

©Copyright 2005

Laura C. Stonehill

Deployment and Background Characterization of the
Sudbury Neutrino Observatory Neutral Current Detectors

Laura C. Stonehill

A dissertation submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Washington

2005

Program Authorized to Offer Degree: Physics

University of Washington
Graduate School

This is to certify that I have examined this copy of a doctoral dissertation by

Laura C. Stonehill

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

Chair of the Supervisory Committee:

John F. Wilkerson

Reading Committee:

John F. Wilkerson

R. G. Hamish Robertson

Stephen D. Ellis

Date: _____

In presenting this dissertation in partial fulfillment of the requirements for the doctoral degree at the University of Washington, I agree that the Library shall make its copies freely available for inspection. I further agree that extensive copying of this dissertation is allowable only for scholarly purposes, consistent with "fair use" as prescribed in the U.S. Copyright Law. Requests for copying or reproduction of this dissertation may be referred to Proquest Information and Learning, 300 North Zeeb Road, Ann Arbor, MI 48106-1346, 1-800-521-0600, to whom the author has granted "the right to reproduce and sell (a) copies of the manuscript in microform and/or (b) printed copies of the manuscript made from microform."

Signature_____

Date_____

University of Washington

Abstract

Deployment and Background Characterization of the
Sudbury Neutrino Observatory Neutral Current Detectors

Laura C. Stonehill

Chair of the Supervisory Committee:
Professor John F. Wilkerson
Physics

The Sudbury Neutrino Observatory (SNO) is a heavy water Cherenkov detector designed to observe ^8B neutrinos from the sun via elastic-scattering interactions with electrons and charged-current and neutral-current interactions with deuterium. SNO is currently taking data in its third phase, in which the detector's capabilities have been enhanced by the addition of the Neutral Current Detectors (NCDs), an array of ^3He proportional counters that detect neutrons liberated in neutral-current interactions. The NCDs were designed and constructed to minimize radioactive impurities that can produce backgrounds. A variety of innovative materials and methods were used in the construction and deployment of the NCD array in order to meet the stringent purity requirements.

It is important to quantify alpha backgrounds in the NCDs and to verify that the impurity levels are within design specifications. Analysis techniques relying on pulse shape characteristics, time coincidences, and energy spectra were used to identify and measure bulk ^{238}U and ^{232}Th chain decays and surface ^{210}Po decays in the nickel NCD bodies. The measured bulk impurities in the NCD array are equivalent to $5.7_{-0.9}^{+1.0}$

picograms ^{232}Th per gram of nickel and $2.8_{-0.8}^{+0.6}$ picograms ^{238}U per gram of nickel. In addition to the 43_{-7}^{+6} alphas per day in the NCD array from these sources, ^{210}Po surface contamination produces 127 ± 2 alphas per day, and another 81_{-2}^{+5} alphas per day are produced by activity in the upper portions of the disequilibrium ^{238}U and ^{232}Th chains. The alpha background studies presented here indicate that the alpha rates in the NCD array are in line with expectations, and will not pose a significant problem to the successful extraction of the neutral-current signal from the NCDs using pulse shape analysis techniques.

TABLE OF CONTENTS

List of Figures	v
List of Tables	ix
Glossary	xi
Chapter 1: Introduction to Solar Neutrinos	1
1.1 Introduction to Neutrinos	1
1.2 The Weak Interaction	3
1.3 The Standard Solar Model	5
1.4 Electron Capture in the CNO Cycle	8
1.5 Overview of Solar Neutrino Experiments	12
1.6 Neutrino Oscillations	16
Chapter 2: The SNO Neutral Current Detectors	18
2.1 The Sudbury Neutrino Observatory	18
2.2 Motivation for the NCD Phase	27
2.3 NCD Signals	31
2.4 NCD Data Acquisition	36

Chapter 3:	NCD Construction and Deployment	39
3.1	NCD Fabrication	39
3.2	NCD Storage and Handling	46
3.3	Pre-deployment Welding	48
3.4	Deployment Hardware	58
3.5	Deployment Procedures	64
3.6	Difficulties during Deployment	71
Chapter 4:	NCD Characterization and Optimization	79
4.1	Cool-down Data	79
4.2	NCD Selection	82
4.3	Deployment Source Tests	85
4.4	Quality Control and Repairs	90
4.5	Pathological Strings	92
4.6	MUX Threshold Optimization	95
Chapter 5:	NCD Pulse Shape Analysis Techniques	107
5.1	Pulse Shape	107
5.2	Pulse Width	114
5.3	Proof-of-principle Pulse Width Analysis	116
5.4	Other Pulse Shape Analysis Techniques	119

Chapter 6:	Alpha Modeling	123
6.1	Alpha Sources	123
6.2	Alpha Energy Simulation	126
6.3	^{210}Po Alpha Simulation	135
6.4	Pulse Width Simulation	139
6.5	Time Coincidence Simulation	142
Chapter 7:	Alpha Background Measurements	146
7.1	Datasets	146
7.2	Width vs. Energy	151
7.3	Energy Spectrum Fits	156
7.4	Time Coincidence	162
7.5	Combined Alpha Background Results	172
Chapter 8:	Implications of Alpha Backgrounds	177
8.1	Converting Alpha Rate to Impurity Level	177
8.2	Comparison to Previous Measurements and Detectors	179
8.3	Statistical Estimate of Alpha Sensitivity in the NCD Phase	181
8.4	Alpha Projection into the Neutron Energy Region	183
8.5	Statistical Alpha Subtraction	186
Chapter 9:	Conclusions	191

Bibliography	200
Appendix A: Data Cleaning	201
Appendix B: String-by-string Alpha Rates	207

LIST OF FIGURES

Figure Number	Page
1.1 The pp chain of nuclear reactions	6
1.2 Standard Solar Model neutrino spectrum	8
1.3 The CNO cycle of nuclear reactions	9
1.4 Solar neutrino spectrum, including the CNO cycle	13
2.1 Artist's drawing of the SNO detector	19
2.2 W-mediated elastic scattering Feynman diagrams	20
2.3 Z-mediated elastic scattering Feynman diagram	21
2.4 Charged current Feynman diagram	22
2.5 Neutral current Feynman diagram	23
2.6 Flux results from the SNO experiment	25
2.7 MSW oscillation parameters	26
2.8 Artist's drawing of SNO with the NCD array	28
2.9 Expected MSW region after NCD phase	30
2.10 Diagram of an NCD string	32
2.11 Idealized neutron spectrum from a ^3He proportional counter	34
2.12 Neutron source data from the NCD array	35

