observable based on the time-profile of the ionization in the counters. The inclusion of this observable in a maximum-likelihood fit increases the potential to distinguish neutrons from backgrounds which are primarily due to alpha-decays. The combination of this new observable with the energy deposited in the counters results in a more accurate determination of the number of neutrons.

The analysis presented in this thesis was limited to one third of the data from the proportional counters, uniformly distributed in time. This limitation was imposed to reconcile different time-lines between the submission of this thesis, a thorough review of this work by the SNO Collaboration and results from an independent analysis that is still underway. Analysis of this reduced data set determined that 398 $\pm$ 29 (stat.) $\pm$ 9 (sys.) neutrons were detected in this reduced data-set. The number compares well to the previous analysis of the data, based only on a spectral analysis of the deposited energy, which determined that 410 $\pm$ 44 (stat.) $\pm$ 9 (sys.) were detected in the same time period. The analysis presented here has led to a substantial increase in the statistical accuracy. Assuming that the statistical

accuracy will increase when the full data set is analyzed, the results from this thesis would bring the uncertainty in the ^{8}B solar neutrino flux to 6.8% down from 8.5% in the previously published results. The work from the thesis is intended to be included in a future analysis of the SNO data and will result in a more accurate measurement of the total flux of solar neutrinos from ^{8}B as well as reduce the uncertainty in the θ_{12} neutrino oscillation mixing angle.

Statement of Originality

The work in this thesis is the original and independent work of the author. This thesis presents an independent analysis of a fraction of the data collected from the Sudbury Neutrino Observatory experiment that was performed by several hundreds of talented collaborators over more than two decades. Participation in the analysis of such an experiment necessarily depended on the work of many people.

The first chapter of this thesis presents an overview of the relevant literature on solar neutrinos and a description of the SNO experiment and does not contain any original material. Similarly, the second and third chapters provide a description of the physics of proportional counters and a description of the NCD array and do not contain any original material either.

The main contribution from the author begins in the fourth chapter where the pulse-shape grid-fitter is introduced. The grid-fitter was designed and optimized by the author with the advice of numerous members of the SNO Collaboration. It was an original idea, presented in this thesis, to introduce the NoverA observable as a means of distinguishing neutrons from alphas. The use of NoverA as an observable in a maximum likelihood analysis was also proposed by the author as a way to improve on the statistical accuracy that can be achieved using energy.

The fifth chapter introduces probability density functions (pdfs) for the energy and

NoverA observables. The methods by which these functions were determined were an original work of this thesis. A novel method for handling systematic uncertainties for pdfs with low statistics is also introduced.

The sixth chapter contains little original material as it is intended as an introduction to the extended likelihood formalism using Markov-Chain Monte-Carlo analysis. It does however contain an outline of the novel way in which systematic uncertainties in the pdfs are handled. The seventh chapter shows various tests of the procedures introduced in this thesis and the eighth chapter applies these procedures to the real NCD data, which has never been done. The last chapter concludes and does thus not include any additional material.

Acknowledgments

I would like to take this opportunity to thank all the people that have made it possible for me to complete my PhD at Queen's University. As I sit here at my new position at Lawrence Berkeley National Labs, it is not without nostalgia that I look back at my time at Queen's.

I have learned a great deal about being a physicist from my supervisors, Art McDonald and Aksel Hallin. I am grateful for all the help and advice that Art has given me as I wrote up my thesis and applied for postdoc positions, as well as his willingness to always take the time to help me, no matter how busy or where in the world he was. I am also grateful for the sense of 'high-standards' that he has given me as a physicist and look forward to applying these as I progress in my career. I am also grateful for the methodologies that I have learned from Aksel, in particular the courage to not back away from examining problems in as much detail as they may require. I feel that working with Aksel has given me the necessary 'rigour' that one needs to be an accomplished experimentalist. I know that I am extremely lucky to have had the chance to work with both of these amazing physicists.

My time at Queen's has been very enjoyable and a great learning experience thanks to so many people. My graduate studies started when Mark Chen and Ian Towner (I like to have 2 supervisors) decided to take me on as a young Master's student. I am very grateful that Mark took me on and introduced me to the SNO group and I have learned a lot about teaching from all the amazing classes that I took with Ian. And then, of course, there all

the amazing graduate students that I have met through the years. I can't even begin to recount all the great times I have had with Mark Kos and Ryan MacLellan; the BBQs, softball games, pig roasts, ski trips, tennis games (I am better than Mark), loosing money gambling on the Internet, winning money gambling on the Internet (at this point we are up and Ryan owes my USD100). I am also grateful for all the good times I have had with Paradorn and would like to thank him for printing my thesis. The numerous discussions I have had with Alex Wright have always kept my mind sharp and it has been great to go through grad school with Alex (and I also learned something about show chickens).

The atmosphere in the SNO group at Queen's has been a fantastic learning environment and a great place in general. Thanks to Peter Skensved for all the computer help and quick resolution of software related matters (even if it wasn't his problem), I loved all the puns from Barry Robertson (my kind of humour), the advice from Hamish Leslie and Hugh Evans, the resolution of hardware related matters by Steve Guillen, the special machine shop training from Chuck Hearns (specific to the manufacturing of pig roasting equipments) and the help in the old days from Jose Maneira who was always willing to point me to google for C++ advice.

The Department of Physics (I prefer that name to the one on my title page) at Queen's has been a great place to do my studies. Thanks to Dick Henriksen for supervising me as an undergrad. I have had a great time TA-ing (although I know many grad students who did not), especially once I started working with Alastair MacLean, from whom I have a learned a lot about teaching and whom I thank for entrusting me with many responsibilities; TA-ing the 106 physics course was one of my favourite things as a graduate student, and I look forward to that component later in my career.

I am also very grateful to the SNO collaboration in general for letting me in and allowing me to participate in such a fantastic experiment. I am not unaware of the privilege that was given to me to work with such a successful and professional collaboration.

I would like to thank my lovely girlfriend, Rebecca, for all her support during the writing

of my thesis. Without her, I would have had a hard time finding the perseverance that was required to finish writing; I am very grateful for the motivation that she has provided me.

Finally, I would like to thank my parents, Claudia and Reg, and my sister, Cristina, for all the support during the many many years. I am extremely grateful for the parents that I was given and it is my hope to make them proud as I continue on. I would like to dedicate this thesis to them as a small token of gratitude for the wonderful life that they have given me.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The analysis in this thesis results in a significant increase in the accuracy of the measurement of the flux of solar neutrinos, compared to previous work. This chapter aims to give a context to the work in this thesis as well as review the current knowledge in the field of solar neutrino physics. It will start with an introduction to the field of solar neutrino physics before introducing the Sudbury Neutrino Observatory and the results from the experiment. The chapter will then finish with an overview of the work in this thesis and highlight its relevance.

1.1 The Aim of this Thesis

This thesis will present an independent analysis of the data from the third phase of the Sudbury Neutrino Observatory (SNO) experiment. The SNO experiment (described in more detail in section 1.3) was an experiment designed to measure the flux of neutrinos coming from the Sun and test the prediction from the standard solar model (SSM). The three phases of the experiment not only showed that electron solar neutrinos change flavour, but also confirmed the theoretical prediction from the SSM and provided accurate measurements of neutrino properties, in particular the 'mixing angle', θ_{12}, and mass-squared difference, Δm_{21}^2, both discussed later in this chapter.

In the third phase of the experiment, an array of 36 strings of proportional counters was deployed to count neutrons produced from solar neutrinos. Previous analysis of the data from these counters only considered the energy deposited by each neutron capture. The number of detected neutrons was inferred by performing a maximum-likelihood fit to the energy spectra of the counters. Alpha emitters from radioactive contaminants in the detector materials provided the main background to neutrons.

The work in this thesis includes additional information from the counters in order to improve the ability to distinguish neutrons from alphas. Namely, the ionization energy as a function of time was recorded for each event in the counters. These 'pulse shapes' often

depend on the nature (neutron or alpha) of the particle that led to the ionization inside the proportional counters. An observable based on these pulse shapes is thus introduced in this thesis and used in conjunction with the deposited energy to decrease the uncertainty in the measured number of neutrons.

A more accurate determination of the number of neutrons detected in the third phase of SNO has a direct impact on the precision of the neutrino oscillation parameters (in particular, θ_{12}) and is the main motivation for the work presented in the following chapters.

The work in this thesis was completed before the SNO collaboration had fully reviewed and fixed the criteria for incorporating this analysis into future publications. An alternative method for analyzing the NCD data is also underway and will be used as a cross-check of the work presented here. Since these two methods were not fully reviewed at the time that this thesis was submitted, it was decided that only one third of the data would be made available for inclusion in the thesis. The restriction in the size of the data set allows for comparisons with previously published work, while maintaining blindness on the statistical accuracy of the new analysis.

1.2   Solar Neutrinos

The field of solar neutrino physics started in 1964 when John Bahcall [1] and Raymond Davis Jr. [2] proposed that solar neutrinos could be detected. John Bahcall was a pioneer of the solar model calculation of neutrino fluxes ([3, 4]) and Ray Davis was the first to conduct a successful experiment to detect solar neutrinos [5]. Neutrino detection had only recently been achieved at that time by Cowan and Reines [6], using anti-neutrinos produced in the Savannah River nuclear reactor, making the prospect of detecting solar neutrinos quite a challenge.

1.2.1 Neutrino Production in the Sun

In 1939, Bethe [7] first published a set of nuclear fusion reactions that he believed were behind the energy production in the Sun. These 'CNO-cycle' reactions (shown in panel (a) of Figure 1.1) use carbon and nitrogen as catalysts to convert the hydrogen (protons) in the Sun into helium (alpha particles). The CNO-cycle contributes at most 2% of the solar energy and most of the fusion reactions occur through the 'pp-chain' [4]. The pp-reactions (panel (b) of Figure 1.1) also convert hydrogen into helium and takes place in three branches.

CNO Solar Fusion Cycle

$$^{12}C + p \rightarrow {}^{13}N + \gamma$$

$$^{15}N + p \rightarrow {}^{12}C + \alpha \qquad \qquad ^{13}N \rightarrow {}^{13}C + e^+ + \nu_e$$

$$^{15}O \rightarrow {}^{15}N + e^+ + \nu_e \qquad \qquad ^{13}C + p \rightarrow {}^{14}N + \gamma$$

$$^{14}N + p \rightarrow {}^{15}O + \gamma$$

p-p Solar Fusion Chain

$$p + p \rightarrow {}^2H + e^+ + \nu_e \qquad p + e^- + p \rightarrow {}^2H + \nu_e$$

$$^2H + p \rightarrow {}^3He + \gamma$$

$$^3He + {}^3He \rightarrow {}^4He + 2p \qquad {}^3He + p \rightarrow {}^4He + e^+ + \nu_e$$

$$^3He + {}^4He \rightarrow {}^7Be + \gamma$$

$$^7Be + e^- \rightarrow {}^7Li + \gamma + \nu_e \qquad {}^7Be + p \rightarrow {}^8B + \gamma$$

$$^7Li + p \rightarrow \alpha + \alpha \qquad \qquad {}^8B \rightarrow 2\alpha + e^+ + \nu_e$$

(a) CNO cycle　　　　　　　　(b) pp-chain
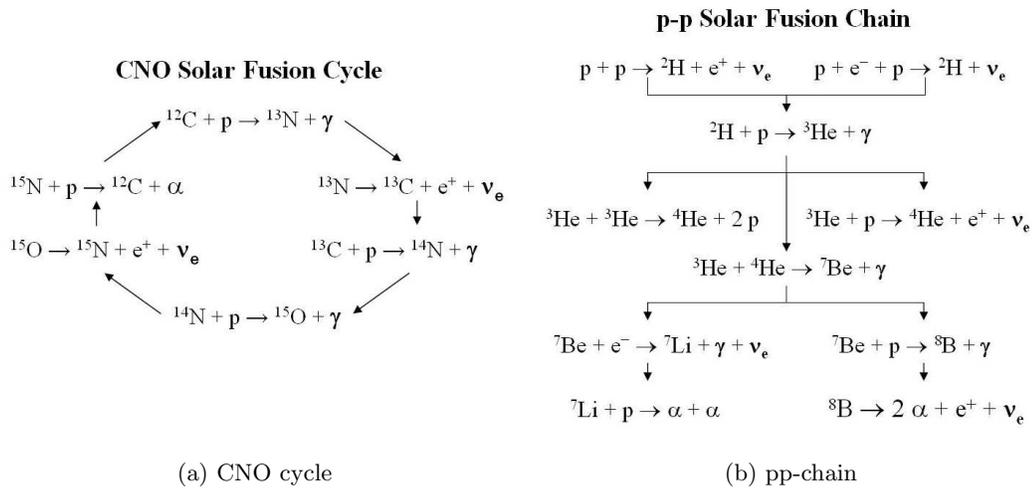
Figure 1.1: Reactions in the CNO cycle (a) and pp-chains (b). Figures taken from [8]

Since the energies released in these fusion reactions are always less than the mass of a muon (or tau), only electron-flavour neutrinos are produced. Figure 1.2 shows the energy spectra of the neutrinos produced by the Sun in these reactions.
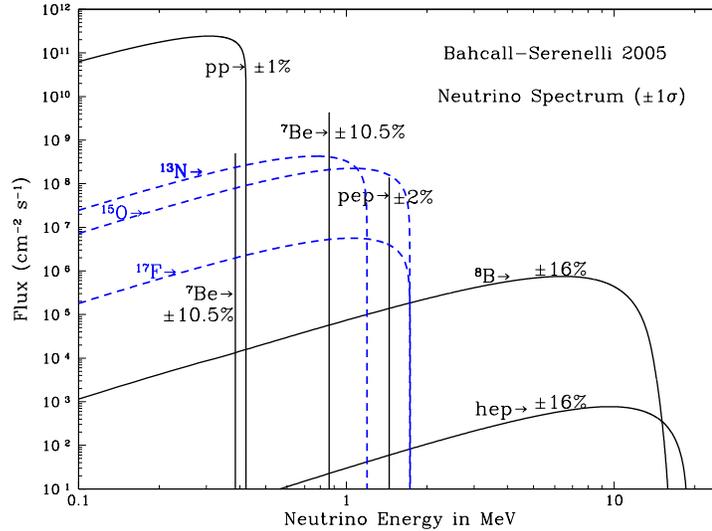
Figure 1.2: Energy spectra of solar neutrinos from the BS2005(OP) solar model. The blue dotted lines show neutrinos from the CNO cycle and the black lines are from the pp-chain. Uncertainties in the flux predictions are also shown. Figure reproduced from [4]

1.2.2 The Solar Neutrino Problem

Ray Davis conducted the first successful experiment to detect solar neutrinos [5]. In order to detect electron-flavour neutrinos, he exploited the inverse beta-capture reaction on chlorine:

$$\nu_e + {}^{37}\text{Cl} \rightarrow {}^{37}\text{Ar} + e^- \tag{1.1}$$

by exposing a tank with 430,000 liters of C_2Cl_4 in the Homestake Mine (South Dakota) 1463 m underground (to shield the experiment from cosmic-ray induced backgrounds). The reaction is sensitive to neutrinos with energies above 0.8 MeV and was thus primarily a test of the flux of neutrinos from the ^{7}Be and ^{8}B reactions in the Sun. Every few months, the argon atoms were chemically extracted from the volume and counted at his laboratory in Brookhaven. The experiment was run from 1967 to 1995 and observed a consistent deficit in neutrinos [9] compared to the standard solar model prediction (summarized in Figure 1.3).

The KamiokaNDE [10] and, later, the Super-Kamiokande [11] detectors also observed a deficit in the number of neutrinos from the Sun. The latter two detectors looked for

Čerenkov radiation in water from electrons scattered by solar neutrinos. Since the scattered electrons point back towards the Sun, the Čerenkov detectors provided the first unquestionable evidence that neutrinos were coming from the Sun. Water Čerenkov detectors are primarily sensitive to electron-flavour neutrinos (by a factor of $\sim$6 over other flavours) and have an energy threshold above $5\,\text{MeV}$ so that they are almost only sensitive to neutrinos from the ^8B decay. KamiokaNDE and Super-Kamiokande (Super-K) both observed a significant deficit in neutrinos compared with the standard solar model prediction.

Two other 'radio-chemical' experiments, SAGE [12] and Gallex [13], used an inverse electron capture reaction on ^{71}Ga and also observed a deficit in the measured number of electron-flavour neutrinos compared to the standard solar model prediction. The gallium experiments were sensitive to neutrino energies down to $0.2\,\text{MeV}$ and were thus able to measure the neutrino flux from the p+p reaction of the pp chain.

This consistent discrepancy between the number of neutrinos measured from the experiments and the number predicted by the solar model was dubbed the 'solar neutrino problem' and was not conclusively resolved until 2001 when the Sudbury Neutrino Observatory published its first results [14]. Figure 1.3 shows a summary of the measured fluxes from the solar neutrino experiments along with the corresponding prediction from the solar model.

1.2.3 Neutrino Oscillations

A hint to the solution of the solar neutrino problem came in 1998, when the Super-K experiment indicated that atmospheric neutrinos could oscillate [16, 17]. Neutrino oscillations are a mechanism first proposed by Pontecorvo [18], in analogy to quark flavour mixing, that allow neutrinos to change flavour.

Neutrino oscillations can occur when there are two different eigenbases to represent neutrino quantum states. In particular, neutrinos are always produced in a 'lepton flavour' eigenstate $|\nu_\alpha>$ (as they are produced together with the corresponding heavy lepton) that
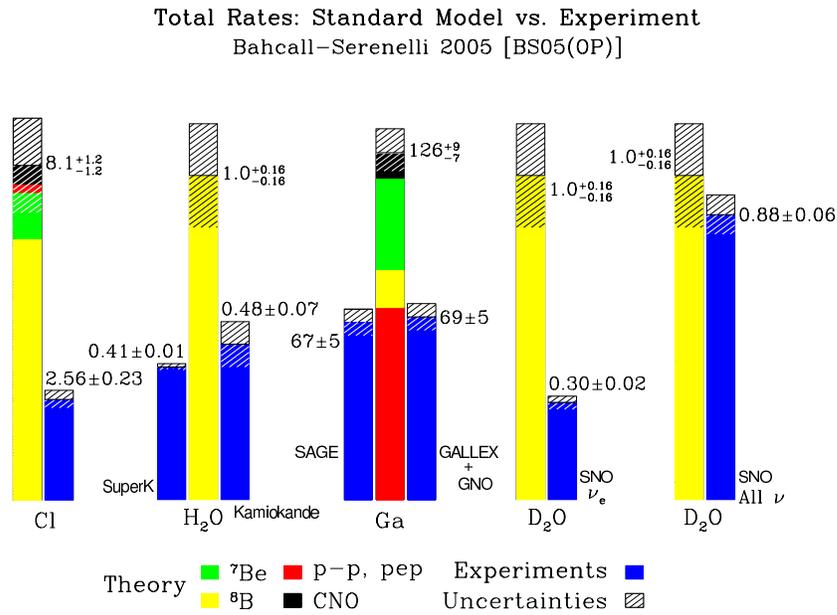
Figure 1.3: Summary of the solar neutrino problem. The theoretical prediction and results from each class of solar neutrino experiment are shown side-by-side. The experiments are: Cl (chlorine, the Homestake experiment), H_2O (the Čerenkov experiments), Ga (the gallium experiments) and D_2O (two results from the SNO experiment described in the next section). The theoretical predictions contain the contributions from different neutrinos depending on the energy threshold of the type of experiment. The blue bars show the experimentally measured fluxes of neutrinos. The last bar from SNO is the 'SNO neutral-current' measurement (the only experiment equally sensitive to all three flavours of neutrinos) which solves the solar neutrino problem, as it agrees with the theoretical prediction. Figure taken from [15].

is not necessarily a mass eigenstate of the Hamiltonian, $|\nu_i>$. Unitarity requires that the two representations be related by a unitary transformation, U:

$$|\nu_\alpha> = U|\nu_i>$$

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{1.2}$$

where α stands for the three flavours e, μ, τ and i stands for the 3 mass eigenstates of the Hamiltonian. The Schrödinger equation is then diagonal in the mass eigenstates of the Hamiltonian:

$$-i\hbar\frac{d}{dt}|\nu(t)> = H|\nu(t)>$$

$$-i\hbar\frac{d}{dt}\begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} = \begin{pmatrix} E_1 & 0 & 0 \\ 0 & E_2 & 0 \\ 0 & 0 & E_3 \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{1.3}$$

For most experiments, neutrino oscillations can be approximated by a two flavour equivalent scenario. In the case of solar neutrinos (see, for example, [19] and [20]), one can model the oscillations as occurring between ν_e and ν_X, where the latter is a linear combination of the muon and tau neutrinos. In this case, the formalism is greatly simplified and the 'mixing matrix', U can be parametrized with a single angle, θ_{12}, the 'solar mixing angle':

$$\begin{pmatrix} \nu_e \\ \nu_X \end{pmatrix} = \begin{pmatrix} \cos\theta_{12} & \sin\theta_{12} \\ -\sin\theta_{12} & \cos\theta_{12} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} \tag{1.4}$$

It is then straightforward to calculate the 'survival probability', P_{ee}, defined as the probability of detecting an electron neutrino at a distance L from a point where an electron neutrino was created, $|\nu(t=0)> = |\nu_e>$:

$$P_{ee} = |<\nu_e|\nu(t)>|^2$$

$$= 1 - \frac{1}{2}\sin^2 2\theta_{12}\left(1 - \cos\left(\frac{\Delta m_{21}^2}{2Ec\hbar}L\right)\right) \tag{1.5}$$

where $\Delta m_{21}^2 = m_2^2 - m_1^2$ is the difference in masses squared between the eigenstates of the Hamiltonian. For an easy-to-follow, complete derivation, see [19].

Thus, the probability of detecting an electron neutrino oscillates with distance, if the neutrino has mass. The frequency of the oscillations depends on the ratio $\Delta m_{21}^2/E$ and the amplitude is determined by the parameter θ_{12}. If there are many neutrinos, produced with a range of different L, with different energies, and detected with a finite energy resolution (as is the case for all solar neutrino experiments), the oscillating part of the survival probability effectively averages out:

$$\bar{P}_{ee} = 1 - \frac{1}{2}\sin^2 2\theta_{12} \qquad (1.6)$$

where the average solar neutrino survival probability, $\bar{P}_{ee}$, is independent of energy.

The case for all three flavours of neutrinos is similar and results in a slightly longer formula [19]. The key difference is that, with three flavours, the most general unitary transformation contains three mixing angles, $(\theta_{12}, \theta_{13}, \theta_{23})$, and potentially, a CP-violating complex phase (δ_{CP}). In addition, the three flavour scenario will contain two different mass-squared differences, Δm_{21}^2 and Δm_{32}^2. Most neutrino oscillation experiments can be analyzed in terms of two flavours of neutrinos, depending on the energy and 'baseline' (L) of the neutrinos. Solar neutrinos have been shown (see [20, 21]) to have a slight sensitivity to the θ_{13} mixing angle which can be well approximated by modifying the two-neutrino survival probability:

$$P_{ee}^{3f} = \sin^4 \theta_{13} + \cos^4 \theta_{13} P_{ee}^{2f} \qquad (1.7)$$

where P_{ee}^{2f} is the survival probability as calculated in the two-neutrino case[1]. The current experimental limits [22] show hints that $\sin^2 \theta_{13} \sim 0.016$ (if not zero), so that this modification is small.

Of particular interest to the solar neutrino problem, neutrino oscillations can explain why

---

[1]One needs to make a slight modification to the P_{ee}^{2f} calculation in the case where matter is present, as shown in [20]

the measured flux of electron neutrinos is smaller than expected, since these can oscillate into the other flavours and avoid detection from experiments looking for electron flavour neutrinos. The atmospheric neutrinos that Super-K measured had energies of several GeV and baselines of order 10 kilometers and thus corresponded to different mixing parameters (θ_{23} and Δm_{32}^2) than solar neutrinos and, therefore, did not provide a conclusive solution to the solar neutrino problem.

Solar neutrinos are also qualitatively different from atmospheric neutrinos since they traverse dense matter (the Sun). Wolfenstein [23], Mikheyev and Smirnov [24] showed that neutrino oscillations are altered by matter because of the different scattering cross-sections for interaction with electrons between electron neutrinos and the muon and tau neutrinos. This, so-called, 'MSW effect' results in an energy dependence of the average solar neutrino survival probability, as shown in Figure 1.4. The MSW effect in the Sun effectively converts electron neutrinos into an almost pure ν_2 eigenstate when these exit the solar matter. The survival probability undergoes a transition around $2\,\mathrm{MeV}$ from a value that is dominated by vacuum oscillations (where the MSW effect is small) to a value that is driven by matter effects.

1.3   The SNO Detector

The Sudbury Neutrino Observatory [25] was an experiment designed to resolve the solar neutrino problem by using heavy water (D_2O) to detect all flavour of neutrinos through the neutral current reaction. The idea of using 1,000 tonnes of heavy water in a deep mine to solve the solar neutrino problem was first proposed by Chen and the SNO Collaboration in 1984 [26]. Construction on the experiment was finished in 1997, it was then commissioned between 1997 and 1999 and data were collected from 1999 to 2006.

The SNO detector was located near Sudbury, Ontario, Canada in the Creighton mine operated by Vale-Inco. The experiment benefited from $2\,\mathrm{km}$ of norite rock overburden
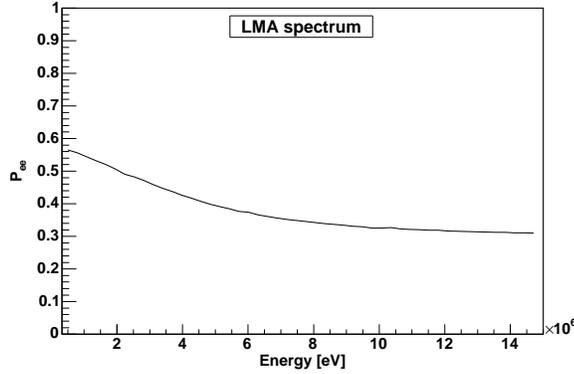
Figure 1.4: Energy dependence of the solar neutrino survival probability shown for the 'LMA' (large mixing angle) choice of neutrino mixing parameters. At low energies the survival probability is not significantly affected by matter. Around 2 MeV, the survival probability undergoes a transition to a value that is the result of the MSW effect.

(6010 m of water equivalent) to shield the detector from cosmic ray muons. The flux of cosmic ray muons at that depth is about 70 per day through the detector volume.

The detector consisted of 1,000 tonnes of heavy water (D_2O) contained in a 12 m diameter acrylic vessel. The heavy water volume was monitored by 9438 8"-diameter Hamamatsu r1408 photo-multiplier tubes (PMTs) held in place on an 17.8 m diameter geodesic support structure. The PMT array provided a 54% photo-cathode coverage of the sphere solid angle. This was all contained within a 22 m diameter cavity (at its widest), 34 m high and filled with 7,000 tonnes of ultra-pure light water (H_2O) to shield the heavy water volume from radioactivity (from the cavity walls and the PMTs). Figure 1.5 shows an artist's rendition of the SNO detector.

1.3.1 Neutrino Reactions in SNO

SNO detected neutrinos by searching for Čerenkov radiation from electrons. The electrons were produced as the result of three reactions involving neutrinos that took place in the
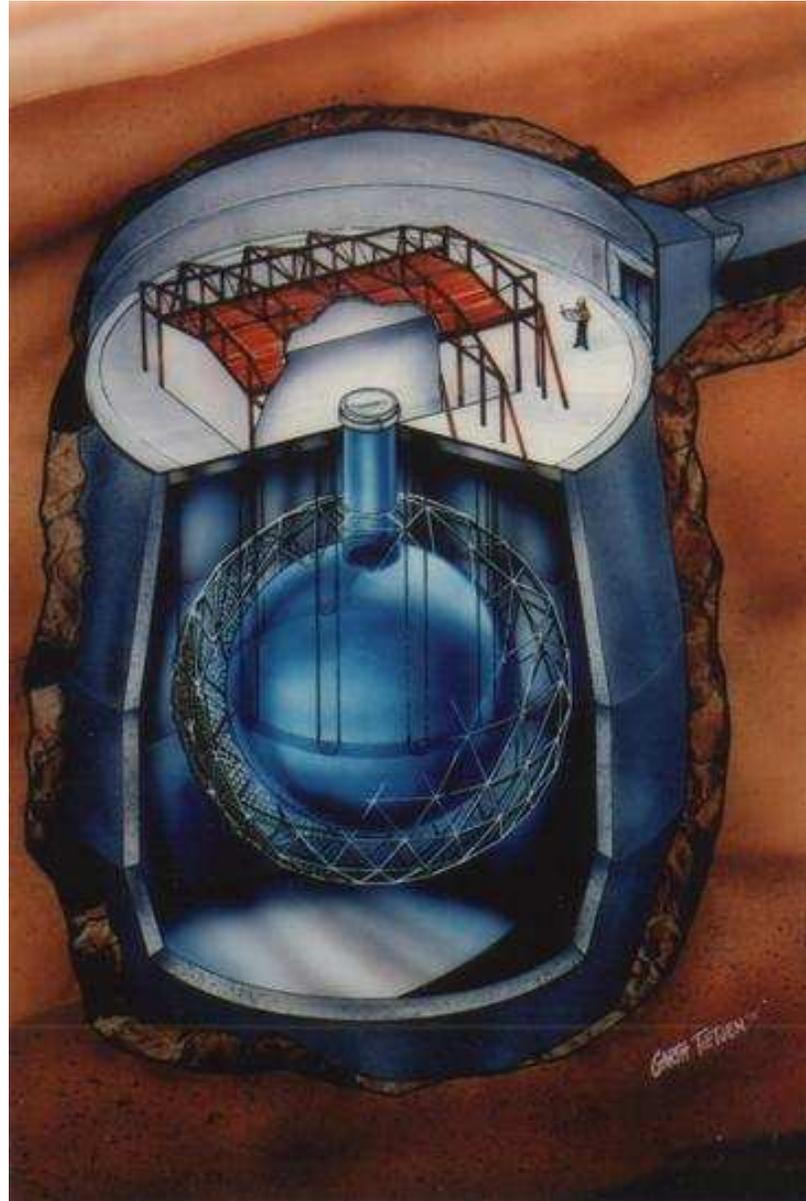
Figure 1.5: An artist's rendition of the SNO detector. The 12 m diameter acrylic vessel contains 1,000 tonnes of heavy water and is supported from the 'deck' with Vectran ropes. The entire cavity containing the 17.8 m diameter photo-multiplier support structure is filled with 7,000 tonnes of light water.

heavy water. The charged-current reaction (CC) is sensitive only to electron flavour neutrinos, ν_e, with a threshold energy of $1.4\,\text{MeV}$.

$$\nu_e + d \to p + p + e^- \tag{1.8}$$

where d is the deuterium nucleus, p are protons and e^- is an electron whose energy is highly correlated with that of the incoming neutrino. Measuring the rate of CC reactions is thus a direct measure of the electron neutrino flux. In addition, by measuring the energy spectrum of the Čerenkov electrons, one can infer the spectrum of the incoming electron-flavour neutrinos.

Heavy water also allows one to detect active neutrinos of all three flavours, ν_x, as they all have an equal cross-section for a neutral-current interaction (NC) with deuterium:

$$\nu_x + d \to p + n \tag{1.9}$$

where a neutron, n, is released, providing the incoming neutrino has an energy above $2.2\,\text{MeV}$. Section 1.3.2 explains how the neutrons were then detected. This NC reaction is a direct measure of the flux of active neutrinos, independent of whether neutrinos oscillate, since it is equally sensitive to all three active flavours. Furthermore, if neutrino oscillations occur, the ratio of the CC/NC rates is a measure of the average solar survival probability and is almost directly proportional to $\sin^2 2\theta_{12}$ for solar neutrinos.

Finally, SNO was also sensitive to an elastic scattering reaction (ES) of neutrinos with electrons:

$$\nu_x + e^- \to \nu_x + e^- \tag{1.10}$$

which is sensitive to all flavour of neutrinos, although the cross-section for electron neutrinos is ~ 6 times greater than it is for the other active flavours. The direction of the out-going electron is correlated to that of the neutrino and, essentially, points back to the Sun. The ES reaction in SNO has low statistics when compared with the Super-K [11] experiment.

1.3.2 The Three Phases of SNO

The SNO experiment was conducted in three phases in order to detect the neutrons from the NC reaction in three different ways. In the first phase (the 'D$_2$O phase') of the experiment [14, 27, 28, 29, 30], Compton-scattered electrons from 6.25 MeV.γ-rays from the neutron captures on deuterium were used to detect neutrons. Čerenkov events from the NC signal were thus indistinguishable from the CC events, as both resulted in the observation of a Čerenkov electron. Thus, the only way to separate NC and CC events was to use a statistical separation based on the expected energy spectra of the two classes of events, which meant that the energy spectrum of the incoming neutrinos could not be measured independently.

In the second phase of the experiment [31, 32], 2000 kg of salt (NaCl) were dissolved into the heavy water. This allowed the neutrons to capture on ^{35}Cl with a much higher cross-section (44 b vs 0.5 mb on deuterium)[33], releasing more energy (8.6 MeV) and, on average, several γ-rays , producing a more isotropic pattern of hits on the PMTs. This resulted in a much higher detection efficiency for neutrons (40.7% vs 14.4% in the first phase), as well as providing an energy-independent way of separating NC and CC signals using the isotropy of the light hitting the PMT sphere. Data from the second phase (the 'salt phase') were then used to measure the energy spectrum of incoming neutrinos. This is of value for testing neutrino oscillation scenarios, in particular to test the MSW distortion of the survival probability from Figure 1.4.

Finally, in the third phase (the 'NCD phase') of the experiment, 36 strings of proportional counters filled with ^{3}He ('Neutral Current Detectors') were deployed in the heavy water to detect neutrons through the reaction:

$$\text{n} +^3\text{He} \rightarrow^3\text{H} + \text{p} + 764\,\text{keV} \tag{1.11}$$

Analysis of the data from these 'Neutral Current Detectors' (NCDs) is the primary concern of the work presented in this thesis and will be discussed at length in the next chapters.

1.4 Results from SNO

This section covers the published results from SNO to date (June 2009). The work in this thesis aims at improving the measurement of the neutrino fluxes from the third phase of the experiment.

1.4.1 The D$_2$O Phase and resolution of the Solar Neutrino Problem

SNO's first results [14] covered data from Nov. 2, 1999 to Jan. 15, 2001 corresponding to 241 live days of data. The time, charge and location of the PMTs were recorded for each event 'trigger', defined as 18 or more PMTs crossing their thresholds in a period of 93 ns .

The data set was reduced by rejecting events that did not have characteristics of Čerenkov radiation. A fiducial volume cut was applied at a radius of 5.5 m to remove external background events from materials outside of the D$_2$O volume. Events with a re-constructed electron kinetic energy below 6.75 MeV were also discarded in order to reduce events from low-energy radioactive backgrounds. This resulted in a final data set with 1,169 events.

The reconstructed energy of events was calibrated using a tagged ^{16}N source of 6.13 MeV γ-rays and verified with a ^{8}Li β source. These sources could be deployed 3-dimensionally within the heavy water. Optical calibrations of the detector were performed by using diffuse light from a pulsed laser source with six different wavelengths ranging from 337 nm to 620 nm . The detector's response to neutrons was calibrated using a ^{252}Cf fission source.

The analysis of the PMT data were done using an extended maximum-likelihood fit for the number of CC and ES events in the data, as well as the number of neutrons. The statistical separation between the different types of signals (CC, ES and neutrons) was performed using, as 'observables', the reconstructed energy of events (Figure 1.6a), the reconstructed radial position of events (Figure 1.6c), and the direction relative to the Sun of the reconstructed Čerenkov events (Figure 1.6b). Probability density functions of

these observables were obtained for each signal by Monte Carlo simulation and verified with calibration data. The resulting fits in the three observables are shown in panels (abc) of Figure 1.6.

SNO's first published results did not include a determination of the number of neutral-current neutrons due to the energy threshold that was chosen to limit contamination from low energy events (including neutrons). The extended maximum-likelihood resulted in 975.4±39.7 CC events, 106.1±15.2 ES events and 87.5±24.7 neutrons. These event numbers can be converted into the flux of solar neutrinos from the ^{8}B decay, by assuming a neutrino energy spectrum undistorted by the MSW effect:

$$\phi_{SNO}^{CC}(\nu_e) = 1.75 \pm 0.07(\text{stat.})^{+0.12}_{-0.11}(\text{sys.}) \pm 0.05(\text{theor.}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$$

$$\phi_{SNO}^{ES}(\nu_x) = 2.39 \pm 0.34(\text{stat.})^{+0.16}_{-0.14}(\text{sys.}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$$

$$(1.12)$$

where the theoretical uncertainty comes from the CC cross-section. The dominant source of systematic uncertainty was due to the energy scale. The ES measurement from SNO was in good agreement with the high precision result from the Super-K experiment at that time:

$$\phi_{SK}^{ES}(\nu_x) = 2.32 \pm 0.03(\text{stat.})^{+0.08}_{-0.07}(\text{sys.}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$$

$$(1.13)$$

A test was then performed for the hypothesis that electron neutrinos had not changed their flavour, by comparing SNO's CC flux (sensitive only to the electron flavour neutrinos) and Super-K's ES flux (sensitive to all flavours). It was found that SNO's CC measurement is inconsistent with the ES measurement from Super-K by 3.3σ. Since the ES reaction is sensitive to other flavours of neutrinos, this is evidence of a non-electron-flavour component

of neutrinos in the solar flux. This was determined to be:

$$\phi^{SNO+SK}(\nu_{\mu\tau}) = 3.69 \pm 1.13 \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$$

(1.14)

which then gives the total flux of solar neutrinos as:

$$\phi^{SNO+SK}(\nu_x) = 5.44 \pm 0.99 \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$$

(1.15)

in excellent agreement with Bahcall's BP2000 solar model prediction [3]:

$$\phi^{BP2000}(\nu_x) = 5.05(1.00^{+0.20}_{-0.16}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$$

(1.16)

A summary of the measured fluxes and the determination of the total flux is shown in panel (d) of Figure 1.6.

This first measurement of CC reactions from SNO combined with Super-Kamiokande's measurement of the ES events had thus shown evidence for a non-electron flavour component in the solar neutrino flux.

In 2002 [27], SNO published the first ^{8}B neutrino fluxes determined using the neutral-current reaction. The data set was obtained from Nov. 2, 1999 to May 28, 2001 comprising 306.4 live days spanning the entire first phase of SNO.

The principal challenge in this analysis was an accurate determination of backgrounds to the neutral-current-induced neutrons. Background neutrons are present in the heavy water from the photo-disintegration of deuterium due to γ-rays with energies above 2.2 MeV from ^{214}Bi and ^{208}Tl (uranium and thorium chain decay daughters). These photo-disintegration neutrons were indistinguishable from neutral-current neutrons and were thus an irreducible background. Water assays were performed in order to constrain the amount of uranium and thorium in the water. These were compared with Čerenkov events (from radioactivity) below the analysis energy threshold and provided consistent constraints on the number of

(a) Fit in reconstructed energy

(b) Fit in direction of event to sun

(c) Fit in events radial position

(d) Total fluxes
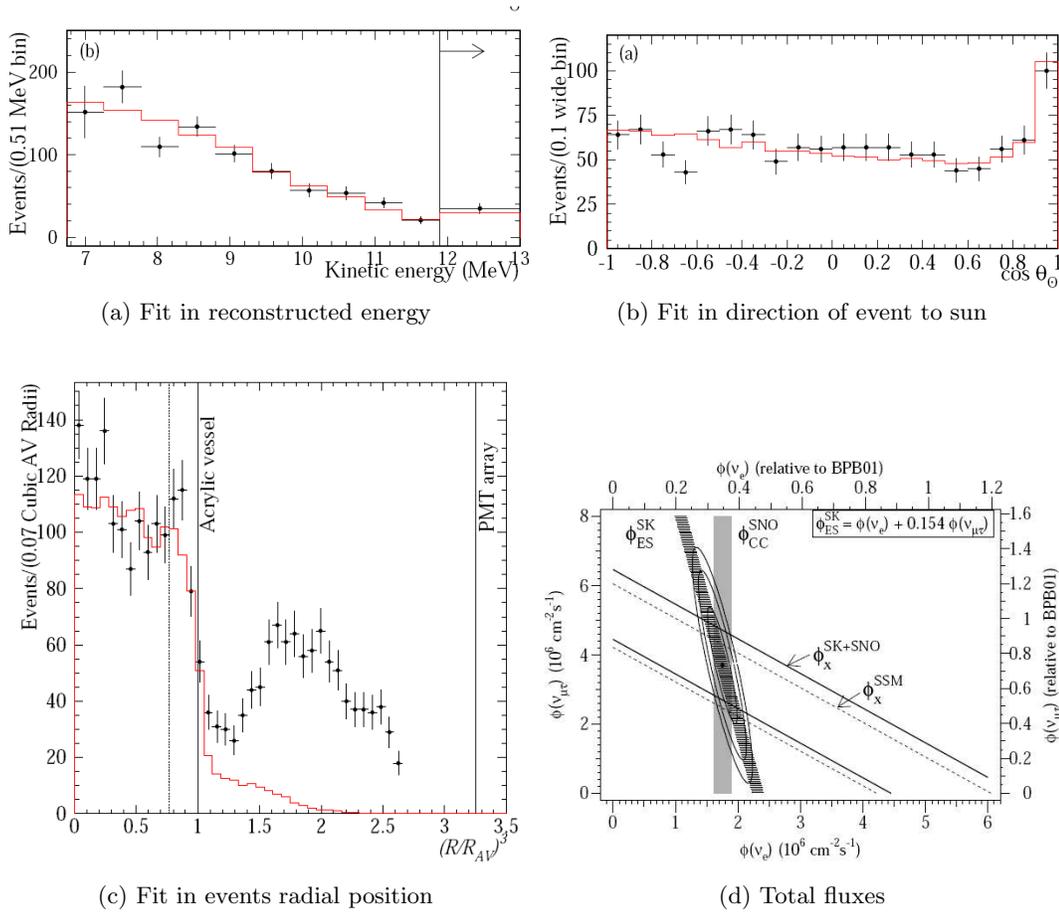
Figure 1.6: Figures summarizing the results from the first published results from SNO. Panels (abc) show the fit to the SNO data in energy, event direction to the sun and event position in the detector, respectively. Panel (d) shows a summary of the inferred neutrino fluxes along with the result from Super-K. The fluxes have been broken into an electron-flavour component (x-axis) and a muon+tau component (y-axis). The CC measurement, sensitive only to electron flavour neutrinos is thus a vertical band, whereas the ES measurement has a slope due to its sensitivity to muon and tau flavour neutrinos. The combination of both measurements (using Super-K's high-precision result) allows the total flux to be determined, and is shown here to agree very well with the standard solar model prediction (dotted lines). Figures taken from [14].

neutrons from radioactive backgrounds.

The reconstructed energy threshold was brought down to $5\,\text{MeV}$ in order to increase the amount of signal from the $6.25\,\text{MeV}$ γ-rays from neutron captures on deuterium. The measured fluxes from this first NC analysis were:

$$\phi_{SNO}^{CC}(\nu_e) = 1.76_{-0.05}^{+0.06}(\text{stat.})_{-0.09}^{+0.09}(\text{sys.}) \times 10^6 \ \text{cm}^{-2}\text{s}^{-1}$$

$$\phi_{SNO}^{ES}(\nu_x) = 2.39_{-0.23}^{+0.24}(\text{stat.})_{-0.12}^{+0.12}(\text{sys.}) \times 10^6 \ \text{cm}^{-2}\text{s}^{-1}$$

$$\phi_{SNO}^{NC}(\nu_x) = 5.09_{-0.44}^{+0.44}(\text{stat.})_{-0.43}^{+0.46}(\text{sys.}) \times 10^6 \ \text{cm}^{-2}\text{s}^{-1}$$

$$(1.17)$$

which included the first direct measurement of the total solar neutrino flux, in excellent agreement with the solar model. The energy scale gave the largest contribution to the overall systematic uncertainty and the neutron detection efficiency gave the next largest uncertainty in the NC measurement. The results from the neutral-current analysis show that the non-electron flavour neutrino component of the flux is non-zero at the 5.3σ level, providing direct evidence that electron neutrinos from the Sun change flavour. This level of significance was strong enough to conclusively solve the solar neutrino problem.

Data from the D_2O phase were also used to investigate possible day-night variations in the neutrino fluxes [28]. A difference in day and night fluxes is expected from the MSW effect as the neutrinos will travel different path lengths in the Earth. Indeed, during the day, the solar neutrinos do not traverse the Earth to reach SNO, whereas at night, they may go through substantial amounts of matter. The MSW effect predicts that the survival probability (the proportion of electron flavour neutrinos, proportional to the CC flux) can be altered by the Earth depending on the value of the neutrino mixing parameters (the mixing angle and mass-squared difference). Detecting any (day-night) asymmetries in the electron flavour flux can test the MSW effect as well as constrain the solar neutrino mixing parameters. A combination of the measured asymmetry from the CC and ES fluxes (fixing the total neutrino flux asymmetry to be zero) gave the day-night asymmetry in the electron

flavour neutrino flux from the first phase, A_{D-N}, as [28]:

$$A^{CC}_{D-N} = 7\% \pm 4.9\%(\text{stat.})^{+1.3}_{-1.2}\%(\text{sys.}) \qquad (1.18)$$

Neutrino oscillations are the accepted mechanism to explain how the electron flavour neutrinos convert into the other flavours. Since the data from SNO confirmed the solar model, data from the other solar neutrino experiments were used to help constrain the solar neutrino survival probability, and thus the mixing parameters. Combining all the available data including the day-night asymmetries from SNO yielded 3σ ranges for the neutrino oscillation parameters [34]:

$$0.27 < \tan^2\theta_{12} < 0.92$$

$$2.4 \times 10^{-5}\text{eV}^2 < \Delta m^2_{21} < 4.7 \times 10^{-4}\text{eV}^2 \qquad (1.19)$$

which were significant enough to excluded regions of parameter space that had been allowed prior to this time. In particular this result disfavored the SMA ('small mixing angle') region of parameter space, for which there was a theoretical prejudice due to the small mixing angles in the quark sector [35]. The favoured region of parameter space has been termed the 'LMA', for 'large mixing angle'. A detailed account of the analysis of the first phase of SNO is presented in [29].

1.4.2   The Salt Phase

In the second phase of SNO, 2 tonnes of salt (NaCl) were dissolved into the heavy water to exploit the capture of neutrons on ^{35}Cl . This improved the ability of SNO to detect neutrons in three ways: first, the cross-section for capture on ^{35}Cl is approximately 40,000 times larger than it is on deuterium, second, the energy released by the subsequent decay of ^{36}Cl is greater (8.6 MeV versus 6.25 MeV) and finally, the 8.6 MeV decay energy is released in several γ-rays (2.5 on average) resulting in several Čerenkov electrons producing light that is more isotropic than for just one electron. The neutron detection efficiency was increased almost three-fold (from 14.4% to 40.7%).

The data from the salt phase was collected from July 26, 2001 to August 28, 2003 and comprised 391 days. An analysis of the data up to October 10, 2002 (254 live days) was published in [31], whereas the full data set, summarized here, was published in [32] with much more detail.

Data reduction was performed in a similar way to the first phase. A fiducial volume cut was applied at a radius of 5.5 m to limit backgrounds that are external to the heavy water and a cut was placed on the reconstructed energy of events at 5.5 MeV to limit contamination from low energy events from radioactive backgrounds. Applying cuts to remove 'instrumental events' (such as micro-discharges within the PMTs) as well as the cuts on the reconstructed observables resulted in a data set of 4,772 candidate neutrino events.

Again, an extended maximum-likelihood 'signal extraction' was performed to determine the number of CC, ES and NC events as well as neutrons from external (to the D_2O) backgrounds. In order to benefit from the increased isotropy of light emitted during neutron captures on ^{35}Cl , a light 'isotropy parameter', β_{14} , was introduced as an observable in the fit:

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

$$\beta_l = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} P_l\left(\cos\theta_{ij}\right) \tag{1.20}$$

where θ_{ij} is the mean pair angle, as viewed from the reconstructed event position, between the N PMTs in an event and P_l is the Legendre polynomial of order l. A larger value of β_{14} corresponds to a less isotropic pattern of hit PMTs.

The neutron response of the detector was again calibrated using a ^{252}Cf fission source and checked with an AmBe neutron source. Both sources were deployed throughout the heavy water volume to characterize the neutron detection efficiency as a function of position.

The maximum-likelihood fit was performed using four observables: reconstructed energy (T_{eff}), reconstructed direction with respect to the Sun ($\cos\theta_\odot$), reconstructed radial

position ($\rho = (R/R_{AV})^3$), and event isotropy (β_{14}). A key difference from the first phase is that multi-dimensional probability density functions were used to account for the correlation of β_{14} with T_{eff} and ρ. In addition, the increased sensitivity to neutrons allowed the use of ρ to constrain the neutron backgrounds from radioactivity produced outside the vessel (by looking for neutron captures near the acrylic, which are distributed differently than the uniform NC neutrons). Figure 1.7 shows the resulting fits in the four observables.

The analysis was carried out in two ways; the first fit constrained the CC energy spectrum to the distribution expected for ^{8}B and the second fit determined the values of the CC signal in 17 energy bins, effectively measuring the energy spectrum of electron flavour neutrinos. The unconstrained fit determined that the data set contained 2,176(78)CC events, 279(26) ES events, 2,010(85) NC events and 128(42) external-source neutrons. Converted into neutrino fluxes, these become:

$$\phi^{CC}_{SNO}(\nu_e) = 1.68^{+0.06}_{-0.06}(\text{stat.})^{+0.08}_{-0.09}(\text{sys.}) \times 10^6 \ \text{cm}^{-2}\text{s}^{-1}$$

$$\phi^{ES}_{SNO}(\nu_x) = 2.35^{+0.22}_{-0.22}(\text{stat.})^{+0.15}_{-0.15}(\text{sys.}) \times 10^6 \ \text{cm}^{-2}\text{s}^{-1}$$

$$\phi^{NC}_{SNO}(\nu_x) = 4.94^{+0.21}_{-0.21}(\text{stat.})^{+0.38}_{-0.34}(\text{sys.}) \times 10^6 \ \text{cm}^{-2}\text{s}^{-1}$$

$$(1.21)$$

The CC/NC ratio, which is proportional to the survival probability, was determined to be:

$$\frac{\phi^{CC}_{SNO}(\nu_e)}{\phi^{NC}_{SNO}(\nu_x)} = 0.34 \pm 0.023(\text{stat.})^{+0.029}_{-0.031}(\text{sys.}) \qquad (1.22)$$

The day-night asymmetry in the CC flux from the un-constrained fluxes was determined to be:

$$A^{CC}_{D-N} = -3.7\% \pm 6.3\%(\text{stat.})3.2\%(\text{sys.}) \qquad (1.23)$$

when the asymmetry in the NC flux is fixed to zero. This data can then be combined with the analysis from SNO's first phase and the other solar neutrino experiments to obtain

limits on the solar neutrino oscillation mixing parameters:

$$\tan^2 \theta_{12} = 0.45^{+0.09}_{-0.08}$$

$$\theta_{12} = 33.9^{+1.8}_{-1.6} \text{ degrees}$$

$$\Delta m^2_{21} = 6.5^{+4.4}_{-2.3} \times 10^{-5} \text{eV}^2 \tag{1.24}$$

which are significantly better determined than they were in the first phase. Results from the KamLAND experiment [36], sensitive to oscillations of anti-neutrinos, can be combined with the data from the solar neutrino experiments (assuming CPT invariance) to further restrict the allowed ranges for the mixing parameters:

$$\theta_{12} = 33.9^{+1.6}_{-1.6} \text{ degrees}$$

$$\Delta m^2_{21} = 8.0^{+0.4}_{-0.3} \times 10^{-5} \text{eV}^2 \tag{1.25}$$

where the precision in the mixing angle has resulted in the literature beginning to quote it in degrees. It is worthy to note that these results significantly ($\sim 7\sigma$) exclude 'maximal mixing', where the mixing angle would be equal to 45 degrees. Figure 1.8 summarizes the results from the salt phase.

1.4.3   The NCD Phase

The third phase of the experiment [37, 38] collected data from November 27, 2004 until November 29, 2006 and consisted of 385 live days. The salt from the previous phase of the experiment had been removed. Thirty-six 'Neutral Current Detectors' (NCDs) were deployed to count the neutrons produced in the heavy water independently from the data collected by the PMTs. The NCDs were strings of cylindrical proportional counters filled with ^{3}He gas which detected neutrons through:

$$\text{n} +^3\text{He} \rightarrow^3\text{H} + \text{p} + 764 \text{ keV} \tag{1.26}$$

The proton (p) and triton (^{3}H) ionize the gas, and the resulting energy deposition (764 keV if both deposit all of their energy) is collected on an anode wire. The current on the anode

(a) Fit for CC energy spectrum, T_{eff}

(b) Fit in direction to sun, $\cos\theta_{\odot}$

(c) Fit in radial position, ρ

(d) Fit in β_{14}
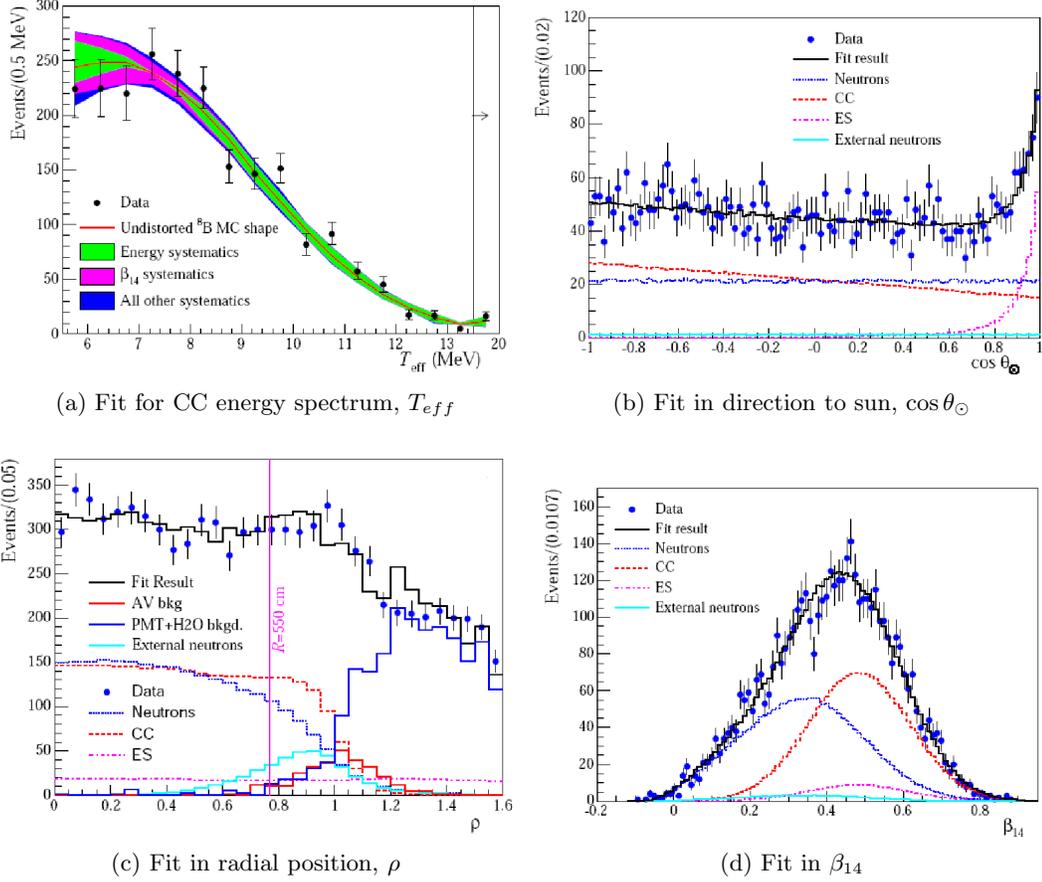
Figure 1.7: Fits over the observables in the salt phase. Panel (a) shows the measured spectrum of CC events and was determined by breaking up the CC flux into 17 energy bins and determining the contribution in each bin. The color bands show the contributions to the systematic uncertainties, whereas the statistical uncertainties are shown on each point. Panel (b) shows the fit in the direction of events with respect to the Sun and places a strong constraint on the amount of ES signal in the data. Panel (c) shows the fit in the radial position of events. The increase in neutron detection efficiency allows ρ to be used to constrain the number of external neutrons. Panel (d) shows the β_{14} distribution of the data and it is clear that this observable provides the ability to distinguish CC and NC events, as can be seen by their different probability distribution. The pdfs for all signal in each observable were determined by Monte Carlo simulation and constrained with calibration data. Figures taken from [32]

(a) Total fluxes

(b) SNO-only neutrino oscillation parameter space
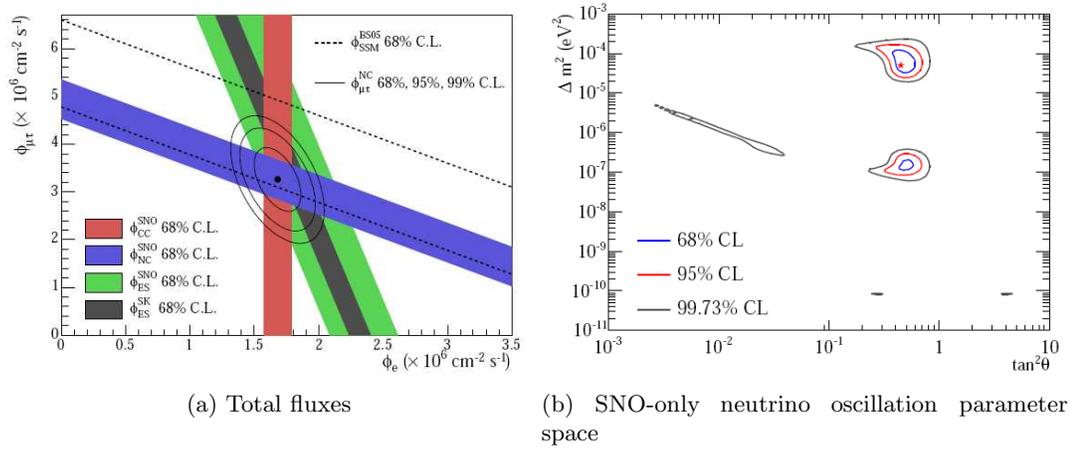
Figure 1.8: Panel (a) shows a summary of the fluxes measured by SNO in the salt phase, as well as Super-K's measurement of the ES reaction. It is again clear that this precision, model independent, measurement of the total solar neutrino flux (blue band, nc signal) is in good agreement with the standard solar model prediction (dotted band). The width of each band is determined by the uncertainties. Panel (b) shows the allowed region in neutrino oscillation parameter space when all of the SNO data is included (first and second phase, day-night asymmetries). Including all neutrino oscillation data then only allows the region in the upper-right to persist at more than 3σ, the so-called 'LMA' solution. Figures taken from [32]

wire as a function of time was read out as a pulse by digitizing scopes. The integral of the charge deposited on the anode for each event was recorded independently and was proportional to the energy deposited in the counter gas. The array of counters is described in detail in chapter 2.