3.1	NCD endcaps	43
3.2	NCD cable bell	49
3.3	Weld fixture	51
3.4	Pre-deployment welds	52
3.5	NCD delay line	54
3.6	The AV with the NCD deployment equipment	59
3.7	The deployment plate	61
3.8	The Remotely Operated Vehicle	63
3.9	Anchor point	69
3.10	The arrangement of the NCD array	72
3.11	The ROV in the neck of the AV	78
4.1	Centroids and widths for each NCD from deployment calibrations	89
4.2	Threshold cut fraction for string 14	97
4.3	MUX threshold amplitude for each string	98
4.4	Rates on each string with different MUX thresholds	101
4.5	Threshold cut fraction for each string	103
4.6	Reoptimized threshold cut fraction for each string	105
5.1	The Bragg curve	108
5.2	Stopping power for a neutron-capture event	109
5.3	Radius as a function of drift time	110

5.4	Sample neutron pulse shapes	111
5.5	Simulated ion tails	113
5.6	Simulated alpha and neutron capture width vs. energy	115
5.7	Pulse width vs. energy plot for neutrino data	117
5.8	Pulse shapes of various NCD events	118
5.9	Simulated proton-triton track	121
6.1	The uranium and thorium decay chains	124
6.2	Simulated NCD wall geometry	129
6.3	Simulated ^{238}U and ^{232}Th alpha wall energies	130
6.4	Energy loss curve for an alpha in nickel	131
6.5	Simulated NCD internal geometry	132
6.6	Simulated ^{238}U and ^{232}Th alpha energy deposition	134
6.7	Effect of space charge on neutron and ^{210}Po simulations	137
6.8	Effect of detector resolution on ^{210}Po simulation	138
6.9	Point of closest approach to the NCD wire	140
6.10	Simulated pulse width vs. energy plot	141
7.1	Impact of data cleaning on energy spectra	149
7.2	Width vs. energy plot of alpha data	151
7.3	Results from width vs. energy analysis	154
7.4	Fits to commissioning and open data	157

7.5	Fit residuals for commissioning and open data	158
7.6	Energy distribution of triples	163
7.7	Distribution of triples and 0.3-second doubles in the NCD array . . .	164
7.8	Energy distribution of 0.3-second doubles	165
7.9	Distribution of doubles in the NCD array	166
7.10	Energy distribution of doubles with fits	168
7.11	Residuals from doubles energy spectral fits	170
7.12	Comparison of all analysis results	174
8.1	Data and simulated spectra with neutrons	187
B.1	^{210}Po rates on each NCD string	208
B.2	Lower uranium chain rates on each NCD string	209
B.3	Lower thorium chain rates on each NCD string	210
B.4	Upper uranium and thorium chain rates on each NCD string	211

LIST OF TABLES

Table Number		Page
1.1	Continuum electron capture correction for CNO cycle	11
1.2	Neutrino fluxes from CNO electron capture	12
1.3	Solar neutrino flux in different detectors	15
2.1	Uncertainties in each phase of SNO	29
3.1	Problems during NCD deployment	74
4.1	NCDs and alpha rates in the optimized array	86
4.2	NCDs and alpha rates in the final array	93
4.3	Discharging NCDs	94
4.4	Low-rate discharging NCDs	95
4.5	Thresholds used in MUX rate tests	100
4.6	“Optimized” and “Reoptimized” MUX thresholds	106
6.1	Alpha energies and branching ratios	128
7.1	Comparison of livetime methods	148
7.2	Impact of non-analyzed strings	150

7.3	Region counts for alpha determination	152
7.4	Results from energy spectrum fits	159
7.5	Results and systematic uncertainties	175
8.1	Conversion to ppt and $\mu\text{Bq/kg}$	178
8.2	Statistical uncertainties from NCD phase	183
B.1	Alpha rates on each string in the NCD array	212

GLOSSARY

PMT: Photomultiplier Tube, a light detector capable of single-photon sensitivity.

PP CHAIN: The proton-proton fusion chain, the dominant hydrogen-burning mechanism in the sun.

CNO CYCLE: The Carbon-Nitrogen-Oxygen cycle of nuclear reactions in stars. It is a sub-dominant hydrogen-burning mechanism in the sun.

SSM: The Standard Solar Model that provides theoretical predictions of solar neutrino fluxes and spectra.

SNP: The Solar Neutrino Problem that arose because solar neutrino experiments detected fewer neutrinos than predicted by solar models.

MSW EFFECT: The Mikheyev-Smirnov-Wolfenstein effect, a density-dependent resonant enhancement of neutrino flavor-change probability.

SNO: The Sudbury Neutrino Observatory, a heavy-water Cherenkov experiment for detection of solar neutrinos.

AV: The Acrylic Vessel that contains the kiloton of heavy water used as a solar neutrino target in SNO.

ES: Elastic-Scattering interaction of neutrinos with electrons, one of the three neutrino interactions in SNO.

CC: Charged-Current interactions of electron neutrinos with deuterium.

NC: Neutral-Current interactions of neutrinos with deuterium.

NCD: Neutral Current Detector, a ^3He proportional counter used in the SNO experiment to detect thermalized neutrons.

PDF: Probability Density Function.

DAQ: Data Acquisition system.

MUX: A Multiplexer used to combine signals from up to 12 NCD channels prior to digitization.

ADC: Analog-to-Digital Converter.

ORCA: The Object-oriented Real-time Control and Acquisition software used to acquire data from the NCDs.

SHARC: The SNO Hardware Acquisition and Real-time Control software used to acquire data from the SNO PMT array.

CVD: Chemical Vapor Deposition, the process used to purify the nickel used in the NCD bodies.

FWHM: The Full Width at Half the Maximum amplitude, a measure of the resolution of a peak.

TDR: Time Domain Reflectometry, a technique used to verify continuity and determine the location of impedance changes in a electrical line.

DCR: The Deck (or Dark) Clean Room above the neck of the SNO AV that provides access to the SNO detector. This space was used to deploy the NCDs.

ROV: The Remotely Operated Vehicle, a remote-controlled submersible used to deploy the NCD array into the SNO detector.

DAC: Digital-to-Analog Converter.

SRIM: The Stopping and Ranges of Ions in Matter software used to calculate energy loss of ionizing particles.

SNOMAN: The SNO Monte Carlo and Analysis package.

PSUP: The PMT Support structure on which the SNO PMTs are mounted.

RNAA: Radiochemical Neutron Activation Analysis, a method of determining the radioactive impurities in a sample.

ACKNOWLEDGEMENTS

I have learned an immense amount from my advisor John Wilkerson. He has taught me that being a physicist involves more than just knowing physics. His commitment to students and his devotion to education make him an excellent mentor. John's ability to take (and make) a joke will be one of the things that I remember most about my graduate career.

Hamish Robertson has more physics knowledge easily accessible in his head than most people will ever know, and I am indebted to him for his willingness to share some of that knowledge with me.

Many thanks go to Peter Doe and Steve Elliott for giving me my first taste of experimental neutrino physics and for introducing me to my husband Minesh.

My advisor-in-law Derek Storm is always willing to answer interesting nuclear physics questions and I appreciate all that I have learned from him.

I am grateful to Sean McGee, Adam Cox, Brandon Wall, and Keith Rielage for making numerous trips to Sudbury a lot more fun and for all the martinis at Ristorante Verdicchio.

I would like to thank my officemates Noah Oblath, Michelle Leber, and Sky Sjue for interesting physics debates, random Google searches, typing contests, and all the bananas. Extra thanks go to Noah for the ROOT and SNOMAN help.

For help with countless little things, many thanks to Kathryn Miknaitis, Jason Detwiler, and Joe Formaggio.

Thanks to Barb Fulton for making life run smoothly and happily at CENPA, and

for the cheerful greetings every day.

I am grateful to Tom Burritt, John Amsbaugh, Allan Myers, Tim Van Wechel, Mark Howe, and Greg Harper for providing numerous technical innovations that made the NCDs possible.

A million thanks to the SNO staff, particularly those of you who are hard at work underground every day, for creating one of the best-run experiments ever, and for making work in Sudbury much more pleasant.

Thank you Pandora for helping with my quantum homework and for providing stress relief when it was most needed.

My eternal gratitude goes to my parents and the rest of my family for supporting me (financially and emotionally) through over 22 years of school.

And above all, thank you to my husband Minesh for always believing in me, helping me, loving me, and taking part of the burden when times got tough.

DEDICATION

To my family.

Chapter 1

INTRODUCTION TO SOLAR NEUTRINOS

1.1 Introduction to Neutrinos

The neutrino started its life as a hypothetical particle. Early interpretation of beta decay experiments implied that energy, momentum, and spin-statistics were not conserved. Rather than discard the cherished conservation principles, in 1930 Wolfgang Pauli simply invented another particle to conserve all of the quantities in the decay. This new particle was postulated to be an uncharged, massless partner to the electron, which Pauli dubbed the neutron. Following James Chadwick's 1932 discovery of what we now call the neutron, Enrico Fermi took up Pauli's hypothesized particle and named it the neutrino (Italian for "little neutral one"). Fermi's detailed theory of beta decay relied on the existence of the neutrino, which was hypothesized to be created spontaneously along with an electron in beta decay:

$$n \rightarrow p + e^{-} + \nu. \quad (1.1)$$

Detection of the neutrino would be possible through inverse beta-decay interactions such as:

$$\nu + p \rightarrow n + e^{+}, \quad (1.2)$$

but the predicted cross-sections were on the order of 10^{-44} cm², implying that a neutrino could penetrate light-years of material without interacting. This made detection

of the neutrino a daunting challenge that was not possible until the advent of atomic bombs and nuclear reactors, which produce copious quantities of neutrinos. After toying with the idea of using an atomic bomb as an intense neutrino source, Frederick Reines and Clyde Cowan set up a neutrino detection experiment near the Hanford nuclear reactor in 1953. In 1955 they moved the experiment to the Savannah River reactor, where they located the detector 11 meters from the reactor core, and 12 meters underground to provide shielding against cosmic rays.

The experimental technique used by Reines and Cowan relied on detecting the emitted positron and neutron in coincidence, as suggested by Equation 1.2. The detector consisted of about 200 liters of water with dissolved CdCl_2 . Liquid scintillator tanks instrumented with photomultiplier tubes (PMTs) were placed above and below the water tank. The positron from the inverse beta decay annihilated with an electron, creating a pair of 511-keV gammas which produced a scintillation signal in the PMTs. In addition, the neutron from the inverse beta decay captured on Cd in the water after several microseconds, producing gammas that also made a scintillation signal. Detection of the prompt annihilation gammas and the appropriately-delayed capture gammas provided a neutrino signature that discriminated against background radiation. In 1956 Reines and Cowan announced the first experimental detection of the neutrino [1], for which Reines was awarded the 1995 Nobel prize.

The neutrinos that Reines and Cowan detected were electron antineutrinos, emitted in decays of fission fragments and capable of initiating interactions with protons to produce neutrons and positrons. In a 1955 experiment at Brookhaven National Laboratory and a more sensitive 1958 follow-up experiment at the Savannah River reactor, Ray Davis showed that these electron antineutrinos were not capable of initiating interactions with neutrons to produce protons and electrons [2, 3]. This indicated that there are at least two states of neutrinos: ones that initiate interactions with protons

(right-handed particles known as antineutrinos), and ones that initiate interactions with neutrons (left-handed particles known as neutrinos). In 1937 theorist Ettore Majorana suggested that since the neutrino is electrically neutral it might be its own antiparticle [4]. In this case the distinction between what are known as neutrinos and antineutrinos is really based on their handedness, not their particle-antiparticle nature. Resolving this question is currently a very active area of experimental neutrino research.

In addition to the electron neutrinos and antineutrinos involved in beta decays, there are two other known flavors of neutrinos. In 1961 the first accelerator neutrino experiment, carried out at Brookhaven National Laboratory, proved that the neutrinos associated with muon production in pion decays were incapable of initiating reactions that produce electrons [5]. This provided evidence for the muon neutrino, a flavor of neutrino distinct from the electron neutrino. After the discovery of the tau lepton in 1975 [6] it was assumed that the associated neutrino would also exist, but it was not experimentally confirmed until the DONUT collaboration at Fermilab detected the tau neutrino in 2000 [7]. Both the muon and tau neutrinos also have associated antineutrinos.