The main background to the neutron signal came from ionizing alpha particles from the decay chains of ^{238}U and ^{232}Th contained within the counter materials as well as from ^{210}Po deposited on the inner surfaces of the counters. In order to characterize the alpha background, four counters filled with ^{4}He gas (insensitive to neutrons) were also deployed. In addition to backgrounds from 'physics events', the data contained a number of instrumental events that were rejected based on their pulse shapes. Six of the NCD strings were rejected from their analysis for reasons discussed in section 2.6.4.

One motivation for the third phase of the experiment was to perform an independent measurement of the neutrons in SNO. The resulting determination of the NC flux would have different systematic uncertainties, providing an independent confirmation of the SSM prediction.

The analysis of the third phase differed from previous phases in several ways. The biggest difference was the inclusion of data from the NCDs in addition to the PMT data. This was handled by constructing a likelihood function that was the sum of a likelihood of the PMT data and a likelihood of the NCD data. Both parts of the likelihood contained the NC signal in common. Another major difference was the use of a Markov-Chain Monte Carlo (MCMC) method [39, 40] (instead of minimization algorithms) to perform the maximum-likelihood analysis. This was preferred due to the large number of parameters, which included parameters to model systematic uncertainties (such as the PMT energy scale).

The PMT likelihood was built using T_{eff} , $\cos\theta_\odot$ and ρ; the same three observables that were used in the first phase. However, the pdfs were built as 3-dimensional functions, to avoid any potential biases from the factorization into three 1-dimensional pdfs. The

data on the PMT side had a 6 MeV threshold applied to it, in order to significantly reduce backgrounds to the neutrino signals. The NCD likelihood was built using the deposited energy of events. The probability density function (pdf) for neutrons was obtained from calibration data, whereas the alpha energy pdfs were derived from a Monte Carlo simulation. In addition to neutrons and alphas, two different instrumental backgrounds were identified by an excess of low energy events in the data from two of the strings. Although these strings were removed from the analysis, energy pdfs for these events were parametrized with analytical functions to allow for the presence of the same type of events in the data. In order to minimize the influence from these events, the NCD data were restricted to events above a threshold of 400 keV. An upper energy cut was placed to keep only events below 1.4 MeV. This allowed 93% of neutrons to be kept in the data, since they peak at 0.764 MeV.

The response of the NCD array to neutrons was measured by dissolving an activated ^{24}Na solution uniformly into the heavy water. The ^{24}Na decay releases 2.754 MeV. γ-rays that can photodisintegrate deuterium and thus result in a uniform neutron source. The capture efficiency for neutrons by the array was determined to be 21.1%. The presence of the NCD array capturing most of the neutrons significantly reduced the possibility for the PMTs to detect neutrons, and the PMT neutron detection efficiency was brought down to 4.9%.

The MCMC method was used to sample the likelihood function in parameter space and obtain posterior distributions for each parameter. The parameter set included not only the signals (neutrino fluxes, backgrounds) but also the systematic uncertainties in several quantities. For example, the energy scale for the PMT energy was floated in the analysis, subject to the constraints obtained from calibration data. This is a philosophically different approach to systematic uncertainties than was employed in the previous phases where data were fit, then re-fit with systematic parameters shifted by their 1-σ values. The 'floating systematics' approach adopted in this phase allowed the data to constrain the systematic parameters beyond the bounds provided by the calibration data.

The joint analysis of the PMT and NCD data yielded $983.4^{+77.0}_{-75.5}$ NC events in the NCD array and $267.0^{+23.8}_{-22.3}$ NC events in the PMT array. It was determined that $1,867.1^{+90.6}_{-100.9}$ CC, $171^{+24.3}_{-22.3}$ ES and $77.0^{+11.6}_{-10.2}$ background neutrons were observed in the PMTs. It was also determined that $185.1^{+24.8}_{-22.3}$ background neutrons were detected in the NCDs along with $6,126.6 \pm 100.6$ alpha+instrumental events. The results of the fits in the PMT observables and the NCD energy are shown in Figure 1.9. The neutrino fluxes from the data in the NCD phase were determined to be:

$$\phi^{CC}_{SNO}(\nu_e) = 1.67^{+0.05}_{-0.04}(\text{stat.})^{+0.07}_{-0.08}(\text{sys.}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$$

$$\phi^{ES}_{SNO}(\nu_x) = 1.77^{+0.24}_{-0.21}(\text{stat.})^{+0.09}_{-0.10}(\text{sys.}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$$

$$\phi^{NC}_{SNO}(\nu_x) = 5.54^{+0.33}_{-0.31}(\text{stat.})^{+0.36}_{-0.34}(\text{sys.}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$$

$$(1.27)$$

and the CC/NC ratio is:

$$\frac{\phi^{CC}_{SNO}(\nu_e)}{\phi^{NC}_{SNO}(\nu_x)} = 0.301 \pm 0.033(\text{total}) \tag{1.28}$$

Including all available solar neutrino data as well as the data from the KamLAND experiment resulted in the best constraints of the mixing parameters to date:

$$\theta_{12} = 33.8^{+1.4}_{-1.3} \text{ degrees}$$

$$\Delta m^2_{21} = 7.94^{+0.42}_{-0.26} \times 10^{-5}\text{eV}^2 \tag{1.29}$$

(a) Fit in NCD energy

(b) Fit in PMT energy

(c) Fit in direction to sun

(d) Fit in radial position
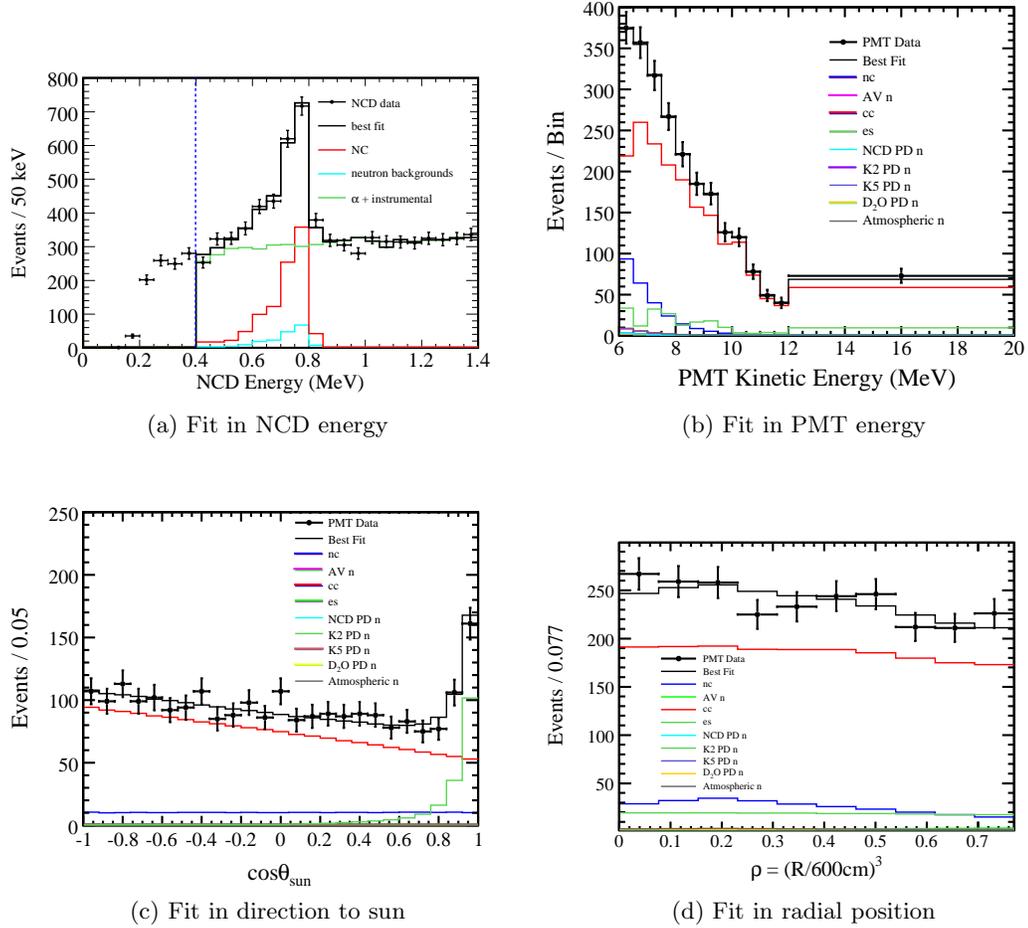
Figure 1.9: Summary of results from NCD phase. Panel (a) shows the fit to the deposited energy of events in the NCD counters. The red pdf corresponds to neutrons and was determined by calibration data. Panel (b) shows the fit in PMT energy, where the data is dominated by CC events. Panel (c) shows the fit in $\cos\theta_\odot$ and panel (d) shows the fit in ρ within the fiducial volume.

1.4.4 SNO's Contribution to Physics

Although SNO has unquestionably solved the solar neutrino problem, it has achieved much more than this initial goal. Indeed, thanks to the SNO experiment, there is now undeniable evidence that electron neutrinos can change flavour. In addition, the uncertainty in the solar neutrino mixing angle, θ_{12}, is approaching 1 degree. This has ushered the community into an area of precision neutrino measurements. Just 10 years ago, the flux of solar neutrinos was uncertain, but SNO has now measured it more accurately than the prediction from the solar model as well as placed stringent constraint on neutrino mixing parameters. The combination of SNO results with KamLAND has also led to an accurate knowledge of the solar neutrino mass-squared difference. The demonstration that solar electron neutrinos change flavour, the resolution of the solar neutrino problem (by confirming the prediction from the solar models) and the accurate measurement of solar neutrino mixing parameters are without doubt the biggest contributions from SNO.

SNO data has been used to set an upper limit [30] in the flux of neutrinos from the 'hep' reaction in the Sun (the highest energy neutrinos in Figure 1.2). The hep flux was determined to be less than $2.3 \times 10^4 \text{cm}^{-2}\text{s}^{-1}$ at 90% confidence level using the data from the first phase (an analysis combining the data from all three phases is underway). Uncertainties in the predicted hep flux are dominated by nuclear matrix element calculations and thus, a measurement of the hep neutrino flux helps to constrain these calculations as well as the standard solar model.

A search for electron anti-neutrinos in the SNO data were performed [41]. These can be detected efficiently as their CC interaction with deuterium (4.03 MeV threshold) releases two neutrons and a positron in coincidence. Searches for anti-neutrinos can limit the possibility for mechanisms that could convert neutrinos into anti-neutrinos. Results provided a 90% confidence upper limit indicating that less than 0.81% of the ^{8}B neutrinos convert into antineutrinos.

The SNO detector was also used to search for nucleon decays [42], which may be signs of grand-unification theories. In particular, the data from SNO was used to search for 'vanishing' neutrons or protons inside ^{16}O nuclei. This can happen if protons or neutrons decay to 'invisible' particles (such as neutrinos), leaving an excited state of the nucleus which subsequently decays with the release of a γ-ray above 6 MeV. SNO's high γ-ray detection efficiency allowed the limits on nucleon lifetimes to be improved to $\tau > 2 \times 10^{29}$ years at 90% confidence level.

The solar neutrino flux has a 1-year period modulation due to the eccentricity of the Earth's orbit. However, an independent analysis [43] of the Super-K data had suggested that the ^{8}B flux also contained a modulation of the same period (9.42^{-1} years) as harmonic of the Sun's rotation period. Since the rotation of the Sun does not influence nuclear reaction rates, a change in neutrino flux with that period would presumably be the signature of a flavour-changing interaction with the solar magnetic field ('new' physics) [44]. A re-analysis of the Super-K data by the collaboration [45] as well as a search for periodicities in the SNO data [46], found no modulation except for that expected from the Earth's orbit.

Finally, the SNO experiment was also used to measure the flux of atmospheric muons through the detector [47]. High-energy muons (>100 MeV release much more Čerenkov radiation in SNO than electrons and are easily detectable. Muons are produced in the atmosphere when cosmic rays induce showers. These 'cosmic ray' muons are seen in SNO as coming from above, with small zenith angles. The cosmic ray showers also produce muon neutrinos (from muon decays), which can reach SNO from all zenith angles (they can originate from atmospheric showers on the other side of the Earth). These muon neutrinos can then interact with matter near the detector and produce a muon that will be seen by SNO. A measure of the number of muons in SNO can thus serve to measure both the muon flux from the atmosphere (by looking at small zenith angles) as well as the muon neutrino flux. Figure 1.10 shows SNO's measured muon flux as a function of zenith angle. Although SNO is much less sensitive to atmospheric neutrinos than Super-K, the flux as a function

of zenith angle favours neutrino oscillations (black boxes versus hashed boxes).
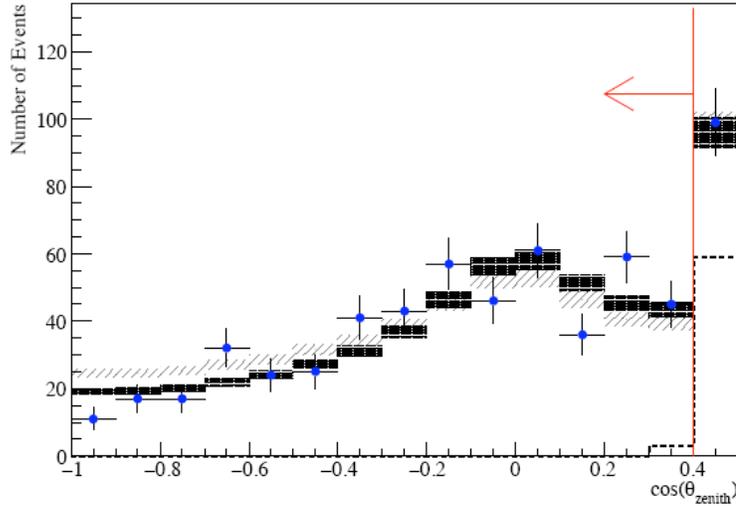
Figure 1.10: Flux of muons at SNO as a function of zenith angle. Zenith angles with $\cos\theta > 0$ correspond to events above the horizon. The dashed line at small zenith angles ($\cos\theta > 0.3$) corresponds to the events from the cosmic ray muons, whereas all other events are from the neutrino-induced flux. The hashed boxes show the prediction for a model with no neutrino oscillations compared to the black boxes for the case where neutrino oscillation take place. The red arrow shows the region below which the muon flux is dominated by neutrino-induced events and the dotted line shows the flux of muons coming directly from the atmosphere. Figure from [47].

1.5 What Can (Still) be Learned from Solar Neutrinos

The previous section highlighted some of the physics results from SNO to date. This section aims to motivate the benefits of a more accurate measurement of the solar neutrino flux and parameters involved in solar neutrino oscillations as well as how this will be accomplished in the upcoming analysis of SNO data (including this thesis).

Recent years have seen the emergence of a 'solar metallicity problem'. Latest measurements of solar metallicities (the abundances of element heavier than helium) have shown lower levels than previously measured [48]. The inclusion of these lower abundances into solar models has led to significant discrepancies with helioseismological measurements which

cannot (to date) be explained by deficiencies in the models. It has been proposed [48] that solar neutrinos can be used to distinguish between high and low metallicity models, as the heavy element abundances have an impact on neutrino fluxes. Previous solar model calculations [3, 15, 4] had large uncertainties ($\sim$18%) on the neutrino fluxes to account for different levels of metallacities. Recent years have seen the emergence [48] of 'high' and 'low' abundance solar model which have smaller uncertainties ($\sim$11%) in the neutrino fluxes. Although the difference in solar neutrino fluxes between these models is highest for CNO neutrinos (44%), the predicted ^{8}B neutrino flux differs by 21% between models so that SNO data can make a significant contribution to this problem.

Through the measurement of the CC/NC ratio, SNO has a direct sensitivity to the survival probability, and hence the θ_{12} mixing angle, which determines the amplitude of the oscillations (equation 1.6). The KamLAND experiment [36] measures the energy spectrum of electron-flavour anti-neutrinos of energies of a few MeV with baselines of order $200\,\mathrm{km}$. Assuming CPT invariance, both experiments are sensitive to the same set of neutrino oscillation mixing parameters, namely the solar oscillation mixing parameters. Since KamLAND measures the energy spectra of neutrinos it is most sensitive to the frequency of the oscillations and hence the mass-squared difference. The combination of solar neutrino experiments and KamLAND, provides constraints on both mixing parameters. Both types of experiment have a slight sensitivity to θ_{13}. Recent results [22] have shown that the best fit values of θ_{12} between the solar (driven by SNO) and KamLAND experiments are slightly different (at levels of low statistical significance), and the difference in this value of θ_{12} can in fact be accommodated by a non-zero value of θ_{13}. Figure 1.11 shows the 'pull' in the best fit θ_{12} between the solar and KamLAND results along with the inferred value of θ_{13}. Thus, a more accurate measurement of θ_{12} from solar experiments could lead to a more robust determination of θ_{13}. It is of important interest to measure this last remaining mixing angle, as it can lead to the measurement of CP-violation in the lepton sector, which could in turn shed some light on the leptogenesis process and whether it can help explain

the matter/anti-matter asymmetry in the universe.
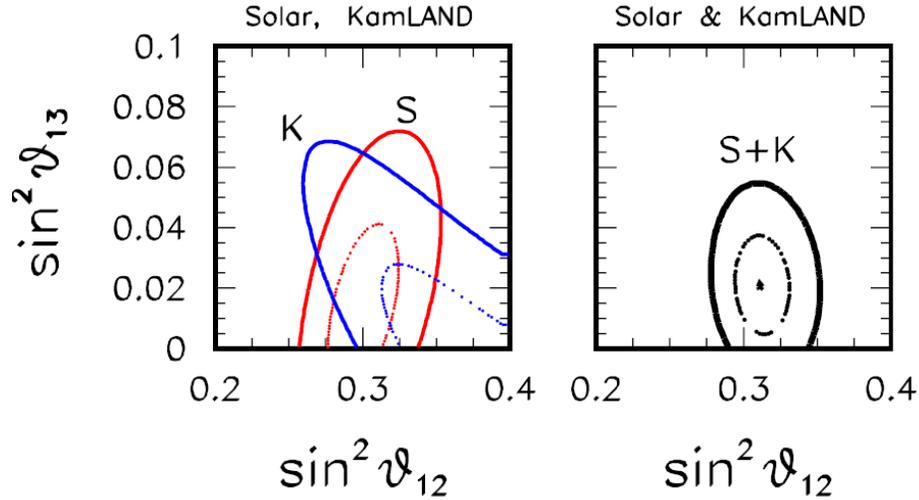
Figure 1.11: Allowed regions of θ_{12}-θ_{13} for solar (red) and KamLAND (blue) data (panel a). The slightly different best fit value of θ_{12} result in a non-zero best-fit value of θ_{13} when the two data sets are combined (panel b). Contour lines are 1σ, the figure was taken from [22].

The energy-dependence of the solar survival probability (Figure 1.4) is not measured to a very good accuracy over the the full energy range. SNO has measured the spectrum at high ($>5\,\text{MeV}$) energies, and the radio-chemical experiments have measured some points at the lower energies. Experiments [49] are starting to probe the 'transition' region, around $2\,\text{MeV}$ where the survival probability goes from its vacuum value to a matter-dominated value. An accurate measurement of the energy spectrum in that region can help to distinguish between the MSW model and other proposals, such as 'mass-varying' neutrinos [50] or 'non-standard-interactions', where, for example, a neutrino magnetic moment alters the MSW interaction with solar matter [51]. Figure 1.12 shows the survival probability that is predicted for certain models as well as the data points that have been taken from various experiments to date. A 'low-energy-threshold-analysis' (leta) of the SNO data is currently underway and aims to extract the energy spectrum down to $4\,\text{MeV}$ using data from the first two phases. In addition to constraining the shape of the survival probability, this analysis will significantly

increase the amount of neutrino data and result in a much more accurate determination of

the mixing angle.
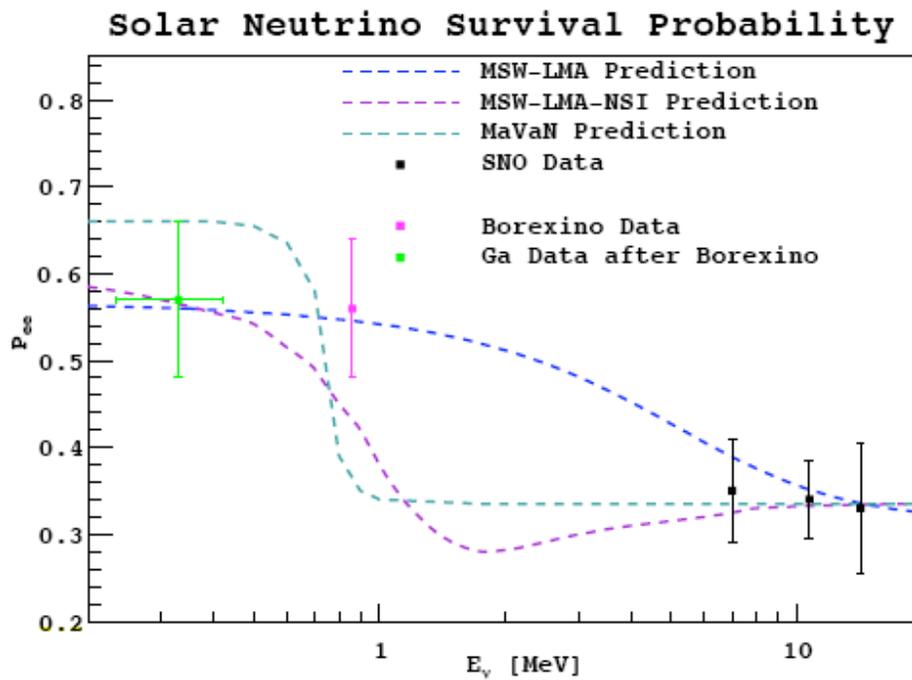
Figure 1.12: The solar neutrino survival probability in the transition region for different models along with data from solar neutrino experiments. Measuring the energy spectrum in the transition region can help to distinguish between models. Figure was taken from [52].

Finally, accurate measurements of θ_{12} can place constraints on theoretical models. There are neutrino mixing models that predict the values of the neutrino mixing angles. Although 'bi-maximal mixing' ($\theta_{12} = \theta_{23} = \frac{\pi}{4}$) is now excluded by SNO, 'tri-bi-maximal' ($\theta_{12} = 35.3$, $\theta_{23} = \frac{\pi}{4}$, $\theta_{13} = 0$)) models are still allowed. Other predictions such as 'quark-lepton' complementarity make use of symmetries to combine neutrino and quark mixing angles to sum to $\frac{\pi}{4}$. Although it is beyond the scope of this thesis to discuss these models, it is clear that an accurate determination of θ_{12} is valuable. For an overview of these neutrino mixing models, see [53].

Future results from the SNO data will include the low energy threshold analysis of the first two phases. In addition, the collaboration is currently working on a '3-phase' analysis

that will include the data from all three phases (as well as the work pursued in this thesis) and will result in the most accurate determination of the solar mixing angle and total ^{8}B solar neutrino flux. The 3-phase analysis will also set new limits on the hep flux. Other experiments that will look at low energy solar neutrinos (in the MSW transition region) are currently underway or will start in the near term. These include Borexino [49] (already running), KamLAND and SNO+ [54], which all use liquid scintillator to look for low energy solar neutrinos. The SNO+ experiment will replace the heavy water in SNO with liquid scintillator.

1.6 Overview of the work in this thesis

This thesis presents an independent analysis of the data from the proportional counters (NCDs) used in the third phase of the SNO experiment. In particular, the ultimate goal of this work is to derive an accurate number of neutrons detected by the counters during the NCD-phase of SNO. An accurate determination of the neutrons detected by the NCD array is directly related to the accuracy in the inferred neutral-current flux of solar neutrinos from the third phase. In addition, a more accurate determination of the neutral-current flux has a direct impact on the precision in the measurement of θ_{12}.

The previous analysis of the NCD phase only used the energy deposited in the counters to fit for neutrons on an alpha background (see panel (a) of Figure 1.9). This thesis will introduce a new variable, 'NoverA', which is based on the shape of scope traces of events in the counters and derived from the comparison with libraries of calibration neutron and alpha pulse shapes. Distributions for the NoverA variable are determined for neutrons and alphas so that this new observable can be used in conjunction with energy in a maximum-likelihood analysis to decrease the uncertainty in the inferred number of neutrons.

Chapter 2 provides a detailed description of the neutral-current detector array and associated electronics. The various calibrations that were performed in order to characterize the

response of the array are also described. Chapter 3 is intended to provide an understanding of the physics behind events in the NCDs. In particular, it motivates why neutron and alpha pulse shapes are expected to look different, so that they may be used in an analysis to differentiate the two types of events. The chapter concludes with a brief description of a Monte Carlo simulation that was designed to model the NCD array.

The original work from this thesis begins in chapter 4, where the 'Queen's Grid Fitter' (QGF) is introduced. QGF uses a library of calibration neutron pulse shapes and a library of calibration alpha pulse shapes. The NoverA variable for NCD events is introduced here and is the result of the comparison of data pulses to these libraries. The NoverA variable is a measure of how a data pulse compares to the two libraries and indicates how neutron-like a given pulse is. This chapter is aimed at motivating the fine details in how the NoverA observable for each event is determined.

Chapter 5 characterizes the probability density functions (pdfs) for neutrons and alphas for the energy and NoverA observables. The distributions and systematic uncertainties for neutrons are determined using calibration data. The pdfs for alphas are determined using the Monte Carlo simulation and the systematic uncertainties are derived from comparison of the Monte Carlo simulation with alpha data from the ^{4}He counters. In addition to neutrons, two additional instrumental-type of events are characterized. Distributions for these events are also introduced and a novel method for handling their systematic uncertainties is presented.

Chapter 6 gives a brief (mathematical) overview of the extended maximum-likelihood technique that will be used to analyze the NCD data using the energy and NoverA observables. Methods for handling systematic uncertainties are presented as well as the Markov-Chain Monte Carlo method that is used in the analysis.

Chapter 7 presents tests of the use of energy and NoverA to determine the number of neutrons. The first part of this chapter presents extensive analysis of Monte Carlo generated data to test for any biases in the determination of the neutron number. The second part of

the chapter presents an analysis of a subset of the real NCD data using the techniques of this thesis. This subset of the data were designed for a 'blind' analysis, so that the techniques presented here are not tuned to 'get the right answer', since the data were already published.

Finally, chapter 8 presents an analysis of one third of the data from the NCD phase using the methods presented in this thesis. The result from this new analysis is compared to the published result. Chapter 9 presents some concluding remarks.

Chapter 2

Neutral Current Detector Characterization

This chapter characterizes the Neutral Current Detector (NCD, hereafter) array and its associated data acquisition system. The array was comprised of 36 (^{3}He +CF$_4$) cylindrical proportional counters sensitive to neutrons and four (^{4}He +CF$_4$) counters (insensitive to neutrons) to measure backgrounds. It was installed in the heavy water after the NaCl was removed, between November 2003 and April 2004, with a commissioning phase until November 2004. The neutrino data were collected from November 27 2004 to November 28 2006.

2.1 Overview of NCD Array

This section describes the physical characteristics of the proportional counters and array. The design was motivated by the need for low radioactive backgrounds in order to detect the ~13 neutral-current-induced neutrons that are produced in the D$_2$O each day.

2.1.1 Neutron Detection Reaction

Thermal neutrons were detected in the ^{3}He -filled cylindrical proportional counters (Figure 2.2) via the reaction:

$$n + {}^3\mathrm{He} \to {}^3\mathrm{H} + p + 764 \ \mathrm{keV} \tag{2.1}$$

with a cross section of 5333 b [33] (compared to 0.5 mb for neutron capture on D_2O and 44 b for capture on ^{35}Cl). The proton and triton are created nearly back-to-back with energies of 573 keV and 191 keV, respectively. They lose their energy by ionizing the gas (the full 764 keV is only deposited providing neither particle collides with the counter wall); the electrons are then collected onto an anode wire. The high electric field at the anode wire causes an avalanche of secondary ionization with a gain of about 220 secondaries per primary electron. The resulting pulse was detected in the data-acquisition system (section 2.2) and the charge in the pulse is proportional to the ionization energy deposited in the gas. The energy spectrum for neutrons from a distributed ^{24}Na calibration source (section 2.4.3) is shown in Figure 2.1; the peak at 764 keV corresponds to the proton and triton depositing their full energy in the counter. The peak is broadened by events where either the proton or triton collides with the wall. The lowest energy events (191 keV) correspond to cases where the proton is immediately absorbed by the wall and only the triton energy is deposited.

2.1.2 Counter Design

This section gives a brief description of the NCD string design; a detailed description can be found in the technical paper [37]. Nickel was chosen as the main material for the counters and was purified (by chemical-vapor-deposition onto aluminum, CVD) to 10^{-12} g/g of ^{238}U and ^{232}Th .

The proportional counters were 2-3 m long with a 5.08 cm outer-diameter and a wall thickness between 305 μm and 533 μm (370 μm nominal). The wall thickness was dictated
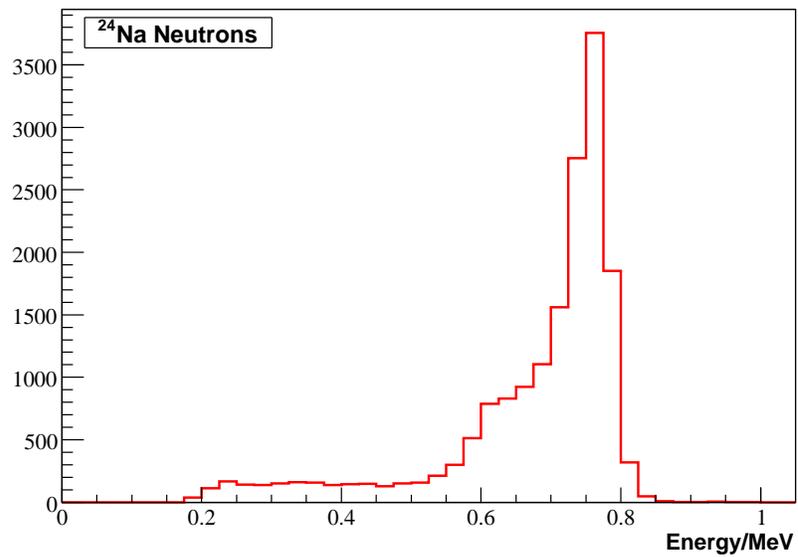
Figure 2.1: Neutron spectrum from a distributed ^{24}Na source calibration (see section 2.4). The peak at 764 keV corresponds to the proton and triton depositing their full energy. The shoulders in the spectrum are due to the proton and triton colliding with the NCD wall and failing to deposit their full energies.

by the need to maintain structural integrity in the heavy water under pressures of 3.2 atm at the bottom of the acrylic vessel. The gas mixture was 85% Helium (^{3}He or ^{4}He) and 15% CF_4 (by pressure) held at 2.5 atm. The anodes were made of 50 μm diameter copper wires.

The cylinders had CVD-nickel endcaps welded 4.3 cm from the end and allowed counters to be electrically connected to each other into strings. A fused-silica feed-through tube insulated the anode wire from the endcap and the counter body. The feed-through tube extended between 2.5 cm and 2.8 cm into the gas volume to create an avalanche-free region where the electric field is otherwise distorted (affecting the gain). The proportional counters were assembled in the underground lab[1] into groups of three or four to make NCD-strings between 9 m and and 11 m in length. A schematic of a string and the connection between counters is shown in Figure 2.2 (a and b).

The anode wires were terminated by open-ended delay lines at the bottom of the strings. The delay lines extended the NCD-strings by about 45 ns , creating a 90 ns delay between current exiting the top of the anode wire and its reflection from the bottom. The delay line was a flat meander circuit etched onto a copper-clad Kapton sheet which was then wrapped around an acrylic core positioned at the bottom of the strings (see Figure 2.2a). The open-ended delay-line was designed for crude longitudinal event position determination, as well as allowing for the full signal charge to be collected (thus improving signal-to-noise characteristics).

The top part of the anode wire was connected to the NCD readout cables by a resistive coupler. The readout cables were responsible for bringing the high-voltage to the detectors as well as collecting the data. The impedance of the detectors is 415 Ω which required a 325 Ω resistive coupler to match the 93 Ω impedance of the readout cables. The readout cable protruded through a standard CVD-nickel endcap (similar to the counter endcaps, except with a larger hole to allow the cable through) into a gap similar to the one between

---

[1]Due to the headroom above the detector, counters had a maximum length of 3 m and were welded together as they were inserted into the heavy water.

counters. The resistive coupler was made of a Teflon cylinder with wire strung back and forth around one side. The wire was connected to a CVD-nickel toothed-ring that slid onto the fused-silica feed-through of the NCD anode wire on one side and the cable connector on the other (see Figure 2.2c). This method of connecting the cables to the detector failed intermittently for some of the strings (and completely for at least one string).

Readout Cable

Nickel Counter Body

Gas Fill
Pinch-Off Tube

^{3}He-CF$_4$ Gas Fill

Anode Wire

Fused Silica Insulator

Delay Line Termination

Vectran Braid

Acrylic ROV Ball

Acrylic Anchor Ball

9 – 11 m

(a) NCD string schematic

Top Counter

Anode Wire

CVD Ni Coupler

Laser Weld

CVD Ni

Spring Ears

1.3 cm

0.5 cm dia.

Endcap

Laser-Welded Lap Joint

^{4}He-Air Endcap Gas

Copper Endplug

Endcap with Pinch-Off Tube

Silica Feedthrough

^{3}He-CF$_4$ Counter Gas

Anode Wire

4.3 cm

0.7 cm

4.3 cm

Bottom Counter

(b) Connection between two counters

Readout Cable

Potting Compound

Acrylic Fill Spacer

Endcap

Laser-Welded Lap Joint

Resistive Coupler

Endcap

Anode Wire

Low-Radioactivity Heat Shrink

9.3 cm

5.6 cm

4.3 cm

Pinch-Off Tube

5.1 cm

(c) NCD and readout cable connection

Figure 2.2: Panel (a) shows the string assembly with the open-ended delay-line on the bottom and the readout cable on the top. The Vectran braid at the bottom held two acrylic spheres; one anchored the string to the acrylic vessel in a specially designed slot and the other was used by a remotely operated vehicle (ROV) to install the string. Panel (b) shows the junction between two counters with the endcaps, the nickel coupler and fused-silica feed-throughs. Panel (c) shows a closeup view of the resistive coupling between the NCD and the readout cable. Figures are taken from [37].

2.1.3 Array Layout

The deployed array consisted of 36 ^{3}He -filled and 4 ^{4}He -filled strings in a 1 m square grid with cylindrical symmetry. The layout is shown in Figure 2.3 and was designed to maximize the neutron detection efficiency while minimizing the shadowing of Čerenkov light to the PMT array. Two-dimensional position reconstruction of neutron events is straightforwardly achieved with this layout. The grid spacing resulted in a 21% capture efficiency for uniformly produced neutrons (section 2.6.2)

The strings were labeled by two different coordinate systems; in one case, each string was assigned a letter to indicate its radial position (see Figure 2.3) and a number to indicate position in the ring, in the other labeling system, strings were simply numbered 0 to 39 (see Table 2.1 for the conversion between labeling systems). This thesis will predominantly identify strings by their number.
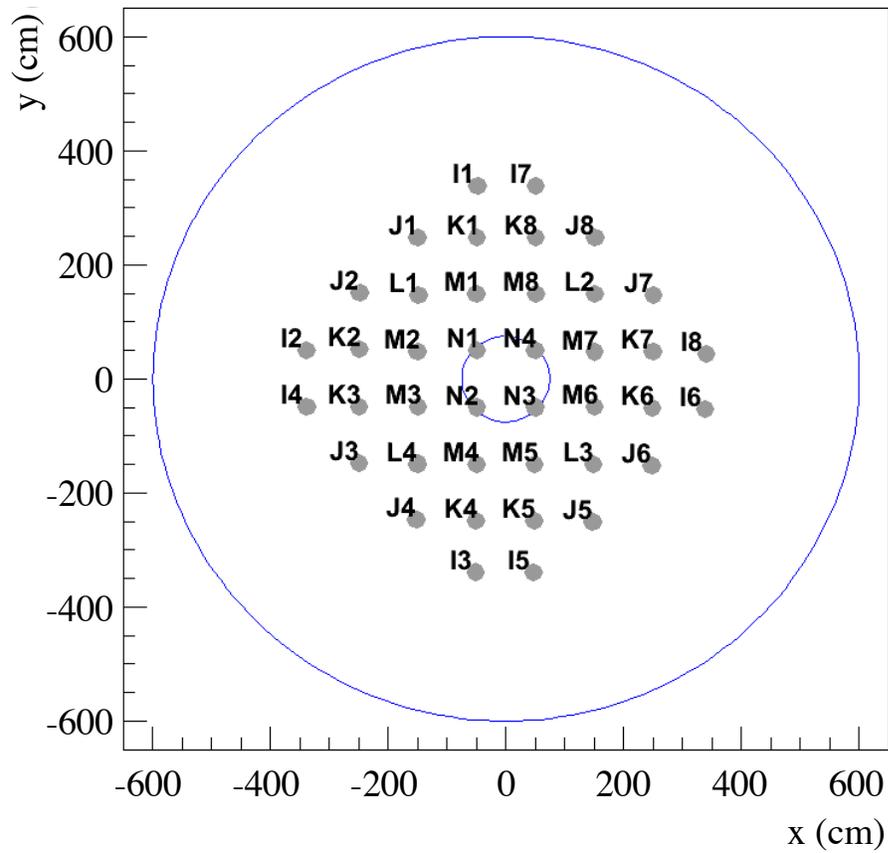
Figure 2.3: NCD array layout projected into the horizontal plane. The outer circle is the equator of the acrylic vessel and the inner circle is the projection of the neck of the vessel. The strings are installed in concentric circles on a square 1 m grid and each ring is labeled by a letter (N for the innermost rings and I for the outermost). Within a ring, different strings are identified by a number. Figure taken from [37]

String Number	Label	String Number	Label
0	N4	20*	I3
1	M8	21	K4
2	K8	22	M4
3*	I7	23	J4
4	J8	24	L4
5	L2	25	N2
6	J7	26	J3
7	M7	27	M3
8	K7	28	K3
9	I8	29	I4
10*	I6	30*	I2
11	K6	31	K2
12	M6	32	J2
13	J6	33	M2
14	N3	34	L1
15	L3	35	J1
16	J5	36	I1
17	M5	37	K1
18	K5	38	M1
19	I5	39	N1

Table 2.1: Conversion between NCD string number and ring-label. The four strings with an asterisk contain the counters filled with ^{4}He .

2.2 Electronics and NCD Data Acquisition

This section describes the flow of data from the readout cable at the top of the NCD and characterizes the various components in the electronics chain. The description is limited to the NCD-side of the electronics and data-acquisition (DAQ, hereafter); the PMT-side is discussed in [25]. A more detailed overview of the NCD electronics can be found in [37] and [55].

The design of the NCD-DAQ is motivated by the desire to be able to discriminate neutron events from background alpha events (from the ^{238}U and ^{232}Th chains) by digitizing the waveforms collected in the counters as well as the ability to collect events from a potential galactic supernova at rates of a few kHz. For this reasons, two different DAQ paths were implemented; the first allowed for the waveform digitization by oscilloscopes (at low event rates, for solar neutrino analysis using particle identification) and the second used custom-designed shaping-ADC boards to record event energies (at kHz event rates, suitable for supernova events).

2.2.1 Electronics Layout

Figure 2.4 shows a diagram of the electronics layout with the fast (shaper) and slow (scopes) DAQ paths. The NCDs were connected to a trans-resistance pre-amplifier which was responsible for the readout (toward the multiplexers), injection of calibration pulses (from the pulser distribution system) and the high-voltage supply of the proportional counters. Pulses from the NCDs were then sent from the multiplexers to the shapers (fast path) and the digitizing scopes (slow path). The two DAQ paths interacted with the SNO Master Trigger Card (MTC)[25], where they were assigned event identification numbers, before being recorded by a Mac computer running OSX and the ORCA (Object-oriented Real-time Control and Acquisition) software which allowed for data acquisition as well as control of the electronics. The following sub-sections will briefly describe the most important components

of the DAQ.
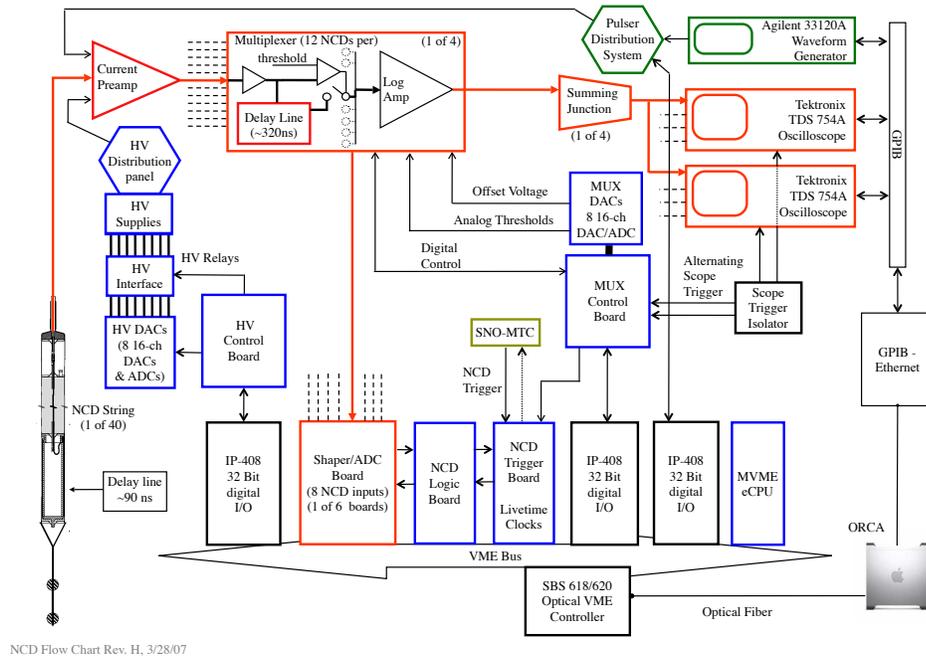
Figure 2.4: Layout of the NCD electronics and data-acquisition systems. Figure taken from [37].

2.2.2 PreAmp

The trans-resistance pre-amplifier was chosen over a charge-amplifier in order to preserve the time profile of the charge collected by the detectors. A two-stage gain was provided by four 2SK152 JFETs followed by an AD8055 operational amplifier and met the design requirement of a 3kHz-45MHz pass-band (determined by the sharpest of physics pulses). The pre-amplifiers (one per NCD) were used to connect the HV supply to the detectors as well as being the injection points for pulses from a waveform generator into the electronics chain (for electronics calibration). The output from the pre-amplifier, $V_{out}(t)$ can be modeled [56] by a high-pass RC circuit:

$$V_{out}(t) = e^{-t/RC} \int_0^t \left[\frac{dV_{in}(\tau)}{d\tau} \right] e^{\tau/RC} d\tau \tag{2.2}$$

where $V_{in}(t)$ is the pulse at the input of the pre-amplifier.

2.2.3 Multiplexers

Four multiplexer boxes (Muxes) were used to send the signals to the digitizing scope readout. The signal from the current pre-amplifiers entered two parallel buffer amplifiers that divided the signal towards the shaper and scope channels. The muxes contained an LT1016 comparator to trigger the digitizing system path. When the signal exceeded the discriminator level it was sent to an AD8307 logarithmic amplifier (logamp) (section 2.2.4). Signals from 12 NCD strings were summed before entering the logamp. The normal trigger rate during data-taking was low enough (0.3 Hz for the entire array) to make it unlikely that more than one event entered the logamp at any time.

The Muxes were controlled by the Mux Control Board (MCB) (itself controlled by the ORCA computer (section 2.2.7)) which is responsible for several functions. The MCB set the discriminator thresholds through an 8-bit digital-to-analog (DAC) card with a resolution on the order of one mili-Volt. The DAC card contained one 8-bit register for each NCD, so that their discriminator thresholds could be controlled individually (they varied by a

few milivolts between NCDs). The MCB was also responsible for sending a 'Mux-Trigger' to the NCD triggering system when the discriminators fired and recorded the pattern of strings that exceeded their threshold in all four Muxes. The MCB triggered the digitizing scopes and controlled the readout of the channel corresponding to the Mux that initiated the trigger. A 320 ns delay line was built into the Muxes to allow the digitizing scopes the time to trigger after the thresholds were exceeded.

2.2.4 Logarithmic Amplifier

The logamps [57] were used in order to extend the range of signals that could be digitized accurately, as well as maintain an amplitude-independent noise in the pulses. The range was dictated by the possible physics-events that could take place in the NCDs, ranging from an inward-going 191 keV triton produced at the inner wall (200 nA current) to an 8.8 MeV ^{212}Po alpha traveling parallel to the wire (30 μA current). This corresponds to a dynamic range of approximately 150. The effect of the logamps on the input pulses was parametrized [56] as:

$$V_{log}(t) = a \times \log\left(1 + \frac{V_{out}(t + \Delta t)}{b}\right) + c \qquad (2.3)$$

where V_{out}(t) is the pulse exiting the pre-amplifiers. The parameters a,b,c and Δt were measured with calibration pulses (described in section 2.3.2). Even though there were only four Muxes (and hence logamps), these parameters were measured for each string, since they each had different pre-amplifiers gains (parametrized by b) as well as for each scope path, totaling 80 different electronic chains and sets of parameters to be determined for each calibration.

2.2.5 Digitizing Scopes

Two Tektronix TDS754A digitizing scopes with four input channels (each) were used to record the pulse waveforms from the NCDs. Each input channel was connected to one

of the four Muxes. The MCB controlled which scope was triggered for the digitization, alternating between scopes when neither was 'busy' or switching to whichever scope was free. Upon triggering, the scope recorded on all four input channels, but the GPIB only read out the channel corresponding to the triggered Mux(es). The signals were digitized at $1\,$GHz and $15\,\mu$s of data were recorded for each Mux-trigger event. In order to obtain a baseline for the pulses, $1500\,$ns before the oscilloscope triggering was recorded in addition to $320\,$ns of the event before the mux-threshold was crossed (to account for the mux delay line). Using the two scopes, the NCD DAQ was able to digitize data at $1.8\,$Hz , which was sufficient for solar neutrino data.

2.2.6   Shapers/ADC

In order to record NCD data at higher rates than in the digitizing path, six eight-channel custom ADC/shaper boards residing in a VME crate were designed. This allowed for the rapid, almost dead-time-less, recording of NCD-pulse charge, and hence event energy. The NCD pulse sent from the Mux was first integrated by four successive operational amplifiers and then divided into three paths; the first path went to an ADC converter (12-bit MAX120CWG), the second path to a differentiating amplifier and the third to a threshold discriminator. The signal from the differentiating amplifier clocked a D flip-flop switch whose D-input was connected to the discriminator. The D flip-flop output then triggered the ADC converter to digitize the integral as well as send a lock-out signal after a $180\,$ns waiting period, so that all coincident channels could be recorded. The digitized event was then read by an embedded CPU (eCPU in the VME crate) and the NCD trigger board was notified to send a 'Shaper Trigger' to the SNO Master Trigger Card.

2.2.7   The ORCA Computer

The Object-oriented Real-time Control and Acquisition (ORCA) software was used to control the hardware thresholds and record the data from the shapers/ADC and scopes. ORCA

was run on a Mac G4 dual-processor computer running OS X. A PCI-to-VME card was used to communicate with the shaper/ADC eCPU, the NCD trigger card, the HV supplies, the pulser distribution system (section 2.3.1) and the Mux Control Board. The data from the scopes was recorded using an Ethernet-to-GPIB adapter that was also used to control the waveform generator of the pulser distribution system (section 2.3.1).

2.2.8 The NCD Triggering System

NCD events were assigned a Global Trigger Identification (GTID) by the NCD Trigger Card (NCD-TC). A register of GTIDs in the NCD-TC was kept synchronized with the SNO Master Trigger Card (MTC), so that NCD and PMT events could be time-ordered by their GTIDs. The MTC recorded event times using (redundant) 10 MHz and 50 MHz clocks.

2.2.9 The NCD Data Stream

The raw data from ORCA for the NCD events (time, channel, trigger type) were written to ZEBRA [58] Data Acquisition/Analysis Bank (ZDAB) files by the 'ncdbuilder' process running on a Mac G5 computer. This was done in parallel to a similar process for the PMT data. A second-level process, 'snobuilder', then merged the two data streams according to event GTIDs into one ZDAB file. The snobuilder process was also responsible for forwarding the events to a network dispatcher connected to the various DAQ monitoring tools.

The data were collected in 'runs' that were approximately 7 hours in length for neutrino data. Each run corresponded to a ZDAB file with the data from both PMTs and NCDs. Calibration runs would typically last less time to keep the size of the files manageable.

2.3 NCD Electronics Calibration

This section briefly describes the calibrations that were regularly performed on the electronics chain.

2.3.1 The Pulser Distribution System

The Pulser Distribution System (PDS) was designed to inject pulses at the NCD pre-amplifier in order to test the gain, linearity, dead-time, and thresholds of the shaper/ADC and scope chains. An Agilent 33120A Waveform Generator was used to create pulses with known shapes to perform these calibrations. In addition to weekly calibrations, a custom pulser was also used to randomly trigger the electronics (during data-taking) in order to measure their associated dead-time at an average rate of $0.01\,\mathrm{Hz}$.

2.3.2 Calibration of the Logarithmic Amplifier

In order to obtain the pulse that was generated from the ionization in the NCDs, the waveforms recorded by the scopes needed to be deconvolved. The main distortion on the pulses was caused by the logarithmic amplification and it was thus important to accurately measure the parameters related to this process. During the data-taking period, 'logamp calibrations' were performed once a week, in which a sine-wave pulse was injected into the electronics and the recorded waveform was then fit using equation 2.3 and the corresponding parameters were extracted for each channel and scope. These parameters were then recorded in SNO data banks so that waveforms could be 'de-logged' during analysis.

2.3.3 Threshold, Gain and Linearity Calibrations

In addition to measuring the effect of the electronics on the pulse shapes, calibrations were also done to monitor the thresholds of the Mux and shaper/ADC paths, as well as the gain and linearity of the shaper/ADCs. The thresholds of both data paths were measured by sending sine-wave pulses of decreasing amplitude and determining when each channel would stop triggering. The varying sine-wave amplitudes tested both the amplitude-threshold of the digitizing path as well as the total charge-threshold of the shaper/ADC channels.

The gain and linearity of the shaper/ADC channels were measured by injecting square-wave pulses of different amplitudes into each shaper/ADC channel. The known amount of injected charge was then compared with the digitized value to monitor the gain and linearity of each shaper channel. Together with neutron source calibrations, one could then set an absolute energy/ADC conversion for the events occurring in the NCDs.

2.3.4 Random Pulser

In addition to the weekly calibrations mentioned above, the PDS was also used to randomly trigger the electronics (during data-taking) in order to measure their associated dead-time at an average rate of 0.01 Hz on a dedicated channel (different than the channels used by the NCDs).

2.4 Neutron Source Calibrations

In order to measure the response of the NCD array to neutrons, several neutron sources described in this section were deployed during the data-taking phase. A manipulator system (described in [25] [37]) allowed for the three-dimensional positioning of various encapsulated calibration sources throughout the heavy water volume.

2.4.1 AmBe Source

Two encapsulated Americium-Beryllium (AmBe) sources, differing by their neutron production rate, were used. The neutrons are produced by the (α,n) reaction of alpha-particles from ^{241}Am on a ^{7}Be target. These sources were deployed monthly and used to monitor the stability of the NCD array. In particular, the AmBe source was used to calibrate the ADC energy spectrum. A summary of the source properties is shown in Table 2.2

Source	Type	Half-life	Source Strength
AmBe-Med	encapsulated	432 yr	23.63(27) neutrons/second
AmBe-Hi	encapsulated	432 yr	68.70(74) neutrons/second

Table 2.2: Properties of the SNO AmBe (medium and high rate) sources. The rate was measured using the PMT array and includes contributions from deuterium photo-disintegration.

2.4.2 Encapsulated ^{252}Cf Source

A ^{252}Cf fission-neutron source encapsulated in an acrylic puck and fixed to a stainless-steel mounting was used in all three phases of SNO to measure neutron detection efficiency. ^{252}Cf can decay by either α-emission or neutron fission. The californium source was used primarily due to the possibility of accurately measuring its absolute neutron production rate [32]. Properties of the source are summarized in Table 2.3.

Source	Type	Half-life	Neutron Multiplicity	Source Strength
^{252}Cf	encapsulated	2.645 yr	3.7676(47)	16.55(08) neutrons/second

Table 2.3: Properties of the SNO ^{252}Cf source. The source strength is a weighted-mean of several measurements [32] and adjusted to June 12, 2001.

2.4.3 ^{24}Na Distributed Source

In order to measure the response of the detector to a uniform source of neutrons, the heavy water was 'spiked' twice with an activated ^{24}Na brine. ^{24}Na beta-decays to ^{24}Mg with the emission of two prompt γ-rays with energies of 2.75 MeV and 1.37 MeV. The 2.75 MeV γ-ray has enough energy to photo-disintegrate deuterium (threshold of 2.2 MeV), effectively producing a uniform source of neutrons, and happens approximately once for every 385 ^{24}Na decays. The ^{24}Na spikes were performed in October 2005 and October 2006. The strength of the ^{24}Na (about 1 neutron detected per second in the array at the beginning of the spike) was measured before the spike using different methods [59], including an encapsulated sample that was counted with the PMT array and compared to ^{252}Cf to set

the absolute rate, as well as with a germanium detector. The ^{24}Na data formed the basis
for measuring the neutron detection efficiency of the NCD array ([38],[59]).

2.5   Sources of Non-Neutron Backgrounds in the NCDs

This section describes sources of events in the NCDs that are not related to neutron capture.
These include alpha-emitters (from ^{238}U and ^{232}Th chains) and spurious electronic events.
There is also the possibility of β and γ decays depositing enough energy in an NCD to
trigger an event.

2.5.1   Alpha Emitters

Low-levels of ^{238}U and ^{232}Th in the bulk of the NCD nickel and anode-wire result in alpha
events in the NCDs. The alpha-particles from decays in the bulk of the metal have their
energy degraded (as a function of the depth of the emitter) and result in a relatively flat
energy spectrum. A more significant source of alpha events in the NCDs comes from the
plating of ^{210}Po (5.304 MeV alpha-decay) on the NCD-wall surface after exposure to air
contaminated with ^{222}Rn . Figure 2.5 shows the NCD spectrum for the entire array over
the NCD phase, the peak from the ^{210}Po decay at 5 MeV and the peak from neutrons at
764 keV are clearly visible. The apparent position of the ^{210}Po peak is below 5.3 MeV, which
is due to space charge effects (section 3.1.4). The rate of alphas in the neutron energy region
(155 to 800 keV) has been estimated [60] to be 16 $\pm$ 1 per day in the 36 ^{3}He-filled strings.

2.5.2   Beta and Gamma Events

Electrons can also ionize the gas and produce events in the NCDs. These β-particles can
come from the decays of elements in the ^{238}U and ^{232}Th chain present in the NCD materials
as well as being the result of Compton scattering of γ-rays from the same chains. Electrons
need to travel long distances in the gas to deposit enough energy and the phase space
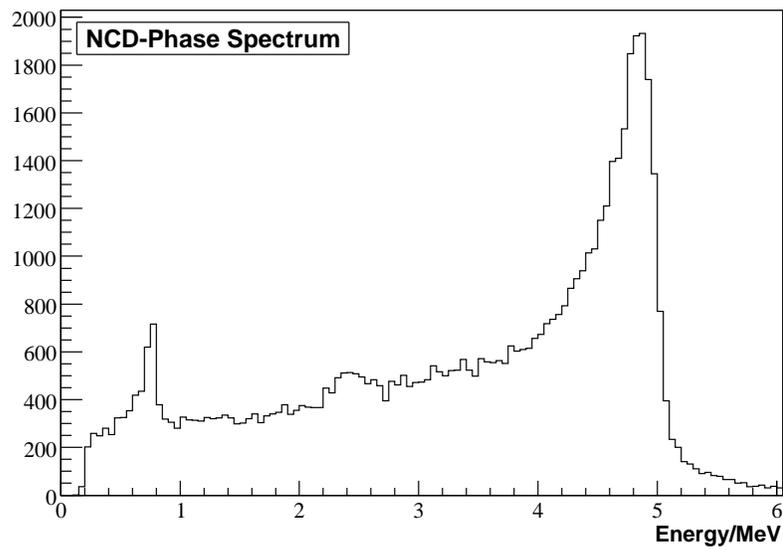
Figure 2.5: NCD spectrum from 0 to 6 MeV. One can see the neutron peak at 764 keV over a background that is primarily due to ^{210}Po . The ^{210}Po peak is shifted down from 5.3 MeV due to space charge effects (section 3.1.4).

for them to do so is thus limited (as they need to travel parallel to the wire while being most likely emitted from the NCD walls). The NCD data has not shown any signs of such contribution in the energy spectrum shown in Figure 2.5

β-particles produced in the gas are more likely to deposit their full energy (from purely geometric considerations). Since the ^{3}He was obtained from tritium decay, care was taken to ensure that the contamination of ^{3}H in the NCDs was small. The end-point energy of the tritium electron is 18.6 keV which is well below the threshold for analysis (191 keV), however, there is the possibility of pileup events (multiple events in one string with a total energy above threshold). As the gas mixture for the NCD was purified, the tritium contamination was reduced by a factor of almost 400 to 2.7 nCi per liter of gas at STP, which resulted in a negligible background.

2.5.3 Electronic Events and Low-Level Data Cleaning

The NCD DAQ also triggered on events that were not the result of ionization in the gas. Such events were generated by thermal noise, micro-discharges, electronic pick-up and HV breakdown in the PMT and NCD systems. Events that did not trigger both the digitizing and shaper/ADC paths were discarded in the analysis. More complex cuts based on the pulse shapes were designed to remove events that had pulse shapes that were inconsistent with ionization events. Pulse-shape cuts were designed in both the the time and frequency domains to remove specific events. The sets of cuts in either domain were seen to remove almost identical fractions of the data. Figure 2.6 shows a class of electronic pulses attributable to discharges in the gas (panel a) as well as the shaper spectrum after successive data-cleaning cuts are applied (panel b). In particular, one notes that the raw data is dominated by a low-energy peak which has been seen to arise mostly from one string (string 27) that appeared to have a much larger rate of 'fork' events (similar to the one in panel a). It was noted that some strings had particular classes of events or higher rates of certain event types (such as string 27 having a large amount of fork events). About half of the strings

in the array were virtually unaffected by data-cleaning cuts, and are sometimes referred to as 'super-clean' strings. The sacrifice of physical data (neutrons) due to the data-cleaning cuts was measured with neutron calibration data and seen to be less than 2%.

| (a) Discharge event | (b) Shaper/ADC spectrum |
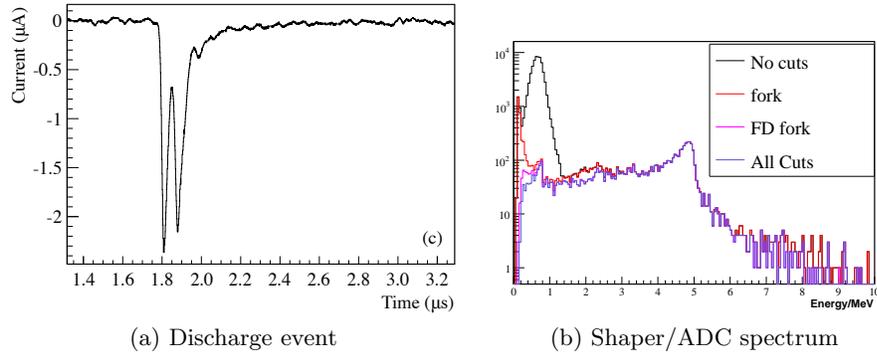
Figure 2.6: Panel a) shows a discharge event ('fork' event) that most likely took place inside an NCD. This class of events is cut relatively efficiently by cuts both in the time and frequency domains. Panel b) shows the shaper/ADC spectrum for a fraction of the NCD data after data-cleaning cuts have been applied. In this case, the raw data is shown in black, the data in red is after applying a time-domain based cut for removing fork events, the magenta line is after the addition of a frequency domain cut for removing fork events and the blue line shows the spectrum after all pulse-shape based data-cleaning cuts have been applied.

2.6 Detector Characterization

2.6.1 Stability of the NCD Electronics

The NCD electronic calibration constants (logamp parameters, thresholds, gain and linearity) were measured weekly during the data-taking phase. Additional calibrations were done during the commissioning-phase over smaller timescales (24 hours) and showed no significant variation of the parameters over such timescales. The calibration constants were saved

in weekly data-banks that were subsequently used to deconvolve the data. Figure 2.7 shows one example of calibration constants measured string-by-string as a function of time during the NCD phase. In this case, 'logamp parameter a' from equation 2.3 has been plotted for ten strings (labeled by their names) as a function of run number using scope 1.

2.6.2 Neutron Capture Efficiency

The neutron capture efficiency for the NCD array is defined as the fraction of neutrons produced uniformly in the heavy water that capture in the detector array. This was measured [62] by a solution of ^{24}Na dissolved and mixed into the detector (section 2.4.3). The capture efficiency, ϵ_{cap}^{NCD} was then determined using:

$$\epsilon_{cap}^{NCD} = f_{non-unif} \times f_{edge} \times \frac{R_{spike}}{A_{spike}} \tag{2.4}$$

where $f_{non-unif}$ is a correction factor to take into account the possible non-uniformity of the ^{24}Na brine and was measured by comparing PMT data to Monte Carlo simulations. The f_{edge} factor corresponds to a Monte Carlo correction that takes into account the possibility of the 2.75 MeV ^{24}Na γ-ray escaping the heavy water without producing a neutron (for ^{24}Na near the acrylic vessel). R_{spike} is the rate of neutrons captured by the NCDs, whereas A_{spike} is the neutron production rate from the brine (both rates corrected to same reference time). Values for these parameters are shown in Table 2.4 along with the capture efficiency for the two ^{24}Na spikes (2005 and 2006) that were carried out. The data from the 2005 spike was used in the published analysis [38] and the 2006 values were used as a verification.

Year	$f_{non-unif}$	f_{edge}	R_{spike}	A_{spike}	ϵ_{cap}^{NCD}
2005	0.9812(21)	0.9702(78)	0.2708(34)n/s	1.240(20) n/s	0.2089(73)
2006	0.9812(21)	0.9702(78)	0.1811(24)n/s	0.838(13) n/s	0.2066(73)

Table 2.4: NCD neutron capture efficiency from ^{24}Na and input parameters. The Monte Carlo correction factors are taken to be the same for the two spikes and the rates are shown in neutrons per second. This table is reproduced from [62]
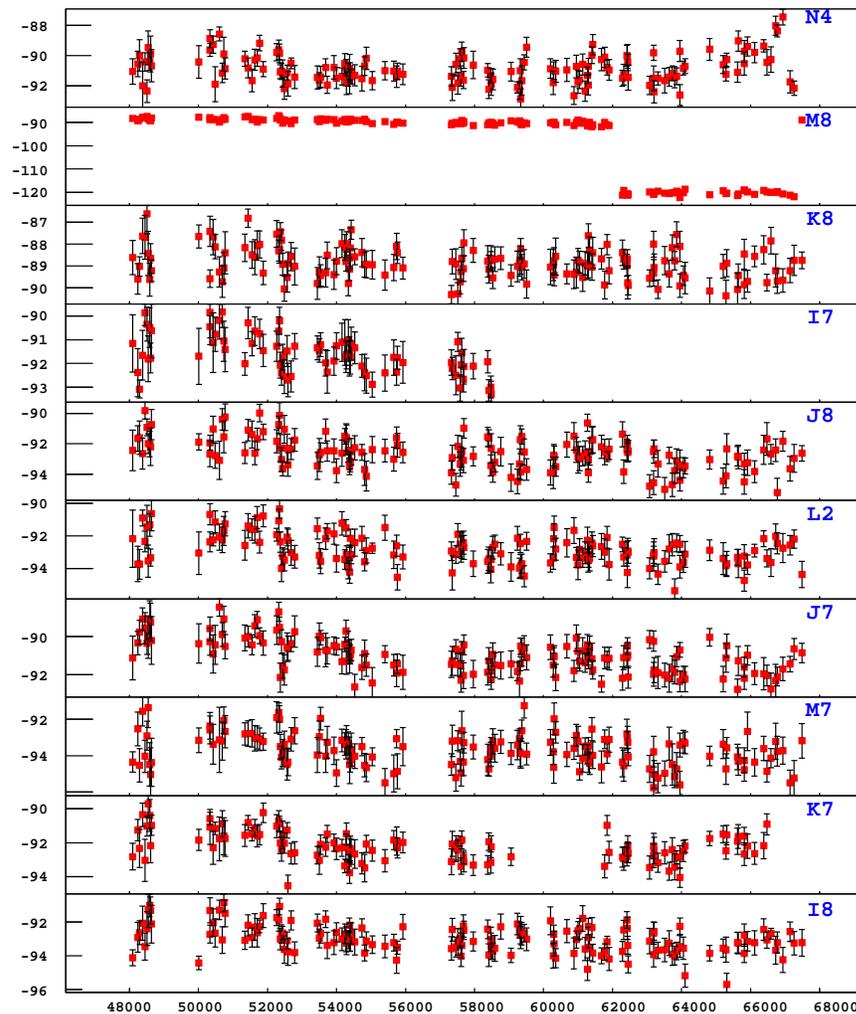
Figure 2.7: Logamp parameter a as a function of run number for ten strings using scope 1. String I7 was turned off half-way during data-taking. Data shown here is representative of the overall quality of electronics calibration. This figure was provided by [61].

2.6.3   Energy Calibration

The absolute scale of the NCD shaper energy spectrum was calibrated monthly using the AmBe sources. These were deployed using the 3-dimensional source manipulator system [25] to 'scan' the NCD array. The energy scales were then defined using the 764 keV neutron peak.

2.6.4   Problematic Strings

Although data-cleaning was shown to be successful at removing spurious events from the NCD data, six strings were singled out due to the poor quality of their data and the impact they would have on the uncertainties in the final analysis. Table 2.5 shows a summary of the strings that were removed from the analysis and the reason for their removal. Two of the strings (0 and 26) were removed due to the presence of unidentifiable events in the neutron window of the energy spectrum. Although these strings were removed, the spectra of their events were retained and used to look for similar events at lower rates in other strings (see section 5.3). String 18 was removed because one of the counters in the strings was observed to be leaking gas, significantly affecting its gain and energy resolution. It was observed that the neutron peak in string 8 shifted back-and-forth during the data-taking phase. Finally, string 1 and 31 had bad resistive coupler connections (Figure 2.2c), confirmed when the strings were removed from the heavy water at the end of the phase. String 1 completely disconnected from the DAQ after a seismic event in the mine and string 31 was seen to intermittently disconnect (for periods of hours), making it difficult to estimate its live-time.

2.7   Summary

This chapter characterized the array of ^{3}He proportional counters that were deployed in SNO to count neutrons. Analysis of the data from these counters is the primary focus of this thesis and motivated a detailed description of the electronics and calibrations that were

String	Reason for Removal
0 (N4)	Low-energy events in neutron window, sometimes in bursts, particular to this string
1 (M8)	Bad resistive coupler connection, difficult to estimate live-time
8 (K7)	Unstable gain, observable shift in neutron peak during calibrations
18 (K5)	Leaking counter, unstable gain and surface neutron source contaminant
26 (J3)	Strange, low-rate events in neutron window. Also, weak resistive coupler connection
31 (K2)	Bad resistive coupler connection, difficult to estimate live-time

Table 2.5: Summary of the ^{3}He-strings that were removed from the NCD phase analysis.

performed for these counters to be operated as effectively as possible.

Chapter 3

Neutral Current Detector Modeling

This chapter aims at characterizing the pulse shapes in the Neutral Current Detectors (NCDs) recorded by the digitizing oscilloscopes. The intention is to motivate the pulse shape analysis (PSA) techniques presented in this thesis and give the reader an understanding of their potential and limitations. The first section will describe the physical processes that generate the pulse shapes, then the characteristics of neutron and α pulse shapes will be examined; the chapter will finish by briefly describing how these pulse shapes are simulated in the SNO Monte Carlo simulation.

3.1   Pulse Shape Generation Processes

This section describes the physical parameters that lead to different pulse shapes in the NCDs. The pulse shape is a record of the amount of charge (electrons) that comes out of the anode wire as a function of time, and is thus determined by the location and orientation of the initial ionization track, the electron drift and avalanche in the gas, the effect of the ions and the electronics chain.

3.1.1 Ionization Energy Deposition

An event in an NCD begins when a charged particle (α-particle or proton + ^{3}H ion pair) leaves an ionization track in the gas. Energy deposition is distributed according to a Bragg curve. Figure 3.1 shows the energy-loss in the NCD gas-mixture (He + CF$_4$) for ^{210}Po alphas, protons and tritons of maximal energy as a function of energy calculated with SRIM [63, 64]. The energy-loss in each step is proportional to the amount of ionization created at that point and hence the amount of charge that is generated at each point on an ionization track. A full energy alpha will have a track that is about 2.5 cm long , a proton will have a track that is around 1.5 cm and the triton 0.3 cm .

(a) Proton

(b) Triton

(c) ^{210}Po α-particle
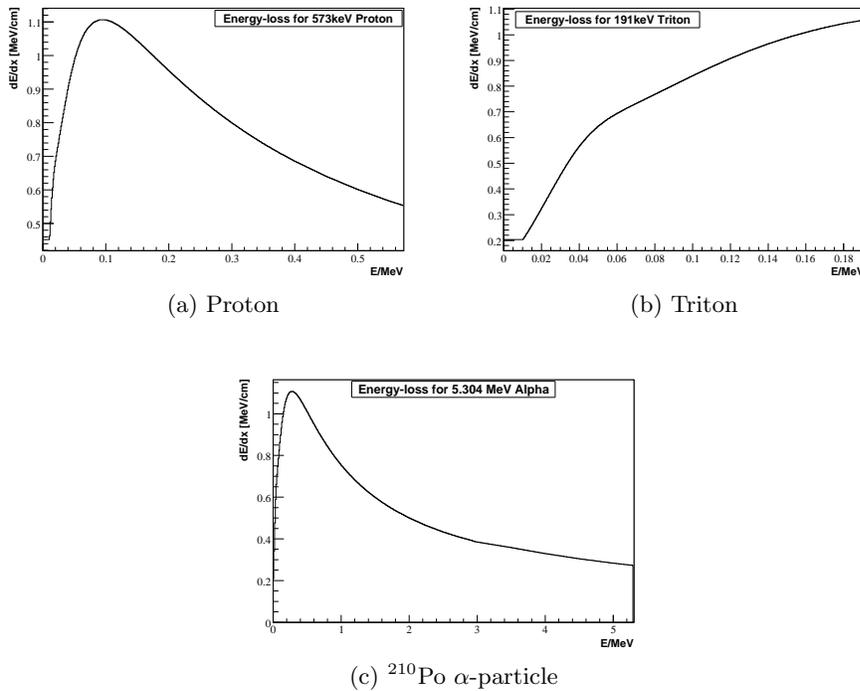
Figure 3.1: Bragg curves for 573 keV protons, 191 keV tritons and 5.304 MeV ^{210}Po α-particles. These show the energy loss in MeV per cm as a function of energy in the NCD gas-mixture (He + CF$_4$). The graphs were made using the SRIM [63, 64] simulation software.

3.1.2 Geometry

The proton and triton from neutron capture have energies of 573 keV and 191 keV respectively, whereas α-particles from the decay of ^{210}Po (the main source of α-particles) have a maximum energy of 5.304 MeV. Compared with the drift speeds of electrons in the NCDs (it takes about 3.5 μs for a electron to drift 2.5 cm from the wall to the anode) one can, to a very good approximation, assume that the track is created instantaneously.

The time-profile of the charge-arrival will then depend primarily on the track orientation and the distribution of charge along it (according to the Bragg curve). Tracks that are parallel to the anode-wire result in a narrow pulse, since all the primary ionization is the same distance to the anode and will all arrive there at the same time. Conversely, tracks that are directed towards the anode will result in the widest pulse shapes. The orientation can also make the primary ionizing particle hit the wall of the NCD and reduce the amount of charge deposited.

3.1.3 Electron Drift and Avalanche

The electrons created along the ionization track will drift and cascade when they reach the region of high electric field, close to the anode. The drift velocity depends on the local electric field and will increase when the electrons are close to the anode. Measurements [65] have shown that the drift time as a function of radius can be parametrized by:

$$t_{drift}(r) = 3.4\mu s \left(\frac{r}{2.54\,cm} \right)^{1.6} \tag{3.1}$$

where 3.4 μs is the maximum drift time (from the wall to the anode) and 2.54 cm is the radius of the counter. As the electrons arrive close to the anode (about 1 anode radius, 25 μm), they have enough energy to further ionize the gas and create an avalanche of secondary ionization. This effective gain was measured [65] to be about 220 at the operating voltage of 1950 V . For the SNO Monte Carlo simulation, the effective gas gain was obtained from string-by-string calibrations that were carried out monthly with the AmBe (section 2.4.1)

source.

3.1.4 Space Charge

Depending on the geometry of the track, space charge effects can alter the effective gas gain. This happens when an avalanche takes place in a region where a previous avalanche has already occurred; the ions from the first occurrence effectively reduce the local electric field thereby reducing the gain for the second avalanche. Space charge effects are the largest for ionization tracks that intersect and are perpendicular to the anode wire, since all the primary electrons will take the same path to the anode. These effects not only affect the overall gain in a pulse, but also the time-profile of the charge-arrival, since later parts of the avalanche will be influenced by the earlier parts. Because of the reduction in gain, space charge results in a decrease of the charge integral for a pulse (as seen in the downshift of the ^{210}Po peak in Figure 2.5). Space charge effects are notoriously difficult to simulate and a model with free (tunable) parameters was used in the official SNO Monte Carlo simulation.

3.1.5 Ion Drift

The electrons collected on the anode are attracted to the ion cloud created by the avalanches so that they cannot immediately leave the NCD. As the ions drift away (much slower than electron drift), the electrons on the anode are slowly allowed to leave the avalanche region, which results in an exponential-like decay in the charge-arrival profile, referred to as the 'ion-tail'. The ion-tail has a large impact on the pulse-shape and washes out some of the physics characteristics mentioned above. In the SNO Monte Carlo, the effect of the ion drift is simulated by the Wilkinson function [66].

3.1.6 Electronics Chain

Once the pulse is generated by the ionization event, it is modified substantially by the electronics. As the charge leaves the avalanche region, some of it travels downwards and

is reflected by the bottom of the NCDs. Propagation along the NCD wire can be modeled with a lossy transmission-line and low-pass RC filter. A slight impedance mismatch causes another reflection at the input of the current pre-amplifier.

The pre-amplifier can be modeled as a combination of low and high-pass RC filters and a gain of 27500 V/A. The chain from the pre-amplifier to the logamp can be modeled by a low-pass RC filter, the log-amp can be described using equation 2.3 and another low-pass RC filter can be used to model the final path to the scope.

3.2 Characteristics from Calibration Data

Using calibration data, one can determine the energy spectra of neutrons and alphas as well as pulse shape characteristics such as pulse width, amplitude and rise-time (the time on the leading edge of the pulse from 10% to 90% of the maximum amplitude). These data can then be used to verify the modelling of pulses. In addition to simply verifying a model, one can also motivate the use of a pulse shape-based analysis to discriminate between neutrons and alphas if these distributions are different.

3.2.1 Energy Spectrum Characterization

Figure 3.2 shows the energy spectra for neutrons from the 2005 ^{24}Na spike as well as the spectrum of alphas taken from the ^{4}He strings. It is clear that these distributions are both a good test of NCD pulse modeling as well as a good tool for extracting the neutron signal from the NCD data. Indeed, the energy spectra were used to determine the number of neutrons in the published analysis of the NCD phase [38].

(a) Neutrons Energy (b) Alphas Energy
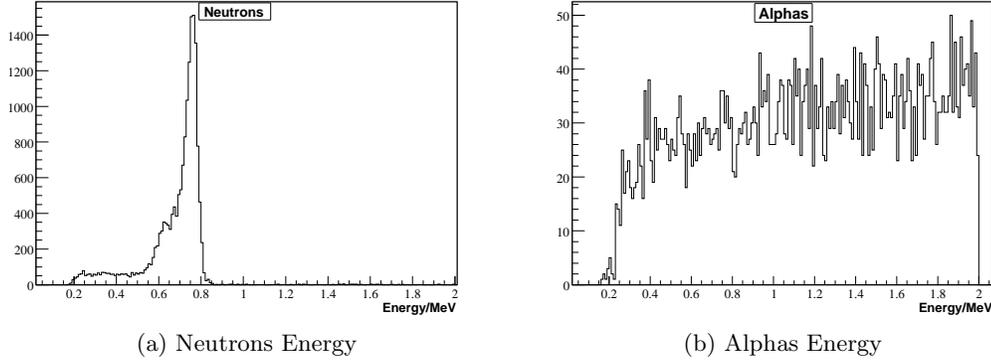
Figure 3.2: Energy spectra for neutrons and alphas. These distributions are useful for testing the modeling of NCD data as well as being the basis for extracting the neutron signal from the NCDs.

3.2.2 Pulse Shape Characterization

Figure 3.3 shows histograms of pulse rise-time, full-width-at-half-max (fwhm) and amplitude for neutrons taken from the 2005 ^{24}Na calibration and alpha particles taken from the ^{4}He-filled strings. It is clear that the average pulse shape between these particles is different and that this information can be used in addition to the energy spectra to extract the neutron signal from the data. The distributions are most different for the rise-time parameter which is also expected to contain the most information about the original ionization that created the pulse; most of the later part of the pulse is generated by the ions leaving the anode wire (the ion tail) and is less characteristic of whether the pulse was created by a neutron or alpha particle.

3.3 The NCD Monte Carlo

In order to better understand the data from the neutral-current detectors, extensive work was done to model the entire process of data acquisition in the NCD phase. A complete Monte Carlo simulation of the NCDs has been produced by the SNO collaboration [67]

(a) Neutron-Rise-time

(b) Alphas-Rise-time

(c) Neutrons-FWHM

(d) Alphas-FWHM

(e) Neutrons-Amplitude
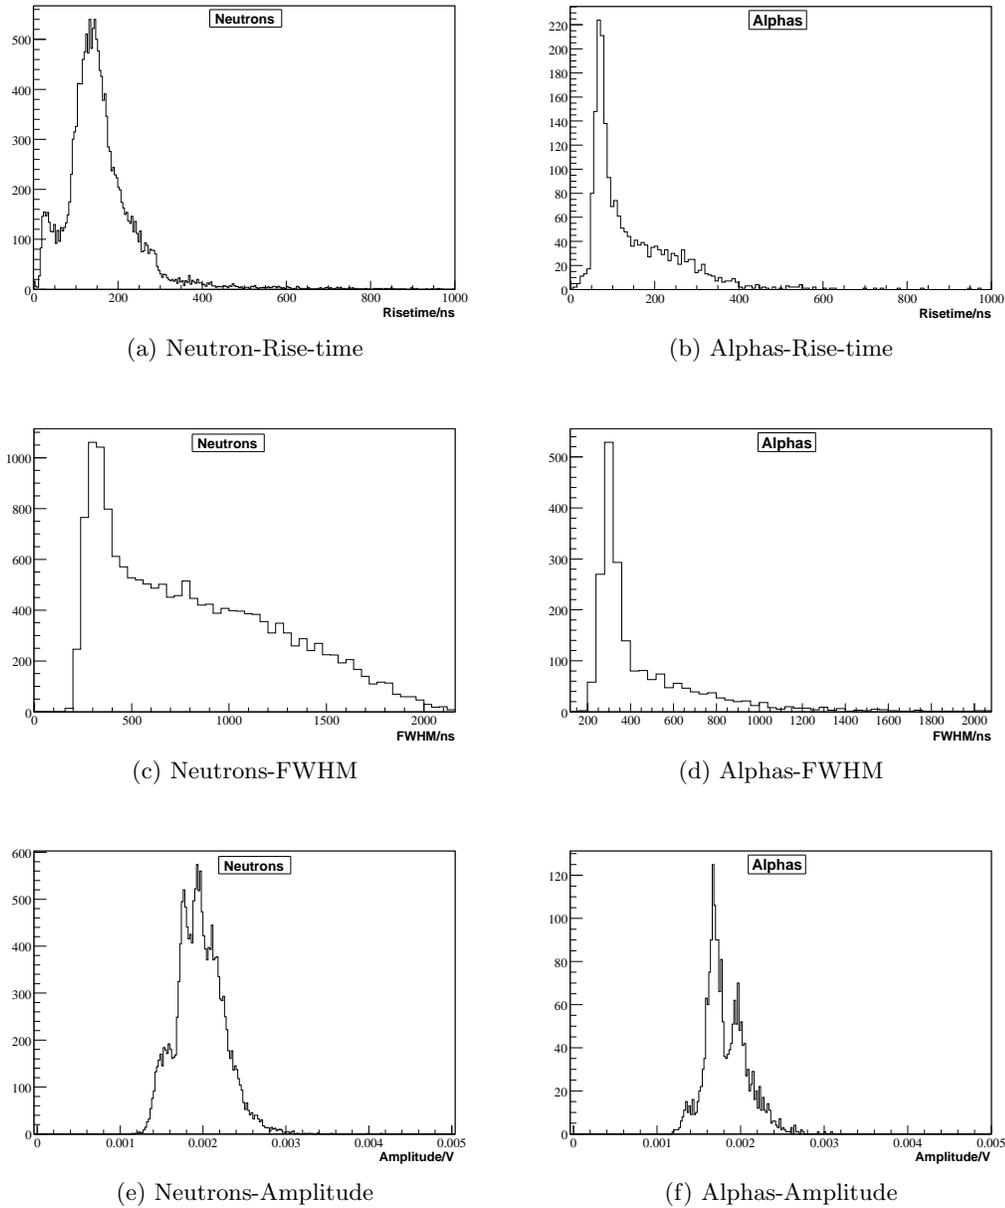
(f) Alphas-Amplitude

Figure 3.3: Pulse rise-time, FWHM and amplitude distributions for neutrons and alphas with energies below 1 MeV. These distributions motivate the idea of using pulse shape characteristics to distinguish between neutron and alpha events in the NCDs as well as to provide a means of testing the NCD Monte Carlo. These pulse shape parameters were measured on normalized pulses to reduce the effects from the differing energy spectra (the amplitudes of the pulses (third row) are thus similar between neutrons and alphas).

[68] and was used to produce probability density functions for the energy spectra of alpha particles used in the published results [38]. This section highlights some of the features of the model used for the Monte Carlo simulation. The SNO Monte Carlo simulation was used for the work in this thesis to model the pulse shapes from alpha particles in order to supplement the data from the ^{4}He strings which is limited statistically. It is thus of particular concern that pulse shapes be modeled adequately.

3.3.1 Description of the Model

Many of the features described in this chapter are implemented in the SNO NCD Monte Carlo simulation. A complete and detailed account of the simulation and verification is given in [67]. This brief section is aimed at highlighting the features that determine pulse shapes, that is, the time profile of charge collected from the NCDs.

SRIM 2003 [63] was used to model the loss of energy from alphas particles in the nickel walls of the NCDs as well as for all particles (alphas, protons and tritons) in the gas region. In all cases, straight ionization tracks were assumed in the gas as the scattering of the ionizing particle did not create any significant effect on the pulse shapes. The tracks were generated in small segments ($\sim 1\,\mu$m) and stopped when the particle lost all of its energy or ran into the wall or anode. Energy straggling was also shown to produce negligible effects on the pulse shapes.

The electron drift towards the anode was modeled from first-principles and verified with the Garfield simulation package [69]. The simulation also included time-resolution effects due to the random nature of the electron drift. The loss of electrons through attachment processes was shown to be negligible.

A few radii away from the anode, the electrons produce avalanches. Average properties of the avalanches are determined from the gain in each NCD string. The gains were measured for each string during AmBe calibration with respect to a reference string and implemented in the Monte Carlo to simulate all strings. The drift of the ions away from the anode wire

is an important factor in the pulse shapes. The ion mobility was modeled and verified with neutron calibrations.

The electronics chain was modeled as described in section 3.1.6. In order to include the effect of the digitization from the scopes, the amplitude in each time bin of the calculated pulse was round-off to the nearest integer. Noise was added to the pulses based on measured frequency spectra of noise taken from calibration runs. Finally, a model for the NCD-DAQ triggering system was also implemented.

3.3.2 Validation of the NCD Monte Carlo Simulation

Numerous tests [67] were performed to validate the Monte Carlo simulation. The neutron model was validated by generating 20,000 simulated ^{24}Na events and compared with the data from the corresponding calibration. This helped to confirm the neutron model as well as the general electronics and ionization simulations and provided the confidence that alpha particles could be simulated appropriately.

The alpha-particle simulation was more difficult to validate in the neutron energy region as the only available data in that region came from the ^{4}He-filled strings, which have low statistics. All other strings could not produce a pure sample of alpha particles, since these would be contaminated with neutrons in that energy range. To gain confidence in the alpha simulation, the Monte Carlo was compared with data at energies above $1.2\,\mathrm{MeV}$, which is free of neutrons. This was achieved by generating 200,000 ^{210}Po and 60,000 ^{238}U and ^{232}Th -chain alpha events on all strings. The (real) data from each string was also fit in order to determine string-by-string proportions of surface (polonium) and bulk (uranium+thorium) contributions. This allowed for a Monte Carlo 'cocktail' data set to be generated that adequately reflected the surface and bulk contributions from each string. This properly 'weighted Monte Carlo' was then compared with the data at energies above $1.2\,\mathrm{MeV}$. The Monte Carlo in the neutron energy region was then verified with the data from the ^{4}He strings. Figure 3.4a shows an example of a fit to the energy spectrum of string

4. Panels (b) and (c) of Figure 3.4 show a comparison of real neutrino data and the Monte Carlo with the weighting factors for all strings taken into account.

(a) Energy fit to string 4

(b) Neutrino data

(c) Monte Carlo
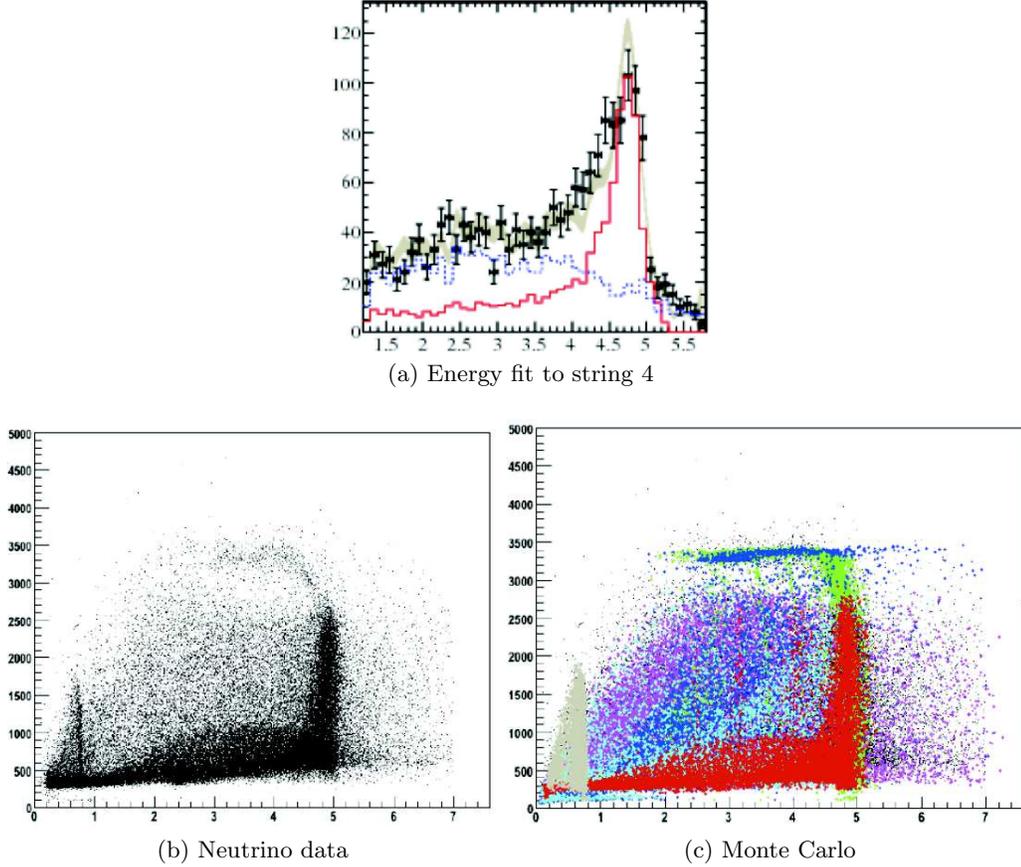
Figure 3.4: Panel (a) shows an example of a fit in energy (MeV) to string 4 data (black) in order to determine the surface (red) and bulk (blue) alpha contributions. Panel (b) shows real data plotted as a function of energy (x-axis, in MeV) and pulse width (y-axis, defined at 40% of the maximum amplitude of a pulse). Panel (c) shows pulse width versus energy (in MeV) for the Monte Carlo, colour coded to show different types of alphas. Green = Wire Po, Blue = Wire U, Cyan = Endcap Nickel Po, Red = Nickel Po, Magenta = Bulk U, Grey =Neutrons. Figures were taken from [67].

Figure 3.5 shows a small sample of the verifications that were performed between the data and Monte Carlo. The top panels compare pulse shape parameters between the neutron Monte Carlo and the data from the 2005 ^{24}Na spike. The bottom panels compare the same parameters for the weighted Monte Carlo with the data from the ^{4}He-strings. In both cases,

the Monte Carlo reproduces the general features visible in the data. The work in this thesis will use a distribution derived from the Monte Carlo alphas (see section 5.2, Figure 5.11b for the 'NoverA' distribution). The particular shape of this distribution will be such that difference between Monte Carlo and data can be accommodated with a simple scale factor and that results in a relative 'robustness' against small disagreements between the Monte Carlo and the data.
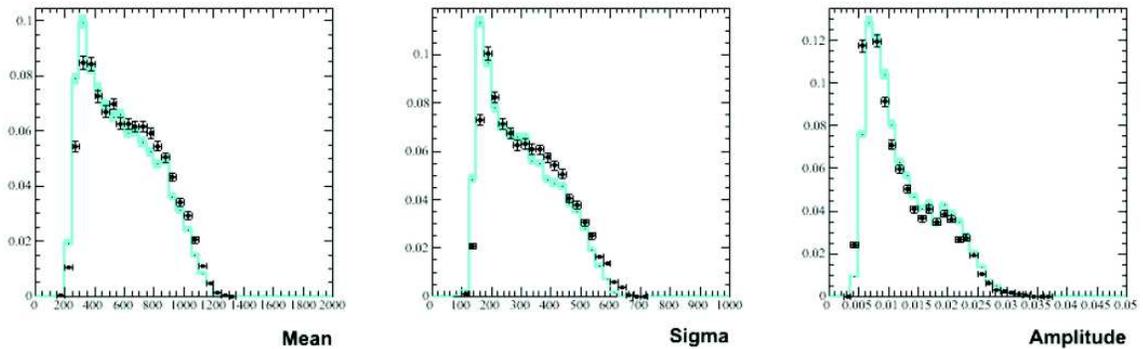
(a) Neutron model validation using ^{24}Na data
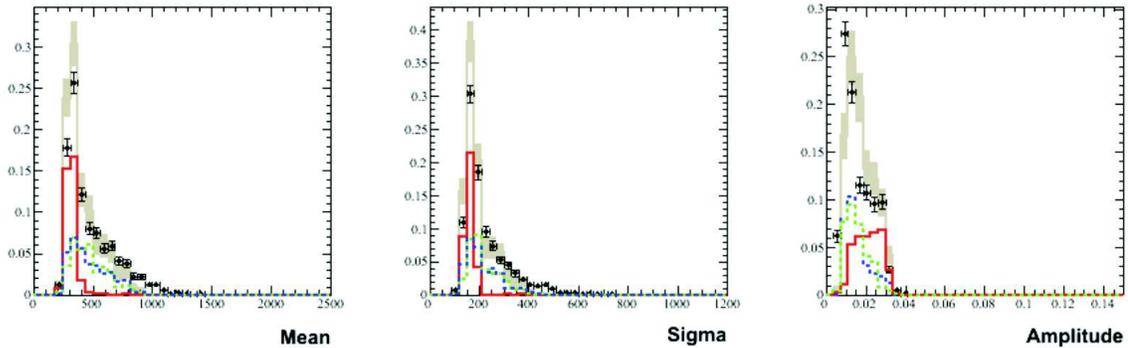
(b) Alpha model validation using ^{4}He data

Figure 3.5: Panel (a) shows the Monte Carlo simulation for neutrons (blue line) compared with data from the 2005 ^{24}Na spike (points with error bars). Three different pulse parameters are shown; the mean, width (sigma) and amplitude. Panel (b) shows the comparison between the same parameters for the alpha model with data from ^{4}He (points with error bars). The weighted alpha simulation (gray) contains alphas from polonium (red), uranium (blue) and thorium (green). Figure was taken from [67].

3.4   Summary

This chapter aimed to give the reader an understanding of the physical processes behind the pulses recorded by the scopes. The main goal was to motivate the idea that pulse shapes for neutrons and alphas look different, and this was shown with simple pulse parameters (such as rise-time and amplitude). Finally, the SNO NCD Monte Carlo was introduced and motivated, as it will be used in this thesis to generate probability density functions for alphas.

Chapter 4

The Queen's Grid Fitter

This chapter describes the main Pulse Shape Analysis (PSA, hereafter) technique developed by this author to distinguish neutrons from alphas in the NCD data. The Queen's Grid Fitter (QGF, hereafter) is designed to select neutron and alpha-like pulses by comparing the data to a library of calibration neutrons (the 'neutron library') and a library of alpha calibration pulses (the 'alpha library'). Neutron pulses are available from the large quantity of neutron calibration data (from the ^{24}Na and AmBe sources) and alpha events can be obtained from the ^{4}He'-filled counters which are insensitive to neutrons. This chapter will introduce the NoverA observable which will be used in conjunction with ADC energy in a maximum-likelihood extraction of the number of neutrons in the NCD data.

4.1   Introduction to Grid Fitting

This section describes how the grid-fitter operates and how one can quantify its effectiveness at distinguishing neutrons from alphas. The basic principle is to compare a data pulse-shape with a library of known shapes (from calibration data) using a chi-squared parameter to measure similarities.

4.1.1 Pulse Normalization

In order to remove amplitude effects and to maximize the influence of the pulse shape in distinguishing neutrons from alphas, pulses are normalized to unit area before comparison. The energy of pulses is used in the comparison between data and library pulses to ensure that the amplitude of a data and library pulse are similar (as the actual amplitude was lost in the normalization). The range over which the pulses are normalized is dynamic and corresponds to the range over which the χ^2 is calculated. The choice of range will be justified in section 4.2.2

4.1.2 Calculating a Chi-Squared Between Two NCD Pulses

The NCD scope events were recorded in histograms (1 ns bins) and are used by the grid-fitter. The chi-squared between data and library pulses is calculated straightforwardly using:

$$\chi^2 = \sum_{i=Min}^{Max} \left(\frac{D_i - L_i}{\sigma_{Data}} \right)^2 \tag{4.1}$$

where D_i (L_i) is the ith bin of the data (library) pulse histogram and σ_{Data} is the 'uncertainty' in the data pulse (defined below). The range (Min, Max) over which the calculation is done will be justified in section 4.2.2. In this scenario, one assumes that the library pulse is exact (the model), although, in reality, it also has associated noise and uncertainty. For this reason, the chi-squared does not satisfy the statistical properties of a rigorous chi-squared and should be considered a pseudo-χ^2 . This thesis will not use the χ^2 as an absolute measure of the probability of being a neutron (or alpha) so this distinction is of no importance.

Time-Shifting

In order to account for artificial time shifts in the zero time of the scope traces arising from the electronics, the pulses are allowed to shift in time during the χ^2 calculation. The pulses are first lined up so that their peaks occur at the same time bin and are then allowed to shift by ± 100 ns until the lowest χ^2 is found (and retained).

Defining the Uncertainty on a Data Pulse

The uncertainty in the data pulse, σ_{Data}, that is used for the χ^2 calculation is determined by measuring the noise rms in the last $3\,\mu$s of the data pulse. The recorded scope-trace is $15\,\mu$s long so that the ion tail has decayed away in the last part of the pulse, as can be seen in the typical neutron pulse shown in Figure 4.1.
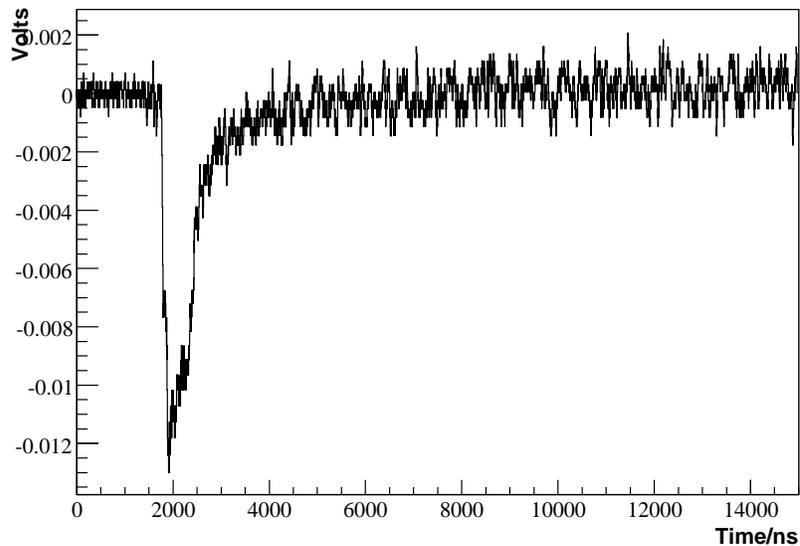
Figure 4.1: Typical neutron pulse from the 2005 ^{24}Na spike. In this case, the pulse is from string 32 and has an energy of 0.75 MeV. The noise in the last part of the pulse is not likely to be biased by the ion tail and is a suitable representation of the uncertainty in each bin.

4.1.3 The Two Libraries of Data Pulses

The QGF technique uses two different libraries of pulses. The first library was designed to contain almost entirely neutrons (obtained from neutron-source calibrations) and data-pulses that find a good match in the neutron library are deemed to be neutron-like. Conversely, the second library was designed to contain events that are alphas (obtained from the ^{4}He strings), and pulses that find good matches in this library are deemed to be alpha-like.

Due to the large amount of calibration neutron data, one has virtually no restriction in the size of the neutron library. This allows one to build a library with a chosen energy spectrum (flat or neutron-like), which will be discussed in section 4.3.2. On the other hand, the limited amount of data from the ^{4}He-strings severely impacts the size of the alpha library unless one adds events that are simulated with the Monte Carlo.

4.1.4 Data sets used for Optimizing the Grid Fitter

In order to measure the performance of the grid-fitter in differentiating neutron events from alpha particles, calibration alpha and neutron data sets will be used throughout this thesis. The neutron data set was taken from the 2005 ^{24}Na spike and consists of 16,659 events isotropically distributed in the detector [59]. The alpha data set was taken from the ^{4}He-strings during the entire SNO-NCD phase (during calibrations and data-taking) and contains 2,112 events distributed in the four ^{4}He-strings. In the rest of the thesis, these data sets will be referred to as 'the' calibration neutron and alpha data sets unless otherwise explicitly mentioned.

4.1.5 Observables from the Grid Fitter

This section introduces the observables that arise when a data pulse is compared to a library using the grid-fitter. These new observables are then associated with each data event, along with energy and other measurable quantities. One can then either place cuts on these

observables or use their associated probability density functions (pdfs) to fit for the amount of signal and background (given the PDFs for signal and background for a given observable).

Best χ^2 to Neutron Library, χ_n^2

When a data pulse is compared to the neutron library, χ_n^2 is the lowest χ^2 that was found by comparison to all the pulses in the library. That is, χ_n^2 is the chi-squared between the data pulse and the best-matching pulse in the library. If χ_n^2 is low, then the data-pulse found a reasonable match in the neutron library; if it is high, then the library is not representative of the pulse.

Best χ^2 to Alpha Library, χ_α^2

Similarly, χ_α^2 is the lowest chi-squared that was found by comparison to the alpha library. A low χ_α^2 is then indicative of a good match in the alpha library .

Figure 4.2 shows an example of the χ_n^2 and χ_α^2 distributions for the calibration neutrons and alphas. The distributions for neutrons and alphas appear to be 'more different' for χ_α^2 than they are for χ_n^2 , indicating that the former is a better discriminator between neutrons and alphas. This can be quantified by considering the integral of the distributions within some range to estimate how they affect the signal (neutrons) to background (alpha) ratio.

χ_n^2/χ_α^2 "NoverA"

One can combine the information from both χ_n^2 and χ_α^2 by using their ratio and it will be shown that this observable is a powerful discriminator between neutron and alphas. Figure 4.3 shows the distributions of χ_n^2/χ_α^2 for neutrons (red) and alphas (black). This observables shows the potential to be used either as a cut on the data or for a PDF extraction.

(a) χ_n^2 Distributions

(b) χ_α^2 Distributions
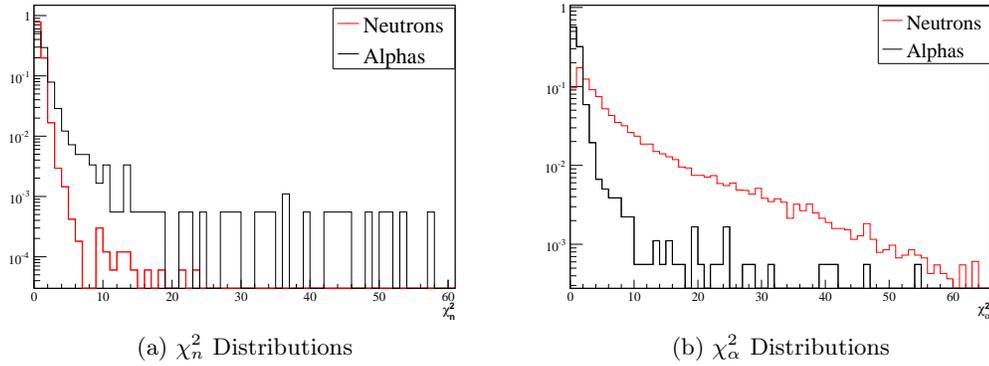
Figure 4.2: Comparison of χ_n^2 and χ_α^2 distributions for neutrons (from ^{24}Na) and alphas (from ^{4}He-strings). The distributions shown here have been normalized to unit area. In the χ_n^2 distribution, 98% of neutrons and 85% of alphas have χ_n^2 less than 2.0. In the χ_α^2 distribution, 28% of neutrons and 90% of alphas have a χ_α^2 less than 2.0. It should be clear from this figure that χ_n^2 and χ_α^2 both have potential as discriminators and that χ_α^2 is more powerful.
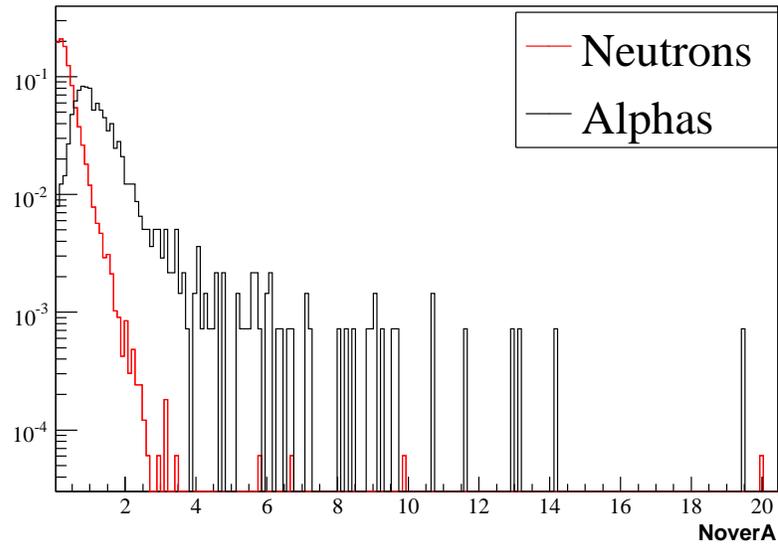
Figure 4.3: Comparison of χ_n^2/χ_α^2 ("NoverA") normalized distributions for neutrons and alphas. 92% of neutrons and 30% of alphas have $0 < \chi_n^2/\chi_\alpha^2 < 0.7$.

4.1.6 Measuring Optimization of the Grid-Fitter

In order to measure the effectiveness of the grid-fitter in doing PSA, the distributions of the variables presented in Section 4.1.5 under different regimes will be considered. In particular, the grid-fitter will be optimized by considering the χ_n^2/χ_α^2 distribution with the aim of making the distributions for neutrons and alphas as different as possible. Additionally, the grid-fitter will also be optimized so that the χ_n^2/χ_α^2 distribution for neutrons is as narrow as possible. This can be quantified by considering the integral of the distributions within a certain range (essentially, measuring the acceptance that the observables would have if used as data-cleaning cut).

4.2 Computational Efficiency Optimization for Grid Fitting

The process of running a grid-fitter can be long and such calculations should be optimized. The number of iterations in calculating the χ^2 with equation 4.1 will clearly have an impact on the amount of CPU-time that is required. In order to minimize the number of iteration in calculating a χ^2 , one can reduce the range over which it is calculated, as well as change the binning of the pulses before the calculation. In addition to minimizing the number of iterations in one calculation, one can reduce the effective size of a pulse library by comparing one data pulse only to pulses in the library that are already similar (decided based on pre-calculated comparators such as energy, pulse-width, pulse rise-time). This section describes the optimization of the various parameters that can reduce the amount of CPU time without significantly affecting the discrimination of the grid-fitter.

4.2.1 Rebinning

Since the data pulses have electronic noise, one is led to consider the possibility of smoothing them. This can be done by either using some smoothing algorithm or rebinning the data (combining several bins into one, which effectively averages the data). In addition to

smoothing the pulse shapes, rebinning also results in a faster calculation of equation 4.1, since there are fewer iterations. Figure 4.4 shows the effect of rebinning a pulse in the neutron data set by a factor of 50, and the smoothing effect is clearly seen. One needs to find an optimal value of rebinning that will not wash away physical features of the pulses. Figure 4.5 shows the the χ_n^2/χ_α^2 distribution for neutrons when the grid-fitter was used with different values of rebinning for the pulses. The distribution gets wider as the pulses are rebinned into coarser resolution times, while the peak moves to the left. A rebinning value of 50 is seen to preserve the distribution while substantially improving the computational efficiency of the algorithm and will be used as the default value in this thesis.

(a) Raw neutron pulse   (b) Neutron pulse rebinned by a factor of 50
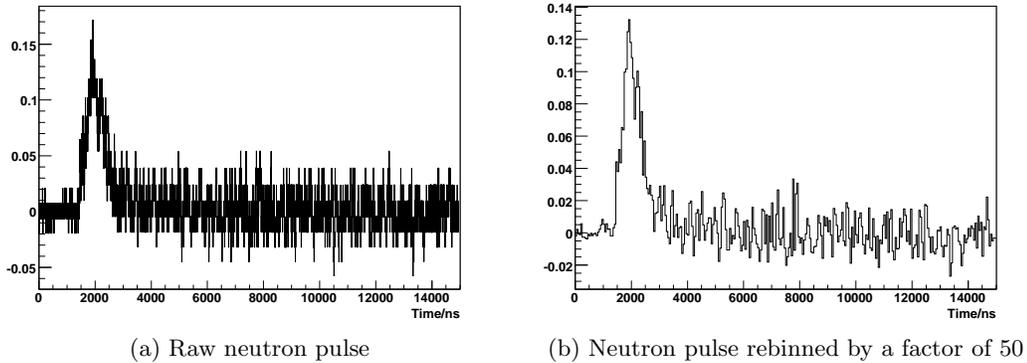
Figure 4.4: Effect of rebinning an NCD pulse. This pulse was taken from the neutron data set. Rebinning by a factor of 50 smooths the pulse while retaining the principal features.

4.2.2 Fit Range

The PSA observables obtained from the grid-fitter will be affected by the range over which χ_n^2 and χ_α^2 are calculated when matching a pulse to the libraries. Clearly, if the range is too small, important characteristics in the pulses can be neglected and the benefits of the pulse shape discrimination could be lost. On the other hand, if the fit range is too long, parts of the pulse that are not characteristic of neutrons or alphas (such as the ion-tail) will be over-weighted and could wash out any PSA discriminating power. It is then natural to
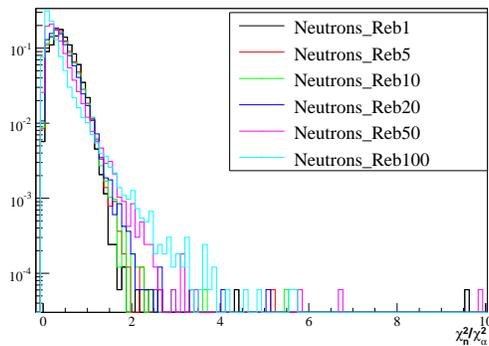
(a) χ_n^2/χ_α^2 distributions for neutrons

Rebinning	Events between 0 and 0.7
1	86%
5	88%
10	88%
20	89%
50	92%
100	93%

(b) Fraction of neutron events between χ_n^2/χ_α^2 =0 and 0.7, for different pulse rebinning

Figure 4.5: χ_n^2/χ_α^2 distribution for calibration neutrons with different values of pulse rebinning. A pulse rebinning of 1 corresponds to the raw pulse, whereas a pulse rebinning of 100 corresponds to summing 100 bins into one, and thus reducing the number of bins in the pulse histogram by that factor. The table on the right shows the integral of the histograms between 0 and 2. One can see that rebinning the pulses by 50 is reasonable. As the rebinning is increased, the width of the χ_n^2 distribution increases and the location of the peak get closer to zero.

optimize the range over which the χ^2 are calculated. In particular, the rising edge of the pulse and peak are expected to be most characteristic of the particle that created it, since those parts are mostly determined by the electrons arriving at the NCD wire. The decay of the pulse amplitude is mostly determined by the ions moving away from the wire and is expected to be less of a discriminator between neutrons and alphas.