1.2 The Weak Interaction

By the 1950s Fermi's theory of beta decay had been expanded and refined to become a comprehensive theory of weak interactions. George Gamow and Edward Teller in 1936 included nuclear spin and proposed that the change in nuclear spin be 0 or ± 1 for allowed transitions [8]. This required either an axial-vector or tensor coupling in addition to Fermi's original vector coupling. Further consideration of the properties of weakly interacting particles led theorists Tsung-Dao Lee and Chen Ning Yang to propose in 1956 that parity conservation (the invariance of behavior under

a coordinate system inversion) might be violated in the weak interaction [9]. In 1956 Chien-Shiung Wu and collaborators performed an experiment looking at the angular distribution of electrons emitted in the beta decay of polarized ^{60}Co [10]. The electrons were preferentially emitted in the direction opposite to the nuclear spin, showing that parity conservation is maximally violated by the weak interaction.

Parity nonconservation and handedness are intrinsically linked, so the picture of the weak interaction that emerged by the late 1950s was one where handedness plays a major role in determining interaction properties. There are two types of handedness in physics, helicity and chirality, which are the same only for massless particles. Helicity is defined as the projection of spin along the direction of motion, which is not boost invariant. Chirality is defined using the projection operators $P_L = (1 - \gamma^5)/2$, which leaves left-handed states unchanged and causes right-handed states to vanish, and $P_R = (1 + \gamma^5)/2$, which picks out right-handed states and destroys left-handed states. If the particle is massless, then the chirality states will be states of definite helicity. If the particle is massive, the chirality will be the same as the helicity when the particle is viewed from a reference frame in which it is relativistic.

In order to fit experimental data and explain the maximal parity violation in the weak interaction, the coupling must be of the “V – A” (vector minus axial vector) form. The Lagrangian interaction term looks like:

$$\bar{\psi}\gamma^\mu(1 - \gamma^5)\psi. \tag{1.3}$$

This structure implies that the weak interaction only acts on particles with left-handed chirality (or on antiparticles with right-handed chirality). Parity violation in weak interactions can be explained by postulating that neutrinos only exist in left-handed helicity states. Right-handed neutrinos (or left-handed antineutrinos) have not been experimentally detected, but the only way to guarantee that neutrinos are always left-handed is to make them massless. This is what is done in the Standard Model

of particle physics, but recent experimental developments have shown that neutrinos are in fact massive.

Dating back to Fermi's original beta-decay theory, the weak interaction was modeled with the exchange of currents, in parallel with electromagnetic theory. The vector bosons that mediate the weak charged-current interaction are the heavy charged W^+ and W^- bosons. In 1971 the Glashow-Weinberg-Salam (GWS) theory of electroweak interactions was shown by Gerard t'Hooft to be renormalizable, which sparked interest in weak neutral currents, a prediction of the GWS theory. Weak neutral currents were discovered in 1973 in a bubble chamber experiment at CERN [11, 12]. Weak charged-current vertices connect a charged lepton (electron, muon, or tau) to the uncharged neutrino of the same flavor through a W boson. Weak neutral currents are mediated by a Z boson and the vertices leave the particle unchanged.

1.3 The Standard Solar Model

Copious quantities of neutrinos are produced in the nuclear reactions that occur in stars. In the sun, the primary reaction mechanism is the pp chain (shown in Figure 1.1), which has the net effect of fusing four protons into helium, producing two positrons and two electron neutrinos, and releasing 26.7 MeV of energy. Another reaction mechanism, the CNO cycle, also occurs in the sun, but produces less energy and fewer neutrinos than the pp chain. The CNO cycle has the same net effect as the pp chain but requires the presence of carbon, nitrogen, and oxygen as catalysts. It dominates over the pp chain at higher stellar temperatures. Because different reactions are involved, the CNO cycle and the pp chain produce different neutrino energy spectra.

The Standard Solar Model (SSM) is a continually evolving, detailed model of the sun that is the basis of theoretical predictions of solar neutrino fluxes and spectra.