Figure 4.6 shows the χ_n^2/χ_α^2 distribution for neutrons where the fit range was varied dynamically both on the rising and falling edges of the pulse. The lowest bin in the fit range was always taken to be at 10% of the maximum pulse amplitude and the highest bin in the fit range was taken when the pulse amplitude had decreased by 10, 20, 40, 60 and 80 percent with respect to the maximum amplitude. The best compromise between calculation speed and efficient use of the pulse shape is achieved when calculating the χ^2 until the pulse has decreased to 40% of its maximal amplitude. If the range is much larger, the distribution becomes too wide.

4.2.3 Grid Reduction Using Pulse Parameters

In order to avoid calculating a χ^2 between a given data pulse and every pulse in the library, the calculation can be chosen to be carried out only if the pulses are deemed similar enough based on pre-calculated criteria, such as pulse width, rise time (the amount of time between when the pulse goes from 10% to 90% of its maximum amplitude), energy (as measured by the shaper/ADC boards) and amplitude. Figure 3.3 from section 3.2 showed the distributions of these 'basic' pulse shape parameters that were used to define 'tolerances' to be met if two pulses are to be fit. The amplitude of the pulses is related to the energy and was not used. It was found that using a 20 ns tolerance (between data and library pulse) on the risetime alone was sufficient to speed up the calculation. From the distributions in figure 3.3 it is clear that the risetime contains more physical information about the pulse than the other basic parameters.
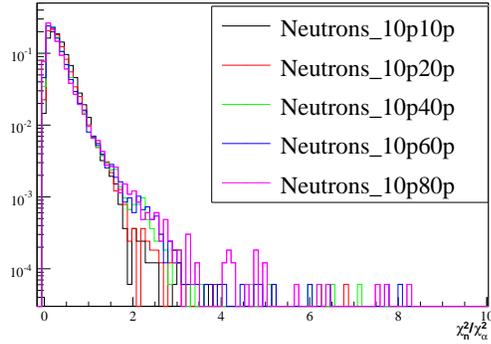
(a) χ_n^2/χ_α^2 distributions for neutrons

Fit range in percent of amplitude	Events between 0 and 0.7
10% to 10%	92%
10% to 20%	92%
10% to 40%	92%
10% to 60%	92%
10% to 80%	92%

(b) Fraction of neutron events between $\chi_n^2/\chi_\alpha^2 = 0$ and 0.7, for different fit ranges

Figure 4.6: χ_n^2/χ_α^2 distribution for calibration neutrons (from the 2005 ^{24}Na spike) fitted with different fit ranges in the grid-fitter. The lowest bin in the fit range was determined dynamically on the rising edge of the pulse at 10% of the maximum amplitude. The highest bin in the fit range (on the falling edge of the pulses) was determined dynamically as the location where the trace reaches 10, 20, 40, 60 and 80% of its maximum amplitude. The table on the right shows the integral of the histograms between 0 and 0.7.

4.2.4 The Energy Cut

Since pulses are normalized to unit area before the χ^2 calculation, information about their amplitude has been lost in the process. Requiring pulses to have similar energies guarantees that the pulses have similar amplitudes. The energy tolerance that is required between pulses has a similar effect in speeding up the calculation as the rise time tolerance. The energy tolerance can, however, introduce a correlation between the χ^2 observables and the pulse energies (for example, with a small energy tolerance and a neutron library with a peak, events with energies close to the neutron peak will be more likely to have a good χ_n^2). Clearly, correlations between PSA observables and energy should be avoided at all costs, since energy will be used as an observable in the data analysis.

Figure 4.7 shows the χ_n^2/χ_α^2 distributions for neutrons and alphas that were fit using an 0.1 MeV (red) and 1 MeV (black) tolerance between data and library pulses (requiring that the library and data pulse be within $\pm$ the tolerance of each other). Although the higher energy tolerance allows pulses to find a better match in a library (in terms of chi-squared), this is offset by the fact that, for example, neutrons will also find better matches to the alpha library. It turns out that the distribution of χ_n^2/χ_α^2 for neutrons is better (narrower and more to the left) with the smaller energy tolerance. This is somewhat compensated by the fact the alpha distribution is slightly better (wider and more to the right) for the 1 MeV tolerance. The difference in shape of the distribution between neutrons and alphas is most different for the 0.1 MeV tolerance; this choice also provides a substantial increase in speed of the algorithm and is used for the remainder of this work.

(a) χ_n^2/χ_α^2 distributions for neutrons

(b) χ_n^2/χ_α^2 distributions for alphas
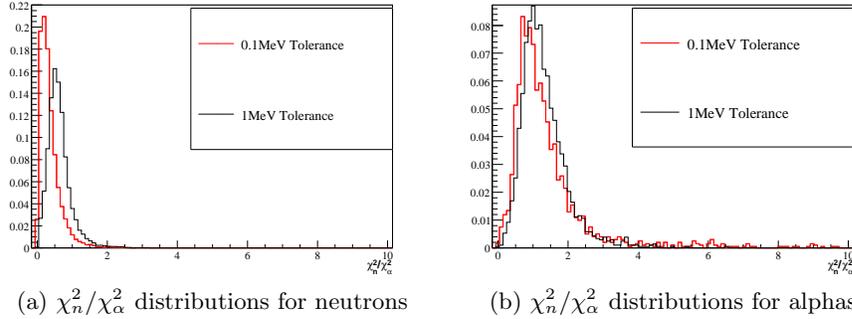
Figure 4.7: χ_n^2/χ_α^2 distributions for neutrons and alphas fitted with different tolerances on the energy between data and library pulses. For neutrons, 92% of events have $\chi_n^2/\chi_\alpha^2 < 0.7$ for an 0.1 MeV tolerance and 72% for a 1 MeV tolerances. This is offset by the fact the for an 0.1 MeV tolerance 25% of alphas are below $\chi_n^2/\chi_\alpha^2 = 0.7$ versus only 15% with the 1 MeV tolerance

4.3 Optimization of the Neutron Library

4.3.1 Choice of the Library Size

The size of the neutron library is an obvious consideration in optimizing a grid-fitter; if it is too small, it is not representative of possible neutron shapes and if it is too large, it will not be computationally efficient. Additionally, the effect of allowing neutrons and alphas to find a good match should be balanced, as a larger library will make it more likely for either type to find a matching library pulse. In this section, the neutron and alpha χ_n^2/χ_α^2 distributions are considered as a function of neutron library size. In each case, the library was generated using data from the 2005 ^{24}Na spike and pulses were selected in order to make the library have a flat energy spectrum (the energy spectrum of the library is discussed in the next section). Figure 4.8 shows the χ_n^2/χ_α^2 distribution for calibration neutrons (from the 2005 ^{24}Na spike) fitted to neutron libraries of different sizes. The distribution becomes narrower as the library size increases (more pulses are finding good matches) until the library contains around 20,000 pulses, which is the size chosen for the rest of this work. Figure 4.9 shows

the χ_n^2 distribution for alphas as well as the integral between 0 and 2 for χ_n^2 . It can be seen that alphas are also more likely to find a better match in a larger neutron library, as anticipated.
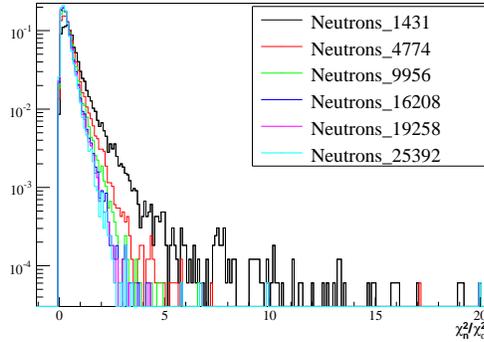
(a) χ_n^2/χ_α^2 distributions for neutrons

Library size	Events between 0 and 0.7
1,431	69%
4,774	81%
9,956	87%
16,208	89%
19,258	92%
25,392	92%

(b) Fraction of neutron events between χ_n^2/χ_α^2 =0 and 0.7, for the different size neutron libraries

Figure 4.8: χ_n^2/χ_α^2 distribution for calibration neutrons (from the 2005 ^{24}Na spike) fitted to neutron libraries of different sizes. As the size of the library increases, more pulses find good matches and the distribution becomes narrower. The table on the right shows the integral of the histograms between 0 and 0.7.

4.3.2   Choice of the Library Energy Spectrum

If the energy spectrum in the neutron library has a peak (as the neutron spectrum in Figure 2.1), any pulse is more likely to find a match if it has an energy close to the peak, simply due to the larger phase space of available pulses. In turn, this can bias the energy spectrum of selected pulses and create a correlation between the pulse-shape parameters and energy. Figure 4.10 shows the energy spectrum of neutrons (panel a) and alphas (panel b) after the use of a PSA variable to cut the data; only events with $\chi_n^2/\chi_\alpha^2 < 0.7$ were kept. The black lines show the spectra when the PSA variables were calculated using a neutron library with a flat energy spectrum and the red curves were calculated using a neutron library with a
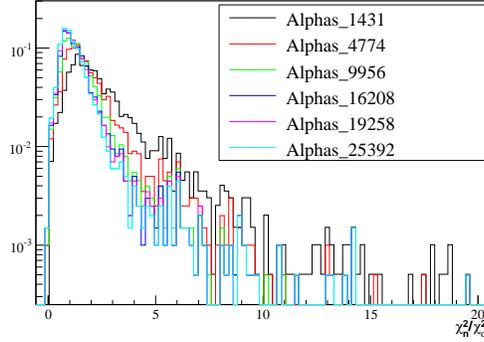
(a) χ_n^2/χ_α^2 distributions for alphas

Library size	Events between 0 and 0.7
1,431	6%
4,774	11%
9,956	17%
16,208	21%
19,258	22%
25,392	25%

(b) Fraction of alpha events between $\chi_n^2/\chi_\alpha^2 =0$ and 0.7, for the different size neutron libraries

Figure 4.9: χ_n^2/χ_α^2 distribution for alphas fitted to neutron libraries of different sizes. As the size of the library increases, more pulses find good matches and the distribution becomes narrower. The table on the right shows the integral of the histograms between 0 and 0.7.

neutron spectrum. It is clear that the spectrum for the alpha particles passing this PSA cut have inherited a neutron peak when the neutron library also contains that peak and an unwanted correlation has been introduced. The neutron library was thus generated with a flat energy spectrum by using the data from the 2005 distributed ^{24}Na calibration source. This was achieved by taking almost all of the pulses in the tail of the energy spectrum and randomly selecting events in the peak.
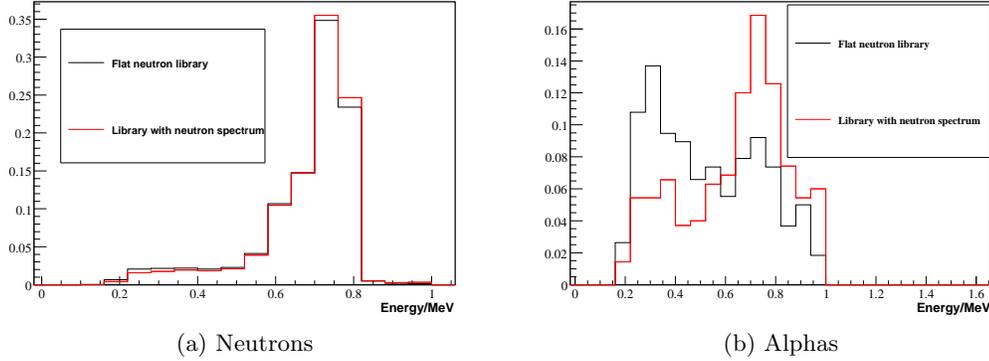
(a) Neutrons   (b) Alphas

Figure 4.10: Normalized spectra for neutron and alphas after a PSA cut selecting events with $\chi_n^2/\chi_\alpha^2 < 0.7$. In one case (black line) the PSA variable was calculated using a neutron library with a flat energy spectrum. One notes that in the case where the neutron library has a neutron spectrum (red line) the energy spectrum of the background (shown here with alpha particles in panel b) is biased towards having a neutron peak.

4.4   Optimization of the Alpha Library

The alpha library is chosen to represent alpha backgrounds in the NCDs as accurately as possible. The events from the ^{4}He strings are an obvious candidate for use as the library. These can potentially be supplemented with events generated by Monte Carlo simulation of alpha particles to increase the statistics in the library.

4.4.1   Choice of Alpha Events

In this section, the use of pulses from the ^{4}He -strings to build the alpha library is motivated by considering the χ_n^2/χ_α^2 distribution. It is shown that the inclusion of Monte Carlo generated pulses into the alpha library does not add any valuable discrimination power to χ_n^2/χ_α^2 . Figure 4.11 shows the χ_n^2/χ_α^2 distribution for neutrons (panel a) and alphas (panel b) compared for two different alpha libraries; the black curve shows the case when the alpha library is made from ^{4}He strings events and the red curve shows the case when the library

has been supplemented by Monte Carlo alpha events[1]. Although the data + Monte Carlo alpha library is more efficient at discriminating alpha events (evidenced by the integral between 0 and 0.7), it is clear the difference in shape of the distribution between neutrons and alpha events is greatest when the alpha library is composed of only events from the ^{4}He strings. The rest of this work thus uses an alpha library that contains only events from the ^{4}He strings.

(a) χ_n^2/χ_α^2 distributions for neutrons   (b) χ_n^2/χ_α^2 distributions for alphas
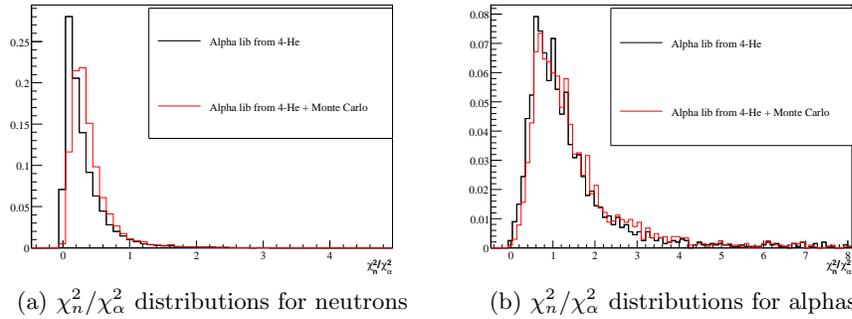
Figure 4.11: χ_n^2/χ_α^2 distributions for neutrons and alphas fitted with different alpha library made from data alphas (black) and including Monte Carlo events (red). For neutrons, 92% of events have χ_n^2/χ_α^2 <0.7 for the data library and 90% for the data + Monte Carlo library. For the alphas, 30% have χ_n^2/χ_α^2 below 0.7 when fitted to the data-only library, versus 23% with the data + Monte Carlo library.

4.5 Summary

This section introduced the use of a grid-fitter (the 'Queen's Grid-Fitter, QGF) to analyze the scope traces from the NCDs. QGF uses a set of neutron scope traces from the 2005 ^{24}Na neutron calibration (the neutron library) and a set of alpha events from the ^{4}He -filled detectors (the alpha library) to use in a comparison with a given data pulse.

The comparison of data pulses to all pulses in a library is done by calculating a χ^2 and

---

[1]The set of Monte Carlo events were obtained from N. Oblath [70], and were generated on a grid in parameter space that is representative of all possible alpha events.

the lowest value for each library (χ_n^2 for the neutron library and χ_α^2 for the alpha library) is retained to qualify the data pulse. A low value of χ^2 is then indicative of the similarity of the pulse with a given library. It was shown that χ_n^2 and χ_α^2 can be efficiently combined into one variable, their ratio, χ_n^2/χ_α^2 ("NoverA"), which was then shown to be effective in discriminating neutrons from alpha events. The NoverA variable was then used to motivate the way in which the QGF fitting is done and will be used in a maximum-likelihood extraction of the number of neutrons in the NCD data.

Chapter 5

Probability Density Functions for NCD Data

In this chapter, the uncertainty in the shape of NCD probability density functions ('pdfs', hereafter) for energy and NoverA is characterized. Both of these quantities have different probability density distributions for different classes of signals. Four classes of signals are considered: neutrons, alphas and two types of pathological events, as determined from the data in strings 0 and 26. This chapter shows how the pdfs and uncertainties in their shape are obtained for use in a maximum-likelihood analysis of the data (described in chapter 6) which will then be the basis for the analysis of the NCD-phase data presented in this work.

The various pdfs will be defined as histograms and the uncertainties on the shapes will be handled by applying scale, shift and resolution functions to the histograms. This method allows the pdf shapes to be distorted continuously and the extent of the distortion will be constrained by calibration data.

5.1 Neutron Probability Density Functions

The neutron probability density functions for ADC energy and NoverA are obtained by considering the ^{24}Na and AmBe neutron calibration data. The pdfs are then defined with their uncertainties so that they agree with all calibration data. Since the data were taken with neutron sources with different rates and at different times during the NCD phase, this method will provide a good test of any such possible influences, as well as characterizing any effect from the position of the neutron sources in the detector.

5.1.1 Energy

Figure 5.1a shows the energy spectra from all AmBe-Hi calibration scans along with the spectra from the 2005 and 2006 ^{24}Na spikes. The AmBe data were collected in 'scans' of the point source that approximate an equal volume weighting of the detector to provide similar calibrations to the ^{24}Na solutions. Statistical error bars are shown for each data set and are seen to be comparable. The overall neutron energy pdf is determined by considering the average of these calibration spectra and is shown in Figure 5.1b. For each bin, the average bin content, $\hat{\mu}$, was found using a simple average over each data set:

$$\hat{\mu} = \frac{1}{N} \sum_{i=0}^{N} \mu_i$$

(5.1)

where μ_i are the bin contents for the N different data sets. In obtaining the average neutron energy pdf, it has been assumed that the ^{24}Na data and the AmBe calibration scans both measure the neutron energy spectrum of the NCD array exposed to a uniform neutron source. It will be seen later in this section that the AmBe data and ^{24}Na agree very well.

In order to determine systematic uncertainties in the shape of the pdf, 'deformation parameters' are applied and then constrained by calibration data. These deformation parameters allow the pdf to change shape by applying a scale, shift and resolution to the

(a) Energy spectra of AmBe-HI+ ^{24}Na

(b) Neutron energy pdf
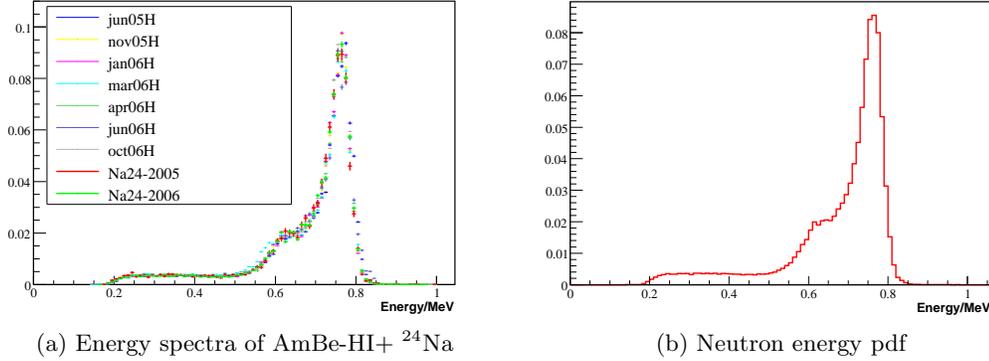
Figure 5.1: Determination of neutron energy pdf. Panel a) shows the ADC energy spectra with statistical uncertainties for all AmBe-Hi scans as well as the two ^{24}Na spikes. The neutron energy pdf is taken as the histogram with the average bin content from the calibration data and is shown in panel b).

x-axis. A function, $f(energy)$, is built by (linearly) interpolating from the average histogram. One can then fit for the scale, shift and resolution function by fitting the function $f(scale * (energy + shift)) \otimes Norm(scale * (energy + shift), resolution)$ to the calibration data (where $Norm(\mu, \sigma)$ is the normal distribution and is convolved with $f()$ to simulate an energy resolution function).

The fit for scale, shift and resolution are performed using a Markov-Chain-Monte Carlo fitter developed by this author and described in chapter 6. Figure 5.2 shows an example of the fitter applied to the energy spectrum of the data from the 2006 ^{24}Na spike where the scale, shift and resolution were determined simultaneously. The data are shown with error bars, the interpolation function is shown in blue and the histogram with the deformation parameters applied (from which the interpolation function was created) is shown in red. The scale, shift and resolution parameters that were determined by this fit (summarized in Table 5.1) do not cause any perceptible deformation.

Since the deformation is very small, one can limit the number of deformation parameters. In particular, the resolution is essentially zero, leaving only scale and shift, which are, in turn, almost exactly correlated. Hence, it is sufficient to fit the neutron energy spectrum

Figure 5.2: Example fit for energy scale, shift and resolution to the data from the 2006 ^{24}Na spike. The data were fit to the function $f(scale * energy + shift) \otimes Norm(scale * energy + shift, resolution)$ using the MCMC fitter described in chapter 6. The data are shown with error bars, the interpolation function is shown in blue and the histogram with the deformation parameters applied (from which the interpolation function was created) is shown in red. The parameters that were determined by this fit do not cause any perceptible deformation to the original pdf (not shown).

with a floating scale, holding shift and resolution both fixed to zero. Table 5.1 shows a comparison of the resulting scale that is fit out, depending on whether or not the shift and resolution are fixed, for the two ^{24}Na spikes (from 2005 and 2006).

The results from the ^{24}Na data can be validated with AmBe calibration data. As the AmBe data were taken more frequently than the ^{24}Na spikes, it can be used to monitor for the time stability of the deformation parameters. In addition, since the AmBe was deployed as a point source, it can be used to check for any variations in the neutron response as a function of position in the detector. The response of the whole detector is, however, better measured by the uniform neutrons from ^{24}Na as this avoids a complicated volume weighing of the AmBe data. Nonetheless, the AmBe data were collected in a way to keep the volume weighing as accurate as possible and provides a valuable consistency check. Each AmBe run from the AmBe-Hi calibration source were also fit for a scale, shift and resolution, as well as for a scale with the resolution and shift fixed to zero.

The fitted scales, shifts and resolution fitted for each run are shown in the left panels of Figure 5.3, whereas the right panels show the data averaged over the scans. Figure 5.4 shows the case where only a scale parameter was fit (resolution and shift were set to zero). The average, $\hat{\mu}$ and uncertainty, $\hat{\sigma}$, for each scan was calculated using

$$\hat{\mu} = \frac{1}{N} \sum_{i=0}^{N} \mu_i$$

$$\hat{\sigma} = \sqrt{\frac{\sum(\mu_i - \hat{\mu})^2}{(N-1)}}$$

(5.2)

where N is the number of runs in each scan and μ_i is the measured parameter in a given run. No systematic variation is seen between scans, indicating that the deformation parameters are stable over time. In addition, it is clear that using only a scale parameter results in a much more stable constraint, as the correlation between the shift and scale is removed.

The values and uncertainties for scale, shift and resolutions from the set of AmBe-Hi

Data	Scale	Shift	Resolution
2005 ^{24}Na	1.0213(55)	-0.0133(13)	0.00030(03)
2005 ^{24}Na	1.0053(05)	fixed	fixed
2006 ^{24}Na	1.0249(39)	-0.0161(29)	0.00023(10)
2006 ^{24}Na	1.0034(03)	fixed	fixed
AmBe-HI	1.0280(110)	-0.0180(61)	0.00047(01)
AmBe-HI	1.0030(50)	fixed	fixed

Table 5.1: Neutron energy scale, shift and resolution uncertainties for the different data. The case when only a scale parameter is more stable and will be retained for describing the systematic uncertainty in the neutron energy pdf. In order to remain conservative, the value from the AmBe data (last line) is used as the neutron energy scale central value and uncertainty.

calibration data are found by creating a histogram of the data from the left panels as shown in Figure 5.5. For each parameter, a value and uncertainty is determined by fitting a normal distribution.

Table 5.1 summarizes the fitted energy scale, shift and resolution for the ^{24}Na and AmBe data. The values for scale, shift and resolution are consistent between the data sets when the three parameters are fit simultaneously. Fixing the shift and resolution parameters to zero results in a more accurate estimate of the scale parameter which, however, disagrees between the two ^{24}Na spikes. Since these parameters represent extremely small deformations of the pdf, the conservative choice of using the value from the AmBe data (which agrees with both ^{24}Na) still maintains a very small overall uncertainty in the neutron energy scale. In the final analysis of the data, the neutron energy pdf will be deformed using only a scale parameter with the constraint of being equal to 1.0030 ± 0.0050.

(a) Fit for energy scale by run

(b) Fit for energy scale by scan

(c) Fit for energy shift by run

(d) Fit for energy shift by scan

(e) Fit for energy resolution by run

(f) Fit for energy resolution by scan
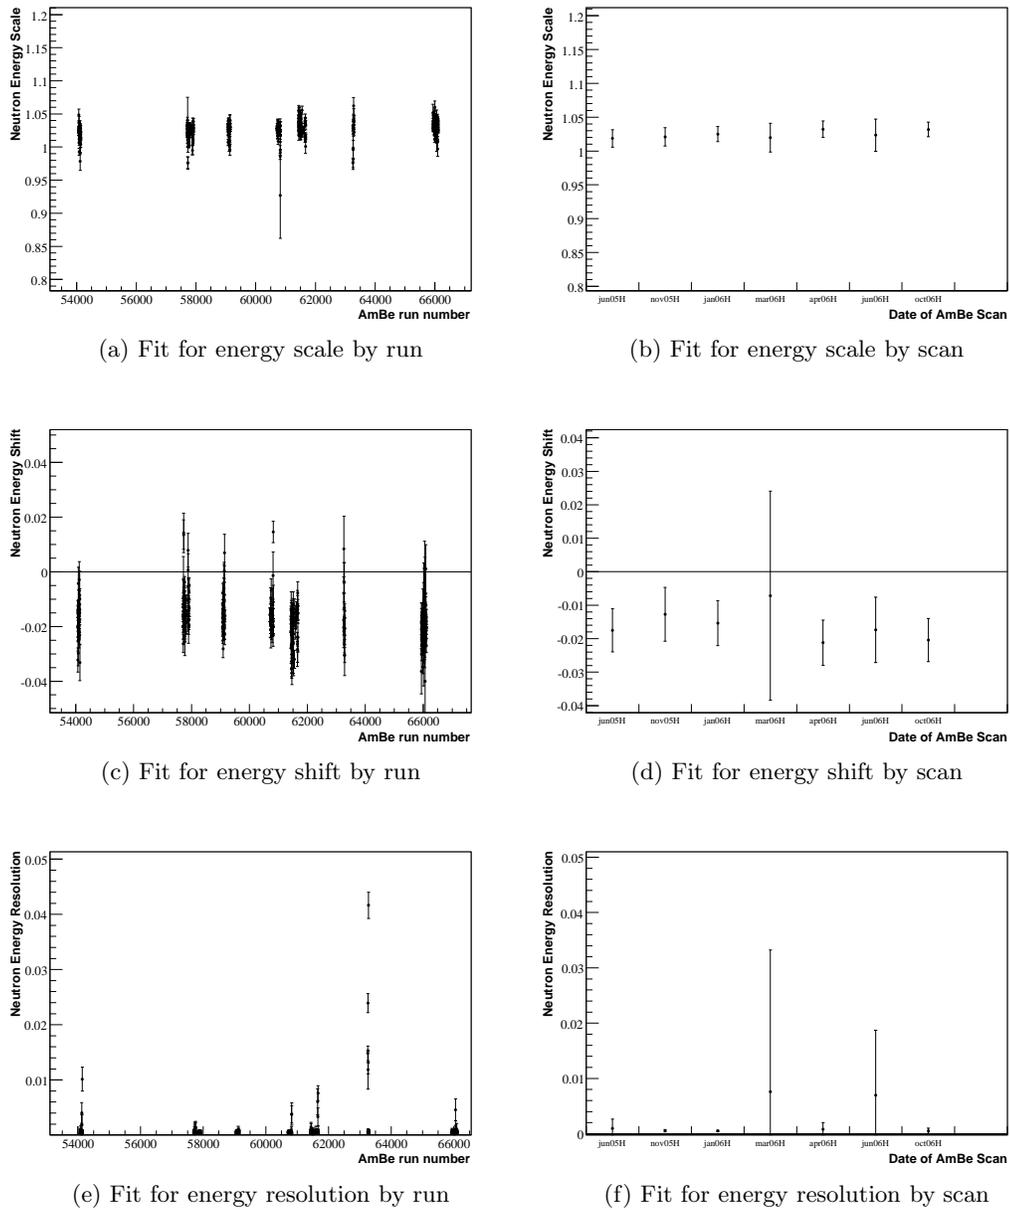
Figure 5.3: Determination of neutron energy scale, shift and resolutions using AmBe-Hi data. The left panels show run-by-run fits and the right-side panels show the parameters averaged over the scans using equation 5.2. These plots corresponds to the results when scale, shift and resolution are determined simultaneously. The grouping of the points on the left-side figures shows points from a same scan.

(a) Fit for only energy scale by run (b) Fit for only energy scale by scan
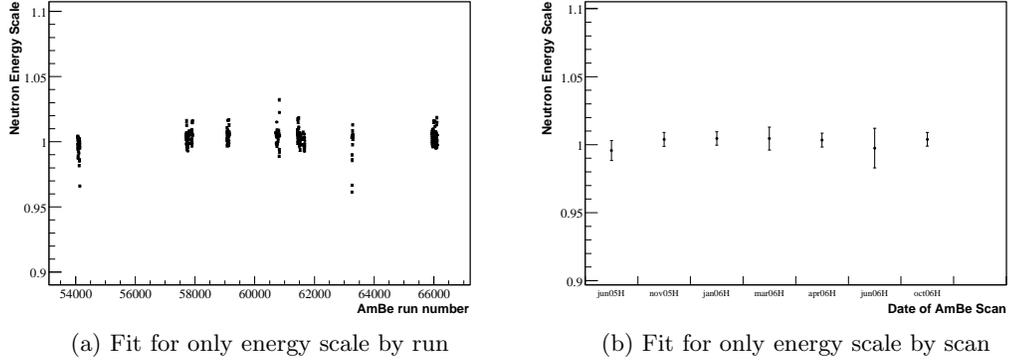
Figure 5.4: Determination of neutron energy scale using AmBe-Hi data. The left panels shows the run-by-run fit and the right-side panel shows the parameters averaged over the scans using equation 5.2. The grouping of the points on the left-side figures shows points from a same scan.

5.1.2 NoverA

This section details how the neutron pdf for NoverA is determined. The NoverA variable for each event corresponds to the χ_n^2/χ_α^2 ratio that was introduced in chapter 4. The neutron pdf for NoverA is determined using the same procedure as was used for the energy pdf. The NoverA pdfs from the different AmBe-Hi scans as well as the ^{24}Na spikes are shown in Figure 5.6a; the average value in each bin was then used to determine the central value of the neutron NoverA pdf, shown in panel b).

The systematic uncertainty in the NoverA pdf shape is also determined by deforming it with scale, shift and resolution parameters. These parameters can again be estimated by fitting the ^{24}Na and AmBe-Hi data. An example of such a fit for the 2006 ^{24}Na data is shown in Figure 5.7; the resulting fitted values are summarized and compared to the case where only a scale parameter is used in Table 5.2. Once again, it is seen that using only a scale is sufficient to properly account for the possible deformation of the neutron pdf.

The fits for NoverA scale to AmBe-Hi data were performed for each calibration run and the resulting fits are shown in Figure 5.8a for each run and in panel 5.8b averaged for

(a) Fitted energy scales

(b) Fitted energy shifts

(c) Fit energy resolutions

(d) Fit energy only scale
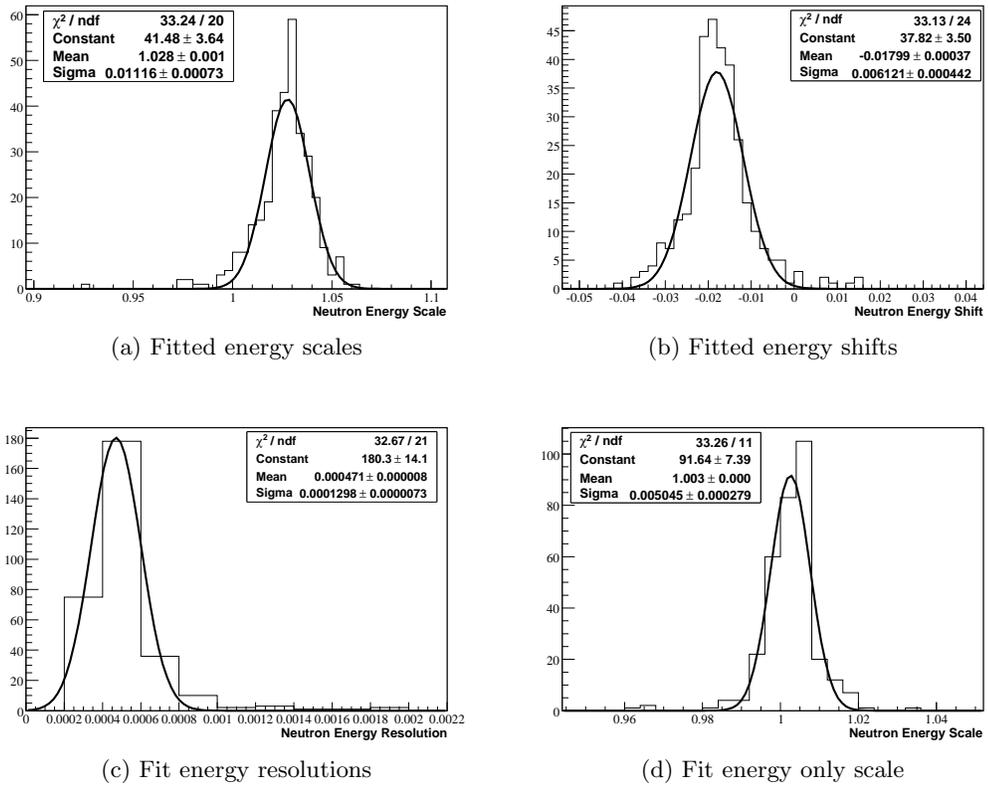
Figure 5.5: Determination of neutron energy scale, shift and resolution uncertainties using AmBe-HI data. The allowed ranges in the deformation parameters are determined by creating a histogram of the data from Figure 5.3 and fitting normal distributions. The ranges of the parameters are summarized in Table 5.1. The first three panels (abc) show the results for the case when all three deformation parameters are fit simultaneously and the last panel (d) shows the case when only a scale parameter was used.

(a) NoverA spectra of AmBe-HI+ ^{24}Na   (b) Neutron NoverA pdf
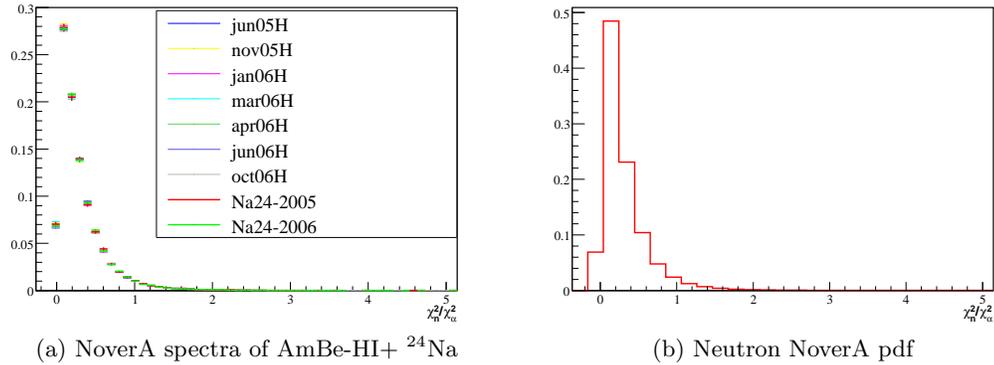
Figure 5.6: Determination of neutron NoverA pdf. Panel a) shows the NoverA pdf with statistical uncertainties for all AmBe-Hi scans as well as the two ^{24}Na spikes. The neutron NoverA pdf is taken as the histogram with the average bin content from the calibration data and is shown in panel b).
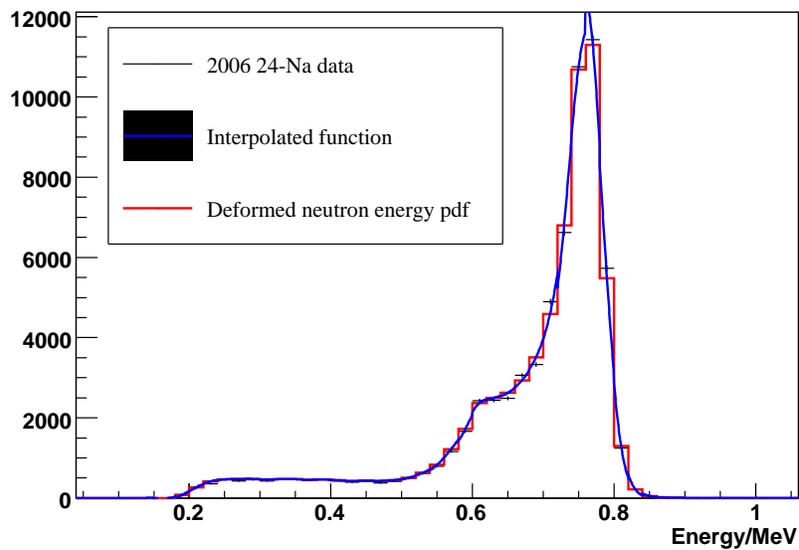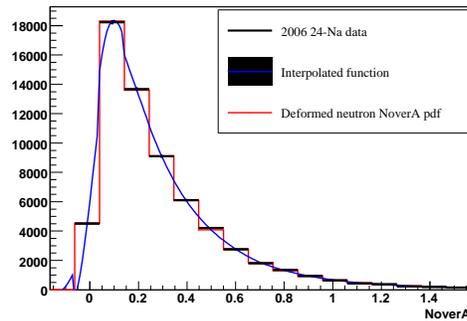

Figure 5.7: Example fit for NoverA scale, shift and resolution to the data from the 2006 ^{24}Na spike. The data were fit to the function $f(scale * (NoverA + shift)) \otimes Norm(scale * (NoverA + shift), resolution)$ using the MCMC fitter. The data are shown with error bars, the interpolation function is shown in blue and the histogram with the deformation parameters applied (from which the interpolation function was created) is shown in red. One should note that the parameters that were determined by this fit do not cause any perceptible deformation.

each scan using equation 5.2. There is no evidence for any systematic shift of the neutron
NoverA scale parameter with time. The resulting value for NoverA scale and its uncertainty
averaged over all the AmBe data is obtained by creating a histogram of Figure 5.8a and
fitting a normal distribution, as shown in Figure 5.9.

(a) Fit for NoverA scale by run   (b) Fit for NoverA scale by scan
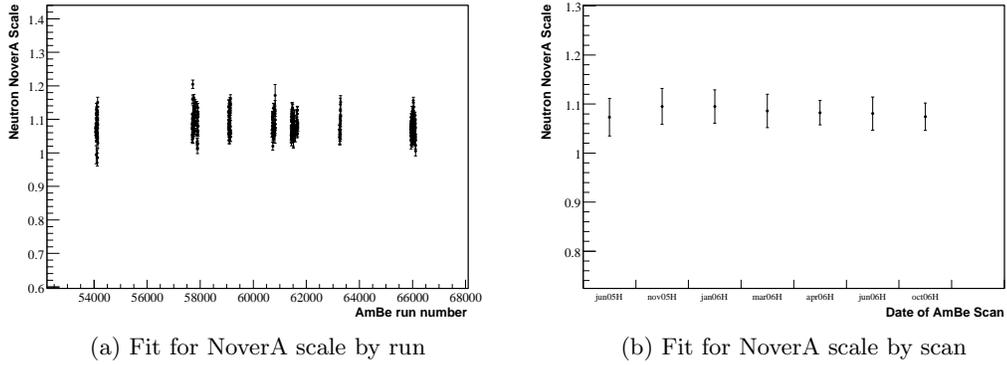
Figure 5.8: Determination of neutron NoverA scale, shift and resolutions using AmBe-HI
data. The left panel shows run by run fits and the right-side panel shows the parameter
averaged over the scans using equation 5.2.

The distribution of Figure 5.9 results in a rather large uncertainty in the NoverA scale
parameter as determined from the AmBe data. This can be attributed to the non-uniform
volume weighting that is a result of the AmBe being a point source. In particular, if the
fitted scale parameter is examined as a function of Z-position in the detector, as in Figure
5.10, there is a clear trend for higher fitted scale for lower Z position. This is consistent
with pulse shapes depending on the location of the event in the NCD string, as this will
result in different reflections off the bottom of the counters. In considering the data from
the AmBe scans, one must then include this as an uncertainty which results in the rather
inaccurate measure of the scale parameter shown in Figure 5.9.

Table 5.2 summarizes the results for the deformation parameters extracted from the
neutron calibration data. Both ^{24}Na data sets agree with each other and the AmBe data
when only a scale parameter is used. The larger uncertainty on the AmBe data is the result
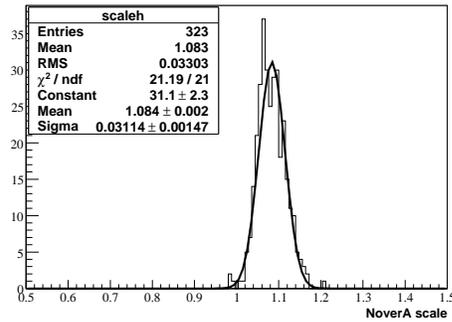
Figure 5.9: Determination of neutron NoverA scale with uncertainties using AmBe-HI data. The allowed ranges in the scale is determined by creating a histogram of the data from Figure 5.8a and fitting a normal distribution. The results are summarized in Table 5.2. There is a clear tail in the distribution to higher values of the scale which increase the uncertainty substantially. This is due to the slight bias from the Z-position of the AmBe source as evidenced in Figure 5.10
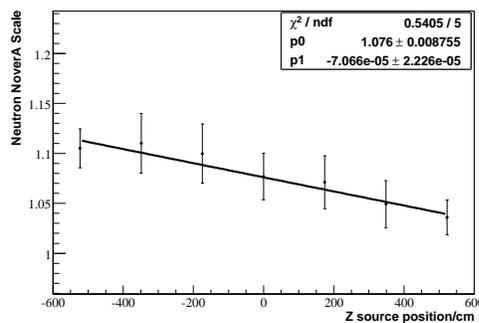
Figure 5.10: Neutron NoverA scale as a function of Z position of the AmBe source. There is a clear trend in higher fitted scale for lower values of Z, consistent with the fact that reflections off the bottom of the NCDs will alter the pulse shapes.

Data	Scale	Shift	Resolution
2005 ^{24}Na	1.016(08)	-0.005(04)	0.0010(02)
2005 ^{24}Na	1.071(07)	fixed	fixed
2006 ^{24}Na	1.049(27)	-0.010(01)	0.0010(01)
2006 ^{24}Na	1.077(03)	fixed	fixed
AmBe-HI	1.084(31)	fixed	fixed

Table 5.2:   Neutron NoverA scale, shift and resolution and uncertainties for the different neutron calibration data. The AmBe data were fit with only a scale deformation parameter and agrees with the ^{24}Na data. The final fit will use the value from the 2005 ^{24}Na data.

of a non-uniform volume weighing of the NCD array and is a conservative estimate. The final fit will thus use only a scale parameter to deform the neutron NoverA pdf and be constrained to be equal to 1.084 ± 0.031, as measured with the AmBe data, which has a larger uncertainty.

5.2 Alpha Probability Density Functions

Due to the relatively small amount of ^{4}He data, the alpha pdfs need to be determined with the help of the Monte Carlo NCD simulation described in chapter 3. The procedure for determining the alpha pdfs with systematic uncertainties will be analogous to the procedure used for the neutrons; in this case, the large amount of neutron calibration data are replaced by Monte Carlo data combined with the ^{4}He data. The uncertainties will thus reflect the variation between the Monte Carlo and data.

The Monte Carlo alpha events can look quite different for different types of alpha sources; for example the energy deposition (and pulse shape) of an alpha decay from a nucleus embedded in the NCD wall is quite different from that of a nucleus deposited on the surface of the anode wire (to take two extremes). The amount of each type of alpha source (eg. surface versus bulk) was determined from a fit to the energy spectrum of the data [67] (described in chapter 3) and these contributions are taken into account when building a 'weighted Monte Carlo' pdf.

The alpha pdfs for energy and NoverA are built by combining the pdfs from the different type of alpha particles. The uncertainties on the pdf are again determined by using scale, shift and resolution parameters that will be constrained by the agreement with the data from the ^{4}He strings. In addition to making the weighted Monte Carlo agree with the ^{4}He data, the allowed ranges for the scale, shift and resolution parameter are further determined by varying systematic uncertainties on the inputs to the Monte Carlo pdfs themselves. The full set of available Monte Carlo data with which to test the pdfs is shown in table 5.3. It includes alpha particles from ^{238}U , ^{232}Th and ^{210}Po deposited on the NCD walls, NCD wires and NCD endcaps. In real data, roughly 99% of the alphas come from the wall of the NCDs, and these are the only ones that are used in the determination of the 'weighted monte-Carlo' alpha pdfs. In addition to the different types of alphas, the data set from table 5.3 also contains wall NCD alphas that were generated with different 'input parameters' to

Label	MC Alpha Type	Label	MC Alpha Type
1	endcap Po	16	Sc Scale Po
2	endcap Th	17	Sc Scale U
3	endcap U	18	E-drift(-) U
4	surface Po	19	E-drift(+) U
5	bulk Th	20	E-drift(-) Th
6	bulk U	21	E-drift(+) Th
7	wire bulk	22	E-drift(-) Po
8	wire Po	23	E-drift(+) Po
9	Depth Po	24	ionmobility(+) Th
10	Depth Th	25	ionmobility(-) Th
11	Depth U	26	ionmobility(+) U
12	Sc gradient Po	27	ionmobility(-) U
13	Sc gradient Bulk	28	^{4}He Data
14	Sc offset Po		
15	Sc offset U		

Table 5.3: Types of alphas generated by the Monte Carlo [67]. Types 1-8 are 'central value' simulations, whereas the others have had 'input parameters' of the Monte Carlo modified. The weighted average of Monte Carlo data contains only types 4-6, which are determined to represent 99% [67] of the data in the neutron energy ranges. Systematic studies will however compare the weighted Monte Carlo pdf to all 27 types of Monte Carlo alphas as well as the ^{4}He data.

the Monte Carlo. These input parameters cover the depth of the alpha emitting materials, parameters describing the electron avalanche (gradient, scale and offset from a Diethorn model) as well as electron and ion drift constants. These are described in detail in [67].

5.2.1 Energy

The energy pdf is determined by using the Monte Carlo data, where the contributions from different types of alphas are weighted appropriately [67]. Energy fits to the different strings in energy ranges above the neutron peak were performed and used to determine string-by-string fractions of the types of alphas. It was shown that the contributions from wire and endcap alphas could safely be ignored for an analysis in the neutron energy region. The weighted Monte Carlo thus only contains contributions from the surface (polonium) and

bulk (uranium and thorium) components.

The uncertainty in the pdf is then modeled as a scale and shift of the x-axis as well as a convolution with a (gaussian) resolution function (in the same manner as the neutrons pdfs were treated). The ranges of the scale, shift and resolution parameters are determined by deforming the Monte Carlo pdf to agree with the ^{4}He data.

Figure 5.11a shows the energy pdfs for the Monte Carlo alphas (types 1-8) as well as the ^{4}He data, panel b) shows the Monte Carlo spectrum when the different contributions are weighted in appropriately, whereas panel c) shows the fit of the weighted Monte Carlo pdf (with a scale, shift and resolution applied) to the ^{4}He data. The deduced values for scale, shift and resolution are summarized in Table 5.4.

The resolution is much smaller than the width of the bins and has no effect, so that it may be ignored. In addition, the scale and shift parameters are very correlated and effectively only trade-off with each other to match the discrepancy at the low energy threshold between the Monte Carlo and the data. For this reason, it is, again, suitable to deform the alpha energy pdf using only one parameter; either a scale or a shift. The shift and scale parameters when fit individually are also shown in Table 5.4. The fitted scale is very far from 1 and does not make much physical sense, whereas the shift parameter is much more convincing.

In order to gain confidence in the constraint of the energy scale or shift deformation parameter, the different types of alphas (table 5.3) that the Monte Carlo group generated can also be examined. The fits for scale and shift on those alpha data sets (types 1-28) are shown in the left panels of Figure 5.12. The ^{4}He data (the last point on the graphs) gives the largest deformation from the Monte Carlo based pdf. This is not surprising as the mismatch appears to come from a different threshold in energy between data and Monte Carlo. The results are summarized in table 5.4.

For the ensemble of Monte Carlo data, the scale, shift and their uncertainties were deduced by fitting a normal distribution to the histogram of the values in the left panels of Figure 5.12. As the shift parameter disagrees less with the Monte Carlo and makes more
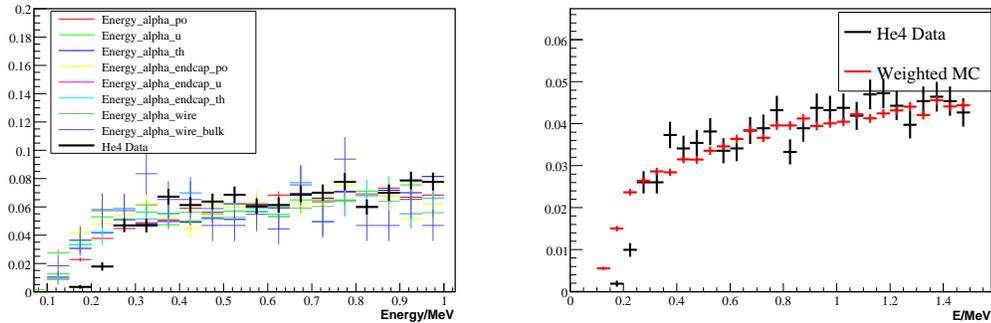
(a) Energy spectra of Monte Carlo alphas and ^{4}He data

(b) Energy spectra of weighted Monte Carlo and ^{4}He data
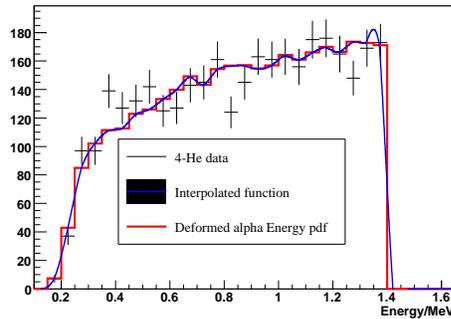
(c) Weighted Monte Carlo Spectrum fit to ^{4}He data

Figure 5.11: Determination of alpha energy pdf uncertainty with ^{4}He . Panel a) shows the energy spectra of the different types of central run Monte Carlo alphas (types 1-8) along with the ^{4}He data, although only the bulk and surface alphas (types 4-6) are used to determine the weighted Monte Carlo pdf. Panel b) shows the spectrum of the weighted Monte Carlo compared to the data from the ^{4}He strings. In panel c), the weighted Monte Carlo pdf is fit to the ^{4}He data using a scale, shift and resolution. The main difference between the Monte Carlo and the data is a shift in the low energy threshold.

sense physically, it will be used to deform the alpha energy pdf. A conservative approach dictates taking a shift parameter than can vary within the range -0.075-0.00 to cover the range from the Monte Carlo and ^{4}He data. Using this range as a 1σ bound, the alpha energy shift parameter constraint will be taken as -0.038 ± 0.038. This seemingly large uncertainty is however not expected to influence the number of neutrons very much, as it essentially only 'slides' the alpha pdf (which is quite flat) along the energy axis, and is the result of a different threshold between the Monte Carlo and the data.
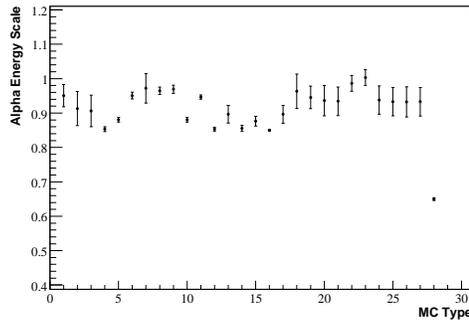
(a) Fits for energy scale by MC type
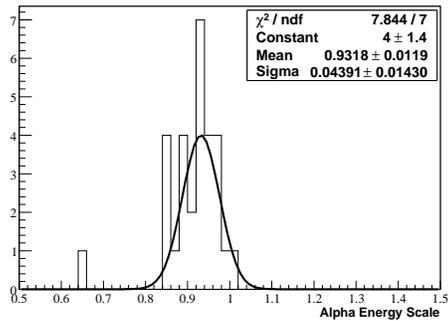
(b) Histogram of fitted energy scale by MC type
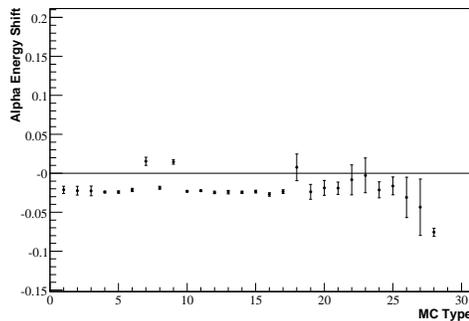
(c) Fits for energy shift by MC type
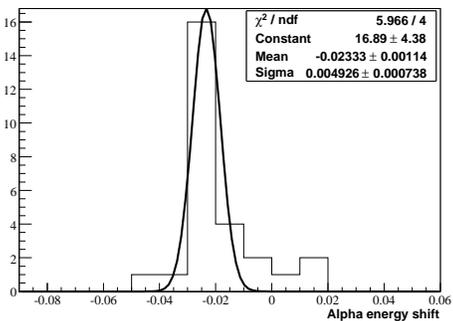
(d) Histogram of fitted energy shift by MC type

Figure 5.12: Determination of alpha energy scale and shift with the ensemble of Monte Carlo data. The points are ordered in the same fashion as in table 5.4. The last point in the graph is from the ^{4}He data. The point from the data disagrees with the Monte Carlo points because of differing energy thresholds between the two. In both cases, the scale or shift parameter were fit as a single deformation parameter. Using only a shift produces a better agreement between the ^{4}He data and the Monte Carlo.

Data	Scale	Shift	Resolution
^{4}He	0.978(48)	-0.070(15)	0.007(11)
^{4}He	0.650(04)	fixed	fixed
^{4}He	fixed	-0.075(05)	fixed
MC ensemble	0.932(44)	fixed	fixed
MC ensemble	fixed	-0.023(05)	fixed

Table 5.4: Alpha energy scale, shift and resolution uncertainties as determined by fitting the ^{4}He data to the weighted Monte Carlo pdf as well as the available ensemble of Monte Carlo alphas. The alpha energy scale and shift from ^{4}He disagree with the Monte Carlo due to a shift in threshold that can be seen in Figure 5.11. Since the shift agrees better, it will be used as the deformation parameter. This results in a large uncertainty in the alpha energy shift, taken to lie between -0.075 and 0 so that it agrees with the ^{4}He data and the Monte Carlo ensemble.

5.2.2 NoverA

The alpha pdf for χ_n^2/χ_α^2 is determined in exactly the same way as the energy pdf. The NoverA pdf for the different Monte Carlo alphas and the ^{4}He data is shown in Figure 5.13a and the weighted Monte Carlo NoverA pdf is shown in panel b). The fit of the Monte Carlo pdf with scale, shift and resolution applied is shown in Figure 5.13c and panel (d) shows the fit with only a scale applied. In this case, it is clear that the NoverA pdf for the Monte Carlo is quite different than it is for the ^{4}He data. The Monte Carlo is considerably narrower than the data, which results in a large resolution to be fit out. However, the resolution, scale and shift end up being very correlated as seen in the 2D likelihoods plotted in Figure 5.14. The shape of the pdf is quite 'simple' and Figure 5.13d shows that a scale parameter is, in fact, sufficient to model the change in shape between the Monte Carlo and the ^{4}He data. The fit quality is slightly worse, as there is less freedom in deforming the pdf but it still produces a good fit and a more accurate constraint on the scale. The fitted values are summarized in Table 5.5.

Again, in order to gain confidence in the measurement of the alpha NoverA scale, the ensemble of Monte Carlo data (types 1-28) is examined. Figure 5.15 shows the fits for the NoverA scale to the ensemble of Monte Carlo data. A large variation in scale is seen
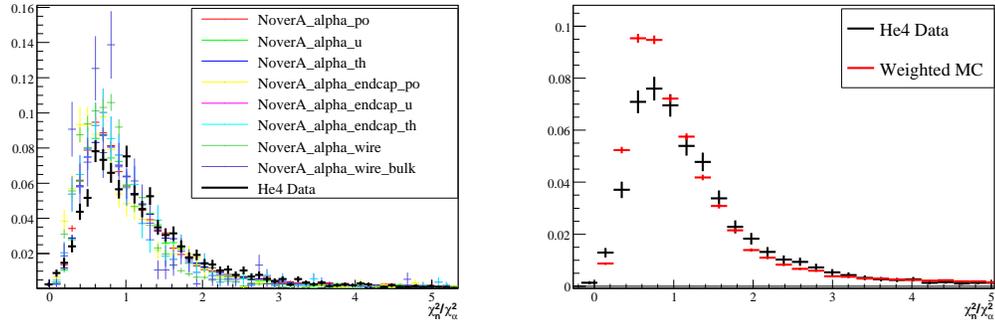
(a) NoverA of Monte Carlo alphas and ^{4}He data

(b) NoverA pdf of Monte Carlo alphas and ^{4}He data
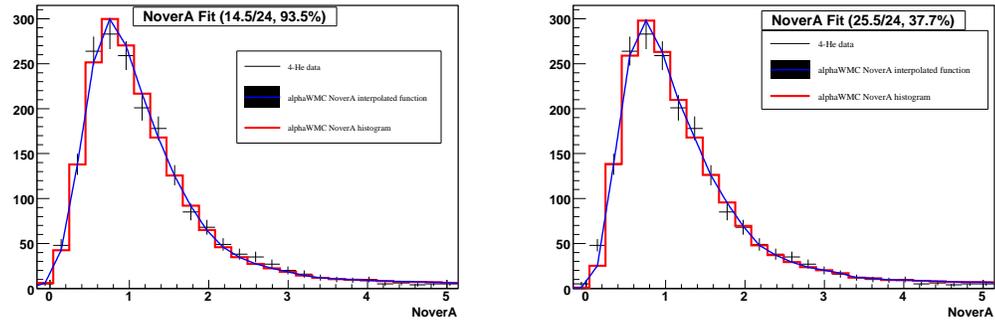
(c) Monte Carlo NoverA fit to ^{4}He data with resolution, scale and shift (shown with goodness of fit)

(d) Monte Carlo NoverA fit to ^{4}He data with only scale (shown with goodness of fit)

Figure 5.13: Determination of alpha NoverA pdf uncertainty. Panel a) shows the NoverA pdf of the different types of alphas modeled in the Monte Carlo as well as the ^{4}He data; the (un-deformed) weighted Monte Carlo alpha pdf and the data from the ^{4}He strings are shown in panel b). Panel c) shows the weighted Monte Carlo pdf fit to the ^{4}He data using a scale, shift and resolution; the correlation between these parameters is seen in Figure 5.14. Panel d) shows the weighted Monte Carlo pdf fit to the ^{4}He data using only a scale. Although the quality of fit in panel (d) is not as good as in panel (c) it is still acceptable and justifies the use of only a scale parameter to deform the alpha NoverA pdf.

(a) scale vs shift

(b) resolution vs scale

(c) resolution vs shift
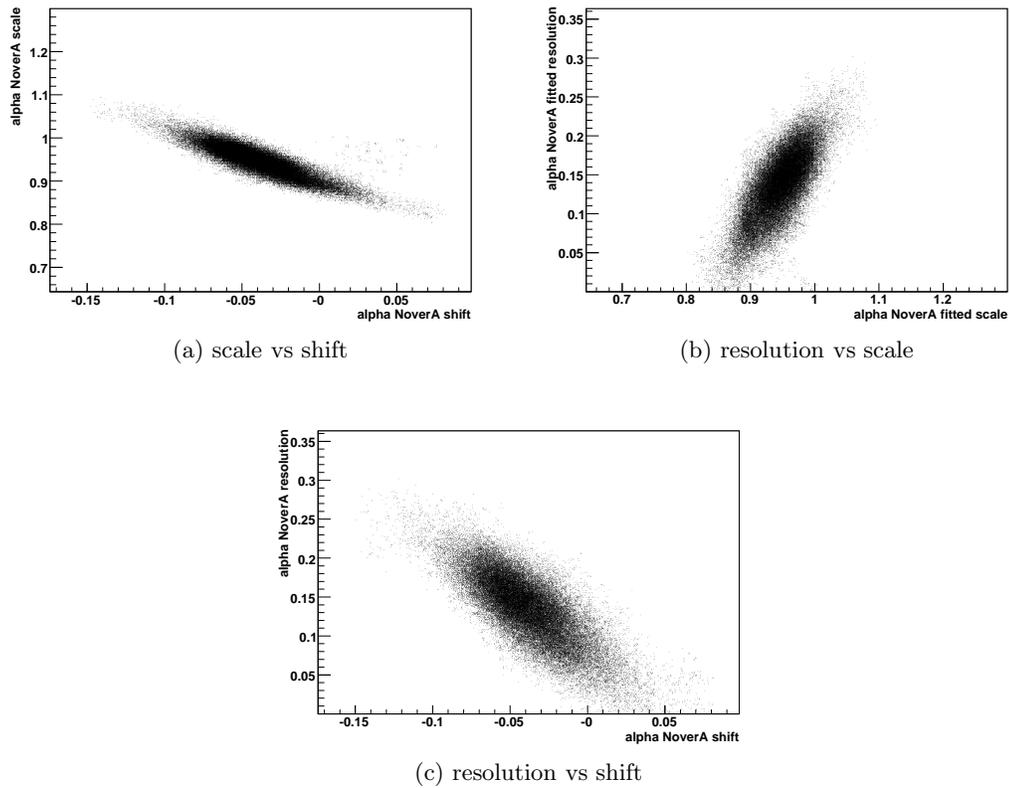
Figure 5.14: These figures show the correlations between the fitted scale, shift and resolution from the fit in Figure 5.13c. These plots are made by creating a histogram of the parameters at each step in the MCMC chain (described in chapter 6) and are thus proportional to the marginalized likelihood as function of the pair of parameters. The obvious slope in the functions reveals a large correlation between the scale, shift and resolution.

although it is consistent with the value from the ^{4}He data. The outliers in Figure 5.15 come from types of alphas that are not a large contribution to the data (such as alpha decays in the endcaps and wire); hence, giving them equal weight in the histogram may seem overly conservative. It should be however noted that, although the Monte Carlo was designed to model pulse shapes, it was not 'vetted' as extensively for that aspect as it was for creating the energy pdf. Oblath [70] nonetheless used the pulse shapes from the Monte Carlo in his doctoral thesis and carried out numerous verifications. It is the philosophy of the present work that allowing the pdf to fluctuate with little constraint on the alpha NoverA scale parameter allows for imperfection in the alpha pulse shape Monte Carlo to be taken into account. Thus, the constraint on the alpha NoverA scale parameter will be taken from the Monte Carlo result, as it agrees with the value from ^{4}He data but has a larger uncertainty. The alpha NoverA scale will be taken as 0.967 ± 0.079

(a) Fits for NoverA scale by MC type   (b) Histogram of fitted NoverA scale by MC type
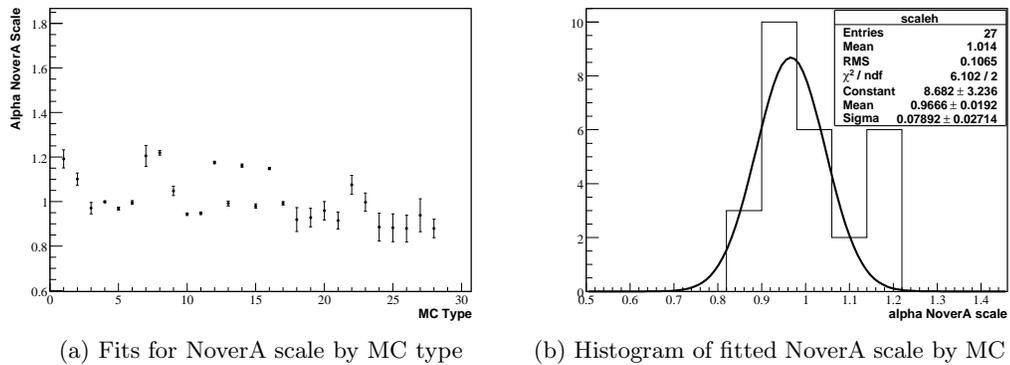
Figure 5.15: Determination of alpha NoverA scale with the ensemble of Monte Carlo data. The points are ordered in the same fashion as in table 5.4. The last point in panel (a) is from the ^{4}He data.

Data	Scale	Shift	Resolution
^{4}He	0.951(33)	-0.042(24)	0.144(37)
^{4}He	0.890(10)	fixed	fixed
MC ensemble	0.967(79)	fixed	fixed

Table 5.5: Alpha NoverA scale, shift and resolution uncertainties as determined by fitting the ^{4}He data to the weighted Monte Carlo pdf. The value of NoverA scale from the MC ensemble is chosen as the constraint for the analysis, as it agrees with the ^{4}He data and yields a conservative estimate.

5.3 Pathological (NNNA) Event Probability Density Function

Two types of events in the NCD data have been identified that do not have the characteristics of neutrons or alphas. These 'non-neutron-non-alpha' (NNNA) events were seen in strings 0 and 26 and identified by a distortion of the energy spectrum of those two strings. Since it has not been possible to determine a physical explanation for these events, it has been assumed that similar events could be present at a lower rate in other strings. For this reason, pdfs in energy and NoverA are determined for both classes of events, which will be referred to as nnna0 and nnna26 events.

The only way of determining pdfs for these events is by comparing the energy spectrum of the bad strings with the spectrum for the rest of the NCD array. If one subtracts the 'normal' spectrum from that of the problematic strings, one is left with the desired pdf. The only issue is where to perform the normalization of the spectra before the subtraction; since the single strings presumably have a different ratios of surface to bulk alpha emitter contamination, the normalization should not necessarily be done, for example, about the 5 MeV ^{210}Po peak.

Figure 5.16 shows the energy spectra of the two problematic strings normalized from 1 MeV to 3 MeV compared with the spectrum from the rest of the NCD array. The NNNA events are clearly visible in both strings as the peaks at low energies, around 0.2 MeV.

(a) String 0 spectrum
(b) String 26 spectrum
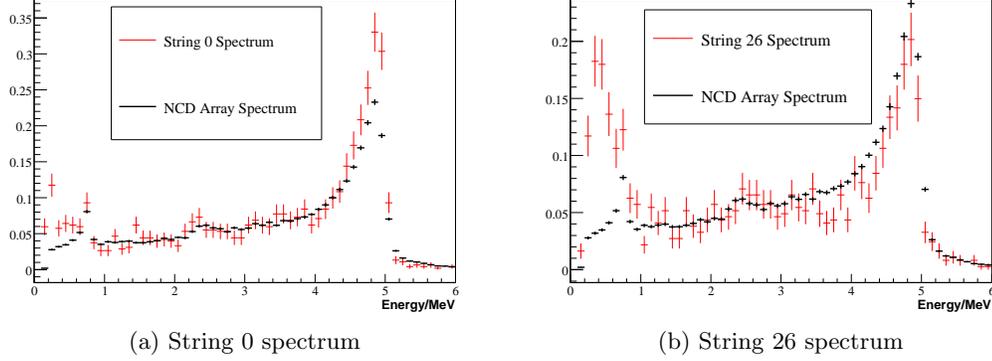
Figure 5.16: String 0(a) and 26(b) energy spectra normalized between 1 MeV and 3 MeV compared with energy spectrum from the 'good' strings in the NCD array. The NNNA events are characterized by the bump in the energy spectra at low energies (around 0.2 MeV). One must be careful in where the spectra are normalized since individual strings are likely to have different amounts of ^{210}Po contamination (the peak at 5 MeV).

5.3.1 Energy

The energy pdfs with uncertainties are determined by subtracting the average NCD array spectrum from the energy spectrum of the offending strings (0 and 26). In order to account for different proportions of ^{238}U , ^{232}Th and ^{210}Po between the offending strings and the rest of the array, the subtraction is performed with different normalizations. The normalizations are taken in 4 different energy ranges that change the weighing of the ^{210}Po peak. This results in 4 different energy pdfs for each string (shown in panel (a) of Figures 5.17 and 5.18). The energy pdf for each string is obtained by taking the average of the four spectra. Since the statistical uncertainties in each bin are large, these are taken into account when calculating the average value in each energy bin, $\hat{\mu}$, and the uncertainty in each energy bin, $\hat{\sigma}$:

$$\hat{\mu} = \frac{\sum \frac{\mu_i}{\sigma_i^2}}{\sum \frac{1}{\sigma_i^2}}$$

$$\hat{\sigma} = \frac{1}{\sqrt{\sum_{i=0}^{N} \frac{1}{\sigma_i^2}}} \tag{5.3}$$

Panel (b) of Figures 5.17b and 5.18 show the resulting energy pdf for nnna0 and nnna26 events with error bars. In order to quantify systematic uncertainties, one does not have the ability to constrain deformation parameters, in the same way as was done for neutrons and alphas, since there is not enough data to impose such a constraint. Most of the uncertainty is the result of the low statistics in each bin.

Barlow and Beeston have resolved the problem of using pdfs with low statistics [71]. Their approach was to treat the bin values as parameters in the likelihood fit that are Poisson distributed about the bin centers. Since the ultimate (correct) values in each bin for the pdfs are not required, an effective integration over the possible (allowed) pdf shapes can be performed. The bootstrap method [72] is a way to perform this integration. In the bootstrap method, the integration over possible pdfs is performed by drawing new pdfs from the original pdf. That is, a new pdf can be built bin-by-bin by choosing a value based on the bin value and error of the original pdf. A novel approach is adopted in this thesis that reconciles the philosophies from both methods. The NNNA pdfs will be constructed as the linear combination of 10 random (bootstrap) pdfs drawn from the original pdf with uncertainties. The likelihood fit will then float the coefficients of the linear combination. The procedure adopted in this thesis is described in more detail in 6.2.3.

5.3.2 NoverA

The NoverA pdf for the two type of NNNA events is obtained in a similar way as energy. Since these NNNA events are distinguishable in energy, the normalizations for the subtraction is done in energy. The NoverA pdfs are then obtained by subtracting the NoverA pdf for the NCD array from those of the offending strings, using the same 4 different normalizations as for the energy pdf. This results in the variations in the NoverA pdf and uncertainties shown in Figures 5.19 and 5.20 for nnna0 and nnna26 events, respectively. Again, an average pdf with uncertainties in each bin is calculated and will be used in the likelihood extraction to draw random pdfs.
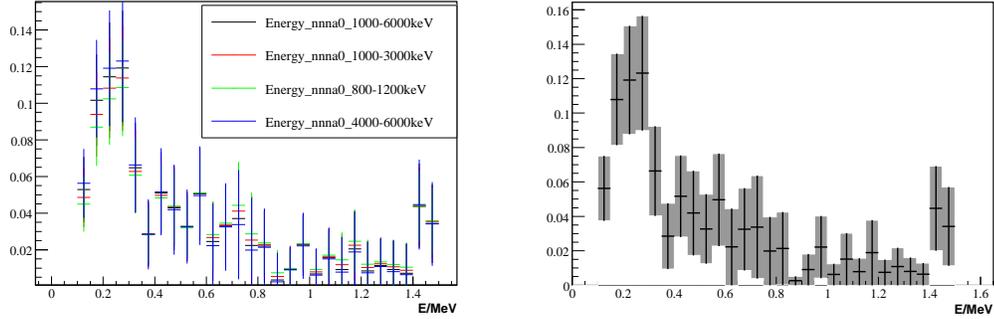
(a) Energy spectrum of nnna0 events for different (b) Energy spectrum of nnna0 events with systematic uncertainty
normalizations

Figure 5.17: Determination of nnna0 energy pdf uncertainty. Panel a) shows the energy spectra (with statistical uncertainties) for string 0 NNNA events for different normalization between the string 0 spectrum and the average spectrum from the NCD array. Panel b) shows the resulting pdf with systematic uncertainties that take into account the variations from panel a) using equation 5.3.
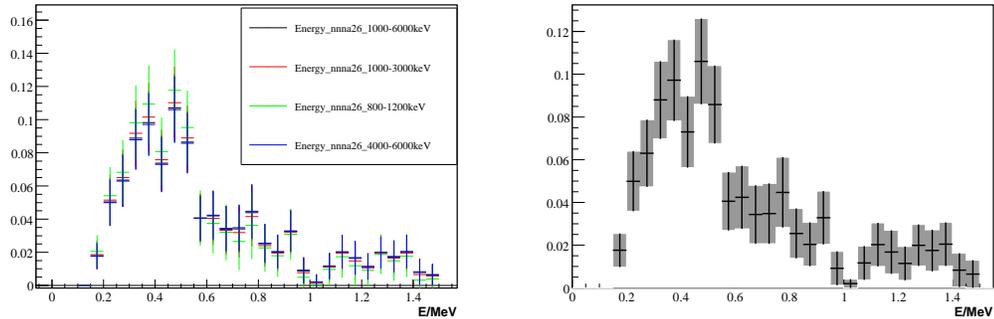
(a) Energy spectrum of nnna26 events for different (b) Energy spectrum of nnna26 events with systematic uncertainty
normalizations

Figure 5.18: Determination of nnna26 energy pdf uncertainty. Panel a) shows the energy spectra (with statistical uncertainties) for string 26 NNNA events for different normalization between the string 26 spectrum and the average spectrum from the NCD array. Panel b) shows the resulting pdf with systematic uncertainties that take into account the variations from panel a).
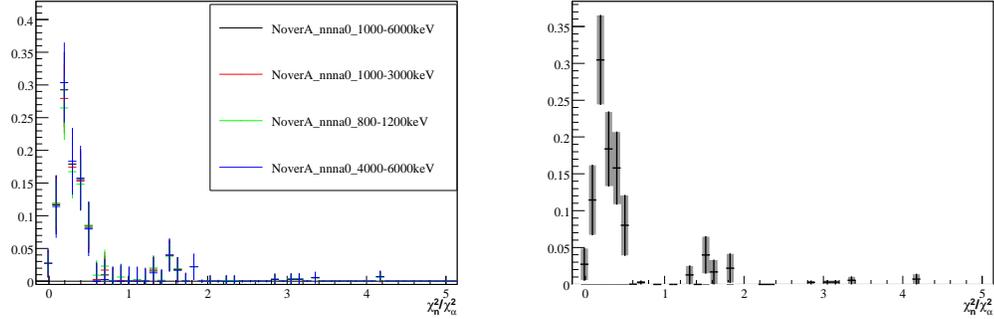
(a) NoverA spectrum of nnna0 events for different normalizations

(b) NoverA spectrum of nnna0 events with systematic uncertainty

Figure 5.19: Determination of nnna0 NoverA pdf uncertainty. Panel a) shows the NoverA distribution (with statistical uncertainties) for string 0 NNNA events for different normalization between the string 0 energy spectrum and the average energy spectrum from the NCD array. Panel b) shows the resulting pdf with systematic uncertainties that take into account the variations from panel a).
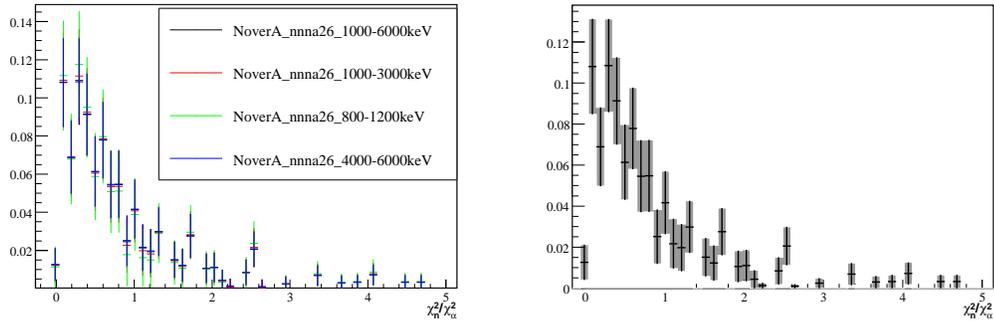
(a) NoverA spectrum of nnna26 events for different normalizations

(b) NoverA spectrum of nnna26 events with systematic uncertainty

Figure 5.20: Determination of nnna26 NoverA pdf uncertainty. Panel a) shows the NoverA distribution (with statistical uncertainties) for string 26 NNNA events for different normalization between the string 26 energy spectrum and the average energy spectrum from the NCD array. Panel b) shows the resulting pdf with systematic uncertainties that take into account the variations from panel a)

.

5.4 Summary

This chapter introduced the technique for determining the pdfs and their systematic uncertainties for the energy and NoverA observables. Four classes of signals (neutrons, alphas, nnna0s and nnna26s) were examined for use in the maximum-likelihood extraction of the NCD data. The average pdfs were built as histograms. For neutrons, different types of calibration data (AmBe and ^{24}Na) were used to determine a spread in the possible shapes of the pdfs, and it was shown that all reasonable shapes can be accommodated by the application of a simple scaling of the x-axis. For alphas, the variation in the pdfs was taken to agree with Monte Carlo events and data from the ^{4}He strings. It was shown that a shift parameter was sufficient to encompass the uncertainty in the shape of the energy pdf and a scale parameter was suitable for the NoverA pdfs. Finally, pdfs were obtained based on pathological events that were detected in strings 0 and 26. The pdfs for these types of events were obtained by comparing the spectra of the offending strings with the average NCD array. Since the uncertainty in these pdfs is primarily due to statistics, the pdfs will be handled in a way that is inspired by a combination of the Barlow-Beeston and bootstrap methods. Figure 5.21 shows a comparison of the energy and NoverA pdfs for all four signals highlighting the differences that neutrons have from the other signals. The difference between neutrons and the other three signals is clearly increased by the additional use of the NoverA observable.
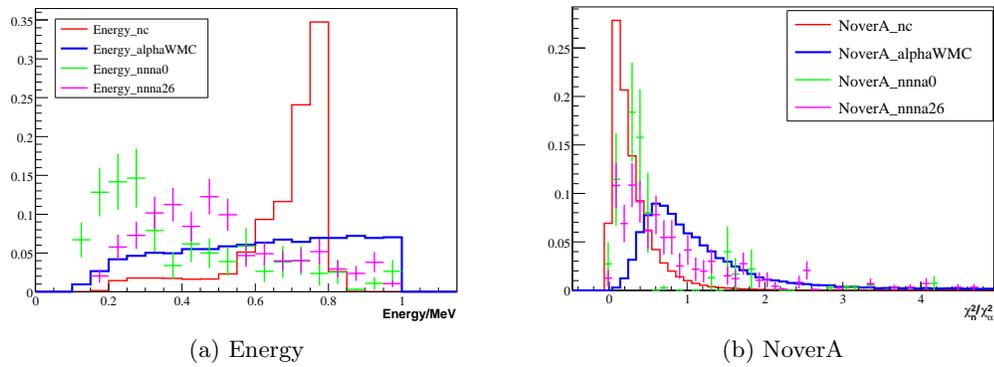
(a) Energy

(b) NoverA

Figure 5.21: Energy and NoverA pdfs with uncertainties for neutrons, alpha, nnna0 and nnna26 events. The difference between neutrons and the other signals in this two-parameter space is seen clearly.

Chapter 6

Description of Markov-Chain-Monte Carlo Fitting

In this chapter, the techniques for parameter determination used in this thesis are introduced. The motivation is to present the framework that will be used to analyze the NCD data and, ultimately, determine the number of neutrons with their uncertainty. It will be shown that a Bayesian approach using a Markov-Chain Monte Carlo (MCMC) is well suited for this classical parameter determination problem. The chapter starts by presenting the framework for parameter estimation using extended maximum-likelihood then introduces the Bayesian approach and the MCMC method. The extended maximum-likelihood function will be built using the pdfs for the energy and NoverA observables that were introduced in the previous chapter.

6.1   Parameter Fitting

The ultimate goal of this analysis is to determine the number of neutrons in the NCD data using the available observables and signals to create a model. The observables that are used in this analysis are energy and NoverA and the signals are, $(n, \alpha, nnna0, nnna26)$: the

number of neutrons, the number of alphas and the number of NNNA events, respectively. These two observables and four signals will be used to model the data by creating a likelihood function for the contribution of each of the four signals, which become the parameters of the problem. In addition to the four parameters that are desired, several 'nuisance' parameters, $\vec{\mu}$, will also be introduced to model systematic uncertainties; the final values of the signals will have to take into account the correct marginalization over these parameters (within their constraints from calibration data, if applicable). Examples of the nuisance parameters are the deformation parameters for the pdfs that were introduced in chapter 5. The vector, $\vec{p}$, will be used to represent the entire set of signal and nuisance parameters.

It should be noted that the 'systematic parameters' will be floated and the data will be used to constrain these. This 'floating systematic uncertainties' approach is a consistent way to handle systematic uncertainties [73] and takes into account the fact that the model for the data ought to be constrained by the data (and not only the calibration data used to make the model). Thus, the deformation parameters for the pdfs, although constrained by the analysis in chapter 5, will be considered here as floating parameters of the model. This section describes how the likelihood function is constructed.

6.1.1 Probability Density Functions

In order to use maximum-likelihood techniques, one requires a way of assigning the probability of a certain signal to have a certain value of an observable. For example, $P_n^E(E|\vec{\mu})dE$ will be defined to be the probability for a neutron to have an energy between E and $E+dE$, where the function depends on the systematic deformation parameters, $\vec{\mu}$. $P_n^{NoverA}(\chi|\vec{\mu})d\chi$ will be the probability of a neutron to have NoverA between χ and $\chi + d\chi$. The functions $P_n^E(E|\vec{\mu})$ and $P_n^{NoverA}(\chi|\vec{\mu})d\chi$ must satisfy the requirements of probability density functions

(pdfs) and must therefore be normalized such that:

$$\int_{AllE} P_n^E(E|\vec{\mu})dE = 1$$

$$\int_{All\chi} P_n^{NoverA}(\chi|\vec{\mu})d\chi = 1 \tag{6.1}$$

$$\tag{6.2}$$

where the limits on the integrals extend over the entire range of the function (as opposed to just the fit range). Since there is no analytical model for the pdfs, these must be determined from data. In chapter 5, histograms to approximate the desired functions were obtained and the probability density functions are then built from these histograms. In order to avoid discontinuities that come from the size of the histogram bins, the pdfs are constructed by using a linear interpolation[1] from the center of the histogram bins, after the histogram has been normalized to have a unit area. The systematic deformation parameters can then be included in the interpolated histogram to give the actual probability density function. Figure 6.1 shows an example of the neutron energy pdf histogram along with the linear interpolated pdf, $P_n^E(E)$ (no deformation parameters applied).

6.1.2 Maximum Likelihood

In a maximum-likelihood fit (see, for example [35], [74], [75]) for the parameters $\vec{p}$ using the data, D, one needs to calculate the likelihood, $\mathscr{L}(D|\vec{p})$, of the data *given* the model (i.e. a choice of the parameters). The data consists of a set of N events which each have a measured value of energy, E_i, and NoverA, χ_i. Because the events are independent of each other, the likelihood of the entire data is the product of the likelihoods, $P(E_i, \chi_i|\vec{p})$, of each event i:

$$\mathscr{L}(D|\vec{p}) = \prod_i^N P(E_i, \chi_i|\vec{p}) \tag{6.3}$$

---

[1]Using a linear interpolation preserves the normalization of the histogram
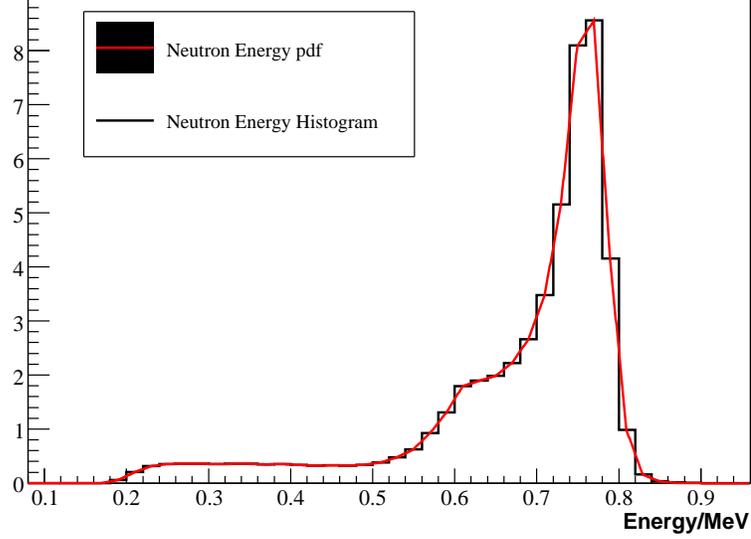
Figure 6.1: Example of interpolation for the neutron energy pdf. The histogram is normalized such that the sum of each bin content times the bin width is equal to 1 and the pdf function is then a linear interpolation between the bin centers.

The likelihood of each event is obtained by adding the probabilities that it is from each of the classes:

$$P(E_i, \chi_i | \vec{p}) = n' \times P_n(E_i, \chi_i | \vec{\mu}) + \alpha' \times P_\alpha(E_i, \chi_i | \vec{\mu})$$

$$+ \, nnna0' \times P_{nnna0}(E_i, \chi_i | \vec{\mu}) + nnna26' \times P_{nnna26}(E_i, \chi_i | \vec{\mu}) \qquad (6.4)$$

where the probabilities for each class depend on the systematic parameters, $\vec{\mu}$ and the dependence on the signal parameters $(n', \alpha', nnna0', nnna26')$ is shown explicitly. One should note that the signal parameters in this expression are fractions (denoted by the apostrophe) constrained such that they sum to 1; for example, the fraction of neutrons in the data, n', is defined as $n' = n/N$. $P_n(E, \chi | \vec{\mu}) dE d\chi$ is then the joint probability of a neutron having an energy between E and $E + dE$ and NoverA between χ and $\chi + d\chi$. In the case where the observables for energy and NoverA are uncorrelated (see section 7.4 for tests of the correlation between energy and NoverA), one can factor the joint probability

$P_n(E, \chi | \vec{\mu})$ into the product of the 1-dimensional pdfs obtained from the histograms of the previous section:

$$P_n(E, \chi | \vec{\mu}) = P_n^E(E | \vec{\mu}) \times P_n^{Nover A}(\chi | \vec{\mu}) \tag{6.5}$$

One then wants to maximize $\mathscr{L}(D | \vec{p})$, with respect to the parameters and find the value of $\vec{p}$ that corresponds to the location of 'maximum-likelihood':

$$\frac{d\mathscr{L}(D | \vec{p})}{d\vec{p}} = 0 \tag{6.6}$$

In practice, it is more convenient to work with the logarithm of the likelihood, $L(D | \vec{p}) = \log(\mathscr{L}(D | \vec{p}))$, which will have the same extrema as the likelihood but is computationally easier to work with:

$$\begin{aligned}
L(D | \vec{p}) &= \log\left(\mathscr{L}(D | \vec{p})\right) \\
&= \log\left(\prod_i^N P(E_i, \chi_i | \vec{p})\right) \\
&= \sum_i^N \log\left(P(E_i, \chi_i | \vec{p})\right) \\
&= \sum_i^N \log\left(n' \times P_n(E_i, \chi_i | \vec{\mu}) + \alpha' \times P_\alpha(E_i, \chi_i | \vec{\mu})\right. \\
&\quad \left. + nnna0' \times P_{nnna0}(E_i, \chi_i | \vec{\mu}) + nnna26' \times P_{nnna26}(E_i, \chi_i | \vec{\mu})\right)
\end{aligned} \tag{6.7}$$

It is then standard practice to work with the negative log-likelihood, $-L(D | \vec{p})$ so that this becomes a minimization problem instead of a maximization.

6.1.3 Extended Maximum-Likelihood

The derivation of the likelihood function assumed that the model for the data predicted an exact number of events, N; that is, of the four signal parameters, $(n, \alpha, nnna0, nnna26)$, only three were independent, since they need to sum to N. One can lift this restriction if one uses the 'extended maximum-likelihood' formalism, where it is assumed that the outcome (i.e. the number of events) of one particular experiment (the one being analyzed) is just one

of several possible outcomes that are Poisson distributed about some mean predicted from the model. The probability, $P^{exp}(N, N_{model})$, of one particular experiment having N events then depends on the number of events in the model, $N_{model} = n + \alpha + nnna0 + nnna26$, according to the Poisson distribution:

$$P^{exp}(N, N_{model}) = \frac{N_{model}^N e^{-N_{model}}}{N!} \tag{6.8}$$

The likelihood of the data, $\mathscr{L}(D|\vec{p})$, must then be multiplied by $P^{exp}(N, N_{model})$, and the (positive) log-likelihood contains an extra term:

$$
\begin{aligned}
L(D|\vec{p}) &= \left[\sum_i^N \log\left(P(E_i, \chi_i|\vec{p}) \right) \right] + \log\left(P^{exp}(N, N_{model}) \right) \\
&= \left[\sum_i^N \log\left(P(E_i, \chi_i|\vec{p}) \right) \right] + N \log(N_{model}) - N_{model} - \log(N!) \\
&= \left[\sum_i^N \log\left(P(E_i, \chi_i|\vec{p}) \times N_{model} \right) \right] - N_{model} \\
&= \left[\sum_i^N \log\left(n \times P_n(E_i, \chi_i|\vec{\mu}) + \alpha \times P_\alpha(E_i, \chi_i|\vec{\mu}) \right. \right. \\
&\qquad \left. \left. + nnna0 \times P_{nnna0}(E_i, \chi_i|\vec{\mu}) + nnna26 \times P_{nnna26}(E_i, \chi_i|\vec{\mu}) \right) \right] \\
&\qquad - (n + \alpha + nnna0 + nnna26)
\end{aligned} \tag{6.9}
$$

where the term $\log(N!)$ has been dropped since it is a constant that will not affect where the maximum-likelihood is located and the fractions from equation 6.7 have been multiplied by N_{model} so that the only parameters left are the actual number of events from each signal ($n = n' \times N_{model}$), and not the fractions. By using the extended maximum-likelihood the sum of the signal parameters is no longer restricted to be N. The last term in equation 6.9 serves as a normalization and restricts each signal from becoming infinitely large (without this term, there would be no absolute maximum, since the likelihood would increase indefinitely with the signals).

6.1.4 Fit Ranges and Normalizations

Care must be taken in correctly normalizing the pdfs when 'fit ranges' are applied to the data; that is, when the likelihood is only calculated for events within a certain range of an observable. For example, it will be of interest to study the number of neutrons depending on the energy range that is chosen to fit the data. In particular, it must be guaranteed that the same number of neutrons is inferred by looking at data from 0.4

,MeV.to 1

,MeV.as for the entire range of energies (so that the fit range may be optimized for the analysis). In order to simplify the derivation, a notation is introduced to rewrite equation 6.9 as:

$$L(D|\vec{p}) = \sum_i^N \log\left(\sum_j^M a_j P^j(\vec{E})\right) - \sum_j^M a_j \tag{6.10}$$

where j is a sum over the M types of signals, a_j ($=\{(n, \alpha, nnna0, nnna26)\}$), and $P^j(\vec{E})$ is the joint probability for a certain type of signal to have a set of observables $\vec{E}$.

It is of interest to maintain the values of each signal, a_j, independent of the cuts that are applied to the data. An example of a cut on the data would be to keep events that have an energy between 0.4

,MeV.and 1

,MeV.as well as NoverA between 0.1 and 5. Cuts are implemented such that events need to satisfy cuts in **all** observables in order to be kept.

Let the (potentially multi-dimensional) cut in observables be defined by keeping events between $\vec{E}_{min}$ and $\vec{E}_{max}$. 'Pass fractions', f_j, can be defined for each class of event:

$$f_j \equiv \frac{\int_{\vec{E}_{min}}^{\vec{E}_{max}} P^j(\vec{E})d\vec{E}}{\int_{All\vec{E}} P^j(\vec{E})d\vec{E}} \tag{6.11}$$

such that

$$a_j^{Cut} = f_j \times a_j \tag{6.12}$$

where a_j^{Cut} is the amount of signal j in the cut data. In the case of uncorrelated observables, the f_j are the product of integrals over each 1-dimensional pdf. If a cut is applied to the data, the sum over N events in equation 6.10 becomes a sum over the number of events that pass the cuts, N^{Cut}, and the a_j should be replaced by a_j^{Cut}. In addition, the probabilities, $P^j(\vec{E})$, in equation 6.10 need to be normalized over the cut range. If the normalizations from equation 6.1 are retained, then probabilities in equation 6.10 need to be replaced by $P^j(\vec{E})/f_j$ (which normalizes them in the cut range) so that the likelihood is now given by:

$$
\begin{aligned}
L(D|\vec{p}) &= \sum_i^{N^{Cut}} \log\left(\sum_j^M a_j^{Cut} P^j(\vec{E})/f_j\right) - \sum_j^M a_j^{Cut} \\
&= \sum_i^{N^{Cut}} \log\left(\sum_j^M f_j \times a_j P^j(\vec{E})/f_j\right) - \sum_j^M f_j \times a_j \\
&= \sum_i^{N^{Cut}} \log\left(\sum_j^M a_j P^j(\vec{E})\right) - \sum_j^M f_j \times a_j
\end{aligned}
\tag{6.13}
$$

and the effect of applying the cuts to the data only modifies the last term and involves a simple calculation to determine the pass fractions, f_j. Thus, in equation 6.13, a_j correspond to the signals within the whole (uncut) data and $P^j(\vec{E})$ are normalized over the entire range.

6.2 The Inclusion of Systematic Uncertainties

In chapter 5, systematic uncertainties in the shape of the pdfs were introduced by allowing the pdfs to be deformed using a set, $\vec{\mu}$, of scale, shift and resolution parameters applied to each pdf. As an example, the case for the neutron energy pdf is worked out here. If the interpolating function from the histogram is denoted by $I(E)$, the neutron energy pdf with

systematic parameters, $P_n^E(E|S_n, O_n, \sigma_n)$ is given by:

$$P_n^E(E|S_n, O_n, \sigma_n) = g\left(S_n(E + O_n), \sigma_n\right) \otimes I\left(S_n(E + O_n)\right)$$

$$= \int_{AllE'} \frac{1}{\sigma_n\sqrt{2\pi}} \exp\left(\frac{-\left(E' - S_n(E + O_n)\right)^2}{2\sigma^2}\right) I\left(S_n(E + O_n) - E'\right) dE'$$

$$(6.14)$$

where S_n is the scale, O_n is the shift (offset) and σ_n is the width of the convolution gaussian, $g(S_n(E + O_n), \sigma_n)$. The convolution integral was explicitly written out in the second line. Once the pdf from equation 6.14 is used in the likelihood calculation, the systematic deformation parameters become variables, just like the signals. The allowed ranges for the parameters that were obtained with the calibration data in chapter 5 will be included when the prior probabilities are discussed in section 6.4.

6.2.1   Deforming the PDFs

Implementing the systematic uncertainties can then be done straightforwardly using equation 6.14 each time the likelihood is calculated for a set of systematic and signal parameters. Each time the pdf is rebuilt, it also needs to be re-normalized. The only issue is the (lengthy) computation of the convolution to regenerate the pdf each time the likelihood function needs to be evaluated. In practice, one can make the convolution reasonably fast using a few simple approximations:

- Instead of integrating over the entire range of the observable, one can limit the integral to within $\pm 3\sigma$ of the central value when generating the convolved pdf

- If the width of the gaussian is small compared to the bin width of the original histogram, the convolution can be ignored

One should note that this method of deforming the pdf does not require rebuilding the pdfs from the data that was used to create it; instead the deformation is, in fact, applied to the interpolating function. In the case where there is a resolution, this is the only way to

proceed. However, it was seen in chapter 5 that all the pdfs can be handled with the use of only a scale or shift parameter. In such a scenario, it is possible to avoid rebuilding the pdf altogether and apply the systematic parameters to the data.

6.2.2 Deforming the Data - No Resolution

Since the deformation of the pdf is applied to the interpolating function (to where it is evaluated), one can bypass the regeneration of the pdf by simply applying the deformation to the data points. This is particularly straightforward in the case where there is no resolution function applied, which is shown in this section.

The idea is that instead of using the deformation parameters to rebuild $P_n^E(E|S_n, O_n, \sigma_n = 0)$, one can just evaluate the un-deformed pdf at a shifted/scaled data point $P_n^E(S_n(E_i + O_n))$. This will however affect the normalization of the pdf, which must be preserved. In the case where the pdf is rebuilt, it can easily be re-normalized at the same time by simply computing its integral. In the case where the deformation parameters are applied to data the normalization can be preserved analytically.

In the case of rebuilding the pdf, the normalization is:

$$\int_{AllE} P_n^E(E|S_n, O_n, \sigma_n = 0)dE = 1 \qquad (6.15)$$

where the integral is over all energies. In the case with the deformation applied to the data, it is required that:

$$\int_{AllE} P_n^E(S_n(E_+ O_n))dE = 1 \qquad (6.16)$$

If one introduces a simple change of variables, $E' = S_n(E + O_n)$, the integral becomes:

$$\int_{AllE} P_n^E(E')S_n dE' = 1 \qquad (6.17)$$

This preserves the normalization of the pdf up to the scale factor, S_n.

Looking back at the likelihood equation 6.13, this results in a few simple changes that can be incorporated by:

- Calculating the f_j with the different limits on the integral in equation 6.11. One should now integrate from $S_n(E_{min} + O_n)$ to $S_n(E_{max} + O_n)$, and use the same likelihood function

- $P_n^E(E|S_n, O_n, \sigma_n = 0)$ should now be replaced with $S_n \times P_n^E(E')$ in the likelihood. As was just shown, the pdf needs to be multiplied by the scale factor to preserve the normalization

6.2.3   Handling the NNNAs (pdfs with low statistics)

Since the pdfs for the NNNA events cannot be deformed continuously with a parameter, they need to be handled differently. In this case, the pdf is a histogram with uncertainties in each bin and any result will have to be integrated over the full range of allowed variations of the pdf. This integration can be done by using a different pdf each time the likelihood is calculated. These pdfs need to be generated such that they cover the range of shapes allowed by the histogram with the errors in each bin. A randomized pdf can be generated by filling each bin with a number drawn from a normal distribution centered about the bin content with a width equal to the uncertainty in that bin.

Figure 6.2a shows an example where 10 different pdfs were randomly generated from the nnna26 energy pdf (from Figure 5.18b of chapter 5), by drawing numbers from a normal distribution. Interpolated pdfs from the random histograms are shown in black, whereas the original histogram with errors is shown in red. It can easily be argued that most of these pdfs are unphysical as they allow big fluctuations between neighboring bins. Any physical process ought to have a smoothly varying pdf, where the bins are correlated. This can be achieved by applying a resolution function (a gaussian smoothing). It is then reasonable to take a resolution (width of the gaussian smoothing function) of the same order as the bin width to smooth out the function. The effect of applying a resolution of the same size as the bin width is shown in panel (b) and leads to a much more convincing set of pdfs.

(a) 10 random pdfs

(b) 10 random smooth pdfs
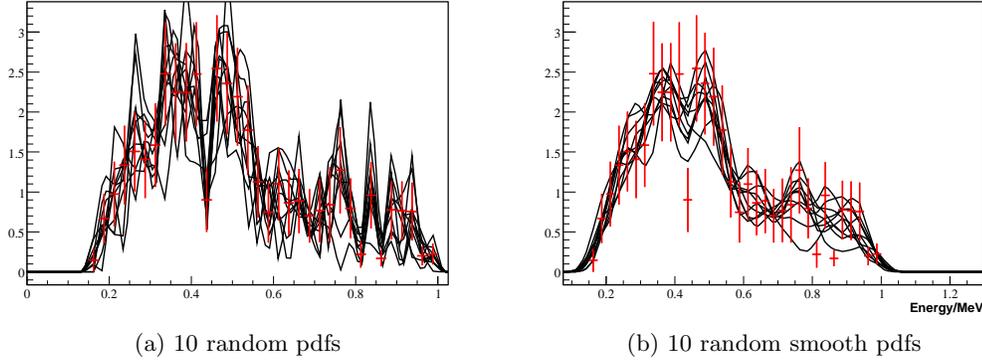
Figure 6.2: 10 randomized pdfs (black) drawn from the average nnna26 energy pdf (red, with bin errors). In panel (a) the 10 black pdfs were interpolated from 10 histograms that were drawn by randomly sampling the bin contents of the (red) average pdf. In panel (b), a resolution of the same size as the bin width was applied to the black pdfs to make them smooth. All curves are normalized.

The rest of this chapter will introduce the Markov Chain Monte Carlo method for performing the fit for the signals. It will be shown that the parameter estimation can be implemented by sampling the likelihood function. Thus, if a different NNNA pdf is used each time the likelihood is estimated, an effective integral over all the possible pdfs will be performed. This can be trivially implemented. The process can be made computationally efficient by generating a finite set (10 or so) of pdfs before the fit and choosing from that small set at random during the computation. This avoids building a random histogram each time the likelihood is evaluated.

The method for handling the systematic uncertainty in the shape of the NNNA pdf can be improved by giving different weights to the random pdfs; of the 10 curves that were generated in Figure 6.2, some will presumably be more representative of the data than others. This can be implemented by creating a new pdf which is a linear combination of the 10 randomized pdfs. The coefficients of the linear combination can then be treated as parameters of the model in the likelihood function. The maximum-likelihood fit will then yield values for the these coefficients which can then be used to generate the 'most likely

pdf', based on the data. In addition, the linear combination method allows the pdf to be deformed continuously as the parameters are varied. The analysis in this thesis will thus use a floating linear combination of random pdfs (drawn from the original histogram with uncertainties) to handle the systematic uncertainty in the shapes of the NNNA pdfs.

6.3 The Bayesian Approach

In this thesis, a Bayesian approach is adopted for analyzing the data. The main philosophy comes from Bayes' theorem, which can easily be derived from simple probability rules [74]:

$$P(X|Y) = \frac{P(Y|X) \times P(X)}{P(Y)} \tag{6.18}$$

where $P(X|Y)$ is the conditional probability of X given Y, and $P(X)$ is the probability of X. The usefulness of Bayes' theorem is that it allows for inversion of the arguments in the conditional probability.

For the purpose of data analysis considered in this thesis, one is interested in the probability, $P(\vec{p}|D)$, of the model parameters, $\vec{p}$, given the data, D. Bayes' theorem then allows that probability to be calculated by using the likelihood of the data given the parameters, $P(D|\vec{p})$. The probability of the parameters, $P(\vec{p})$, which is often called the 'prior', depends on the knowledge that was available before the experiment. Such knowledge can be embodied by the simple proposition that the number of neutrons cannot be negative, or more complicated propositions obtained from calibration data. Finally, the term $P(D)$ is the probability of the data and is sometimes given the name 'evidence'; however, since it does not depend on the model parameters, it only serves the purpose of a normalization constant[2].

---

[2]One should note that this term cannot be ignored when different models are considered, for example when comparing models with a different number of parameters. However, in this case, only the parameters *within* the model are varied and it follows that 'evidence' can be ignored [76]

Using the notation from this thesis, Bayes' theorem can be written as:

$$P(\vec{p}|D) \propto P(D|\vec{p}) \times P(\vec{p}) \qquad (6.19)$$

where the normalization from the evidence has been dropped. Equation 6.19 can then be interpreted as giving the probability of the parameters after analyzing the data. Philosophically, this is a very powerful statement that provides a guideline for analyzing data; it embodies the idea that the previous knowledge about the parameters, $P(\vec{p})$, is improved by using the data (in the form of the likelihood function) and resulting in the 'posterior' knowledge about the parameters, $P(\vec{p}|D)$.

The goal is then to determine $P(\vec{p}|D)$ as a function of $\vec{p}$ to obtain the value of the parameters as well as their uncertainties. The standard approach (in the literature) has often been to use a minimizing algorithm to search for the minimum of the negative log-likelihood and then search the nearby area of parameter space to determine 1σ errors on the parameters. The Bayesian approach adopted in this thesis will be philosophically different in that the goal will be to determine the posterior distribution functions for each parameter and infer the values of the parameters from those 'posteriors'. For example, if one is interested in the number of neutrons, the quantity of interest will be $P(n|D)$ which is given by *marginalizing* (i.e. integrating) the function $P(\vec{p}|D)$ over all parameters:

$$P(n|D) = \int n \times P(\vec{p}|D)d\vec{p} \qquad (6.20)$$

Section 6.5 introduces the MCMC method for obtaining the posterior probability distributions and interpreting them.

6.4   The Prior Distributions

The prior distributions will be used to quantify the knowledge about the model parameters that was available before analyzing the data. There are essentially two categories of prior knowledge that are available when applied to the problem considered in this thesis:

- **Signal Priors**: This should embody the prior knowledge about the amount of signals (neutrons, alphas, NNNAs) in the data. In the most 'ignorant' approach, one can, at least, say that these signals should not be negative. In that case, one assigns a 'flat prior' where all amounts of signals are equally probable but any negative signal gives a probability of zero.

- **Systematics Priors**: This part of the prior knowledge is meant to include calibration data. In this case, looking back at chapter 5, it was shown that the deformation parameters for the pdfs can be constrained by calibration data to lie within certain ranges. This part of the prior distribution is sometimes included in the likelihood and referred to as a 'constraint' or 'penalty function'; the inclusion of the information as a Bayesian prior is perhaps more appealing, philosophically.

It will be assumed that the parameters in the prior distribution are uncorrelated so that $P(\vec{p})$ can be factored:

$$P(\vec{p}) = \prod_i^M P(p_i) \tag{6.21}$$

as the product of the priors on the M parameters.

Although one could, in principle, make estimates of the amounts of signals contained in the data, the only prior will be to impose that signals cannot be negative. It is clear that by using the previous results from SNO and other solar neutrino experiments to fix the flux of neutrinos, one could constrain the number of neutrons that are produced in the detector. However, this would make the result of the present analysis correlated with the previous results and is thus avoided here. One should however note that this is exactly the method that should be used to combine the results of the different experiments or phases of SNO.

The prior probability for the deformation parameters will be modeled using a normal distribution. The ranges determined in chapter 5 are then used to define the $\pm 1\sigma$ span of the normal distribution. If parameter, p_i is restricted to lie between p_{min} and p_{max}, then

$P(p_i)$ is given by:

$$P(p_i) = Norm(p_{avg}, \sigma_i) \tag{6.22}$$

where $Norm()$ is the normal distribution, $p_{avg} = 1/2(p_{min} + p_{max})$ and $\sigma_i = 1/2(p_{max} - p_{min})$. The posterior distribution is then the product of the likelihood with all the prior distributions. Throughout the remainder of the thesis, unless specified, the words 'likelihood' and 'posterior distribution' will be used loosely, in the sense that they will refer to the product of the likelihood and the priors. This simplification in terminology will bridge the gap between a Bayesian approach and other methods of analyzing data. Although the approach used in this thesis is technically Bayesian, references to 'fitting for parameters' and 'minimizing the likelihood' will be made.

6.5   Markov Chain Monte Carlo 'Fitter'

This section presents the Markov Chain Monte Carlo method that will be used for the parameter estimation problem that was posed in the beginning of the chapter. The solution has historically been obtained by using minimization techniques that are widely discussed in the literature [77]. However, these minimization algorithms often suffer from limitations that cannot always be resolved by increasing computing power. Of course, of primary concern is the ability to properly sample parameter space in the case where there is a large number of parameters. Then, there is the fear of being 'stuck' in local minima of the likelihood function and, finally, the fact that the likelihood space is not smooth, due to the likelihood function being constructed using real data with real statistical fluctuations. Minimizing techniques, which are based on estimating local derivatives to find a minimum, can easily fail when the likelihood function is not smooth.

It will be argued that the MCMC technique presented here will eliminate at least some of the concerns that arise from minimization algorithms. The trade-off for more accurate parameter estimation will be at the expense of CPU needs. It will be however seen that,

with the advent of reasonably sized computer farms (100 CPUs), computational issues are irrelevant.

6.5.1 Monte Carlo Techniques

One should recall that the ultimate goal is to evaluate the posterior distribution function for neutrons, $P(n|D)$, given by equation 6.20:

$$P(n|D) = \int n \times P(\vec{p}|D)d\vec{p}$$

That is, one needs to integrate the posterior distribution for the parameters, $P(\vec{p}'|D)$, over all parameters. The posterior distribution is itself the product of the likelihood function, $P(D|\vec{p})$, and the prior distribution functions, $P(\vec{p})$. The bottom line is that one needs to compute a very complicated integral in order to obtain $P(n|D)$:

$$P(n|D) = \int n \times P(D|\vec{p})P(\vec{p})d\vec{p} \tag{6.23}$$

The difference with a minimization technique is stressed once again; whereas minimizing algorithms try to minimize the function $-\log\left(P(D|\vec{p})P(\vec{p})\right)$, the Bayesian approach requires one to map out the function and integrate it. This is a much greater task (computationally) but also has a greater payoff: one obtains the posterior distribution for the parameter of interest with all correlations (with other parameters) properly integrated

Monte Carlo integration has been around for a long time ([78], [75]) and can be implemented easily. The idea is to use Monte Carlo simulation to calculate the expectation value of the number of neutrons, $\langle n \rangle$:

$$\langle n \rangle = \int n \times P(\vec{p}|D)d\vec{p} \tag{6.24}$$

using the basic theorem of Monte Carlo integration:

$$\int n \times P(\vec{p}|D)d\vec{p} \approx V\frac{1}{N}\sum_{i=0}^{N-1} n_i \times P(\vec{x}_i|D) \tag{6.25}$$

where N independent points, $\vec{x}_i$, have been drawn uniformly from a volume, V, that encloses all of parameter space. This simple technique is however inefficient, since it can easily spend

a large amount of time sampling points in parameter space that do not contribute to the integral. The obvious remedy to this problem is to draw the $\vec{x}_i$ from the distribution $P(\vec{p}|D)$ itself, which will then make the process efficient. This modification is known as 'importance sampling', but has the caveat that drawing random, independent numbers from a complicated, multi-dimensional function is difficult in itself. In the case of importance sampling, where the $\vec{x}_i$ are drawn from $P(\vec{p}|D)$, the expectation value $\langle n \rangle$ is trivially given by:

$$\langle n \rangle = \frac{1}{N} \sum_{i=0}^{N-1} n_i \tag{6.26}$$

6.5.2 The Metropolis-Hastings Algorithm

The Metropolis-Hastings algorithm ([39], [40]) is a way to generate a sample of points drawn from the distribution $P(\vec{p}|D)$. This is achieved by obtaining a sequence of $\vec{x}_i$ that are not independent but are, however, properly distributed according to $P(\vec{p}|D)$, which is sufficient to evaluate equation 6.26.

The algorithm to generate the sequence of $\vec{x}_i$ uses a 'proposal distribution', $q(\vec{x}_{i+i}, \vec{x}_i)$, which generates the next point in parameter space, $\vec{x}_{i+1}$ using the current point, $\vec{x}_i$. In this thesis, the proposal distribution will be taken as a multivariate gaussian such that each component of $\vec{x}_{i+1}$ is drawn from a gaussian about the corresponding component of $\vec{x}_i$ using a different gaussian width for each component, $\vec{\sigma}$.