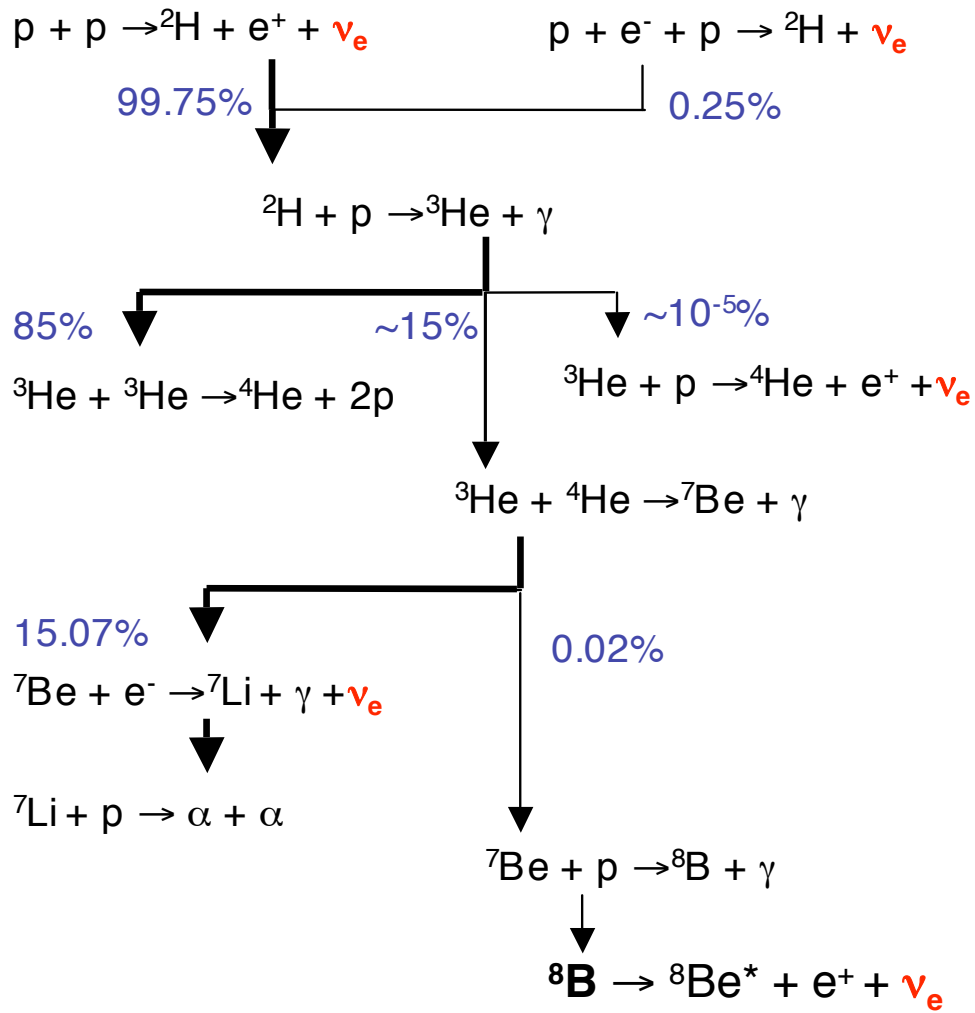


Figure 1.1: The pp chain of nuclear reactions in stars.

The SSM assumes that the sun was homogeneous before hydrogen burning began and that it is currently in hydrostatic equilibrium, with the pressures inside the sun exactly balancing gravity to prevent collapse. The energy released in nuclear reactions inside the sun is assumed to radiate and convect outwards. Starting with these assumptions, a detailed model of the sun's evolution has been developed that reproduces the current features of the sun. Many different inputs are required to define this model, such as nuclear reaction rates and energies, solar abundances of various elements, an equation of state relating the pressure and density in the sun, radiation transport, and the opacity of the solar plasma to photons. In addition to solar neutrino fluxes and spectra, the SSM produces testable predictions of the speed of sound in the convective zone near the solar surface. These sound-speed predictions can be compared with helioseismology data, which measures the frequency spectrum of pressure oscillations in the outermost layers of the sun. In general, the agreement between helioseismology data and SSM predictions has been good, although recent predictions based on reduced heavy element abundances disagree with helioseismology data [13]. The SSM continues to be refined as experiments determine the various input parameters more precisely.

The predicted SSM neutrino energies range up to almost 20 MeV, and the total flux is 6.48×10^{11} neutrinos/cm²·s reaching earth. The neutrino spectrum predicted by the most recent SSM [13] is shown in Figure 1.2. The lowest-energy, highest-flux *pp* neutrinos come from the dominant first reaction in the pp chain, the fusion of two protons. The *pep* line is from the sub-dominant first reaction in the chain, which has an electron on the input side instead of producing a positron. The two ⁷Be lines are from electron capture on ⁷Be, which can populate two different states of ⁷Li, producing two different neutrino energies. The low-flux, highest-energy *hep* neutrinos come from the reaction of ³He with a proton. The high-energy ⁸B neutrinos come

from the beta decay of ^8B . Also shown are neutrino spectra from beta decays of ^{13}N , ^{15}O , and ^{17}F in the CNO cycle.




Figure 1.2: The Standard Solar Model neutrino spectrum. Figure is from [13].

1.4 Electron Capture in the CNO Cycle

The driving component for nuclear burning in the sun is the pp fusion chain. However, it is predicted that a portion of the solar neutrino flux also comes from the CNO cycle [14] (shown in Figure 1.3). The CNO reaction products that have been shown to produce significant neutrino fluxes are the β^+ decays of ^{13}N , ^{15}O , and ^{17}F , but an additional source of neutrinos not included in Figure 1.2 is electron capture on




Figure 1.3: The CNO cycle of nuclear reactions in stars.

those isotopes. Electron capture produces a mono-energetic line 1.022 MeV above the endpoint of the β^+ continuum. The electron capture processes that occur in the CNO cycle involve the following reactions:



Capture from both free and bound-state electrons was calculated in [15], and is described here. If the electron capture process is dominated by bound electrons, it is possible to relate the electron capture flux directly to the β^+ decay flux [16]. At solar temperatures and densities, however, one must take into account the contribution from both bound and continuum electrons. The ratio between electron capture rates in the sun and laboratory measurements is given by [17]:

$$R \equiv \frac{\lambda_{\text{sun}}}{\lambda_{\text{lab}}} = n_e \frac{|\psi(0)_{\text{sun}}|^2}{2|\psi(0)_{\text{lab}}|^2}, \quad (1.7)$$

where n_e is the solar electron density and the atomic wave functions ψ are given by:

$$|\psi(0)_{\text{lab}}|^2 = \frac{1}{\pi} Z^3 \kappa(Z) \quad (1.8)$$

$$|\psi(0)_{\text{sun}}|^2 = \exp\left(-\frac{Z\beta}{R_D}\right)(\omega_c + \omega_b). \quad (1.9)$$

Here Z is the nuclear charge, $\kappa(Z)$ is a correction term applied to the pure Coulomb field of $4Z^3\alpha^3$, as tabulated in [18], $\beta \equiv \frac{1}{kT}$ is expressed in units of $\hbar = e = m_e = 1^1$, and T is the solar temperature. The factors ω_c and ω_b are continuum and bound state electron density ratios at the nucleus for Coulomb-distorted waves relative to plane waves. Also included is a weak solar plasma screening correction which depends on

¹With this convention, the unit of energy is the Hartree energy $m_e c^2 \alpha^2$ and the unit of length is the Bohr radius a_0 .

the Debye radius (R_D) [19]. The continuum and bound state electron density ratios are given by [20]:

$$\omega_c = \left\langle \frac{2\pi\eta}{1 - e^{-2\pi\eta}} \right\rangle \quad (1.10)$$

$$\omega_b = \pi^{\frac{1}{2}} (2Z^2\beta)^{\frac{3}{2}} \sum \frac{1}{n^3} \exp\left(\frac{Z^2\beta}{2n^2}\right), \quad (1.11)$$

where $\eta = Z/v$ is the inverse velocity averaged over the electron Maxwell-Boltzmann distribution.