The Metropolis-Hastings algorithm for generating the sequence of $\vec{x}_i$ goes as follows:

- Given a starting position in the chain, $\vec{x}_i$

- A proposed point, $\vec{x}_{i+1}$, is generated using the proposal distribution

- The 'Metropolis acceptance ratio', r, is calculated, $r = Min\left(1, \frac{P(\vec{x}_{i+1}|D)}{P(\vec{x}_i|D)}\right)$, using a ratio of likelihoods

- The new point is accepted with a probability equal to r (which can be done by generating a random number between (0,1) and checking that it is smaller than r)

- These steps are repeated until the user is satisfied with the size of the sample of $\vec{x}_i$

The sequence of numbers $\vec{x}_i$ satisfy the properties of a Markov Chain (any point only depends on the previous point) and the process is thus called Markov Chain Monte Carlo integration. Although it is beyond the scope of this thesis to provide a rigorous proof that the Metropolis-Hastings algorithm produces a set of $\vec{x}_i$ that are distributed according to $P(\vec{p}|D)$, it is enlightening to provide a simplified overview. For a rigorous treatment, see [75].

One can imagine the MCMC as a random walk in parameter space. The chain starts at some point and then chooses a random point which it always accepts if the likelihood there is better ($r = 1$, since $P(\vec{x}_{i+1}|D) > P(\vec{x}_i|D)$), and sometimes (with a probability equal to r) accepts if it is worse. At least intuitively, it should make sense that the chain will tend to move into regions of parameter space where $P(\vec{p}|D)$ is larger and thus produce more $\vec{x}_i$ in that region.

An important concept is that the Markov Chain needs to reach a 'stationary' state; if the chain is started in a region of low likelihood it will tend to walk more in a particular direction (towards higher likelihood). Thus, the first 'few' steps in the MCMC are often discarded as 'burn-in' when posterior distributions are generated so that the chain can reach this stationary state.

One can introduce the 'transition kernel', $T(\vec{x}_{i+1}, \vec{x}_i)$, as the probability that the Metropolis-Hastings algorithm will draw and accept the proposed point, $\vec{x}_{i+1}$. This is simply given by:

$$T(\vec{x}_{i+1}, \vec{x}_i) = q(\vec{x}_{i+1}, \vec{x}_i) \times r(\vec{x}_i, \vec{x}_{i+1}) \tag{6.27}$$

which is the probability of generating a point multiplied by the probability of accepting

that point. It is then easy to show that the chain satisfies the 'detailed balance equation':

$$P(\vec{x}_i|D) \times T(\vec{x}_{i+1}, \vec{x}_i) = P(\vec{x}_i|D) \times q(\vec{x}_{i+1}, \vec{x}_i) \times r(\vec{x}_i, \vec{x}_{i+1})$$

$$= P(\vec{x}_i|D) \times q(\vec{x}_{i+1}, \vec{x}_i) \times Min\left(1, \frac{P(\vec{x}_{i+1}|D)}{P(\vec{x}_i|D)}\right)$$

$$= q(\vec{x}_{i+1}, \vec{x}_i) \times Min\left(P(\vec{x}_i|D), P(\vec{x}_{i+1}|D)\right)$$

$$= q(\vec{x}_i, \vec{x}_{i+1}) \times P(\vec{x}_{i+1}|D) \times Min\left(\frac{P(\vec{x}_i|D)}{P(\vec{x}_{i+1}|D)}, 1\right)$$

$$= P(\vec{x}_{i+1}|D) \times T(\vec{x}_i, \vec{x}_{i+1})$$

$$\Rightarrow P(\vec{x}_i|D) \times T(\vec{x}_{i+1}, \vec{x}_i) = P(\vec{x}_{i+1}|D) \times T(\vec{x}_i, \vec{x}_{i+1}) \tag{6.28}$$

where the proposal distribution is taken as symmetric $(q(\vec{x}_{i+1}, \vec{x}_i) = q(\vec{x}_i, \vec{x}_{i+1}))$. If one integrates equation 6.28 with respect to $\vec{x}_i$, one obtains

$$\int P(\vec{x}_i|D) \times T(\vec{x}_{i+1}, \vec{x}_i)d\vec{x}_i = \int P(\vec{x}_{i+1}|D) \times T(\vec{x}_i, \vec{x}_{i+1})d\vec{x}_i$$

$$= P(\vec{x}_{i+1}|D) \times \int T(\vec{x}_i, \vec{x}_{i+1})d\vec{x}_i$$

$$= P(\vec{x}_{i+1}|D) \tag{6.29}$$

Thus, it is shown that if $\vec{x}_i$ is drawn from the desired distribution, then all subsequent points, $\vec{x}_{i+1}$, will also be from that distribution. Thus, the MCMC process generates a set of points that are distributed according to $P(\vec{x}_i|D)$, as advertised.

6.5.3 The Posterior Distribution

The Metropolis-Hastings algorithm that has been shown here allows one to create a sample of points in parameter space that are distributed according to the posterior distribution. Obtaining the posterior distribution for, say, neutrons, can then be done simply by making a histogram of the neutron-component of the set of points, $\vec{x}_i$, from the parameter space. The MCMC method is then a very straightforward way of obtaining the posterior distribution for all of the parameters. The histogram can be normalized to unit area, if one desires a true probability density function.

6.5.4   Determining Parameter Values and Uncertainties

Once the posterior distribution has been obtained, one really has freedom in how to present these results. The most rigorous option is to show the posterior distribution. In terms of the number of neutrons, the mean number of neutrons can be calculated using equation 6.26:

$$\langle n \rangle = \frac{1}{N} \sum_{i=i_{min}}^{N-1} n_i$$

where the n_i are just one component of the set of $\vec{x}_i$ and i_{min} has been adjusted to the 'burn-in' length.

In this thesis, the number of neutrons (and all other parameters), will be obtained by fitting a gaussian to the histogram of the n_i. The fit range of the gaussian will be determined to be $\approx \pm 10$ bins around the highest point in the distribution. The mean of the gaussian will be taken as the number of neutrons and σ will be taken as the symmetric uncertainty. This is consistent with the assumption of a parabolic log-likelihood near the extremum. In order to obtain asymmetric uncertainties, the histogram will be re-fit to a gaussian on the positive (negative) sides with a fixed mean to determine the positive (negative) asymmetric uncertainties. If this procedure fails, the parameter and uncertainty are then taken as the mean and rms of the distribution.

Figure 6.3 shows an example of fitting a posterior distribution to determine the value of the parameter with uncertainties. The black line shows the gaussian used to determine the mean, the blue(red) gaussian uses that mean to determine the positive(negative) asymmetric uncertainties.

6.5.5   The Auto-Correlation Function

The set of points that are generated using the Metropolis-Hastings algorithm is correlated, so that care must be taken when applying the the basic theorem of Monte Carlo integration. In particular, one needs to make sure that the Markov-Chain has 'forgotten' it's starting
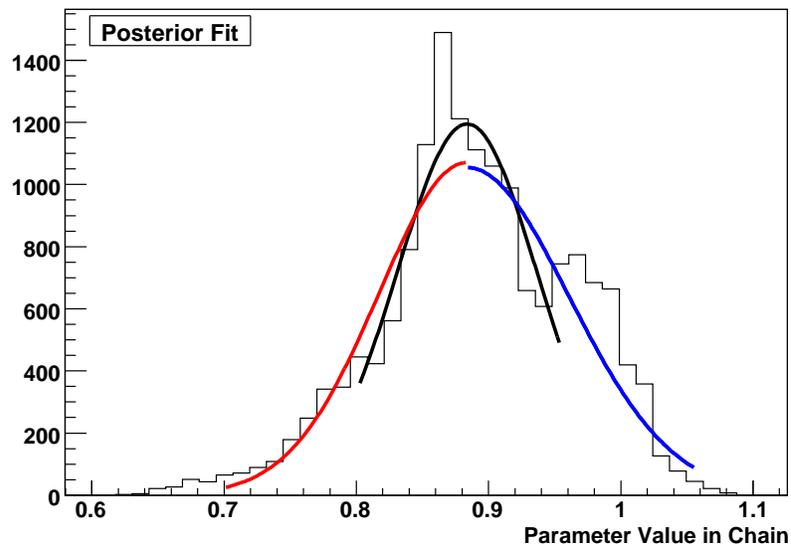
Figure 6.3: Example of fit to a posterior to determine the parameter value and uncertainties. First, the range of the central gaussian (black) is determined by a 12 bin range centered about the maximum entry. The black gaussian is then fit to determine the mean and width of the distribution which are quoted as the parameter value and symmetric uncertainty. That mean is then used to fit two more gaussians (red and blue) with fixed mean and floating width for the asymmetric uncertainties over a 12 bin range (with one end determined by the mean). This particular case shows that the asymmetric uncertainties can account for skewed posteriors.

point in parameter space and has indeed reached a stationary state. This can be verified by introducing the 'Auto-Correlation Function', ACF, for a parameter, p, which depends on the mean value of the parameter, $\bar{p}$:

$$ACF(p, h) = \frac{\sum (p_i - \bar{p})(p_{i+h} - \bar{p})}{\sqrt{\sum (p_i - \bar{p})^2}\sqrt{\sum (p_{i+h} - \bar{p})^2}} \tag{6.30}$$

where h is called the 'lag'. The ACF is essentially a measure of how far a point p_{i+h} must be from another point p_i in the chain to not be correlated. The ACF is an important measure of the MCMC performance and should be monitored to check that the chain has a reached a stationary state, independent from its starting position in parameter space.

Figure 6.4b shows an example of the ACF for a MCMC chain with four parameters. The auto-correlation function is similar to a decaying exponential [75], so it is reasonable to state that the correlation has been lost when the ACF is smaller than $\sim 1/e$. One can see that the parameters forget about each other after about 100 steps. This is an indicator of the amount of time before the chain reaches a stationary state. Also, if one were to create a sample of points from the chain, taking every hundreth point, then one would end up with an essentially uncorrelated sample from the posterior distribution. Panel a) of figure 6.4 shows the values of one of the parameters as a function of step in the chain. In this particular case, the parameter was given an initial value of 100 and one can clearly see the 'burn-in' phase of the chain in which the parameter systematically increases. Once the chain has reached an area of high likelihood in parameter space, the value of the parameter oscillates around its most likely value (about 950) and the histogram for the posterior distribution can be generated. One should note that it is a coincidence that the ACF reaches zero at approximately the same time as the chain burns in; both these conditions should always be checked (in particular, the chain can often appear to have burned-in long before the auto-correlation function is close to zero). The ACF for each parameter also needs to be checked as different parameters will take different amounts of time to lose their (auto)correlation.

(a) Parameter as a function of step

(b) ACF as a function of lag
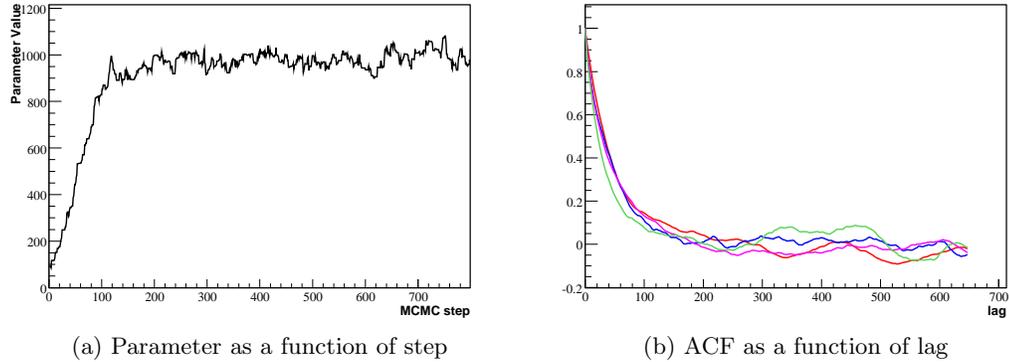
Figure 6.4: Example of Auto-Correlation Function (b) and evolution of a parameter in the chain (b). The ACF indicates a lag of about 100 before the parameter forgets its value. Panel a) shows the parameter evolving along the chain from its starting value of 100 to an area around 950. The burn-in length is, coincidentally, also around 100.

6.5.6 Acceptance Tuning

The most delicate part of the MCMC process is the choice of the proposal distribution, $q(\vec{x}_{i+i}, \vec{x}_i, \vec{\sigma})$. In this thesis, a multivariate (symmetric) gaussian will be used. This leaves one with the choice of the σ for each parameter. Intuitively, it should be obvious that this choice will have a significant impact on the performance of the chain; if the chain is in the region of maximum-likelihood and all the σ are large, there is a very small probability that any new choice of parameters will be accepted (the Metropolis acceptance ratio will be very small) and the chain will take a long time to properly sample parameter space. On the other hand, if the σ are too small, the chain will move around parameter space very slowly and will run a greater risk of being stuck in a local minimum.

Although there is no exact prescription for determining the proper 'widths' for the proposal distributions, it has been argued that one should aim for an acceptance ratio around 23%, to have an efficient process [79]. One can imagine various algorithms to carry out the tuning of the parameter widths or simply tune them 'by hand', which, unfortunately, is often the fastest method.

The main caveat of the MCMC method is the tedious adjustment of these widths. Not only do they influence the efficiency of the method, but also its likelihood of getting 'stuck' in a local minimum; although, in theory, the chain cannot get stuck in local minima if it runs for long enough, in practice, this can happen for finite length chains. The MCMC method is, however, impervious to discontinuities in the likelihood space, since there is no need to worry about derivatives and these discontinuities can just be smoothed-out by 'jumping across them'.

6.6 Summary

This chapter started by introducing the framework for determining the number of neutrons in the NCD data using an extended maximum-likelihood formalism that can accommodate variable fit ranges as well as systematic deformation parameters. A novel method for implementing the uncertainty in the shapes of the NNNA pdfs was introduced. The Markov Chain Monte Carlo method was then presented as a way of implementing a Bayesian approach to data-analysis using the likelihood function.

Chapter 7

Data Analysis Tests with Monte Carlo and Blind Data

This chapter is aimed at demonstrating that the MCMC techniques presented in chapter 6, as well as the choice of signals and observables, are suitable for analyzing the NCD data and extracting an accurate number of neutrons. In the first part of this chapter, the MCMC fitter is applied to fake data (Monte Carlo events generated from the pdfs) to test for correlations between observables and biases in the resulting fits. The MCMC analysis procedure is then 'tuned' to give the smallest uncertainty in the number of neutrons. In the second part of this chapter, the technique is applied to a fraction of the real NCD data, that has been blinded in order to perform these tests. The blinding scheme involved the removal of an unknown fraction of the data as well as the addition of an unknown amount of 'muon-follower' events[1].

---

[1] These are events that are close in time to vetoed muon events. This type of occurrence is generally rich in neutrons created from spallation

7.1 Generating Monte Carlo Data

In order to test for biases and correlations between the signals, several sets of Monte Carlo data were generated. The data sets were made using the pdf histograms that were obtained in chapter 5. For each signal, sets of events with energies and NoverA distributed according to the (1-dimensional) pdfs were randomly generated using CERN's ROOT analysis package [80] and the TH1F::GetRandom() method. In order to examine possible correlations between observables, limited data were generated by sampling from 2-dimensional histograms (energy vs NoverA, using ROOT's TH2F::GetRandom2() method) for neutrons and alphas.

In all data sets, unless otherwise specified, the MC samples were generated with 2000 neutrons, 8000 alphas, 100 nnna0 and 200 nnna26 events.

7.2 Extracting Signals from Monte Carlo Data Using Energy

This section examines the use of the energy observable alone to determine the neutron number. The effect from the choice of the energy range, binning and the contribution from the systematic uncertainties are characterized.

7.2.1 Without Systematic Uncertainties

In order to examine the effects of fit range and binning of the pdfs on statistical accuracy, it is useful to perform the fits without any systematic uncertainties.

The Effect of Fit Range and Binning

It is clear that a larger fit range should produce greater statistical accuracy. However, if one limits the data to events with higher energies (above $0.4\,\mathrm{MeV}$), the contribution from the NNNA events is greatly diminished (as these peak at lower energies). Since the NNNA pdfs have the largest uncertainties on them, it may be advantageous to work in a regime where their contribution is low. For that reason, the effect of a smaller energy range is also

considered. It is expected that changing the energy range will also impact the uncertainty on the NNNA contributions, as these are best constrained at lower energies.

Two different energy ranges will be considered; on the one hand, a 'full energy range' (from 0.2 MeV to 1 MeV) and, on the other, a 'short' energy range, (0.4 MeV to 1 MeV) to limit the number of NNNA events.

In addition to the fit ranges, two choices of binning for the pdfs will also be examined; 50 keV and 25 keV bins. The smaller bin size should improve statistical accuracy, whereas the larger bin size will wash out some of the allowed deformations of the pdfs and could improve the systematic uncertainties.

Table 7.1 shows the resulting signals that are fit out from a set of Monte Carlo data using energy with the two different binnings and fit ranges. In terms of neutrons, the number is stable with respect to these changes and the statistical uncertainty is lowest when the finer binning is used. The main effect of the fit range is in the number of alphas and their uncertainty. With the smaller energy range, the fitter has trouble constraining the NNNA signals, which can then trade off with the number of alphas as well as increase the uncertainties on the corresponding signals. From these fits, it is concluded that the smaller energy binning (25 keV) combined with the larger energy range provides overall better constraints on all signals. The uncertainty on the number of neutrons is virtually unaffected by the fit range.

Figure 7.1 shows the resulting fits in energy for the two energy ranges (panels (a) and (b)) and the finer binning, as well as the correlation between the alpha signal and the nnna26 signal (panels (c) and (d)). The correlation is obtained by creating a two-dimensional histogram of the signals generated by the MCMC. Hence, this two-dimensional histogram is proportional to the marginalized likelihood as a function of the two parameters and is, in fact, the two-dimensional analog of the posterior distribution.

(a) Full Range Fit

(b) Small Range Fit

(c) Alpha and nnna26 correlation, full range fit   (d) Alpha and nnna26 correlation, small range fit
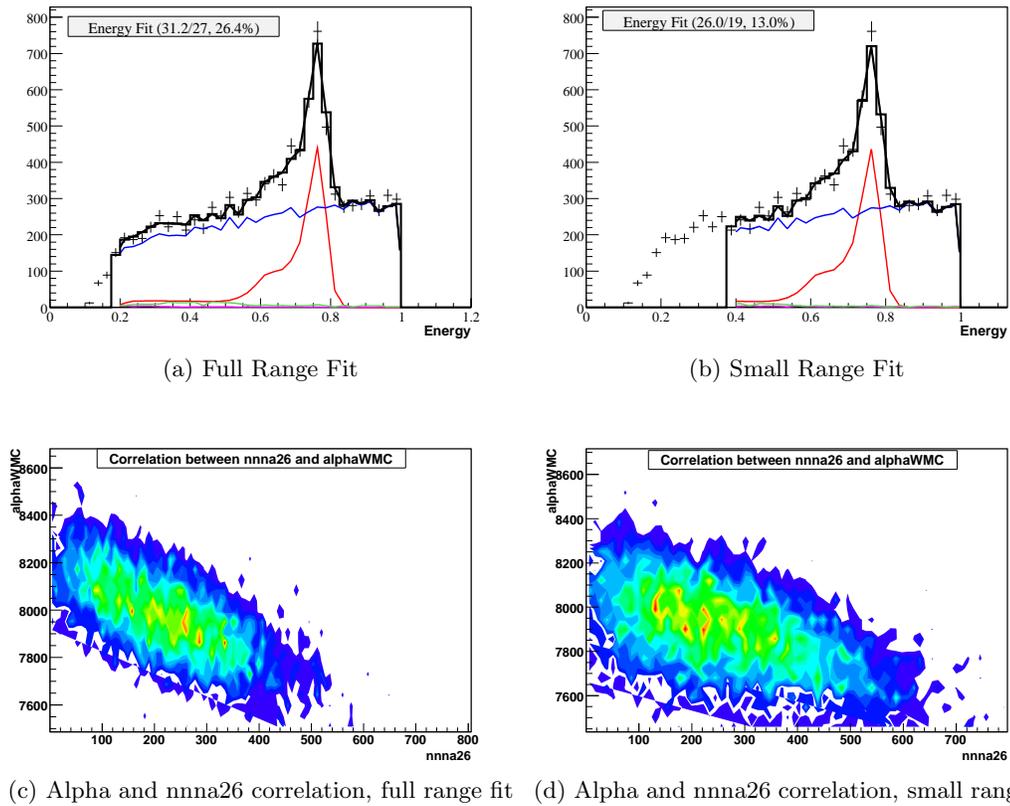
Figure 7.1: Comparison of energy fits for full range (a) and small range (b). Panels (c) and (d) show the correlation between alphas (y-axis) and nnna26 events (x-axis) for the two ranges. The fit over the larger range places a better constraint on the alphas and nnna26 events, but these remain correlated.

50 keV bins	0.2 MeV 1 MeV	0.4 MeV 1 MeV
neutrons	2033(93)	2034(91)
alphas	7932(199)	7790(202)
nnna0	117(102)	59(139)
nnna26	154(152)	392(195)
Fit Quality	13.5/11 (26.5%)	11/7 (14%)
25 keV bins	**0.2 MeV 1 MeV**	**0.4 MeV 1 MeV**
neutrons	2014(79)	1995(79)
alphas	7960(176)	7902(195)
nnna0	53(76)	122(157)
nnna26	234(114)	228(167)
Fit Quality	31.2/27 (26.4%)	26/19 (13%))

Table 7.1: Energy fit to MC data for different fit ranges and binning performed using the MCMC fitter. The (symmetric) uncertainties and mean values are obtained from a Gaussian fit to the posterior distribution centered about the maximum. The quality of each fit is shown as a chi-squared and probability as calculated between a histogram of the data and histogram that sums the pdfs from the different signals, as shown in panels (a) and (b) of Figure 7.1.

7.2.2   With Systematic Uncertainties

This section examines the impact of the deformation parameters for the pdfs on the uncertainties in the extracted signals. The Monte Carlo data will be generated with and without the deformation parameters applied and the fit will be carried out with and without the prior distributions imposed on the deformation parameters. These tests will quantify how well the fitter can extract the values of the applied deformation parameters, as well as test the strength of the constraint supplied by the prior distributions from the calibration data. The effects from neutrons, alphas and NNNA events will be applied sequentially. Table 7.2 summarizes the energy deformation parameters that were obtained in chapter 5 as the mean and spread of a gaussian distribution:

Including Neutron Energy Scale

The effect of the neutron energy scale is considered first. Two different data sets were generated: in one case the data has no scale applied to it and in the other case, the data has

Signal	scale	shift	resolution
neutrons	1.003(05)	fixed	fixed
alphas	fixed	-0.038(38)	fixed

Table 7.2: Summary of energy deformation parameters for neutrons and alphas. The parameters are given a central value and uncertainty from the ranges determined in chapter 5. These are designed to be applied as prior distributions using gaussians with mean and sigma given by the central values and uncertainties.

a energy scale of 1.05 applied only to the neutrons (inconsistent with the prior). This data is then fit with and without the priors applied to the neutron energy scale. No deformation parameters are floated for the other signals. Table 7.3 summarizes the results for the fits. It is seen that, with both data-sets, the prior has a very small effect on the uncertainty in the number of neutrons which increases slightly with no prior. The prior does constrain the number of alphas and including the prior results in a more accurate estimate of the number of alphas. In the second case, where the prior disagrees with the actual scale that was used, the fitted scale comes out slightly wrong and impacts the fit quality. The uncertainty on the number of neutrons appears to have increased slightly with the inclusion of a neutron energy scale (from ∼79 to ∼83), but it is difficult to make a statistically sound estimate of the effect with just one Monte Carlo data set. The point of this demonstration was to show that the neutron number uncertainty does not 'blow up', even when no prior constraint is imposed.

Including Alpha Energy Shift

Similar tests can be performed to test the influence of the alpha energy shift on the number of neutrons and the associated uncertainty. In this case, the neutron energy scale is floated with the prior constraint and the fake data is generated with a neutron energy scale fixed to the central value of the prior applied to the (fake) neutrons. Two data sets are then created, with and without an alpha energy shift applied. In the case where the data is generated with an alpha energy shift, a value of -0.07 MeV is chosen (consistent with the value inferred

Undeformed Data	With Prior	No Prior
neutrons	1988(83)	1989(84)
alphas	8008(138)	7921(185)
nnna0	35(29)	32(40)
nnna26	216(114)	264(119)
neutron energy scale	1.004(2)	1.005(2)
Fit Quality	31/27 (27%)	31/27 (27%)
Deformed Data	**With Prior**	**No Prior**
neutrons	2045(86)	2059(90)
alphas	7917(113)	7869(140)
nnna0	67(50)	28(44)
nnna26	271(65)	285(109)
neutron energy scale	1.044(02)	1.053(03)
Fit Quality	48/27 (1%)	33/27 (20%)

Table 7.3: Energy fit to MC data with neutron energy scale systematic. In the first row, the data were generated with no scale and then fit for a neutron scale that comes out essentially equal to 1. The first column imposes a prior of 1.003(5) on the neutron scale, whereas the second column lets the scale float freely, which does not affect the uncertainty on the extracted number of neutrons. In the second row, the data were generated such that the neutrons have a 1.05 scale applied to them.

from the ^{4}He data). The data sets are then fit with and without the prior constraint on the alpha energy scale. The results are summarized in Table 7.4. The inclusion of the alpha energy shift deformation parameter is seen to have no effect on the uncertainty in the number of neutrons, although the use of the prior constraint on the alpha energy shift seems to increase the uncertainty in the number of alphas.

Including NNNA Energy Systematic Uncertainties

The NNNA systematic uncertainties are included as described in section 6.2.3. That is, from a histogram with uncertainties in each bin (the original energy pdf), 10 new histograms are created by randomly sampling each bin with a normal distribution. In each step of the MCMC, a linear combination of those 10 pdfs is built; the parameters in that linear combination are then floated. At the end of the fit, an effective NNNA pdf has been measured, by building a linear combination from the parameters deduced from the fit.

Undeformed Data	With Prior	No Prior
neutrons	2016(82)	2012(80)
alphas	7980(185)	8026(170)
nnna0	114(102)	93(82)
nnna26	296(65)	282(116)
neutron energy scale	1.007(2)	1.007(2)
alpha energy shift	-0.011(1)	-0.011(1)
Fit Quality	42/27 (3%)	42/27 (3%)
Deformed Data	**With Prior**	**No Prior**
neutrons	2083(80)	2082(80)
alphas	7860(158)	7912(144)
nnna0	66(78)	87(52)
nnna26	199(98)	100(66)
neutron energy scale	1.005(2)	1.006(2)
alpha energy shift	-0.066(2)	-0.066(2)
Fit Quality	26/28(55%)	27/28 (53%)

Table 7.4: Energy fit to MC data with alpha energy shift systematic. In both rows, data were generated with a neutron energy scale set to 1.003 and the neutron energy scale was floated with a prior. In the first row, the alpha energy shift was floated with and without a prior. In the second row, the Monte Carlo data were generated with an alpha energy shift of -0.07 MeV (consistent with the value determined from ^{4}He data).

In this section, the fake data sets were generated by sampling from the NNNA average pdfs as well as from the neutron and alphas (as in the previous sections). The neutron energy scale is fixed to 1.003 in the generated data and is floated with a prior. Similarly, the alpha energy shift is fixed to -0.038 and floated with its prior. The interest is to show that letting the NNNA pdfs float this way does not result in any large increase of the neutron uncertainty.

Table 7.5 shows fit results for the signals when the NNNA pdfs are floated in this manner using either 5 or 10 random pdfs in the linear combination. The first row shows the case where only the nnna0 pdf is floated, the second row shows the case for nnna26 and the third row shows the result when both NNNA energy pdfs are floated. The last row corresponds to the scenario that will be used in the energy part of the NCD fit to the data. Figure 7.2 shows the pdfs that were determined from the fit (red) compared to the input pdf (black) for the two different NNNAs for the case where both pdfs were floated as a linear combination of 10 random pdfs (the last row, second column of Table 7.5).

The number of random pdfs that are chosen do not significantly impact the number of neutrons. Increasing the number of random pdfs from 5 to 10 seems to increase the uncertainty on the extracted number of NNNAs but slightly decrease the uncertainty in the number of alphas. The analysis in the rest of this thesis will use 10 randomized energy pdfs for each type of NNNA as this will result in a better integration over the possible shapes.

7.2.3   Biases and Correlated Signals in the Energy Fit

It is important to understand any bias in the signals that could result from using energy as an observable in the maximum-likelihood analysis. This might happen if one type of signal 'trades off' with another, because of having similar pdfs. Indeed, it has just been seen that the NNNA pdfs are similar enough that the extracted numbers are not very accurate. Of particular interest are any possible biases in the neutron number as any trade-off between other observables has no impact on the neutrino flux measurement.

Floating nnna0	5 random pdfs	10 random pdfs
neutrons	2073(84)	2059(83)
alphas	7843(178)	7831(178)
nnna0	97(77)	96(109)
nnna26	244(111)	271(126)
neutron energy scale	1.006(2)	1.006(2)
alpha energy shift	-0.037(2)	-0.037(2)
Fit Quality	24/28 (70%)	24/28 (69%)
Floating nnna26	**5 random pdfs**	**10 random pdfs**
neutrons	2075(86)	2070(86)
alphas	7873(175)	7886(165)
nnna0	108(108)	152(116)
nnna26	171(147)	200(140)
neutron energy scale	1.006(2)	1.006(2)
alpha energy shift	-0.037(2)	-0.038(3)
Fit Quality	24/28 (67%)	25/28 (65%)
Floating nnna0 and nnna26	**5+5 random pdfs**	**10+10 random pdfs**
neutrons	2079(90)	2072(86)
alphas	7873(202)	7839(182)
nnna0	143(102)	127(108)
nnna26	185(121)	228(143)
neutron energy scale	1.006(2)	1.005(2)
alpha energy shift	-0.037(2)	-0.037(3)
Fit Quality	25/28(65%)	26/28 (55%)

Table 7.5: Energy fit to MC data including the systematic uncertainties in the shapes of the NNNA pdfs. The uncertainties are implemented by building linear combinations of either 5 or 10 random pdfs drawn from the average NNNA pdfs. The first row shows the fitted signals when only the nnna0 pdfs is floated, the second row shows the case when only the nnna26 pdf is floated and the last row shows the case when both are floated. It is seen that using 10 random pdfs for each NNNA results in a larger uncertainty in the number of NNNA, but has virtually no effect on the number of neutrons.

(a) measured nnna0 energy pdf　　　　　(b) measured nnna26 energy pdf
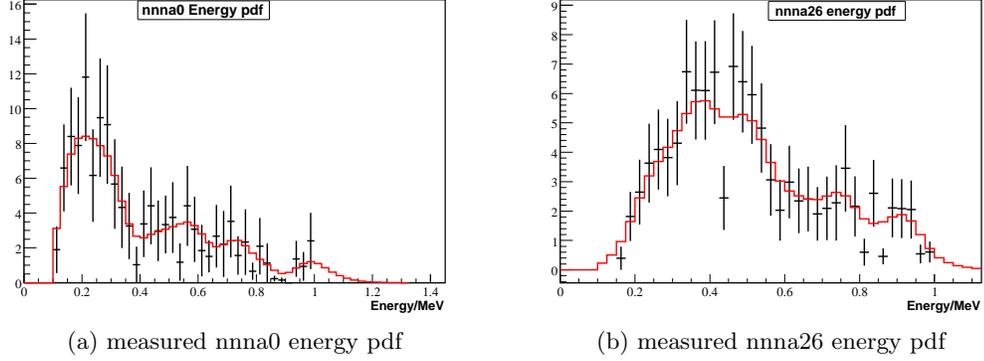
Figure 7.2: Example of NNNA energy pdf determined from fit to fake data. The black histograms corresponds to the pdfs that were used to generate the fake data and were determined in section 5.3. The red lines are the linear combination of 10 histogram randomly drawn from the black histogram. The coefficients of the linear combination were determined by the MCMC. Panels (a) and (b) show the case for nnna0 events and nnna26 events, respectively. These pdfs correspond to the fit in the second column for the third row of Table 7.5.

In order to study biases from the energy observable, 100 sets of Monte Carlo data were generated in the same way as they were for Table 7.5. A neutron energy scale and alpha energy shift were applied (equal to the central values of their prior) and all 100 data sets were fit by floating the neutron and alpha parameters with a prior and the NNNA pdfs were floated with 10 random pdfs each.

Each data set was generated with a Poisson-distributed number of events, N_{MC}^i, from each signal, i, distributed about a mean, λ^i. Neutrons were generated about a Poisson mean of 2000, alphas about a mean of 8000, nnna0s about 100, and nnna26 about a mean of 200. The panels in Figures 7.3 and 7.4 show the 'bias' for each signal plotted as a function of Monte Carlo data-set, along with a histogram projecting the data on the y-axis. The bias for each signal, B^i, was defined with an uncertainty, δB^i, as:

$$B^i = \frac{N_{Fit}^i - \lambda^i}{\lambda^i}$$

$$\delta B^i = \frac{\delta_{Fit}^i}{\lambda^i} \tag{7.1}$$

where N_{Fit} is the number of the signal fitted out with uncertainty δ^i_{Fit}. The bias is thus a measure of the relative difference between the Poisson mean of the distribution and the numbers that were determined in the fit. The uncertainties in the bias show whether the uncertainty determined from the extended maximum-likelihood fit are consistent with the data being Poisson distributed. The projections of the biases on the y-axis show that the fit overestimates the number of neutrons by $\sim 3.1\%$ and underestimates the number of alphas by $\sim 1.2\%$. This can be taken into account as either a correction or an uncertainty.

(a) neutron energy bias

(b) neutron energy bias projection

(c) alpha energy bias

(d) alpha energy bias projection
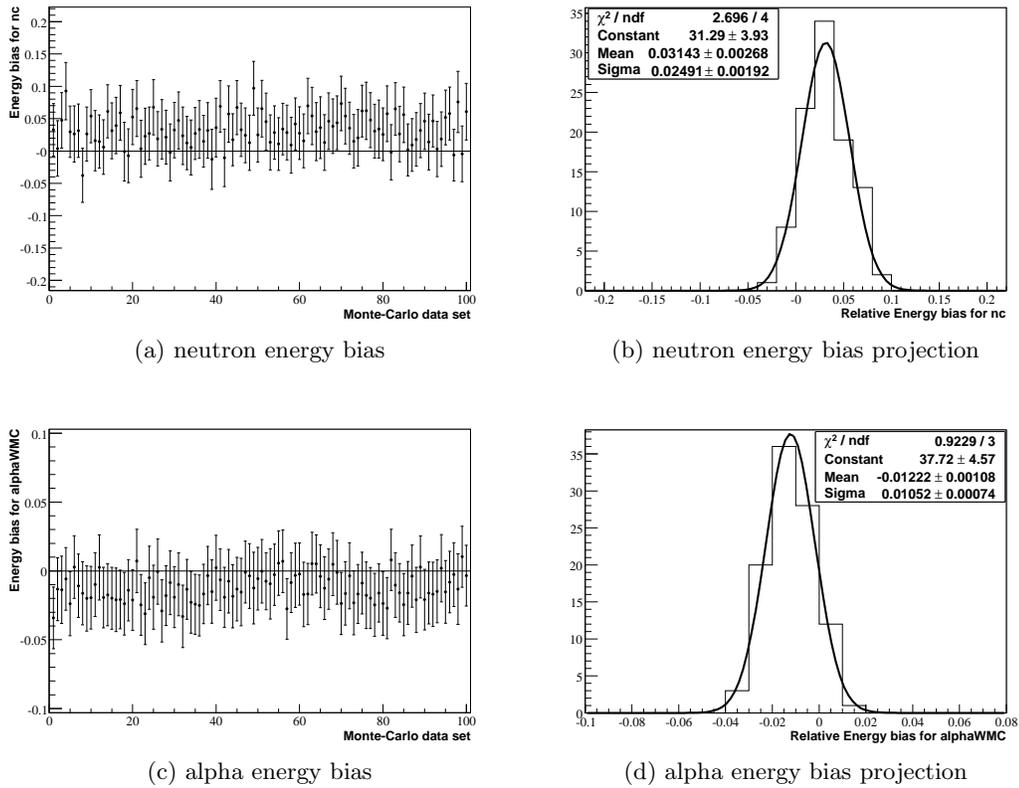
Figure 7.3: Biases (and projections) for neutron and alpha signals when using energy as an observable plotted as a function of the 100 Monte Carlo data sets. It is clear that for neutrons and alphas, there appears to be a systematic pull away from the number of events that were put in each data set, as seen in Figure 7.5.

The panels in Figure 7.5 show the distributions of the 'pull', p^i, for each signal defined
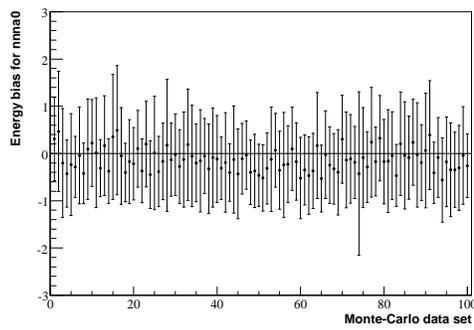
(a) nnna0 energy bias
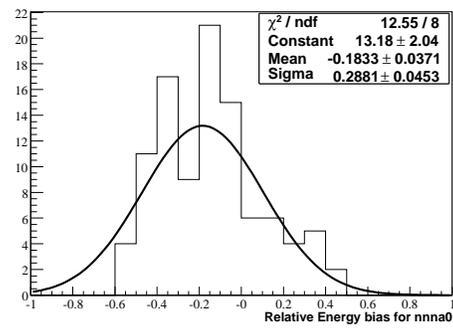
(b) nnna0 energy bias projection
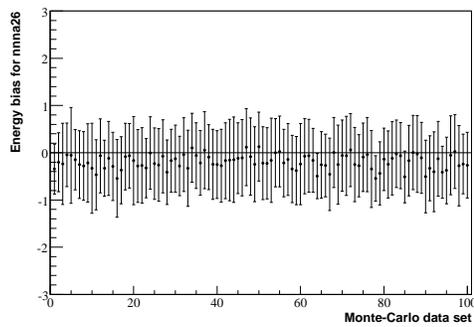
(c) nnna26 energy bias
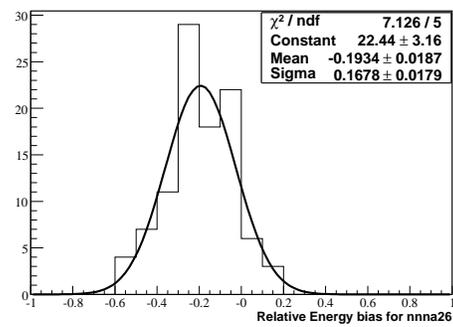
(d) nnna26 energy bias projection

Figure 7.4: Biases (and projections) for the NNNA signals when using energy as an observable plotted as a function of the 100 Monte Carlo data sets.

as:

$$p^i = \frac{N^i_{Fit} - \lambda^i}{\delta^i_{Fit}} \tag{7.2}$$

where λ^i is the mean of the Poisson distribution used to generate the Monte Carlo events for that signal. The pull is then expected to follow a normal distribution about 0 with a standard deviation of 1, if the fit and uncertainties are determined correctly. If the mean of the distribution is not 0, there is an indication that the results determined from the fit are 'pulled' towards a certain direction. The value of the mean is in units of the error on the measurement. For example, if the mean of the pull distribution is 0.5, the N^i_{Fit} are $0.5\delta^i_{Fit}$ from the λ^i. The width of the pull distribution is a measure of the accuracy of the error determined by the fit. If the standard deviation of the pull distribution is less than 1, the fit has overestimated the uncertainties on the fitted number of neutrons and conversely if the width is larger than 1.

Again, it is clear from the pull distributions, that the number of neutrons is slightly pulled towards higher values and the number of alphas towards lower values. This suggests a correlation between the number of neutrons and the number of alphas that are determined in the fit. Figure 7.6 shows the 2-D likelihood for neutrons and alphas for the fit from the second column of the third row of Table 7.5, and the slight correlation between the two signals is seen clearly. The pull can be taken into account as an uncertainty when the data is fit.

(a) neutron energy pull

(b) alpha energy pull

(c) nnna0 energy pull

(d) nnna26 energy pull
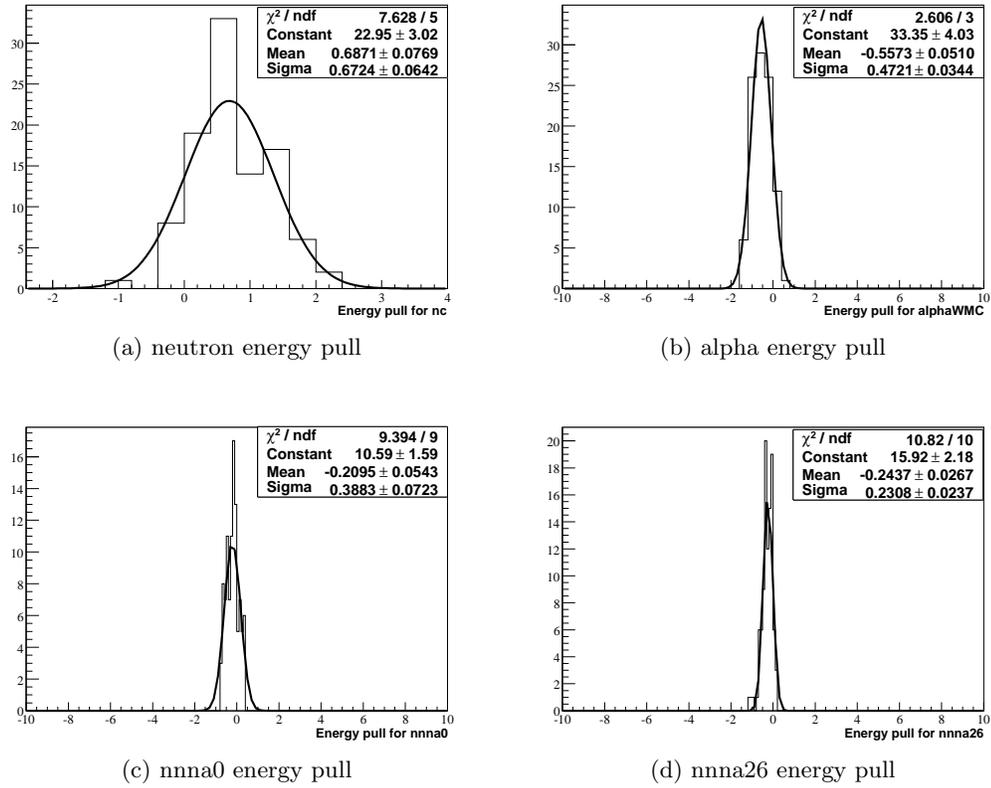
Figure 7.5: Pull plotted for each signal as determined from a sample of 100 Monte Carlo data sets. The data sets were fit using the energy observable. Neutrons and alphas show evidence for a slight pull as the distributions are not centered about 0.
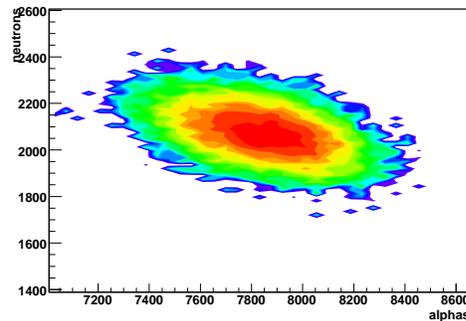
Figure 7.6: Correlation between neutrons and alphas in the energy fit from the second column of the third row of Table 7.5. The correlation is show by plotting the 2-dimensional likelihood as a function of neutron and alphas.

7.3 Extracting Signals from Monte Carlo Data Using NoverA

This section examines the use of NoverA as the observable in a likelihood extraction of the signals, in the same way as was done for energy. Sets of 'fake' data are generated using the pdfs so that biases and uncertainties in the signals can be examined. The systematic uncertainties from neutrons, alphas and NNNAs will be added progressively so that their individual effect on the number of neutrons can be understood. The Monte Carlo data sets will again comprise of 2000 neutrons, 8000 alphas, 100 nnna0s and 200 nnna26s.

7.3.1 Without Systematic Uncertainties

Table 7.6 shows the extracted signals when the Monte Carlo data set is fit in NoverA with different bin sizes (0.1 and 0.2 bin widths) and fit ranges (0 to 5 and 0 to 10). The fit with the smallest bin width and largest fit range is the most accurate (smallest uncertainties) and will be chosen for the rest of the analysis. The neutron number uncertainties are larger than in the case where energy was used, because neutrons do not look 'as different' in NoverA as they did in energy. There also appears to be a significant downward bias in the number of neutrons. This will however disappear when the NNNA systematics are taken into account and is the result of a statistical fluctuation in the average NNNA pdfs (the nnna26 pdf has a statistically high bin content in the same place as the neutron pdf and results in a neutron bias).

Figure 7.7 shows the resulting fit using the small bin size and large range (panel a) and the neutron posterior distribution (panel b). The posterior highlights the fact that care must be taken when quoting a parameter value from the posterior, as the peak of the histogram is closer to the actual number of neutrons that is inferred from the Gaussian fit (also seemingly resulting in a downward bias in the number of neutrons).

0.1 bin	0 - 10	0 - 5
neutrons	1932(93)	1922(122)
alphas	7962 (122)	7946(138)
nnna0	153(86)	140(126)
nnna26	268(135)	210(140)
Fit Quality	76/90(90%)	44/45(51%)
0.2 bin	**0 - 10**	**0 - 5**
neutrons	1928(128)	1984(99)
alphas	8035(120)	8049(124)
nnna0	160(105)	113(76)
nnna26	190(163)	76(118)
Fit Quality	41/45(64%)	20/21(55%)

Table 7.6: NoverA fit to MC data for different fit ranges and binning. The MC data were fit using a bin width of 0.1 and 0.2 and fit ranges from 0 to 5 and 0 to 10 in NoverA. The smaller bin width and larger fit range results in the smallest uncertainties and is chosen for the analysis.

(a) NoverA Fit

(b) neutron posterior
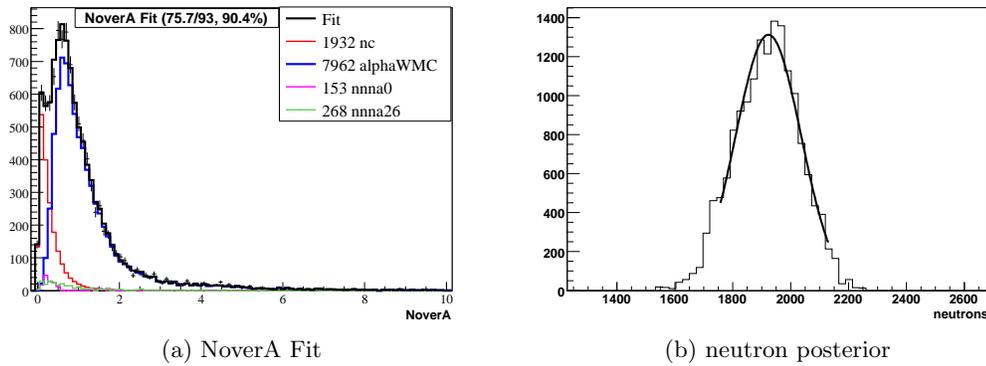
Figure 7.7: Panel (a) shows the fit to MC data for NoverA using the small (0.1) bin width and large fit range (0-10) of Table 7.6. Panel (b) shows the resulting neutron posterior distribution and highlights the fact that care needs to be taken in quoting a number from a posterior distribution. In this case, the Gaussian fit resulted in a smaller number of neutrons than the peak of the distribution and leads to an apparent downward bias in the number of neutrons.

Signal	scale	shift	resolution
neutrons	1.084(31)	fixed	fixed
alphas	0.967(79)	fixed	fixed

Table 7.7: Summary of NoverA deformation parameters for neutrons and alphas. The parameters are given a central value and uncertainty from the ranges determined in chapter 5. These are designed to be applied as prior distributions using gaussians with mean and sigma given by the central values and uncertainties.

7.3.2   With Systematic Uncertainties

The following sections characterize the effect of including the systematic uncertainties in the shape of the pdfs. In chapter 5, it was shown that the uncertainty in the shape of the NoverA pdf for both neutrons and alphas could be accommodated using a scale parameter. Table 7.7 summarizes the constraints that were obtained on the neutron and alpha NoverA scale. The Monte Carlo data sets in this section will be produced in the same way as they were when energy was examined in the previous sections. Data will be generated with and without deformation parameters applied and will then be fit with and without the priors applied to the deformation parameters.

Including Neutron NoverA Scale

The effect of the neutron NoverA scale is examined by fitting with and without the prior constraint. Data with and without a neutron NoverA scale of 1.05 applied will also be generated to test the impact of the prior constraint on the neutron number uncertainty. Table 7.8 shows the results from the four possible combinations. The neutron uncertainty is measurably affected by the inclusion of a neutron NoverA scale and goes from ∼93 to ∼123. It is also apparent that the neutron NoverA scale is not determined very accurately by the fitter.

Undeformed Data	With Prior	No Prior
neutrons	1838(123)	1913(107)
alphas	7983(121)	7908(130
nnna0	193(111)	100(90)
nnna26	247(138)	372(141)
neutron NoverA scale	1.015(0.09)	1.001(8)
Fit Quality	77/94(89%)	73/91(91%)
Deformed Data	**With Prior**	**No Prior**
neutrons	1811(113)	1819(164)
alphas	8000(135)	7901 (127)
nnna0	195(108)	119(102)
nnna26	266(147)	440(140)
neutron NoverA scale	1.103(18)	1.084(43)
Fit Quality	76/94(92%)	72/89(91%)

Table 7.8: Effect of neutron NoverA scale on fits to MC data with a prior (first column) and no prior (second column) applied to the neutron NoverA scale. In the second row, the MC data were generated with a neutron NoverA scale of 1.05.

Including Alpha NoverA Scale

In this section, the effect of the alpha NoverA scale is included together with that from the neutron NoverA scale. Data were generated with a neutron scale of 1.084 applied and fit with and without the prior on the alpha scale. In addition, a data set with an alpha scale of 0.89 (consistent with the ^{4}He data) was also generated and analyzed in the same way. In all cases, the neutron scale was constrained with the prior from Table 7.7. The results from the fits are summarized in Table 7.9.

The inclusion of the alpha NoverA scale uncertainty has facilitated the process of determining the neutron NoverA scale, which results in overall smaller neutron number uncertainties. In the case where the data has been deformed, the neutron uncertainty is decreased even more. This can be explained by the fact that deforming the alphas results in a wider NoverA alpha distribution which makes them 'more-different' from neutrons, and results in the neutron being extracted better.

Undeformed Data	With Prior	No Prior
neutrons	1962(110)	1897(102)
alphas	8029(128)	7975(128)
nnna0	143(98)	175(71)
nnna26	189(149)	299(104)
neutron NoverA scale	1.082(23)	1.114(11)
alpha NoverA scale	0.999(02)	0.999(02)
Fit Quality	81/94(84%)	77/95(90%)
Deformed Data	**With Prior**	**No Prior**
neutrons	2036(90)	2024(82)
alphas	8019(126)	8002(129
nnna0	80(64)	93(87)
nnna26	153(104)	197(134)
neutron NoverA scale	1.093(09)	1.081(11)
alpha NoverA scale	0.889(02)	0.889(02)
Fit Quality	77/94(90%)	77/94(90%)

Table 7.9: Effect of alpha NoverA scale on the NoverA fit to MC data. The neutron scale was fixed to 1.084 in the data and was floated subject to the prior constraint. The data were fit with and without the prior constraint applied to the alpha NoverA scale. The data in the second row was generated with an alpha scale of 0.89, similar to the value determined for the ^{4}He data (chapter 5). The decrease in the neutron uncertainty seen in the second row is due to the alphas in the fake data having a scale applied and thus 'looking more different' from neutrons than they did in the first row.

Including NNNA NoverA Systematic Uncertainties

The uncertainties in the NNNA NoverA pdfs are handled analogously to the case for energy. This section shows the effect of using the linear combination of randomized pdfs on the fit for the signals. The fake data set was generated with neutron and alpha NoverA scales set to their central prior values from Table 7.7. The NNNA events were generated from the average pdfs. Table 7.10 summarizes the results from the study carried out using either 5 or 10 random pdfs. Figure 7.8 shows an example of the resulting pdfs that were determined for the NNNAs for the fit with 10 random pdfs each. In the case where the systematic uncertainties are applied to only one type of NNNA, the use of 10 pdfs increases the uncertainty in the number of neutrons compared to when only 5 pdfs are used. However, when both types of NNNAs have their pdfs floated, the use of 10 pdfs results in an overall smaller uncertainty in the number of neutrons. The final analysis will thus use the linear combination of 10 pdfs for each type of NNNA.

(a) nnna0   (b) nnna26
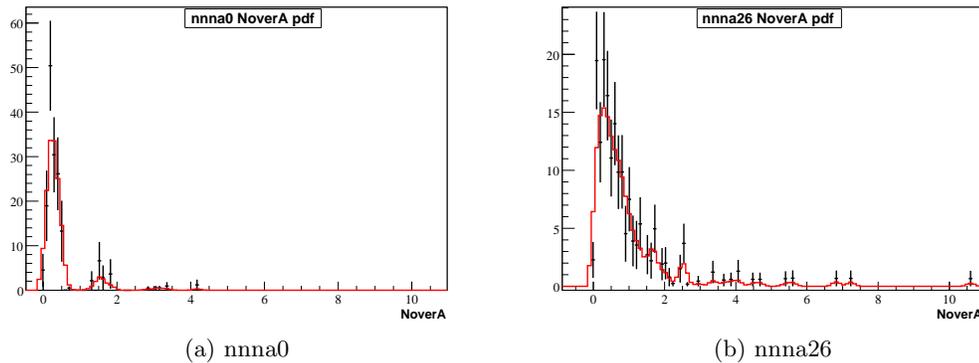
Figure 7.8: Example of NNNA NoverA pdf determined from fit to fake data. The black histograms corresponds to the pdfs that were used to generate the fake data and correspond to the histograms determined in section 5.3. The red lines are the linear combination of 10 histogram randomly drawn from the black histogram. The coefficients of the linear combination were determined by the MCMC. Panels (a) and (b) show the case for nnna0 events and nnna26 events, respectively. These pdfs correspond to the fit in the second column for the third row of Table 7.10.

Floating nnna0	5 random pdfs	10 random pdfs
neutrons	1993(114)	1922(169)
alphas	7909(132)	7920(152)
nnna0	36(29)	48(50)
nnna26	361(127)	205(225)
neutron NoverA scale	1.098(31)	1.112(25)
alpha NoverA scale	0.966(02)	0.966(02)
Fit Quality	80/94(84%)	83/94(79%)
Floating nnna26	**5 random pdfs**	**10 random pdfs**
neutrons	1989(80)	1941(95)
alphas	7978(121)	8032(121)
nnna0	70(66)	107(143)
nnna26	185(149)	199(145)
neutron NoverA scale	1.080(27)	1.120(12)
alpha NoverA scale	0.966(02)	0.966(02)
Fit Quality	81/94(84%)	81/94(83%)
Floating nnna0+nnna26	**5+5 random pdfs**	**10+10 random pdfs**
neutrons	1926(165)	1979(106)
alphas	7963(130)	7984(143)
nnna0	58(80)	99(78)
nnna26	258(148)	149(130)
neutron NoverA scale	1.112(31)	1.101(09)
alpha NoverA scale	0.966(02)	0.966(02)
Fit Quality	82/94 (82%)	79/94(86%)

Table 7.10: Influence of the NNNA pdf uncertainties on the NoverA fit to Monte Carlo data. The data were generated with neutron and alpha NoverA scales set to the central values from their prior. The neutron and alpha scales were then floated subject to their prior.

7.3.3   Biases and Correlated Signals

The biases and pulls are used again here to characterize any biases on the inferred signals from the use of NoverA as an observable. 100 fake Monte Carlo data sets were generated, with neutron NoverA scale and alpha NoverA scale set to their mean values from Table 7.7. The number of events for each signal were distributed about the same Poisson means as they were for energy. The scale parameters (neutrons and alphas) were floated with the constraint from the priors and the NNNA pdfs were fit as the linear combination of 10 randomized pdfs for each signal.

Figures 7.9 and 7.10 show the biases for the 4 signals as a function of the fake data set along with projections on the y-axis and figure 7.11 shows the corresponding pulls. These plots show that neutron have a slight tendency (1.1%) to fit out low. The nnna0 events have the largest bias and pull.
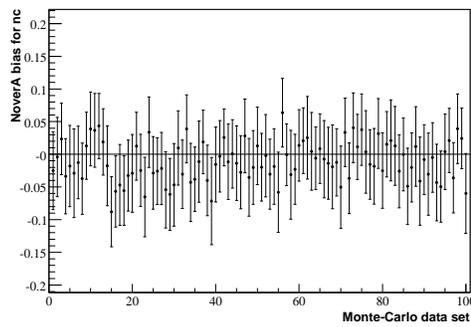
(a) neutron NoverA bias
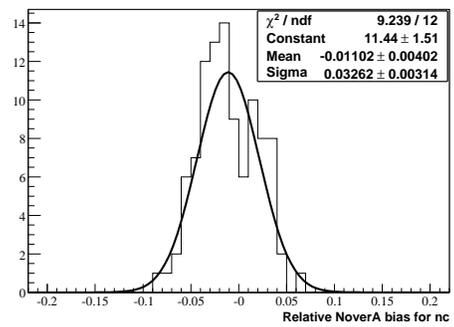
(b) neutron NoverA bias projection
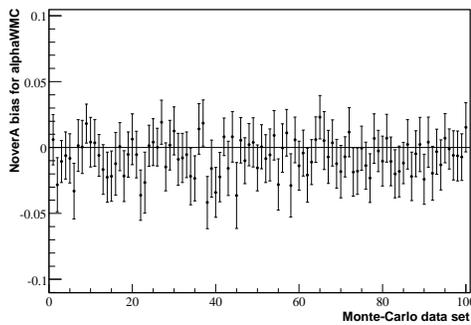
(c) alpha NoverA bias
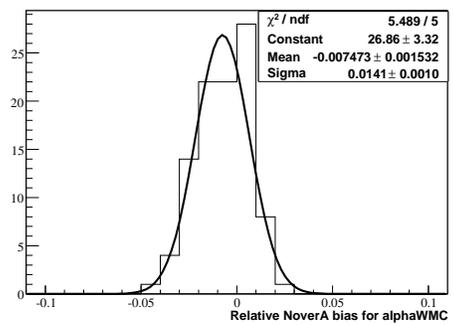
(d) alpha NoverA bias projection

Figure 7.9: Biases and projections for the neutron and alpha signals when using NoverA as an observable plotted as a function of the 100 Monte Carlo data sets.

(a) nnna0 NoverA bias

(b) nnna0 NoverA bias projection

(c) nnna26 NoverA bias

(d) nnna26 NoverA bias projection
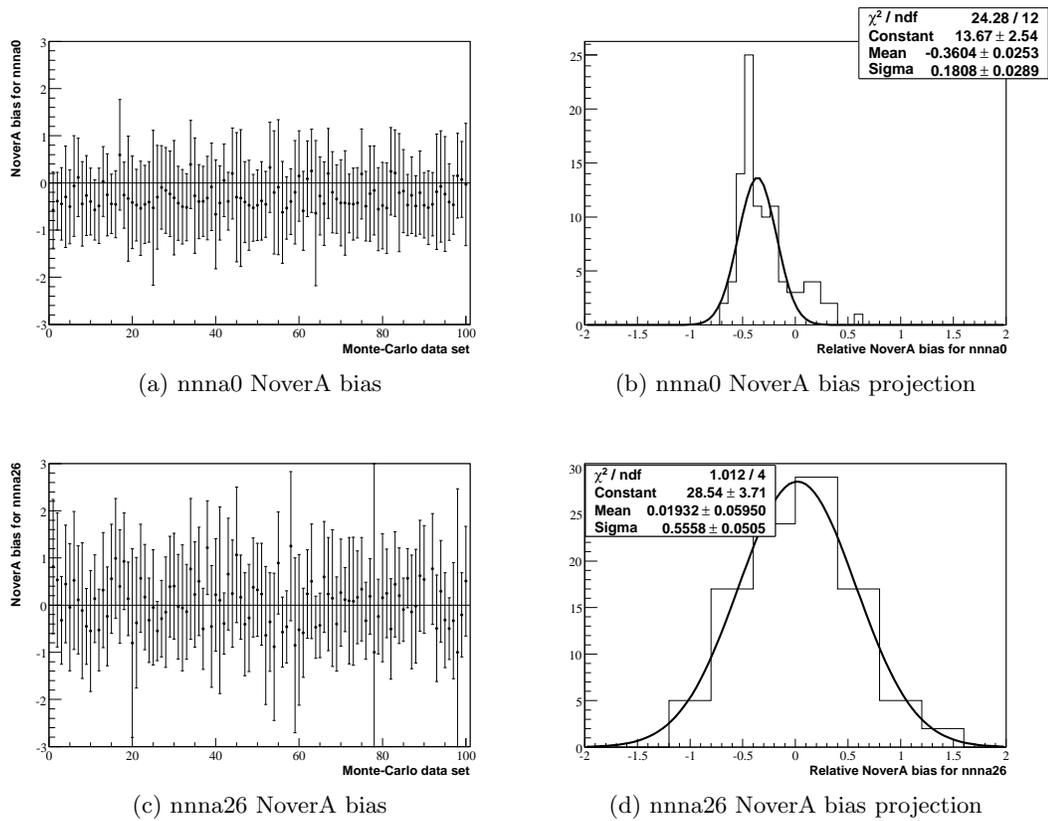
Figure 7.10: Biases and projections for the NNNA signals when using NoverA as an observable plotted as a function of the 100 Monte Carlo data sets.

(a) neutron NoverA pull

(b) alpha NoverA pull

(c) nnna0 NoverA pull

(d) nnna26 NoverA pull
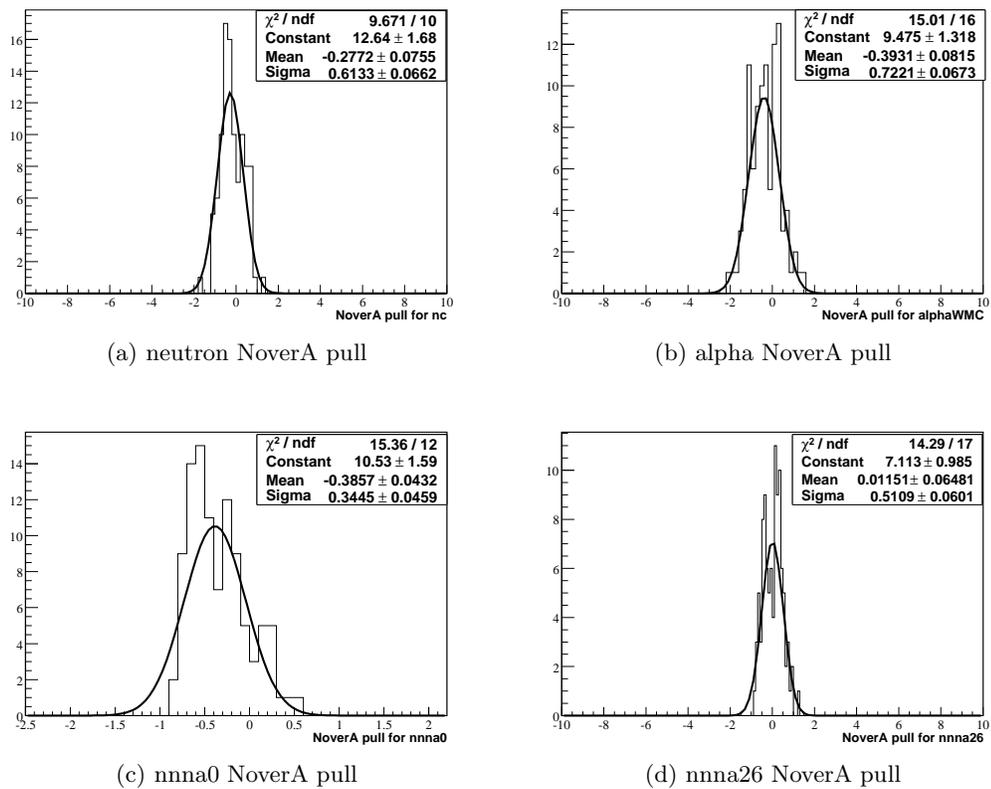
Figure 7.11: Pull plotted for each signal as determined from a sample of 100 Monte Carlo data sets. The 100 data sets were fit using the NoverA observable and floating all systematic parameters.

7.4   Correlations Between Energy and NoverA

This section justifies the factorization of the joint pdfs for energy and NoverA for neutrons and alphas. This factorization will allow the use of two 1-dimensional pdfs (in energy and NoverA) instead of a single 2-dimensional pdf for each signal when the fit is performed using both energy and NoverA. If, for example, energy and NoverA were correlated, the joint probability can not necessarily be factorized as was suggested in equation 6.5 from chapter 6.

In order to validate this factorization, 2-dimensional pdfs need to be obtained for the signals. This is easily achieved for neutrons, by using the data from the ^{24}Na spike. A 2-dimensional pdf for alphas can be obtained from the ^{4}He-strings with reasonable statistics. However, the statistics in the NNNA events do not allow for 2-dimensional pdfs to be generated. Since fits to data show that the number of NNNA events is low, these are ignored in this validation to a very good approximation.

Figure 7.12a shows the 2-dimensional distribution of energy and NoverA that is obtained for neutrons from the 2005 ^{24}Na data. Panel (b) shows the same pdf as obtained by the product of the two 1-dimensional distributions and looks qualitatively similar. Figure 7.13 shows the case for alphas, as deduced from ^{4}He data. The 2-dimensional alpha pdfs disagree the most at higher energies.

Validating the factorization of the pdfs was be done by creating data sets from the 2-dimensional pdfs and then fitting those with the 1-dimensional pdfs. It was found that limiting the energy fit range to 0.9 MeV removes any bias in the inferred amount of neutrons. Figure 7.14 shows the bias and pull distribution for neutrons as deduced from 100 sets of Monte Carlo data. The Monte Carlo data were generated by 'throwing' events from the 2-dimensional pdfs for neutrons and alphas (panels (a) of Figures 7.12 and 7.13) using ROOT's TH2F::GetRandom2() method. The data set contained neutrons distributed about a Poisson mean of 2000 and alphas distributed about 8000. Fitting the data using the same

(a) 2-d pdf for ^{24}Na neutrons

(b) Product of two 1-d pdfs for ^{24}Na neutrons
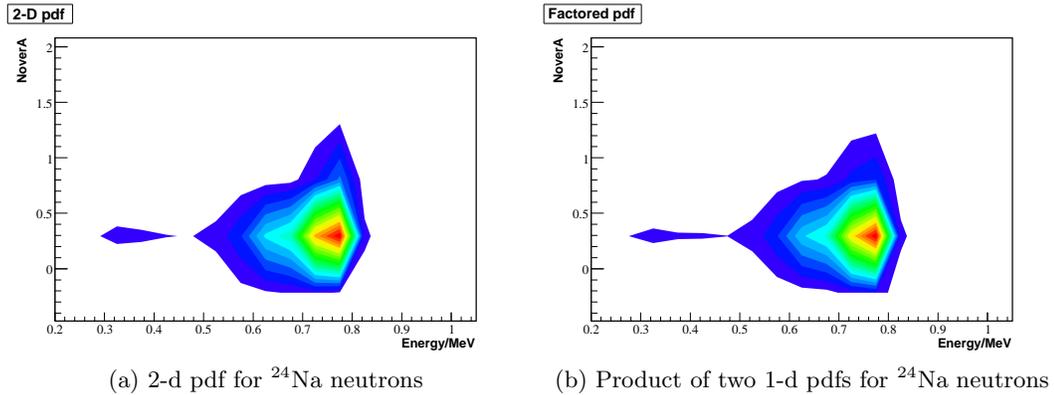
Figure 7.12: Two-dimensional distribution of energy and NoverA for neutrons as measured with data from the 2005 ^{24}Na spike. Panel (a) shows the 2-dimensional (normalized) distribution and panel (b) shows the product of the two 1-dimensional distributions for energy and NoverA.

(a) 2-d pdf for ^{24}Na neutrons

(b) Product of two 1-d pdfs for ^{24}Na neutrons
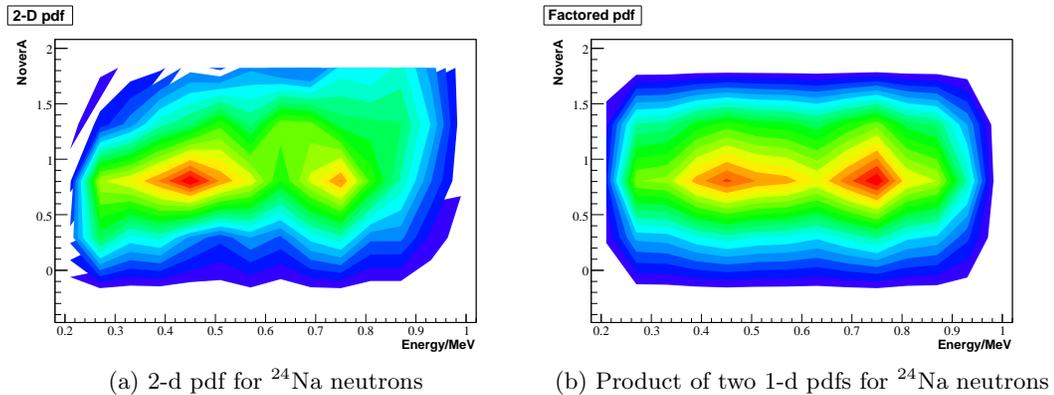
Figure 7.13: Two-dimensional distribution of energy and NoverA for alpha, measured with data from the ^{4}He-strings. Panel (a) shows the 2-dimensional (normalized) distribution and panel (b) shows the product of the two 1-dimensional distributions for energy and NoverA.

2-dimensional pdfs resulted in no bias. The 1-dimensional pdfs that were used to fit the data were the same neutron and alpha pdfs used throughout this thesis. It was found that building 1-dimensional pdfs from the ^{24}Na and ^{4}He data did not change the results for this bias study, even though those pdfs are closer to what was used to generate the fake data sets. This robust result validates the use of factored 1-dimensional pdfs to represent the joint probabilities for neutrons and alphas.

| (a) neutron bias | (b) neutron pull |
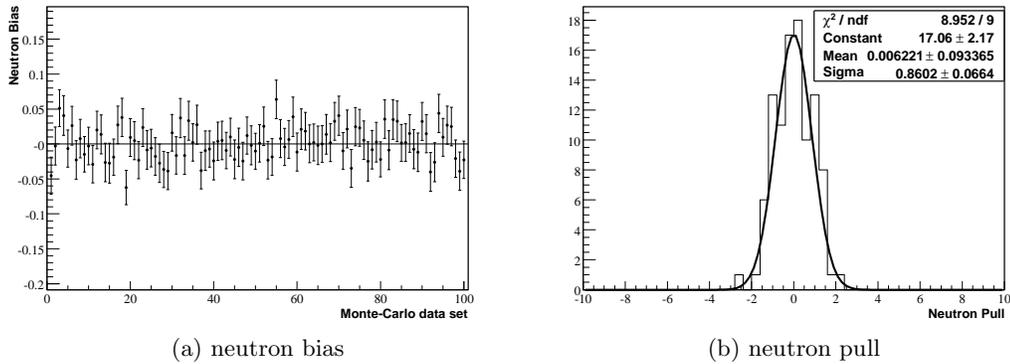
Figure 7.14: Bias and pull for neutrons from 100 Monte Carlo 'fake' data sets. The fake data were thrown from the 2-dimensional pdfs for neutrons and alphas determined in Figures 7.12 and 7.13. The MCMC fit was then done using the 1-dimensional pdfs for neutrons and alphas from this thesis (as opposed to 1-d pdfs generated from the same data). There is no significant bias or pull in the number of neutrons, thus justifying the factorization of the (energy, NoverA) joint probabilities into 1-dimensional pdfs. It was found that the energy range in the fit needs to be limited to 900 keV in order for this factorization to be valid.

7.5 Extracting Signals from Monte Carlo Data Using Energy and NoverA

This section considers the combined use of energy and NoverA to determine the number of neutrons in the NCD data. It is expected that the overall uncertainty in the extracted number will be smaller than if only one of the observables were used. The addition of the NoverA observable to energy in a maximum-likelihood analysis is the original idea that

neutrons	2005(59)
alphas	7849(136)
nnna0	85(49)
nnna26	350(78)
Energy Fit Quality	13/24(97%)
NoverA Fit Quality	80/94(85%)

Table 7.11: Fits to Monte Carlo data using Energy+NoverA and no systematic uncertainties. Using the combination of observables has resulted in a substantially smaller uncertainty in the number of neutrons.

motivated the work in this thesis. The influence of the energy and NoverA systematic uncertainties is explored here along with any possible biases in using the combination of observables. This study will again be carried out by analyzing Monte Carlo data sets generated from the 1-dimensional pdfs.

7.5.1 Without Systematic Uncertainties

Data with no scale parameters applied to it were generated and analyzed with the 1-dimensional pdfs with no systematic uncertainties. This is done here in order to later understand the effect of including the systematic uncertainties in the shapes of the pdfs. Table 7.11 shows the signals and uncertainties that were fit by using energy and NoverA as observables. Figure 7.15 shows the resulting fits to the Monte Carlo data. The data were fit with $0.25\,\text{keV}$ bins in energy from $0.2\,\text{MeV}$ to $0.9\,\text{MeV}$ and in 0.1 bins in NoverA from 0 to 10. As anticipated, the uncertainty on the number of neutrons improves substantially (from $\sim$80 to $\sim$60) compared to the fits in energy or NoverA alone.

7.5.2 With Systematic Uncertainties

The effect of the systematic uncertainties from energy and NoverA are first considered independently, then in conjunction.

(a) Energy Fit   (b) NoverA Fit
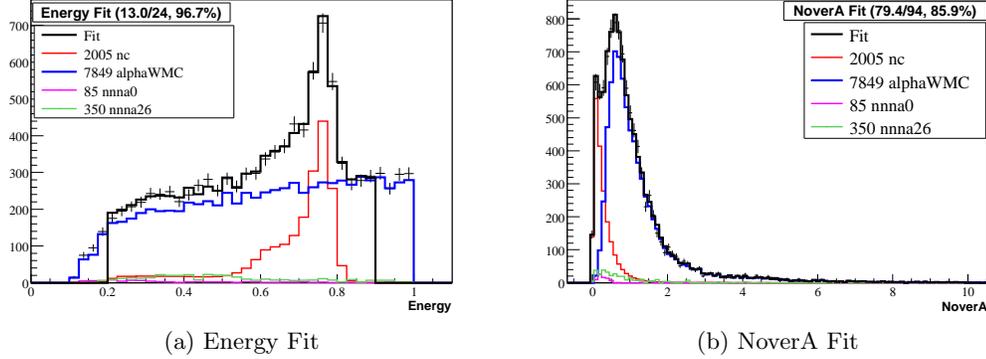
Figure 7.15: Energy (a) and NoverA (b) fits to Monte Carlo data with the fit carried out jointly over the two observables.

With Energy Systematic Uncertainties

Monte Carlo data were generated with the neutron and alpha energy deformation parameters set to their central values. The data were then fit with the energy deformation parameters floated with their prior constraints and the NNNA energy pdfs were built as the linear combination of 10 randomized pdfs. No systematic uncertainties were applied to the NoverA pdfs and the NoverA part of the data had no systematic uncertainties applied. Table 7.12 shows the results from the fit. The inclusion of the energy systematic uncertainties does not impact the neutron uncertainty. This demonstrates that the combination of both observables results in a stable way of estimating the number of neutrons.

With NoverA Systematic Uncertainties

The effect of the NoverA systematic uncertainties is studied independently by fitting a Monte Carlo data set where only the NoverA pdf uncertainties are floated. The data were generated with neutron and alpha NoverA scales set to their central values from their prior constraints. The data were then fit with the NoverA deformation parameters floated subject to their priors and the NoverA NNNA pdfs floated as the linear combination of 10 random

neutrons	1988(59)
alphas	7807(136)
nnna0	60(53)
nnna26	432(101)
neutron energy scale	1.000(02)
alpha energy shift	-0.040(03)
Energy Fit Quality	18/24(78%)
NoverA Fit Quality	118/94(5%)

Table 7.12: Fits to Monte Carlo data using Energy+NoverA and including energy systematic uncertainties. The data were generated with energy deformation parameters applied and set to the central values from their priors. The fit was then done by floating the energy deformation parameters subject to their priors and the NNNA pdfs were floated with 10 components. The NoverA systematic parameters were set to their un-deformed values in the data and were not floated. Adding the systematic uncertainties has no effect on the neutron uncertainty, but has affected the quality of the NoverA fit.

pdfs (each). The energy deformation parameters were not floated and set to un-deformed values to generate the data. The results of the fit are show in Table 7.13. The NoverA systematic uncertainties also result in no significant changes in the neutron uncertainty.

With Energy+NoverA Systematic Uncertainties

Finally, the full effect of the pdf systematic uncertainties on a fit in energy and NoverA is studied. A Monte Carlo data set was generated with all (energy+NoverA) deformation parameters set to the central values from their priors. The fit was then done by floating all deformation parameters subject to their priors and the NNNA pdfs were (each) floated with 10 randomized pdfs. This results in a fit with 4(signals)+ 4(deformation parameters)+40(linear combination parameters)=48 parameters[2]. The uncertainty in the number of neutrons is not increased substantially compared to the case with no deformation on the pdfs. Including the uncertainties in both observables has now allowed for 'good fits' in both observables.

---

[2]It should be clear to the reader why the MCMC method was chosen over minimizing algorithms

neutrons	1999(59)
alphas	7804(145)
nnna0	45(60)
nnna26	417(112)
neutron NoverA scale	1.095(02)
alpha NoverA scale	0.967(02)
Energy Fit Quality	13/24(97%)
NoverA Fit Quality	94/94(49%)

Table 7.13: Fits to Monte Carlo data using Energy+NoverA and including NoverA systematic uncertainties. The data were generated with NoverA deformation parameters applied and set to the central values from their priors. The fit was then done by floating the NoverA deformation parameters subject to their priors and the NNNA pdfs were floated with 10 components.

neutrons	2000(61)
alphas	7852(154)
nnna0	97(67)
nnna26	358(124)
neutron energy scale	1.000(02)
alpha energy shift	-0.040(3)
neutron NoverA scale	1.091(22)
alpha NoverA scale	0.968(02)
Energy Fit Quality	19/24(77%)
NoverA Fit Quality	92/94(75%)

Table 7.14: Fits to Monte Carlo data using Energy+NoverA and including all systematic uncertainties. The data were generated with deformation parameters for energy and NoverA applied and set to the central values from their priors. The fit was then done by floating all deformation parameters subject to their priors and the NNNA pdfs were floated with 10 components for each observable and signal.

7.5.3 Biases and Correlated Signals

Again, biases and pulls are examined in the same way as for energy and NoverA. 100 fake data sets were generated with the deformation parameters set to their central prior values. The data sets were then fit by floating all deformation parameters using their prior constraints and the NNNA pdfs were each floated as the linear combination of 10 randomized pdfs. The amount of each signal in the data sets was generated according to Poisson distributions.

Figures 7.16 and 7.17 show the biases for each signal with their projections and Figure 7.18 shows the corresponding pulls. It is clear that the use of both observables removes any bias or pull in the number of neutrons that was seen with either observable. This encouraging result will lead to no additional uncertainty in the number of neutrons other than determined from the posterior distribution. However, there is clearly a large bias in the number of alphas that trade off with the nnna26 events, which show the largest relative bias.

(a) neutron Energy+NoverA bias

(b) neutron Energy+NoverA bias projection

(c) alpha Energy+NoverA bias

(d) alpha Energy+NoverA bias projection
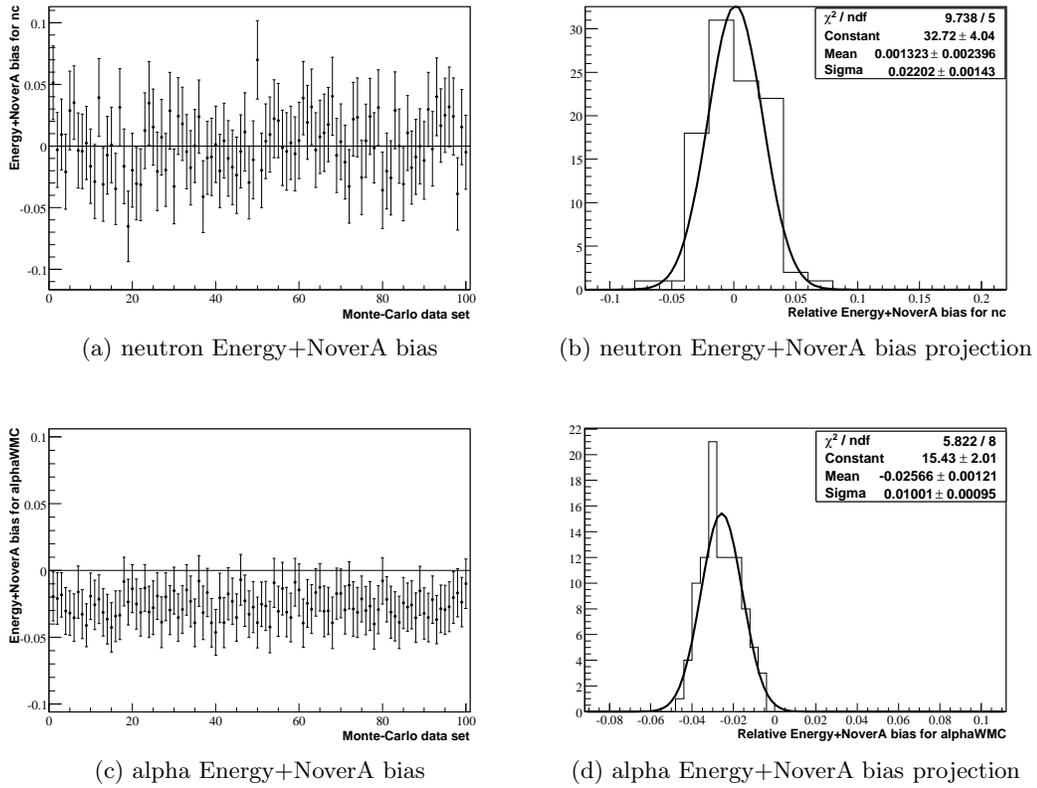
Figure 7.16: Biases and projections for the neutron and alpha signals when using Energy+NoverA as an observable plotted as a function of the 100 Monte Carlo data sets. There is no bias in the number of neutrons when both observables are combined. There is however clear evidence for a large trade off between alphas and the nnna26 events. The corresponding pulls are shown in Figure 7.18.

(a) nnna0 Energy+NoverA bias

(b) nnna0 Energy+NoverA bias projection

(c) nnna26 Energy+NoverA bias

(d) nnna26 Energy+NoverA bias projection
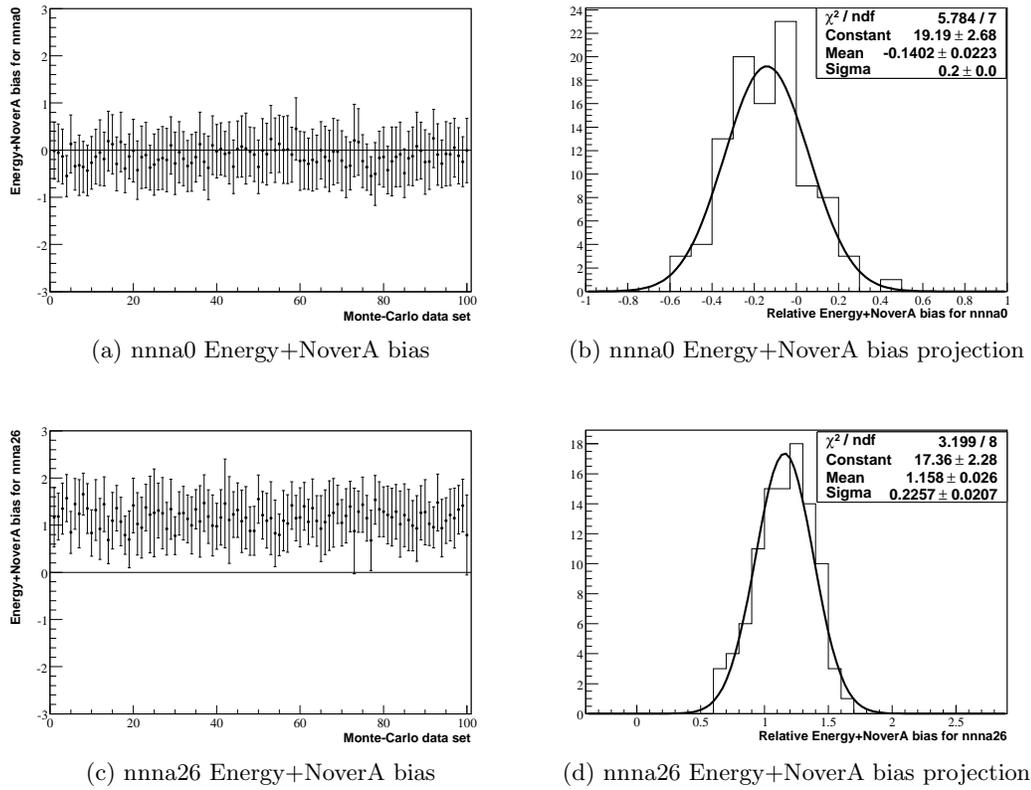
Figure 7.17: Biases and projections for the NNNA signals when using Energy+NoverA as an observable plotted as a function of the 100 Monte Carlo data sets. There is no bias in the number of neutrons when both observables are combined. There is however clear evidence for a large trade-off between alphas and the nnna26 events. The corresponding pulls are shown in Figure 7.18.

(a) neutron Energy+NoverA pull

(b) alpha Energy+NoverA pull

(c) nnna0 Energy+NoverA pull

(d) nnna26 Energy+NoverA pull
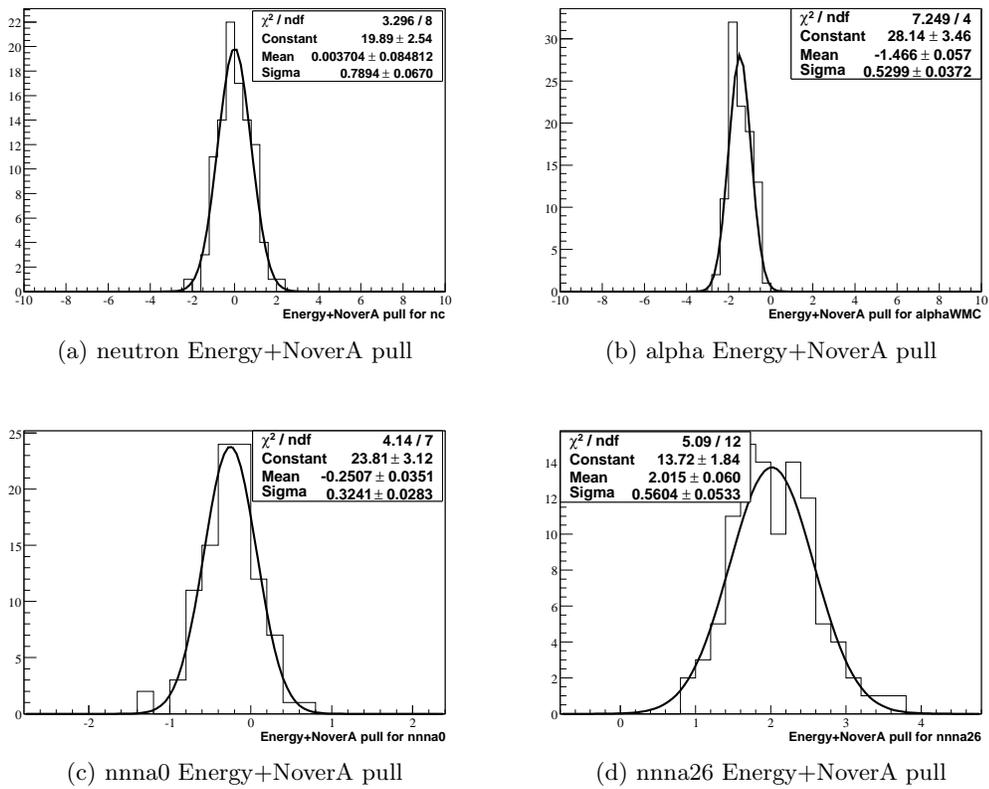
Figure 7.18: Pull plotted for each signal as determined from a sample of 100 Monte Carlo data sets that were fit in Energy and NoverA. Neutrons are consistent with no pull. Alphas and nnna0 have a large negative pull and clearly trade-off with the nnna26 events.

7.6 Signal Extraction on Blind Data

The tests that have been presented so far only served to demonstrate the 'separation power' of the energy and NoverA observables as well as the impact of the pdf uncertainties on the extracted number of neutrons. The last section of this chapter demonstrates that these observables can indeed be used to analyze the real data that was taken with the Neutral Current Detectors.

In order to test this 'signal extraction' on real data, a blind data set was created. This data were generated by removing an (unknown) fraction of events as well as adding in an unknown number of neutrons from muon spallation events. This 'Wilkerson-Elliott' blinding procedure was implemented in a similar fashion to the data for the currently published analysis [38]. At the time the analysis presented in this thesis was performed (June 2009), the data from the NCD phase had already been published; the data set discussed here is thus the result of a 're-blinding' of the previously 'open' data and was implemented in order to avoid any bias resulting from the knowledge of the results from the published analysis. The next chapter will present the results on a data set that only has a statistical blindness applied. This thesis chapter was written before 'un-blinding' and was part of documentation submitted to the SNO collaboration to request the removal of the blindness scheme.

7.6.1 Just Energy

It is relevant to consider the fit of the blind data carried out in the individual observables. It is of particular interest to determine the number of neutrons using only the energy as this observable has already been used ('vetted') in the published analysis [38]. It should be noted that the energy fit discussed in this work is somewhat different than in the published analysis, specifically in the way that the backgrounds are handled.

Table 7.15 shows the result of the fit to the blind data using only energy as an observable. The first column shows the results with no systematic uncertainties applied to the shape

	Without Systematics	With Systematics	No Prior
neutrons	285(33)	284(35)	271(37)
alphas	925(91)	908(107)	860(115)
nnna0	14(20)	32(39)	67(86)
nnna26	161(58)	185(75))	479(195)
neutron energy scale	fixed	1.001(04)	0.996(06)
alpha energy shift	fixed	-0.047(15)	-0.276(140)
Energy Fit Quality	25/24 (39%)	21/24 (66%)	34/24(9%)

Table 7.15: Fit to blind neutrino data using an energy fit with 0.25 keV bins from 0.2 MeV.to 0.9 MeV (fitting up to 1 MeV does not change any of the results). The first column shows the results without the systematic uncertainties in the shape of the energy pdfs. In the second and third columns, the NNNA energy pdfs were floated as the linear combination of 10 random pdfs. The second column was done with priors applied to the systematic parameters and no priors were applied in the the third.

of the pdfs, the second column shows the results when the energy systematic uncertainties are included with a prior and the third column shows the case when no prior is included. The fit was done from 0.2 to 0.9 MeV using 0.25 keV bins. Extending the range to 1 MeV changes the number of neutrons with uncertainties (and priors) included to 286(31), which results in a slightly improved uncertainty. Limiting the range to 0.9 MeV was done so that effect of including the NoverA observable could be directly compared. The pdf uncertainties have little influence on the accuracy in the number of neutrons. They however play a role in changing the uncertainties in the other signals. Panel (a) of Figure 7.19 shows the resulting fit in energy (with systematic uncertainty and priors) of the blind data as well as the energy pdfs that were thus determined for the NNNA events. The case with no prior applied results in a very large alpha energy shift (300 keV) which is clearly an indication of a failure in the algorithm. This is due to the relative flatness of the alpha and NNNA pdfs which can trade-off with each other. In this case, the alpha pdf was shifted out of the range and the nnna26 signal was increased to compensate. Panel (b) of Figure 7.19 shows the correlation between the alpha energy shift with the nnna26 signal and panel (c) shows that the neutron number is virtually independent of this failure to determine the alpha energy shift.
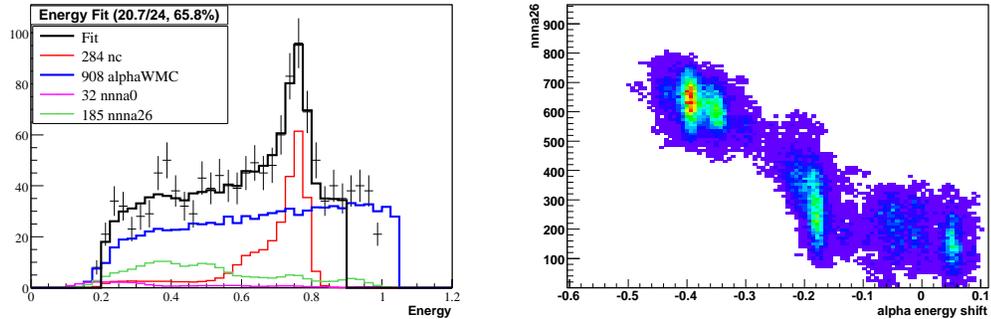
(a) Energy only fit with priors

(b) Alpha energy shift correlation with nnna26 in fit with no prior.
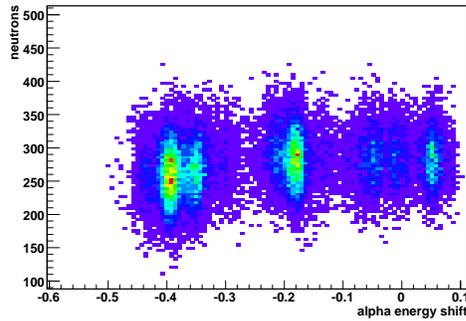
(c) Alpha energy shift correlation with neutrons in fit with no prior.

Figure 7.19: Panel (a) shows the fit to the blind neutrino data in energy using systematic uncertainties in the shape of the pdfs with prior constraints applied. The resulting fit has also determined the shape of the NNNA energy pdfs. Panels (b) and (c) show the correlations of the nnna26 and neutron signal with the alpha energy shift in the case where no prior was applied. The lack of prior makes the alpha energy shift difficult to determine because it can trade off with the nnna26 signal, but the number of neutrons is largely unaffected.

	Without Systematics	With Systematics	No Prior
neutrons	274(40)	299(43)	320(39)
alphas	973(44)	962(39)	956(41)
nnna0	63(46)	40(55)	29(36)
nnna26	19(28)	40(33)	39(36)
neutron NoverA scale	fixed	1.049(27)	0.986(16)
alpha NoverA scale	fixed	0.938(22)	0.929(05)
NoverA Fit Quality	99/94(34%)	98/94(38%)	96/94 (41%)

Table 7.16: Fit to blind neutrino data using NoverA. The first column shows the results without the systematic uncertainties in the shape of the NoverA pdfs. In the second and third columns, the NNNA NoverA pdfs were floated as the linear combination of 10 random pdfs.

7.6.2 Just NoverA

Table 7.16 shows the signals determined by fitting the blind data using only the NoverA observable from 0 to 10 in bins of 0.1. This results in a less accurate estimate of the number of neutrons compared to the energy fit, due to the fact that neutron pdf is 'less different' in NoverA than it was in energy. Also, the NoverA fit results in considerably less nnna26 events than the energy fit. The inclusion of the systematic uncertainties results in a substantial change in the number of neutrons, suggesting that NoverA is not particularly good at constraining neutrons when used by itself. This was indeed already seen in the Monte Carlo tests.

The NoverA distribution however provides a strong constraint on the number of alphas in the data and results in a smaller uncertainty in alphas than for the energy fit. It is this added constraint on the number of alphas that leads to the overall smaller uncertainty in the number of neutrons when both observables are combined.

7.6.3 Using Energy and NoverA

Table 7.17 summarizes the results for fits to the blind data when both observables are used. As anticipated, the use of both observables has significantly improved the accuracy in the
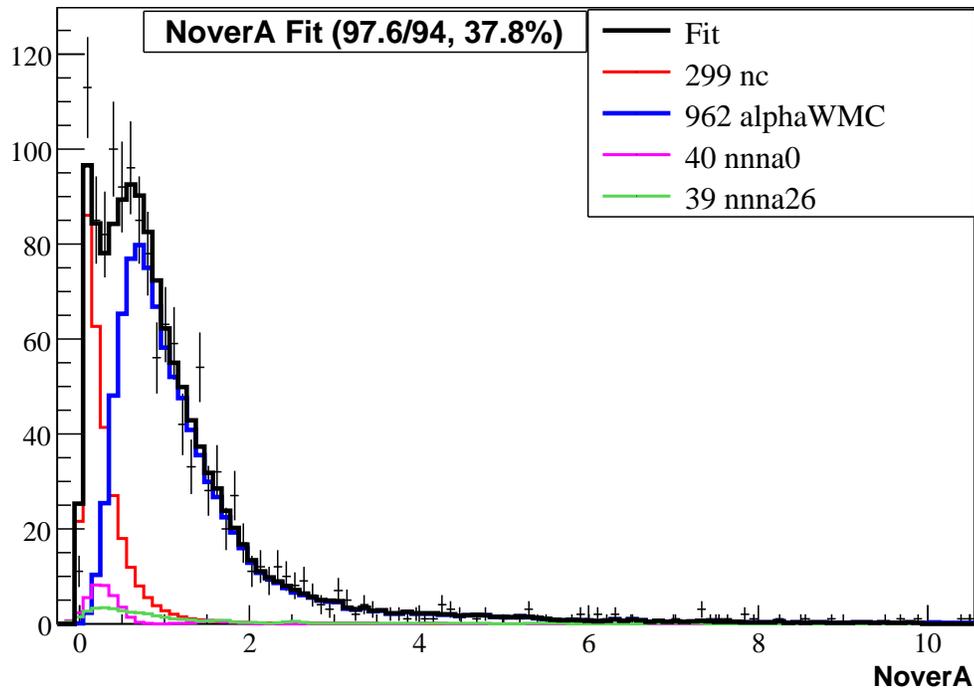
Figure 7.20: Fit to the blind neutrino data in NoverA using systematic uncertainties (with priors) in the shape of the pdfs. The resulting fit has also determined the shape of the NNNA NoverA pdfs.

number of neutrons. It is clear that the fit in energy provides a stronger constraint on the number of neutrons, but the fit in NoverA constrains the alpha signal and thus results in an overall smaller uncertainty in the number of neutrons. The measured number of neutrons in this blind data set from fitting the posterior distribution is determined to be $274.3^{+23.0}_{-21.1}$ when priors are taken into account. Removing the constraints from the prior results in an increase in the neutron uncertainty indicating that the priors were inconsistent with the data. It is evident that the alpha energy shift is responsible for this discrepancy. A conservative approach dictates that the number of neutrons quoted from this data set should be taken from the fit with no prior: $278.8^{+27.5}_{-25.3}$

Some 'diagnostic plots' of the fit to the blind data conclude this chapter. All plots were done for the case where prior constraints were imposed on the parameters. Figure 7.21 shows the result of the fit in energy and NoverA, the negative log-likelihood, as well as the auto-correlation functions for 'some' of the parameters in the fit (the signals, the neutron and alpha deformation parameters and 4 of the linear combination coefficients for the NNNA pdfs). Figures 7.22 and 7.23 show the posterior distributions that were used to determine the values and uncertainties of the parameters in Table 7.17. The Markov-Chain had a length of 200,000 steps where 40,000 steps were thrown out as 'burn-in', resulting in 160,000 iterations to estimate the posterior distributions. The auto-correlation functions show that parameters have forgotten their state after a lag of $\sim$5,000, justifying the burn-in decision.

(a) Energy Fit

(b) NoverA Fit

(c) Negative Log-Likelihood

(d) ACF for signals

(e) ACF for neutron-alpha systematics

(f) ACF for NNNA linear combination parameters
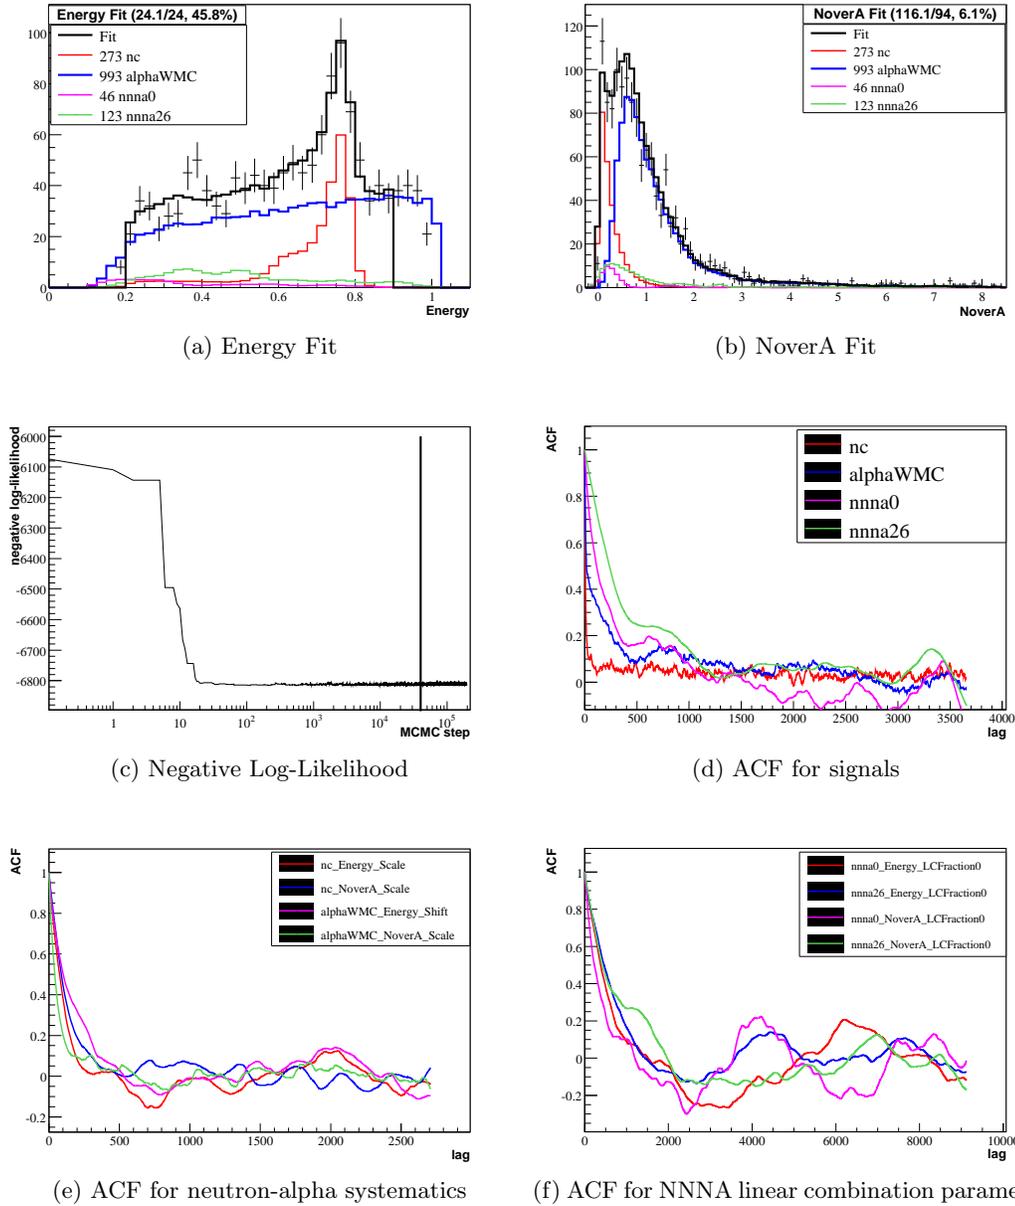
Figure 7.21: 'Diagnostic' plots for the energy+NoverA fit to the blind data. Panels (a) and (b) show the fit in energy and NoverA. Panel (c) shows the negative log-likelihood as a function of step in the MCMC, the vertical line shows where the burn-in cut was placed (at 40,000 of 200,000 steps). Panels (d) and (e) show the auto-correlation function for the signals as well as the neutron and alpha deformation parameters. Panel (f) shows the ACF for 4 of the linear combination parameters that are used to build the NNNA pdfs.

	Without Systematics	With Systematics	No Prior
neutrons	280(22)	274(22)	279(28)
alphas	953(49)	993(63)	957(63)
nnna0	35(21)	46(30)	40(33)
nnna26	143(44)	123(49)	119(56)
neutron energy scale	fixed	1.000(03)	0.998(04)
alpha energy shift	fixed	-0.018(32)	0.034(38)
neutron NoverA scale	fixed	1.078(32)	1.07(12)
alpha NoverA scale	fixed	0.991(09)	0.993(08)
Energy Fit Quality	28/24(27%)	24/24(46%)	23/24(53%)
NoverA Fit Quality	110/93(11%)	116/94(6%)	110/94(12%)

Table 7.17:   Fit to the blind neutrino data using both energy and NoverA. The second column includes systematic uncertainties in the shape of the pdfs.

(a) Neutron posterior $274.3^{+23.0}_{-21.1}$

(b) Alpha posterior $993.5^{+59.5}_{-61.9}$

(c) nnna0 posterior $45.7^{+34.0}_{-27.7}$

(d) nnna26 posterior $122.8^{+58.9}_{-47.2}$
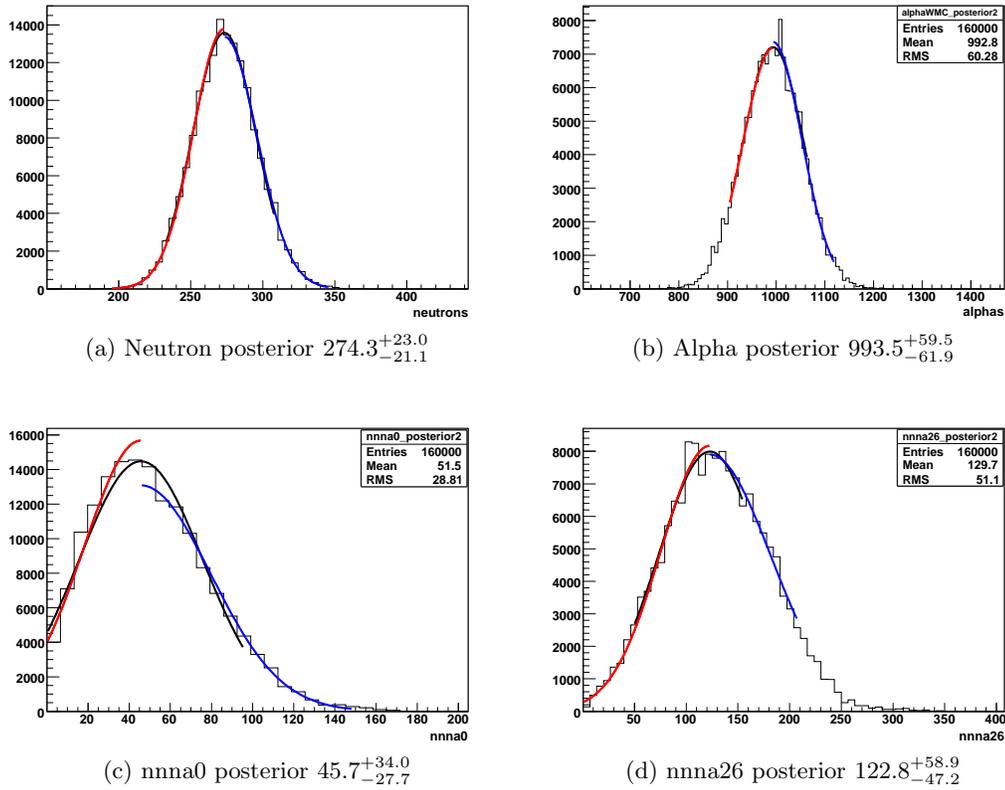
Figure 7.22: Posterior distributions for the signals from the fit to the blind data in energy+NoverA. The histograms are shown together with the fits to normal distributions that were used to determine the parameter means and uncertainties.
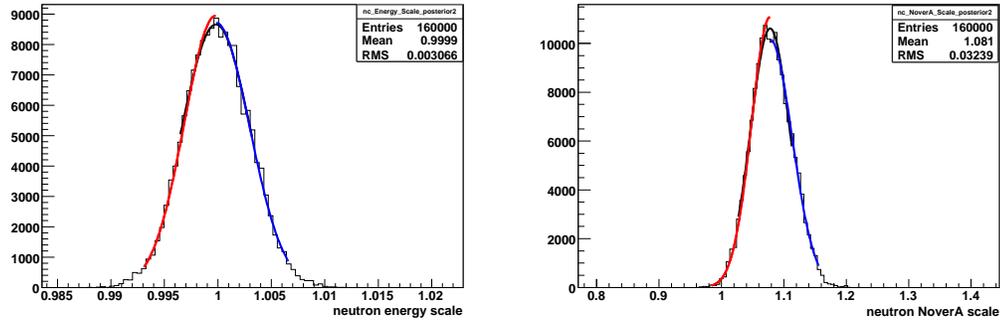
(a) Neutron energy scale posterior $0.9998^{+0.0032}_{-0.0030}$

(b) Neutron NoverA scale posterior $1.0780^{+0.0353}_{-0.0298}$
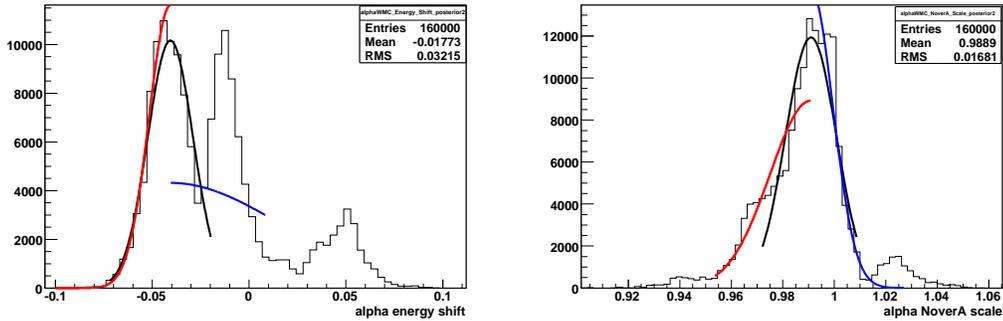
(c) Alpha energy shift posterior - mean+rms

(d) Alpha NoverA scale posterior $0.9909^{+0.0086}_{-0.0159}$

Figure 7.23: Posterior distributions for the neutron and alpha systematic parameters from the fit to the blind data in energy+NoverA. The alpha energy shift posterior is multi-modal, so that the quoted value was not based on the fit to a normal distribution, but rather, using the mean and rms of the posterior distribution. It should be noted that this 'multi-modality' in the parameter is not an indication that the fit for the other parameter has not converged. It is a reflection of the fact that the likelihood has several local minima as a function of that parameter. The MCMC chain has, however, integrated (marginalized) over this distribution when determining, for example, the number of neutrons.

Chapter 8

Final Analysis Using the One-Third Data Set

This chapter uses the extended maximum-likelihood fit developed in the previous chapters to determine the number of neutrons in the reduced NCD data set. This data set contains exactly one third of the data set that was published and the results presented here are thus directly comparable. The MCMC will be used to carry out this analysis using the energy and NoverA observables presented in this thesis. Results are then compared to the analysis that was published previously as well as the standard solar model prediction.

8.1 The One-Third Data Set

To date, no publication of the NCD data has included information from pulse shapes to determine the number of neutrons in the data. The analysis presented in this thesis is one of two methods being pursued by the SNO Collaboration that incorporate pulse shape information. A blindness scheme (described in the previous chapter) was thus imposed on the data in order to avoid influencing these two methods with the results from the published analysis. These pulse-shape based analysis techniques are intended to be used by the SNO

Collaboration in a future publication that will cover all three phases of the experiment and are, as such, subject to review by a SNO Pulse Shape Topical Committee. At the time that this thesis was submitted, the committee has not had time to fully review both methods and it was decided that the blindness would not be lifted on the full data set before this thesis was submitted. However, in order to allow the work in this thesis to have a 'physics result', it was decided to partially lift the blindness and retain only a 'statistical blindness' scheme. This has led to the creation of a data set that includes only one third of the data from the NCD phase; the data set was created by selecting every third event in the NCD phase, so that it is uniformly distributed in time and representative of the entire data set. Analysis of this 'reduced data set' then allows statistical and systematic effects to be fully defined and compared with the results from the previous analysis. In addition, it is also possible to infer the expected uncertainties that this method will have once the full data set is analyzed, so that the impact of this new technique can be quantified.