The electron density ratios are evaluated at both a fixed point in the solar core (R_0), and integrated over the entire solar volume (R_∞). The fixed point used is 0.057 of the solar radius, where the ^{13}N , ^{15}O , and ^{17}F fluxes peak. At this location, the temperature is 1.48×10^7 K, the Debye radius is 0.45, and the density is 5.32×10^{25} atoms/cm³ [21]. The effect of the full integration on the fluxes is small for the nuclei of interest ($\sim 3\%$ for ^{13}N , and less than 1% for ^{15}O and ^{17}F). The total correction due to continuum electron capture is shown in Table 1.1. The relative K-shell/L-shell occupancy for ^{13}N , ^{15}O , and ^{17}F are all greater than 90% [18]. Capture of both K- and L-shell electrons has been included here. For evaluation of the electron capture rate with accuracy of a few percent the radiative corrections should be included (see,

Table 1.1: The fraction of bound state electrons in the solar core, the atomic wave function at the nucleus in the sun, and the total correction to the electron capture rate. Both fixed point (R_0) and volume-integrated (R_∞) ratios are shown. ^7Be is shown for comparison.

Element	$\omega_b/(\omega_c + \omega_b)$	$ \psi(0)_{\text{sun}} ^2$	R_0	R_∞
^7Be	0.302	3.76	0.858	0.804
^{13}N	0.662	11.08	0.419	0.403
^{15}O	0.749	16.14	0.400	0.398
^{17}F	0.818	23.75	0.406	0.405

for example, [22]).

Table 1.2 shows the expected total rate of neutrinos from K-shell and continuum electron capture processes, assuming the solar burning cycle is dominated by pp fusion as in the SSM. The major contribution to the uncertainties on the electron capture fluxes comes from the uncertainties on the SSM β^+ decay fluxes [21]. The neutrino flux from these electron capture sources is of the same order as the ^8B flux, though at lower neutrino energies. The solar neutrino spectrum, including the CNO electron capture neutrino lines, is shown in Fig. 1.4. There is an additional electron-capture branch for ^8B decay [23], but its total flux is $1.3 \text{ cm}^{-2} \text{ s}^{-1}$, too small to appear on the graph.

Table 1.2: Neutrino fluxes from CNO electron capture. The final electron capture flux takes into account the correction for capture of continuum electrons (R_∞). The CNO cycle is assumed to be at the level dictated by the SSM [21].

Element	SSM β^+ decay flux ($\text{cm}^{-2} \text{ s}^{-1}$)	(EC/ β^+ decay) _{lab}	EC flux ($\text{cm}^{-2} \text{ s}^{-1}$)
^{13}N	$5.48 \times 10^8 \left(\begin{smallmatrix} +0.21\% \\ -0.17\% \end{smallmatrix} \right)$	1.96×10^{-3}	4.33×10^5
^{15}O	$4.80 \times 10^8 \left(\begin{smallmatrix} +0.25\% \\ -0.19\% \end{smallmatrix} \right)$	9.94×10^{-4}	1.90×10^5
^{17}F	$5.63 \times 10^6 \left(\begin{smallmatrix} +0.25\% \\ -0.25\% \end{smallmatrix} \right)$	1.45×10^{-3}	3.32×10^3

1.5 Overview of Solar Neutrino Experiments

The first solar neutrino detector, for which Ray Davis received the 2002 Nobel prize, was a 615-ton tank of perchloroethylene dry cleaning fluid located in the Homestake gold mine in South Dakota. The experiment was located underground since the 4850-foot rock overburden dramatically reduced cosmic-ray backgrounds in the detector.




Figure 1.4: Solar neutrino flux at 1 A.U., including electron capture in the CNO cycle. Line fluxes are in $\text{cm}^{-2} \text{s}^{-1}$ and spectral fluxes are in $\text{cm}^{-2} \text{s}^{-1} \text{MeV}^{-1}$. The pp chain and the CNO β^+ decay fluxes are from [24], at the levels calculated in [21].

In Davis's experiment, neutrinos participated in inverse beta-decay interactions with ^{37}Cl in the fluid, producing ^{37}Ar and an electron. The argon was chemically extracted from the perchloroethylene approximately once every two months and the decay of ^{37}Ar ($T_{1/2} = 35$ days) back to ^{37}Cl was detected in a proportional counter. The energy threshold for the inverse beta-decay reaction on ^{37}Cl is 0.813 MeV. Davis's experiment was sensitive primarily to neutrinos from ^8B decay in the sun, but was also sensitive to other neutrinos from the pp chain and the CNO cycle above the energy threshold.

The motivation for Davis's experiment was provided by the very property of neutrinos that makes them so difficult to detect: their incredibly small interaction cross-sections. Unlike photons, whose mean free path in the sun is a few millimeters, neutrinos typically stream directly out of the sun without interacting. Ray Davis and John Bahcall recognized in the early 1960s that detection of neutrinos from the stellar core could provide verification of the hypothesis that the sun's energy is produced by nuclear fusion reactions [25, 26]. Starting in 1968 Davis's experiment did indeed detect solar neutrinos [27], verifying solar fusion, but it raised a larger puzzle than the one it was designed to solve.

Davis's experiment saw only about a third of the neutrinos predicted by Bahcall's SSM. This discrepancy became known as the Solar Neutrino Problem (SNP), and plagued physicists for almost 35 years. It was assumed at first that the solar models were wrong, since the ^8B neutrino flux has a T^{25} dependence on the solar temperature so it is extremely difficult to correctly predict neutrino fluxes. But increasing levels of detail in the SSM and impressive agreement with helioseismology data gradually persuaded many skeptics that the model was not at fault. Davis's experiment produced consistent results for over 30 years, which continued to fall short of the theoretical flux predictions.

In order to find the solution to the SNP, physicists built several different types of solar neutrino detectors. The Russian American Gallium Experiment (SAGE) [28] and Gallex (later renamed GNO) [29, 30] are radiochemical detectors, like Davis's experiment, but with gallium instead of chlorine as the target. Neutrino interactions with ^{71}Ga produce ^{71}Ge and an electron. The ^{71}Ge is chemically extracted from the gallium and counted in a proportional counter. The benefit of gallium is the 0.232-MeV energy threshold for neutrino reactions, which allows neutrinos from pp fusion, the initial reaction in the pp chain, to be detected in addition to the higher-energy

neutrinos to which Davis’s chlorine experiment was sensitive. The radiochemical neutrino detectors rely on charged-current interactions for neutrino detection, thus they are only sensitive to electron neutrinos.

In addition to radiochemical detectors, there are the water Cherenkov detectors Kamiokande and Super Kamiokande (SuperK) [31]. A neutrino scatters off an electron in a tank of water, imparting enough energy to the electron to make it radiate Cherenkov light, which is detected by photomultiplier tubes that line the walls of the tank. A Cherenkov detector provides real-time detection and directional sensitivity, unlike radiochemical experiments, but has a much higher energy threshold of about 5 MeV. The water Cherenkov detectors detect neutrinos via elastic scattering off electrons, thus they are primarily sensitive to electron neutrinos (via charged-current or neutral-current scattering) with reduced sensitivity to muon and tau neutrinos (via neutral-current scattering only). None of these experiments have seen the total predicted solar neutrino flux, as shown in Table 1.3.

Table 1.3: Fraction of SSM neutrino flux detected by various different solar neutrino detectors.

Experiment	Dates	Detection Reaction	Exp./Theory
Chlorine (127 ton)	1970 – 1995	$^{37}\text{Cl} + \nu_e \rightarrow ^{37}\text{Ar} + e^-$	0.34 ± 0.03
Kamiokande (680 ton)	1986 – 1995	$\nu_x + e^- \rightarrow \nu_x + e^-$	0.54 ± 0.08
SAGE (23 ton)	1990 –	$^{71}\text{Ga} + \nu_e \rightarrow ^{71}\text{Ge} + e^-$	0.55 ± 0.05
Gallex/GNO (12 ton)	1991 – 2004	$^{71}\text{Ga} + \nu_e \rightarrow ^{71}\text{Ge} + e^-$	0.57 ± 0.05
SuperK (22 kiloton)	1996 –	$\nu_x + e^- \rightarrow \nu_x + e^-$	$0.451^{+0.017}_{-0.015}$

1.6 Neutrino Oscillations

Solar neutrino data point to the neutrinos themselves as the culprit in the SNP. Not only do all solar neutrino experiments see a deficit of neutrinos, but experiments sensitive to different neutrino energies see different deficits, implying that whatever is happening to the solar neutrinos is energy dependent. A possible explanation for the SNP is neutrino oscillations, first postulated in 1969 by Bruno Pontecorvo and Vladimir Gribov [32]. Neutrinos interact in weak eigenstates, which are defined by the charged lepton that is also involved in the interaction, but they propagate in mass eigenstates. If neutrinos have non-zero mass, and the weak eigenstates of neutrinos are mixtures of the mass eigenstates, then a neutrino produced in a particular weak eigenstate (such as an electron neutrino) could transform into a different weak eigenstate (a muon or tau neutrino) as it travels from the sun to the earth. Thus a detector that is only sensitive to electron neutrinos, but not to muon or tau neutrinos, would detect fewer neutrinos than expected.

Within the neutrino oscillation model, in the approximation that only two neutrino flavors need to be considered, the probability of a neutrino created as an electron neutrino changing flavors can be expressed as:

$$P(\nu_e \rightarrow \nu_{\mu\tau}) = \sin^2(2\theta) \sin^2(1.27\Delta m^2 L/E_\nu). \quad (1.12)$$

In this equation, θ is the mixing angle between the neutrino flavor and mass eigenstates, Δm^2 is the mass-squared difference between neutrino mass eigenstates in eV^2 , L is the distance the neutrino has travelled in meters, E_ν is the neutrino energy in MeV, and 1.27 accounts for factors of \hbar and c in these units. Data from solar neutrino experiments with different E_ν provide information about the mass-squared differences between neutrino mass eigenstates as well as the size of the neutrino mixing angle.

Drawing on analogy to the quark sector, where mixing between weak eigenstates

and mass eigenstates is quite small, the bias in the physics community while neutrino oscillation theory was being developed was that neutrino mixing would also be small. As is evident in Equation 1.12, it is impossible to achieve large flavor transformation for small values of the mixing angle. Thus neutrino oscillations were not looked upon favorably as the solution to the SNP, since small mixing angles could not explain the large reduction in neutrino flux observed by the various experiments.

An extension of the oscillation theory known as the MSW effect was developed by Mikheyev, Smirnov, and Wolfenstein [33, 34] and includes effects of matter enhancement of the neutrino flavor-change probability. When an electron neutrino passes through matter it interacts with electrons via the weak charged current as well as the weak neutral current. Muon and tau neutrinos cannot participate in charged-current interactions since there are no muons or taus present in normal matter. This adds an effective potential energy that only the electron neutrinos feel, which can lead to a resonant enhancement of the flavor change probability for certain densities of matter. Incorporating the MSW effect into the neutrino oscillation theory allows even small mixing angles to produce large flavor change probabilities. As it turned out, the neutrino mixing angles are quite large, contrary to the early theoretical bias, but matter effects are still required to explain solar neutrino data.

For certain mass-squared differences and matter-enhanced effective mixing angles, the solar neutrino data are consistent with neutrino oscillations explaining the observed deficit. However, this is only circumstantial evidence for neutrino flavor change – an appearance measurement of muon or tau neutrinos from the sun would give definitive proof of flavor change. An experiment capable of detecting all three neutrino flavors with equal sensitivity could provide this, as well as measure the total active neutrino flux from the sun. This motivated the construction of the Sudbury Neutrino Observatory.

Chapter 2

THE SNO NEUTRAL CURRENT DETECTORS

2.1 The Sudbury Neutrino Observatory

The Sudbury Neutrino Observatory (SNO) is a heavy-water Cherenkov detector located 6800 feet (6010 meters-water-equivalent) underground near Sudbury, Ontario in Canada [35]. The SNO laboratory is housed in the INCO Creighton mine, an active nickel mine. The detector (shown in Figure 2.1) consists of one kilotonne of ultra-pure heavy water (D_2O) enclosed in a spherical transparent acrylic vessel (AV) 12 meters in diameter. This is surrounded by a geodesic support structure holding 9456 photomultiplier tubes (PMTs), which provide 60% coverage. All of this is immersed in seven kilotonnes of ultra-pure H_2O , which shields the heavy water from radioactivity originating in the rock walls of the cavity.