8.2 The Published Analysis

A neutrino flux analysis of the NCD data has already been carried out and published [38]. This section briefly describes how this work was done and how it can be compared to the results in this thesis.

The published analysis used a Markov-Chain Monte Carlo method to determine the solar neutrino flux from the data and is fully described in [81, 82]. The likelihood function was built using the data from both the PMT and NCD data.

The PMT data were fit using energy, $\cos\theta_{\odot}$ and event radial position as observables. The PMT pdfs were generated from Monte Carlo and contained contributions from neutrons, ES events and CC events. The ES and CC events were fit as separate signals in each of 13 PMT energy bins. The neutrons were further divided into contributions from: NC, acrylic vessel photo-disintegration, D_2O neutron backgrounds, NCD neutron backgrounds (from

materials in the NCDs), K2 neutron backgrounds (from a localized source located on string 31), K5 neutron backgrounds (from a localized source located on string 18) and atmospheric neutrino-induced neutron backgrounds. All neutron backgrounds were constrained with priors determined from other analyses [83]. This resulted in 33 different signals in the PMT data.

In addition to the signals, 12 systematic deformation parameters were floated in the PMT-side of the likelihood subject to priors obtained from calibration data. These parameters included an energy scale and resolution, shift parameters for reconstructed positions (x, y, z), an overall scale parameter for reconstructed positions, two parameters to model a z-position resolution and one PMT angular resolution parameter to deform the ES $\cos\theta_\odot$ pdf. The resulting fits over the PMT observables were summarized in chapter 1 and illustrated in Figure 1.9.

The NCD data were fit using ADC energy as the only observable. Signals included neutrons, alphas and events from the NNNA strings. The neutron pdf was obtained from the 2005 ^{24}Na calibration and was allowed to float with a resolution parameter. The neutron signal was further broken up into the same contributions as the neutron signal in the PMTs. The NNNA signals were modeled as skew gaussians subject to a prior constraint on the mean. The alpha (central) pdf was determined from the Monte Carlo in the same way as was done in this thesis (chapter 5). However, the central alpha energy pdf was allowed to float by deforming it with the addition of 8 different polynomials to simulate 8 different variations to the inputs of the Monte Carlo. The amplitude, α_i, of the deformation from each parameter was then floated. The effective (deformed) alpha pdf, $P^{sys}(E)$, was built from the central alpha pdf, $P^{Central}(E)$, as follows:

$$P^{sys}(E) = P^{Central}(E) + \sum_i^7 \alpha_i \times s_i(E) \tag{8.1}$$

where priors on the s_i were obtained by deforming the central Monte Carlo pdf to make it agree with pdfs generated from Monte Carlo simulations with different input parameters.

The s_i then model the variation between the central simulation and the ones with different parameters [67]:

$$^{210}\text{Po alpha depth variation: } s_0 = -2.06 + 6.58E - 6.56E^2 + 2.11E^3$$

$$\text{Bulk alpha depth variation: } s_1 = -0.0684 + 0.0892E$$

$$\text{Drift time variation: } s_2 = -0.131 + 0.252E - 0.117E^2$$

$$\text{Avalanche width offset variation: } s_3 = -0.0541 + 0.0536E$$

$$\text{Avalanche gradient offset variation: } s_4 = -0.0138$$

$$\text{Ionmobility variation: } s_5 = -0.00930$$

$$\text{Po/bulk fraction variation: } s_6 = -0.00405 + 0.0386E$$

$$\text{Data cleaning systematic: } s_7 = 0.861 - 2.77E + 2.72E^2 - 0.870E^3$$

$$(8.2)$$

This results in a large number of degrees of freedom for the alpha pdf and is similar to floating a polynomial to represent the background (alpha + NNNA), subject to the constraints available from calibration data (the ^{4}He -strings and analysis of the NNNA strings).

The number of neutrons from the NC interaction of neutrinos were constrained to be the same in the NCD and PMT sides of the likelihood. This was achieved by floating the flux of neutrinos and using the PMT and NCD neutron detection efficiencies (determined from calibration data) to convert to a number of events seen in each type of instrumentation. Each component of the neutron backgrounds in the PMT was obtained by converting the number seen in the NCDs using the ratio of expected number of events seen in the PMTs to the number seen in the NCDs.

Results from the previously published analysis are difficult to compare directly to the results in this thesis since they determined the neutrino flux and not the number of neutrons in the NCDs directly. The conversion from the flux into a number of neutrons includes systematic uncertainties that are not directly inherent to the analysis of the NCD data.

Section 8.4 describes how a neutrino flux can be converted into a number of detected neutrons in the NCDs. This conversion involves systematic uncertainties that arise from the conversion of neutrino flux into a number of neutrons (detection efficiencies, cross-sections) as well as uncertainties related to the number of background neutrons that are in the data (neutrons from radioactivity). The process of determining a 'number of neutrons detected in the NCDs' is further complicated by the fact that the systematic uncertainties in the conversion from neutrino flux to neutrons were also floated in the published analysis.

Since the previous analysis determined the neutrino flux with statistical and systematic uncertainties, this number must be converted into a number of neutrons with the corresponding contributions. Table 8.5 from section 8.4 shows how a neutrino flux, Φ, can be converted into the number of detected neutrons using a multiplicative factor, X, and the number of background neutrons, B. In order to obtain the number of neutrons with uncertainties that are representative of the analysis of the NCD data and not the conversion from neutrino flux, a few assumptions must be made to remove the uncertainties from the conversion between neutrino flux and neutron number:

- The uncertainty in the multiplicative factor (the conversion between neutrino flux and number of detected NC neutrons) can be removed in quadrature from the quoted systematic uncertainty in the neutrino flux.

- The uncertainty in the number of background neutrons can be removed in quadrature from the systematic uncertainty in the total number of neutrons.

Although it may seem that these systematic uncertainties can be removed in quadrature from the total systematic uncertainty that was quoted, this is not necessarily true because these were floated in the MCMC. In theory, the resulting marginalization results in an equal or smaller effect than addition in quadrature. It was found that the results from the fit determined the uncertainties in the multiplicative factor to be 3.308% compared with 3.334% before the fit and the uncertainty in the number of background neutrons

was determined to be 12.8%, compared with 13.7% before the fit. The changes in the uncertainties being smaller than 1%, it can be safely assumed that these uncertainties can be removed in quadrature from the neutrino flux systematic uncertainty.

The previously published analysis [38] reported the neutrino flux measured with the PMTs and NCDs to be:

$$\phi_{NCD+PMT}^{NC}(\nu_x) = 5.54 \pm 0.32(\text{stat.}) \pm 0.35(\text{sys.}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1} \qquad (8.3)$$

where the asymmetric uncertainties were 'symmetrized' here by taking the average of the positive and negative uncertainties. The 'statistical' uncertainty in the flux was determined by running the MCMC with no systematic parameters. The systematic uncertainty was then determined by running the MCMC with systematic parameters and removing in quadrature the previously determined statistical uncertainty. Any change in the mean value of the flux from the inclusion of the systematic uncertainties was ignored. It should be noted that this method of 'separating out' the systematic uncertainty will not necessarily lead to a robust determination of its value and really depends on how one wishes to define 'systematic uncertainty'. This is because the part that is defined as the statistical uncertainty also includes contributions that are due to the method that is employed to determine the flux, and therefore, of a 'systematic' nature.

Since this analysis included data from the PMTs, it is not directly comparable to the results in this thesis. For this reason, the author of this thesis has used the official software package that was used for the published analysis to determine the same flux by ignoring the data from the PMTs. The resulting fit in NCD energy for the published analysis and the one that was run without the PMTs are compared in Figure 8.1 and seen to be qualitatively very similar. The result that was obtained for the neutrino flux without using the PMT data is:

$$\phi_{NCD}^{NC}(\nu_x) = 5.45 \pm 0.39(\text{stat.}) \pm 0.22(\text{sys.}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1} \qquad (8.4)$$

and it should be noted that the 'balance' between statistical and systematic uncertainties

has changed between this flux and the published one, even though the total uncertainties are similar. This is an illustration of the caution that should be taken when the total uncertainty returned by the MCMC is interpreted and divided into statistical and systematic components.

These neutrino fluxes can now be converted into a number of neutrons with statistical and systematic uncertainties using the multiplicative factor and the number of background neutrons. Removing the uncertainties from the multiplicative factor and background neutrons in quadrature from the systematic uncertainty results in:

$$N_{NCD+PMT} = 1,246 \pm 60(\text{stat.}) \pm 61(\text{sys.})$$

$$N_{NCD} = 1,229 \pm 76(\text{stat.}) \pm 27(\text{sys.})$$

$$(8.5)$$

where $N_{NCD+PMT}$ and N_{NCD} are the number of detected neutrons determined with and without the inclusion of the PMT data, respectively. N_{NCD} is directly comparable to the results in this thesis since it only depends on data from the NCDs and is not biased by any pull from the observables on the PMT side of the analysis. In order to compare N_{NCD} to results from this thesis, it is assumed that the statistical uncertainty will scale by $\sqrt{3}$ and the systematic uncertainty will scale by 3 when the number is scaled down to a value for the one-third data set:

$$N_{NCD}^{\frac{1}{3}} = 410 \pm 44(\text{stat.}) \pm 9(\text{sys.}) \pm 45(\text{tot.})$$

$$(8.6)$$

where the total uncertainty was obtained by adding the statistical and systematic uncertainties in quadrature.

(a) Fit from published result   (b) Fit without PMT data
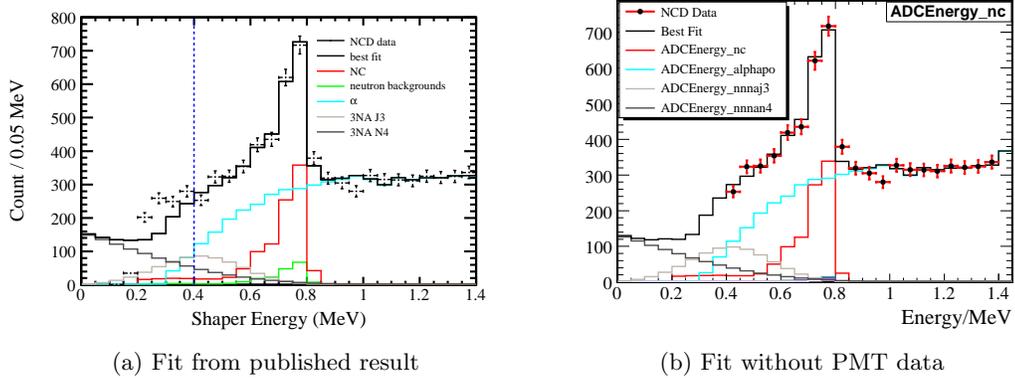
Figure 8.1: Comparison of fit to NCD phase NCD energy spectrum with (a) and without (b) the PMT data. Panel (a) is the fit that was used in the publication, whereas panel (b) is a fit performed by the author of this thesis using the same software but removing the contributions from the PMT data.

8.3 The New Analysis

This section presents the same analysis of the NCD data as was done on the blind data, in section 7.6. In this case one third of the data set that was used for the publication [38] is analyzed for the number of neutrons. Systematic uncertainties are applied in the same way as they were for the blind data set (floating subject to their priors, 10 randomized components for each NNNA pdf).

8.3.1 Using Energy

The NCD data is first re-analyzed using only a fit in the energy observable so that the result from the energy fit can be compared with the published result. This comparison is useful since the method presented here differs significantly from what was previously done (particularly in the way that the backgrounds to the neutron signal are handled). Table 8.1 shows the signals that were extracted from the NCD data with the systematic uncertainties in the energy pdfs applied. The result presented here is virtually unaffected by systematic uncertainties and the uncertainty is thus expected to scale with statistics.

	Without Systematics	With Systematics	No Prior
neutrons	411(39)	406(40)	401(41)
alphas	1490(89)	1505(90)	1482(93)
nnna0	18(25)	122(94)	0(87)
nnna26	65(59)	33(37)	30(49)
neutron energy scale	fixed	1.005(03)	1.008(04)
alpha energy shift	fixed	-0.035(41)	-0.074(97)
Energy Fit Quality	30/24 (19%)	31/24(15%)	35/24(6%)

Table 8.1: Fit to one-third neutrino data using energy. The first column shows the resulting numbers determined with the pdfs fixed to their un-deformed shapes and all systematic parameters fixed. The second column allowed the pdfs to float with deformation parameters constrained with the prior information and the last column allowed the pdfs to float with no constraints from the priors. The quality of the fit is expressed as the chi-squared per number of degrees of freedom and is shown with the corresponding probability.

	Without Systematics	With Systematics	No Prior
neutrons	384(48)	387(50)	351(56)
alphas	1393(49)	1387(58)	1413(61)
nnna0	29(37)	25(34)	40(41)
nnna26	81(59)	43(76)	37(59)
neutron NoverA scale	fixed	1.056(22)	1.243(32)
alpha NoverA scale	fixed	0.986(16)	0.990(14)
NoverA Fit Quality	80/94(86%)	81/94(84%)	77/94(90%)

Table 8.2: Fit to one-third neutrino data using NoverA.

8.3.2   Using NoverA

To provide a comparison with the analysis that was done on the blind data, it is of interest to consider the results when only NoverA is used as an observable. Table 8.2 shows the signals that were determined from a fit in NoverA. Similarly to the case that was seen on the blind data set, it appears that NoverA by itself has trouble in constraining the number of neutrons, especially if the systematic parameters are allowed to float with no priors, as seen in the third column.

8.3.3 Using Energy and NoverA

Finally, the main result of this thesis is presented here. Table 8.3 shows the extracted signals on the one third NCD data set when both energy and NoverA are used. As anticipated, the uncertainty in the number of neutrons has decreased as compared to the fits using only one observable. Allowing the systematic parameters to float with no priors does not change the number of neutrons or the uncertainty. Removing the priors only has a small effect on the alpha energy shift, which was already shown to be difficult to constrain in the blind data. A key difference between this data set and the blind data is that the fit quality in NoverA has improved substantially (72% here vs 12% in the blind data) whereas the quality of the fit in energy has decreased somewhat (from 53% to 25%).

The number of neutrons determined in the one third data set is thus $398.3^{+30.7}_{-29.4}$. The uncertainty is almost unaffected by the inclusion of the systematic parameters. This section concludes with a set of diagnostic plots resulting from the fit in both observables using prior constraints on the systematic parameters (second column of Table 8.3). Figure 8.2 shows the resulting fits in energy and NoverA, the negative log-likelihood as well as the auto-correlation function for some of the parameters. Figures 8.3 and 8.4 show the posterior distributions for the signals and systematic parameters, respectively. The MCMC was run with 200,000 steps and the first 40,000 were discarded as burn-in.

(a) Energy Fit

(b) NoverA Fit

(c) Negative Log-Likelihood

(d) ACF for signals

(e) ACF for neutron-alpha systematics

(f) ACF for NNNA linear combination parameters
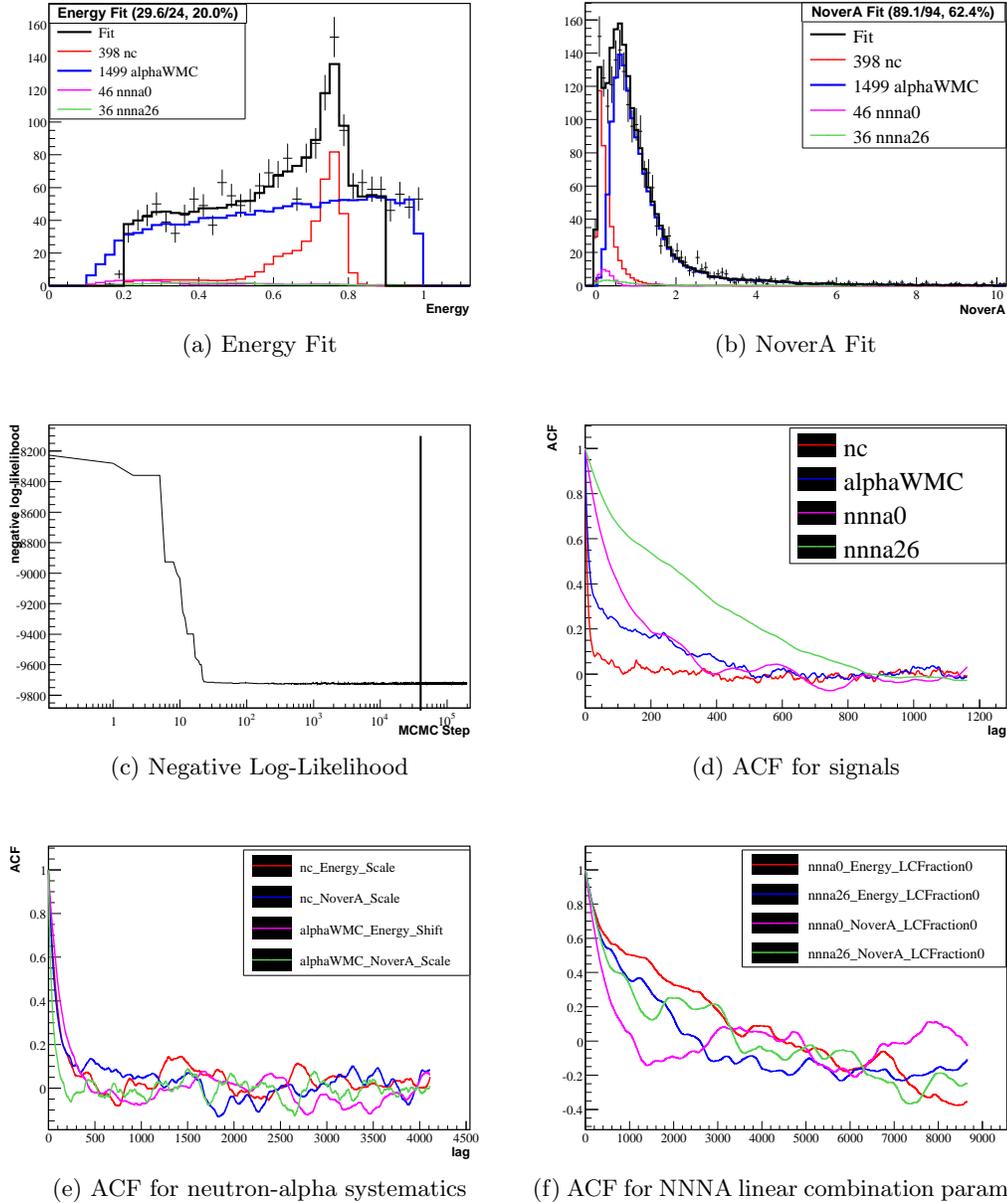
Figure 8.2: 'Diagnostic' plots for the energy+NoverA fit to the one-third data. Panels (a) and (b) show the fit in energy and NoverA. Panel (c) shows the negative log-likelihood as a function of step in the MCMC, the vertical line shows where the burn-in cut was placed (at 40,000 of 200,000 steps). Panels (d) and (e) show the auto-correlation function for the signals as well as the neutron and alpha deformation parameters. Panel (f) shows the auto-correlation function (ACF) for 4 of the linear combination parameters that are used to build the NNNA pdfs.

	Without Systematics	With Systematics	No Prior
neutrons	409(29)	398(30)	398(30)
alphas	1447(58)	1499(64)	1471(59)
nnna0	44(24)	46(31)	38(31)
nnna26	92(59)	36(37)	41(42)
neutron energy scale	fixed	1.005(03)	1.007(03)
alpha energy shift	fixed	0.010(27)	0.033(22)
neutron NoverA scale	fixed	1.082(23)	1.095(22)
alpha NoverA scale	fixed	1.044(10)	1.046(09)
Energy Fit Quality	33/24(9%)	30/24(20%)	28/24(25%)
NoverA Fit Quality	84/94(76%)	89/94(62%)	85/94(72%)

Table 8.3: Fit to one-third neutrino data using energy+NoverA.
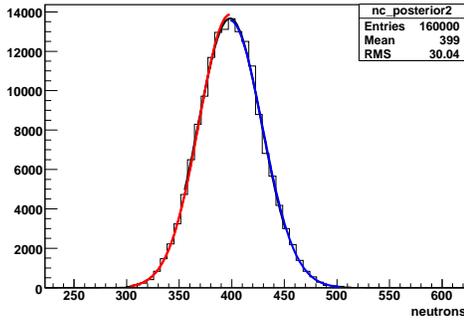
(a) Neutron posterior $398.3^{+30.7}_{-29.4}$
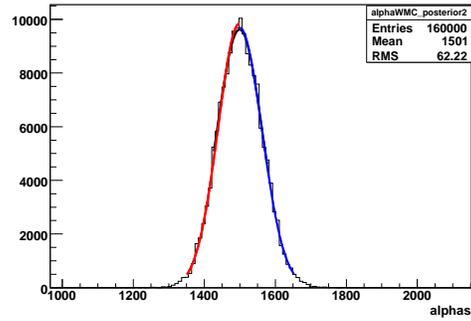
(b) Alpha posterior $1499.3^{+63.0}_{-60.6}$
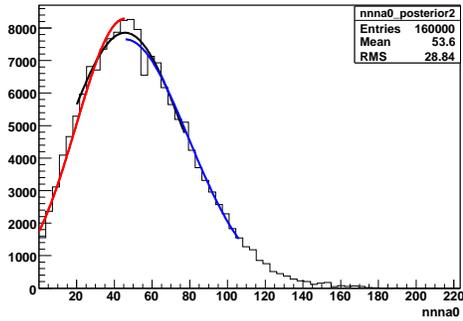
(c) nnna0 posterior $45.7^{+33.5}_{-25.8}$
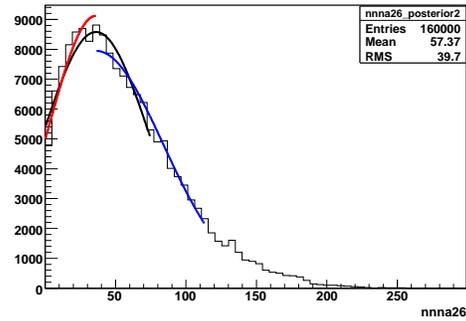
(d) nnna26 posterior $36.1^{+47.7}_{-32.8}$

Figure 8.3: Posteriors for the signals as determined by fitting one third of the NCD data set.
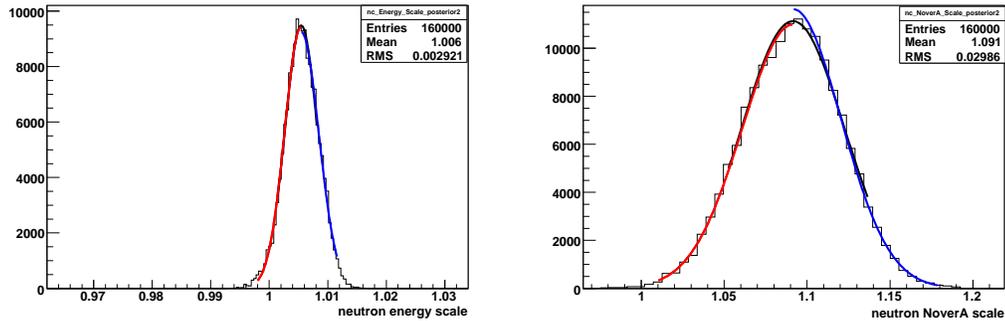
(a) Neutron energy scale posterior $1.0055^{+0.0030}_{-0.0028}$

(b) Neutron NoverA scale posterior $1.0915^{+0.0290}_{-0.0307}$
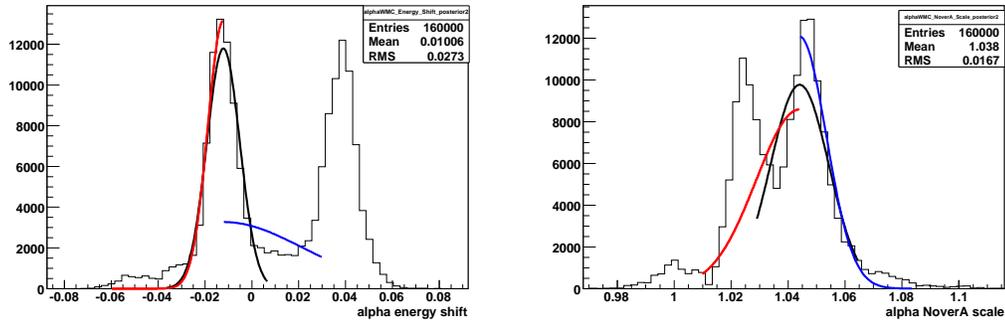
(c) Alpha energy shift posterior - mean+rms

(d) Alpha NoverA scale posterior $1.0441^{+0.0094}_{-0.0153}$

Figure 8.4: Posterior distributions for the systematic parameters for the fit of the one third data using energy and NoverA. It is again noted that the alpha energy shift is not constrained very well, as was found in the blind data set. However, the MCMC algorithm properly marginalizes over this multi-modal posterior distribution.

8.3.4   Comparison with Published Result

The results from this thesis are now compared to $N_{NCD}^{\frac{1}{3}}$, the number of neutrons determined in section 8.2 for the case where the PMT data were excluded. In order to obtain statistical and systematic contributions to the uncertainties from the results in this thesis, a comparison was made between the number of neutrons extracted with and without systematic uncertainties. The difference in quadrature is identified with the systematic uncertainty. Table 8.4 shows a summary of the number of neutrons determined on the one-third data set using energy, NoverA and energy+NoverA with statistical and systematic contributions included. The numbers are compared with $N_{NCD}^{\frac{1}{3}}$. The systematic uncertainty in the the NoverA-alone number also includes a contribution to take into account the discrepancy between the number of neutrons that was obtained with and without the prior constraint (the uncertainty was increased by adding in quadrature half of the difference in neutron numbers between the case with and without priors).

Table 8.4 shows that the number of neutrons determined with the methods in this thesis are in good agreement with the number that was published. A comparison of the central values shows that the number of neutrons agrees within the (uncorrelated) systematic uncertainties of each method. The fit using only energy from this thesis is in excellent agreement with that performed for the published analysis and yielded comparable uncertainties. Furthermore, it is demonstrated that the use of energy+NoverA has resulted in a substantial decrease in the total uncertainty in the number of neutrons. The gain has been obtained due to the statistical accuracy of the method being presented in this thesis which has increased because of the inclusion of a second observable. The total uncertainty has gone down from 10.7% (using the energy analysis of the published data) to 7.5% using the energy+NoverA analysis on the one-third data set. The next section extrapolates the results from this thesis to the full data set and determines the impact on the solar neutrino flux determination.

Source	Neutrons	Stats	Sys.	Total
$N_{NCD}^{\frac{1}{3}}$	410	44	9	44 (10.7%)
Energy Alone	406	39	9	40 (9.9%))
NoverA Alone	387	48	23	53 (13.7%)
Energy+NoverA	398	29	8	30 (7.5%)

Table 8.4: Comparison of neutron numbers obtained in this thesis with an analysis similar to the one that was published except for the exclusion of the PMT data (first row).

8.4   Conversion Between Number of Neutrons and Neutrino Flux

This section introduces the inputs and the method by which the conversion between the solar neutrino flux and the number of neutrons detected in the NCDs can be achieved. These are then used to convert between the number of detected neutrons and the neutrino flux. The number of neutrons produced in SNO can be calculated by the SNO Monte Carlo simulation package [25] using a neutrino flux. The number of neutrons produced per day in SNO from solar neutrinos is determined to be 12.963 ± 0.065 neutrons/day [83] for a neutrino flux of $5.145 \times 10^{-6} \mathrm{s}^{-1} \mathrm{cm}^{-2}$. The number depends on the neutral-current (NC) cross-section for neutrinos on deuterium, the number of deuterium targets and is integrated over the live-time of the detector including effect from the eccentricity of Earth's orbit. This number can then be scaled according to different predictions for the neutrino flux, or used to determine the neutrino flux from a number of neutrons detected per day in SNO.

Neutrons from the NC reaction are produced uniformly throughout the heavy water volume. The probability for uniformly produced neutrons to capture in the NCD array was determined using the data from the distributed 2005 ^{24}Na calibration source and was found to be 0.211 ± 0.007 ([38, 62]). The live-fraction of the scopes along with the efficiency of the data-cleaning cuts (including the removal of the six bad strings) results in a $92.3 \pm 0.4\%$ [83] probability for captured neutrons to be detected within the full energy range ($0.2\,\mathrm{MeV}$ - $1\,\mathrm{MeV}$). The live-time of the NCD phase was determined to be 385.15 ± 0.14 days [84].

Contribution	Value	Uncertainty
Predicted Flux $[10^{-6}\text{s}^{-1}\text{cm}^{-2}]$	Φ	%
Neutrons produced per day	$\times \frac{12.963 \pm 0.065}{5.145}$	0.5%
Number of live days	$\times\ 385.17 \pm 0.14$	0.036%)
Capture Efficiency	$\times\ 0.211 \pm 0.007$	3.318%
Detection Efficiency	$\times\ 0.923 \pm 0.004$	0.433%
Multiplicative Factor (X)	$\times 188.9978 \pm 6.395$	3.383%
Detected background neutrons in 385 days (B)	$+\ 198.76 \pm 25.5$	12.8%
Total number of neutrons detected in NCDs	$\Phi\times\text{X}+\text{B}$	

Table 8.5: This table shows how to convert a neutrino flux, Φ, into a number of predicted neutrons detected in the NCD phase, in the full energy range, by using a multiplicative factor, X, and the number of background neutrons, B.

Finally, the number of detected neutrons will include those produced from the solar neutrinos as well as background neutrons from radioactivity and atmospheric neutrinos. The number of detected background neutrons (including capture and detection efficiencies) for the entire NCD phase was determined [83] to be 198.76 ± 25.5 (detected neutrons within the full energy range) . All these contributions are summarized in Table 8.5, which shows how a neutrino flux can be converted into a total number of neutrons detected in the array for the entire NCD phase.

In order to estimate the impact of the results from this thesis on the measurement of the neutrino flux, the neutron numbers from this thesis are scaled up by a factor of three. Neutron numbers for the entire NCD phase can then be converted into neutrino fluxes using the procedure outlined in Table 8.5. An 'expected' uncertainty in the total number of neutrons is introduced in order to estimate the increased statistical accuracy that would be obtained by performing the analysis on the entire data set. This is achieved by scaling the statistical uncertainty by $\sqrt{3}$ and the systematic uncertainty by 3, similarly to what was done to scale down the number from the published analysis. A 'conservative' uncertainty was obtained by assuming that the total uncertainty scales by a factor of 3 and is not affected by the increase in statistics.

Table 8.6 shows the neutrino fluxes inferred from the methods presented here, compared

with predictions from standard solar models (BP2000 [3], BS2005(OP) [4], BPS2008(GS and AGS)[48]), the previously published analysis and the analysis similar to the published analysis that excluded the PMT data, to which the results from this thesis are expected to be most similar. The BP2000 solar model is the solar model that was current when SNO's first published results came out, BS2005(OP) was the current solar model at the time that the NCD data were published (and contained new opacity models) and BPS2008 (GS and AGS) are the latest solar models that have been split into a high and low metallicity model, respectively. This division has allowed for the uncertainty in the neutrino flux to be reduced, since it does not contain a contribution from the knowledge of the metal abundance. The high metallicity model is consistent with helioseismological measurements, but inconsistent with recent abundance measurements. The low metallicity model includes the new abundances but disagrees with the solar acoustic measurements. It is clear that SNO's measured neutrino flux is approaching an accuracy that can allow to distinguish between these different solar models and contribute to the resolution of the 'solar metallicity problem'.

As previously observed, the main benefit from this analysis is the substantial improvement in the statistical accuracy of the neutrino flux. The systematic uncertainties are not affected very much as they are dominated by the knowledge of the neutron background and the multiplicative factor. With the expected scaling of the statistical uncertainty, the results presented in this thesis would reduce the total uncertainty in the ^{8}B neutrino flux from 8.5% to 6.8% . It is difficult to make a precise statement about the central value of the neutrino flux obtained from the analysis in this thesis, due to the statistical inaccuracy.

Source	Neutrino Flux $10^{-6}\mathrm{s}^{-1}\mathrm{cm}^{-2}$	(stat.)	(sys.)	(total)
BP2000	5.05			0.91(18%)
BS2005(OP)	5.69			0.91(16%)
BPS2008(GS) (high metallicity)	5.94			0.65(11%)
BPS2008(AGS)(low metallicity)	4.72			0.52(11%)
Published	5.54	0.32	0.35	0.47(8.5%)
Publication Analysis, no PMTs	5.45	0.39	0.22	0.45(8.3%)
Energy (conservative uncertainty)	5.39	0.62	0.25	0.66(12.2%)
Energy (expected uncertainty)	5.39	0.36	0.25	0.42(7.8%)
Energy+NoverA (conservative uncertainty)	5.27	0.46	0.24	0.50(9.5%)
Energy+NoverA (expected uncertainty)	5.27	0.27	0.24	0.36(6.8%)

Table 8.6: . Comparison of neutrino fluxes from this analysis with the previously published analysis and the SSM predictions. The number from this thesis have been calculated with an expected uncertainty that assumes that the statistical accuracy would improve by a factor of $\sqrt{3}$. Results from this thesis should be compared to the value obtained from the analysis that is similar to the published analysis but excludes the PMTs (fourth line). In all cases, the systematic uncertainties are dominated by the contributions from the background neutrons and the multiplicative factor.

8.5 Summary

This chapter presented an analysis of one third of the data from the NCD phase. It was shown that the combination of energy and NoverA in a maximum-likelihood analysis resulted in a smaller uncertainty in the number of detected neutrons compared to the previous publication of this data. The uncertainty on the neutrino flux from this new analysis is expected to be 6.8% compared to the 8.5% uncertainty on the previously published result.

Chapter 9

Conclusion

The aim of this thesis was to provide an independent analysis of the data from the pro-portional counters from the third phase of SNO. This analysis differed significantly from previous work as pulse shape information was used to distinguish neutrons from alphas and instrumental backgrounds (NNNA events, 'non-neutron-non-alpha'). The work in this the-sis built on the previous use of shaper energy in a maximum-likelihood analysis to include a pulse-shape observable in the likelihood.

After introducing the NCD array, electronics and calibration (chapter 2), this thesis covered the physics of pulses in proportional counters to motivate the use of pulse shapes to distinguish neutrons from other events (chapter 3).

The Queen's Grid Fitter (QGF) was then characterized as a means of measuring the 'NoverA' observable for each event (chapter 4). The probability density functions of NoverA were then obtained for neutrons, alphas, and NNNAs in chapter 5. Systematic uncertainties in these distributions were handled in different ways, depending on the amount of data available for determining the pdfs. In the case where many data were available (such as for neutrons and alphas), average pdfs were determined and then deformed with either a scaling or shifting of the x-axis. The scaling and shifting parameters were allowed to float in the fits, but were constrained by calibration data for neutrons and a combination of calibration

and Monte Carlo data for alpha particles.

The pdfs for the NNNA type of events had large uncertainties in each bin due to the small amounts of data that were used to generate them. The work in this thesis presented a novel method for dealing with the systematic uncertainties arising from these 'low statistics' pdfs. This was achieved by defining the pdfs as histograms with uncertainties in each bin. A set of 10 'new' pdfs were then drawn from the histograms with uncertainties by randomly sampling a value in each bin and creating a new pdf. In order to introduce bin to bin correlations (as would be expected for a physical process), these pdfs were smeared with a gaussian resolution function. The resulting 10 pdfs were then used in a linear combination and the coefficients of that linear combination were floated as free parameters. This novel technique allowed for the integration over possible pdf shapes as well as the determination of a best fit pdf from the data. The problem of pdfs determined from low statistics can often be a problem in particle physics when Monte Carlo data is used to model rare events. In those cases, even large amounts of Monte Carlo data may not have the desired statistics to model the distributions for rare processes and the techniques presented here may be of use.

In chapter 6, the details of an extended maximum-likelihood analysis were introduced. It was shown that a Bayesian approach using the Markov-Chain Monte Carlo (MCMC) technique was suitable for performing an analysis of the NCD data using energy and NoverA as observables. Test of the MCMC-based analysis were then performed on sets of Monte Carlo data to test for any biases (chapter 7). The tests were then extended to a blind data set and showed that the new technique, which integrated NoverA into the likelihood, gives consistent results compared to an analysis that used only energy but results in a smaller uncertainty in the number of detected neutrons.

In chapter 8, the extended maximum-likelihood analysis was applied to one third of the data from the proportional counters in the NCD phase. This analysis used both energy and NoverA as observables to determine the number of detected neutrons in the NCDs. The

analysis presented here determined that the one-third data set contains 398 ± 29 (stat.) $\pm$ 9 (sys.) neutrons. This can be directly compared to an analysis that is analogous to the published result but which excluded data from the PMTs and determined that the one-third data set contains 410 ± 44 (stat.) ± 9 (sys.) neutrons, where systematic uncertainties are uncorrelated between the two analyses. The addition of the NoverA observable has resulted in a substantial decrease in the statistical uncertainty in the number of neutrons and is the main result from this thesis. The work presented here has thus conclusively shown that the use of the pulse shapes from the proportional counters results in a more precise measurement of the number of detected neutrons in the NCD phase of the SNO experiment.

Assuming that the statistical uncertainty scales by a factor of $\sqrt{3}$ the analysis from this thesis alone results in a 6.8% uncertainty in the total ^{8}B solar neutrino flux, compared with 8.5% uncertainty quoted in the previously published result [38]. This result will have an impact on further constraining solar models and in particular, can help solve the 'solar metallicity problem', where the ^{8}B neutrino flux differs by 21% between high and low metallicity solar models. Increased accuracy in the NC flux of the NCD phase will lead to a better determination of the CC flux and the CC/NC ratio, which will result in a more accurate value of θ_{12}. A better knowledge of θ_{12} will lead to a more robust determination of θ_{13}, which will set the scale for CP-violation in the lepton sector. Phenomenological models that predict the values of the mixing angles will also be further constrained. The inclusion of the analysis presented here to a combined analysis of all three phases of SNO will lead to even more accurate results and, ultimately, the most precise value of the total active neutrino flux and θ_{12} ever determined.

Finally, the last chapter of this thesis provided some conclusions and a brief summary, much like this one.

Bibliography

[1] J. N. Bahcall. "Solar Neutrinos I. Theoretical." *Phys. Rev. Lett.*, **12** (11), 300–302, 1964. 3

[2] R. J. Davis. "Solar Neutrinos II. Experimental." *Phys. Rev. Lett.*, **12**, 303–305, 1964. 3

[3] J. N. Bahcall, M. H. Pinsonneault, and S. Basu. "Solar models: Current epoch and time dependences, neutrinos, and helioseismological properties." *Astrophys. J.*, **555**, 990, 2001, `astro-ph/0010346`. 3, 17, 33, 211

[4] J. N. Bahcall, A. M. Serenelli, and S. Basu. "New solar opacities, abundances, helioseismology, and neutrino fluxes." *Astrophys. J.*, **621**, L85, 2005, `astro-ph/0412440`. 3, 4, 5, 33, 211

[5] R. J. Davis, D. S. Harmer, and K. C. Hoffman. "Search for neutrinos from the sun." *Phys. Rev. Lett.*, **20**, 1205–1209, 1968. 3, 5

[6] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. "Detection of the free neutrino: A Confirmation." *Science*, **124**, 103–104. doi:10.1126/science.124.3212.103, 1956. 3

[7] H. Bethe. "Energy Production in Stars." *Phys. Rev.*, **55** (11), 103–103, 1939. 4

[8] A. Wright. "Ph.D. Thesis, Queen's University.", 2009. 4

[9] B. T. Cleveland et al. "Measurement of the solar electron neutrino flux with the Homestake chlorine detector." *Astrophys. J.*, **496**, 505–526. doi:10.1086/305343, 1998. 5

[10] Y. Fukuda, T. Hayakawa, K. Inoue, K. Ishihara, H. Ishino, S. Joukou, T. Kajita, S. Kasuga, Y. Koshio, T. Kumita, K. Matsumoto, M. Nakahata, K. Nakamura, K. Okumura, A. Sakai, M. Shiozawa, J. Suzuki, Y. Suzuki, T. Tomoeda, Y. Totsuka, K. S. Hirata, K. Kihara, Y. Oyama, M. Koshiba, K. Nishijima, T. Horiuchi, and K. Fujita. "Solar Neutrino Data Covering Solar Cycle 22." *Phys. Rev. Lett.*, **77** (9), 1683–1686. doi:10.1103/PhysRevLett.77.1683, 1996. 5

[11] J. P. Cravens et al. (Super-Kamiokande). "Solar neutrino measurements in Super-Kamiokande-II." *Phys. Rev.*, **D78**, 032002. doi:10.1103/PhysRevD.78.032002, 2008, 0803.4312. 5, 13

[12] J. N. Abdurashitov, T. J. Bowles, M. L. Cherry, B. T. Cleveland, R. Davis, S. R. Elliott, V. N. Gavrin, S. V. Girin, V. V. Gorbachev, T. V. Ibragimova, A. V. Kalikhov, N. G. Khairnasov, T. V. Knodel, K. Lande, I. N. Mirmov, J. S. Nico, A. A. Shikhin, W. A. Teasdale, E. P. Veretenkin, V. M. Vermul, D. L. Wark, P. S. Wildenhain, J. F. Wilkerson, V. E. Yants, and G. T. Zatsepin. "Measurement of the Solar Neutrino Capture Rate by SAGE and Implications for Neutrino Oscillations in Vacuum." *Phys. Rev. Lett.*, **83** (23), 4686–4689. doi:10.1103/PhysRevLett.83.4686, 1999. 6

[13] W. Hampel et al. (GALLEX). "GALLEX solar neutrino observations: Results for GALLEX IV." *Phys. Lett.*, **B447**, 127–133. doi:10.1016/S0370-2693(98)01579-2, 1999. 6

[14] Q. R. Ahmad et al. (The SNO Collaboration). "Measurement of the charged current interactions produced by B-8 solar neutrinos at the Sudbury Neutrino Observatory." *Phys. Rev. Lett.*, **87**, 071301. doi:10.1103/PhysRevLett.87.071301, 2001,

`nucl-ex/0106015`. 6, 14, 15, 18

[15] J. N. Bahcall. "Solar models and solar neutrinos: Current status." *Phys. Scripta*, **T121**, 46–50. doi:10.1088/0031-8949/2005/T121/006, 2005, `hep-ph/0412068`. 7, 33

[16] The Super-Kamiokande Collaboration. "Measurement of the Flux and Zenith-Angle Distribution of Upward Through-Going Muons in Kamiokande II+III." *Phys. Rev. Lett.*, **81**, 2016, 1998. 6

[17] The Super-Kamiokande Collaboration. "Evidence for an Oscillatory Signature in Atmospheric Neutrino Oscillations." *Phys. Rev. Lett.*, **93**, 101801, 2004. 6

[18] B. Pontecorvo. "Inverse beta processes and nonconservation of lepton charge." *Sov. Phys. JETP*, **7**, 172, 1958. 6

[19] R. Martin. "Numerical methods and approximations for the calculation of solar neutrino fluxes in three flavours focused on results for the Sudbury Neutrino Observatory." UMI-MR-15302. 8, 9

[20] C. S. Lim, K. Ogure, and H. Tsujimoto. "How precisely can we reduce the three-flavor neutrino oscillation to the two-flavor one only from delta(m(12)**2)/delta(m(13)**2) ¡ approx. 1/15?" *Phys. Rev.*, **D67**, 033007, 2003, `hep-ph/0210066`. 8, 9

[21] T. K. Kuo and J. T. Pantaleone. "THREE NEUTRINO OSCILLATIONS AND THE SOLAR NEUTRINO EXPERIMENTS." *Phys. Rev.*, **D35**, 3432, 1987. 9

[22] G. L. Fogli, E. Lisi, A. Marrone, A. Palazzo, and A. M. Rotunno. "Hints of $\theta_{13} > 0$ from global neutrino data analysis." *Phys. Rev. Lett.*, **101**, 141801. doi:10.1103/PhysRevLett.101.141801, 2008, `0806.2649`. 9, 33, 34

[23] L. Wolfenstein. "Neutrino oscillations in matter." *Phys. Rev.*, **D17**, 2369, 1978. 10

[24] S. Mikheyev and A. Y. Yu. Smirnov. *Yad. Fiz.*, **42**, 1441, 1985. 10

[25] J. L. Boger et al. (The SNO Collaboration). "The Sudbury Neutrino Observatory." *Nuclear Instruments and Methods in Physics Research*, **A449**, 172, 2000. 10, 48, 55, 63, 209

[26] H. H. Chen. "Direct approach to resolve the solar-neutrino problem." *Phys. Rev. Lett.*, **55**, 1534–1536. doi:10.1103/PhysRevLett.55.1534, 1985. 10

[27] Q. R. Ahmad et al. (The SNO Collaboration). "Direct Evidence for Neutrino Flavor Tranformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory." *Phys. Rev. Lett.*, **89**, 011301, 2002. 14, 17

[28] Q. R. Ahmad et al. (The SNO Collaboration). "Measurement of day and night neutrino energy spectra at SNO and constraints on neutrino mixing parameters." *Phys. Rev. Lett.*, **89**, 011302. doi:10.1103/PhysRevLett.89.011302, 2002, `nucl-ex/0204009`. 14, 19, 20

[29] B. Aharmim et al. (The SNO Collaboration). "Measurement of the nu/e and total B-8 solar neutrino fluxes with the Sudbury Neutrino Observatory phase I data set." *Phys. Rev.*, **C75**, 045502. doi:10.1103/PhysRevC.75.045502, 2007, `nucl-ex/0610020`. 14, 20

[30] B. Aharmim et al. (The SNO Collaboration). "A search for neutrinos from the solar hep reaction and the diffuse supernova neutrino background with the sudbury neutrino observatory." *Astrophys. J.*, **653**, 1545–1551. doi:10.1086/508768, 2006, `hep-ex/0607010`. 14, 30

[31] S. N. Ahmed et al. (The SNO Collaboration). "Measurement of the Total Active ^8B Solar Neutrino Flux at the Sudbury Neutrino Observatory with Enhanced Neutral Current Sensitivity." *Phys. Rev. Lett.*, **92**, 181301, 2004. 14, 21

[32] B. Aharmim et al. (The SNO Collaboration). "Electron energy spectra, fluxes, and day-night asymmetries of ^8B solar neutrinos from the 391-day salt phase of SNO." *Phys. Rev. C*, **72**, 055502, 2005. 14, 21, 24, 25, 56

[33] S. F. Mughabghab, M. Divadeenam, and N. E. Holden. *Neutron Cross Sections*. Academic Press, New York, 1981. 14, 40

[34] J. N. Bahcall, M. C. Gonzalez-Garcia, and C. Pena-Garay. "Before and after: How has the SNO neutral current measurement changed things?" *JHEP*, **07**, 054, 2002, `hep-ph/0204314`. 20

[35] C. Amsler et al. "Particle Data Group: Reviews, Tables and Plots." *Physics Letters*, **B667** (1), 2008. 20, 126

[36] S. Abe et al. (KamLAND). "Precision Measurement of Neutrino Oscillation Parameters with KamLAND." *Phys. Rev. Lett.*, **100**, 221803. doi:10.1103/PhysRevLett.100.221803, 2008, `0801.4589`. 23, 33

[37] J. F. Amsbaugh et al. "An array of low-background ^{3}He proportional counters for the Sudbury Neutrino Observatory." *Nucl. Instrum. Meth.*, **A579**, 1054–1080. doi:10.1016/j.nima.2007.05.321, 2007, `0705.3665`. 23, 40, 44, 46, 48, 49, 55

[38] B. Aharmim et al. (The SNO Collaboration). "An 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. doi:10.1103/PhysRevLett.101.111301, 2008, `0806.0989`. 23, 57, 61, 69, 72, 186, 196, 200, 202, 209, 215

[39] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. "Equation of State Calculations by Fast Computing Machines." *The Journal of Chemical Physics*, **21** (6), 1087–1092. doi:10.1063/1.1699114, 1953, URL `http://link.aip.org/link/?JCP/21/1087/1`. 26, 141

[40] W. K. Hastings. "Monte Carlo sampling methods using Markov chains and their applications." *Biometrika*, **57** (1), 1970. 26, 141

[41] B. Aharmim et al. (The SNO Collaboration). "Electron antineutrino search at the Sudbury Neutrino Observatory." *Phys. Rev.*, **D70**, 093014. doi:10.1103/PhysRevD.70.093014, 2004, `hep-ex/0407029`. 30

[42] S. N. Ahmed et al. (The SNO Collaboration). "Constraints on nucleon decay via 'invisible' modes from the Sudbury Neutrino Observatory." *Phys. Rev. Lett.*, **92**, 102004. doi:10.1103/PhysRevLett.92.102004, 2004, `hep-ex/0310030`. 31

[43] P. A. Sturrock. "Analysis of Super-Kamiokande 5 Day Measurements of the Solar Neutrino Flux." *The Astrophysical Journal*, **605** (1), 568–572, 2004, URL `http://stacks.iop.org/0004-637X/605/568`. 31

[44] "The simplest resonant spin-flavour solution to the solar neutrino problem." *Nuclear Physics B*, **595** (1-2), 360 – 380. ISSN 0550-3213. doi:DOI:10.1016/S0550-3213(00)00546-0, 2001. 31

[45] J. Yoo et al. (Super-Kamiokande). "A search for periodic modulations of the solar neutrino flux in Super-Kamiokande-I." *Phys. Rev.*, **D68**, 092002. doi:10.1103/PhysRevD.68.092002, 2003, `hep-ex/0307070`. 31

[46] B. Aharmim et al. (SNO). "A search for periodicities in the B-8 solar neutrino flux measured by the Sudbury Neutrino Observatory." *Phys. Rev.*, **D72**, 052010. doi:10.1103/PhysRevD.72.052010, 2005, `hep-ex/0507079`. 31

[47] B. Aharmim et al. "Measurement of the Cosmic Ray and Neutrino-Induced Muon Flux at the Sudbury Neutrino Observatory.", 2009, `0902.2776`. 31, 32

[48] C. Pena-Garay and A. Serenelli. "Solar neutrinos and the solar composition problem.", 2008, `0811.2424`. 32, 33, 211

[49] The Borexino Collaboration. "Measurement of the solar 8B neutrino flux with 246 live days of Borexino and observation of the MSW vacuum-matter transition.", 2008, `0808.2868`. 34, 36

[50] V. Barger, P. Huber, and D. Marfatia. "Solar mass-varying neutrino oscillations." *Phys. Rev. Lett.*, **95**, 211802. doi:10.1103/PhysRevLett.95.211802, 2005, `hep-ph/0502196`. 34

[51] A. Friedland, C. Lunardini, and C. Pena-Garay. "Solar neutrinos as probes of neutrino - matter interactions." *Phys. Lett.*, **B594**, 347. doi:10.1016/j.physletb.2004.05.047, 2004, `hep-ph/0402266`. 34

[52] D. D'Angelo. "Results from the Borexino experiment after 192 days of data-taking." Talk given at Miami 2008, Fort Lauderdale, December 2008. URL `http://server.physics.miami.edu/~cgc/Miami2008/DAngelo2008.ppt`. 35

[53] A. Y. Smirnov. "Neutrino-2008: Where are we? Where are we going?" *J. Phys. Conf. Ser.*, **136**, 012002. doi:10.1088/1742-6596/136/1/012002, 2008, `0810.2668`. 35

[54] M. Chen. "The SNO Liquid Scintillator Project." *Nuclear Physics B - Proceedings Supplements*, **145**, 65–68. ISSN 0920-5632. doi:DOI:10.1016/j.nuclphysbps.2005.03.037. NOW 2004, 2005. 36

[55] G. A. Cox et al. "Sudbury Neutrino Observatory neutral current detectors signal readout system." *IEEE Trans. Nucl. Sci.*, **51**, 2227–2230. doi:10.1109/TNS.2004.836148, 2004. 48

[56] A. G. Cox and M. H. Wehrenberg. *Characterization of the transformations of NCD events signals.* Internal SNO Document, 2004. 50, 51

[57] R. G. H. Robertson and T. Van Wechel. *Characteristics of AD8307 logarithmic amplifier.* Internal SNO Document SNO-STR-05-003, 2005. 51

[58] CERN Program Library Long Writeup Q100/101, CERN, Geneva, Switzerland. 53

[59] K. Boudjemline et al. "The calibration of the Sudbury Neutrino Observatory using uniformly distributed radioactive sources." 56, 57, 80

[60] L. C. Stonehill. "Deployment and background characterization of the Sudbury Neutrino Observatory Neutral Current Detectors." UMI-31-98856. 57

[61] B. A. VanDevender, A. Cox-Morbrand, and W. J. Internal SNO document, ECA calibration webpage. 62

[62] G. Guillian. "The Determination of the Neutral Current NCD Capture Efficiency with ^{24}Na Data." Internal SNO Document, 2008. 61, 209

[63] J. F. Ziegler. "SRIM 2003." Software Package, www.srim.org, 2003. 66, 72

[64] J. F. Ziegler, J. P. Biersack, and U. Littmark. *The Stopping Range of Ions in Matter.* Pergamon Press, 1985. 66

[65] A. Hime. "Detector operating parameters and constraints." Internal SNO document. 67

[66] D. H. Wilkinson. *Ionization Chambers and Counters.* Cambridge University Press, 1950. 68

[67] B. Beltran, J. Monroe, N. S. Oblath, G. Prior, K. Rielage, R. G. H. Robertson, and H. S. Wan Chan Tseung. "NCD MC Readiness Report for the First NCD Paper." Internal SNO document, 2008. 70, 72, 73, 74, 75, 108, 109, 198

[68] B. Beltran et al. "A Monte Carlo simulation of the Sudbury Neutrino Observatory proportional counters." *Nuclear Instruments and Methods in Physics Research*. In preparation-Internal SNO document for now, 2008. 72

[69] R. Veenhof. "Garfield Software Package." Www.cern.ch/garfield, 2007. 72

[70] N. Oblath. "Ph.D Thesis, University of Washington.", 2009. 93, 116

[71] R. Barlow and C. Beeston. "Fitting using finite Monte Carlo samples." *Computer Physics Communications*, **77**, 219–228, 1993. 119

[72] B. Efron. "Bootstrap Methods: Another Look as the Jackknife." *Ann. Statist.*, **7** (1), 1–26, 1979. 119

[73] S. Oser. *Why floating systematics is mandatory.* Internal SNO document, MAN-6VF2FA. 125

[74] D. S. Sivia and J. Skilling. *Data Analysis: A Bayesian Tutorial.* Oxford University Press, 2006. 126, 136

[75] P. C. Gregory. *Bayesian Logical Analysis for the Physical Sciences.* Cambridge University Press, 2005. 126, 140, 142, 146

[76] P. Davis and D. H. Kenyon. *Of Pandas and People: The Central Question of Biological Origins.* Foundation for Thought and Ethics, 1989. 136

[77] F. James. "MINUIT - Function Minimization and Error Analysis - Reference Manual.", 2000, URL `http://wwwasdoc.web.cern.ch/wwwasdoc/minuit/minmain.html`. 139

[78] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical recipes in C (3rd ed.): the art of scientific computing.* Cambridge University Press. ISBN 0521880688, 2007. 140

[79] G. O. Roberts, A. Gelman, and W. R. Gilks. "Weak Convergence and Optimal Scaling of Random Walk Metroplolis Algorithms." *Ann. Appl. Probab.*, 110–120, 1997. 147

[80] R. Brun and F. Rademakers. "ROOT - An Object Oriented Data Analysis Framework." *Nuclear Instruments and Methods in Physics Research*, **A389**, 81–86. Proceedings AIHENP'96 Workshop, Lausanne, Sep. 1996. See also http://root.cern.ch/, 1997. 150

[81] B. Jamieson. "SNO NCD Phase Signal Extraction on Unblinded Data with Integration over Systematic Nuisance Parameters by Markov-Chain Monte-Carlo - Version 1.1." Internal SNO document, 2008. 196

[82] B. Jamieson (For The SNO Collaboration). "Measurement of the Solar Neutrino Flux with an Array of Neutron Detectors in the Sudbury Neutrino Observatory.", 2008, 0810.3760. 196

[83] The SNO Collaboration. "SNO NCD Phase Analysis Unidoc." Internal SNO Document, 2008. 197, 209, 210

[84] J. Detwiler, N. McCauley, K. Rielage, G. Tesic, and F. Zhang. "Livetime for the Full NCD Data Set." Internal SNO Document, 2008. 209