The SNO experiment detects solar neutrinos, primarily from 8B decay in the pp chain. Neutrino interactions in the heavy water produce Cherenkov light cones that are detected by the PMTs surrounding the AV. The number of PMTs that are hit in each event is approximately proportional to the energy of the event. In many cases, the pattern of hit PMTs forms a ring characteristic of the Cherenkov light cone. The position and direction of each event is reconstructed by the timing of the light reaching each PMT in the event and the positions of the hit PMTs. The signal of about ten detected neutrino events per day has to be separated from about 20 Hz of background events caused by instrumental effects, radioactivity, and cosmic-ray muons. This is done with a series of data-cleaning cuts based on the pattern of PMTs that trigger in




Figure 2.1: An artist's rendition of the SNO detector, showing the AV, the PMT support structure, the control room, and the clean room above the neck of the AV.

each event. Regular calibrations of the SNO electronics and the detector's response to various types of signals are required to interpret the data and study uncertainties.

Three different neutrino interactions are possible in heavy water. All three active neutrino flavors can undergo elastic-scattering (ES) interactions with electrons in the heavy water:

$$\nu_x + e^- \rightarrow \nu_x + e^-, \quad (2.1)$$

the same interaction used by Kamiodande and Super Kamiokande. Feynman diagrams for ES interactions mediated by the W and the Z, respectively, are shown in

Figures 2.2 and 2.3. For muon and tau neutrinos, only Z exchange is permitted because there are no muons or taus present to couple to the neutrinos via the W. For electron neutrinos, both W exchange and Z exchange are permitted, and these two processes interfere destructively [36]. The cross-section for Z exchange is considerably smaller than that for W exchange, so despite the destructive interference in the cross-section for ν_e elastic scattering, it is about six times larger than the cross-section for ν_μ or ν_τ elastic scattering. Thus the ES interaction has reduced sensitivity to ν_μ and ν_τ compared to ν_e .




Figure 2.2: Feynman diagrams for an electron neutrino participating in an elastic-scattering interaction with an electron, mediated by a W boson. The s-channel (left) and the t-channel (right) amplitudes both contribute. In these diagrams, time increases to the right.

In SNO, the relativistic outgoing electron from the ES interaction produces a Cherenkov light cone that is detected in the PMT array. An interesting and useful feature of the ES interaction is its distinct directional dependence. The scattered





Figure 2.3: Feynman diagram for a neutrino of any flavor participating in an elastic-scattering interaction with an electron, mediated by a Z boson. In this diagram, time increases to the right.

electrons are strongly peaked in the forward direction, along the direction of motion of the neutrino. Thus the direction of the electrons can be used to verify the solar origin of the neutrinos and can help distinguish the ES signal from other signals.

Electron neutrinos can also participate in charged-current (CC) interactions with deuterium [37]:



mediated by a W boson, as shown in Figure 2.4. For muon or tau neutrinos of solar energies, CC interactions cannot occur because the neutrinos are not energetic enough to create a muon or tau. Like in the ES interaction, the relativistic electron from the CC interaction produces a Cherenkov light cone that is detected in the PMTs. The electron carries away much of the initial neutrino energy, less the 1.44-MeV threshold

and a correction for the proton recoil. Thus the measured energy spectrum of the electrons from the CC interaction directly reflects the neutrino energy spectrum.




Figure 2.4: Feynman diagram for an electron neutrino participating in a charged-current interaction with the neutron in deuterium. In actuality, the W couples to one of the d quarks in the neutron, not to the entire neutron. In this diagram, time increases to the right.

The unique feature of the SNO experiment is its ability to detect all three active neutrino flavors with equal sensitivity [38]. This is possible via neutral-current (NC) interactions of neutrinos with deuterium:

$$\nu_x + {}^2\text{H} \rightarrow p + n + \nu_x, \quad (2.3)$$

which takes place through the exchange of a Z boson, as shown in Figure 2.5. The signature of this interaction is the free neutron produced, so a primary goal in SNO is the detection of free neutrons. This is done with a different detection medium in each of the three phases of the SNO experiment: pure heavy water, salt, and

Neutral Current Detectors (NCDs). In all three phases, the neutrons thermalize before capturing, so no information about the neutrino energy or direction is preserved in the NC signal. The CNO electron capture flux calculated in Section 1.4 produces 0.4 ^{15}O neutrino NC events per year and 0.01 ^{17}F events per year, which contribute a small model-dependent background to the ^8B flux measurement.




Figure 2.5: Feynman diagram for a neutrino of any flavor participating in a neutral current interaction with deuterium. In actuality, the Z couples to one of the quarks in the deuterium nucleus, not to the entire nucleus. In this diagram, time increases to the right.

In the first phase of the SNO experiment, which ran from November 1999 to May 2001, neutrons liberated in NC interactions were observed by their capture on deuterium, which releases a 6.25-MeV gamma. The gammas Compton-scatter electrons, imparting enough energy for these electrons to produce Cherenkov light that is detected in the PMTs. The events detected in this phase of SNO were statistically separated into NC, CC, and ES using maximum-likelihood fits to probability density

functions (PDFs) of the expected energy, angle relative to the sun, and radial position in the AV.

During the second phase, from July 2001 to August 2003, chlorine was added to the heavy water in the form of ultra-pure table salt (NaCl). A neutron capture on ^{35}Cl produces up to three photons with 8.6 MeV total energy. The neutral-current signal was enhanced by the 44-barn neutron-capture cross section on ^{35}Cl , compared to the 0.0005-barn cross section on deuterium. The higher energy of the gammas shifted the neutral-current signal further above low-energy radioactive backgrounds compared to the pure heavy water phase. In addition, the multiple gammas produced a more isotropic signal than a single gamma or the single Cherenkov cone from CC or ES interactions, so an isotropy parameter could be added to increase the power of the statistical separation.

By using the isotropy of NC events in the salt phase of SNO, along with the radial position, the angle relative to the sun, and the energy, it was possible to statistically separate the NC, CC, and ES signals without assuming that the CC spectrum was that expected from solar ^8B neutrinos. This allowed a model-independent measurement of the energy spectrum of the neutrinos involved in the CC interaction. It also allowed a search for differences between the measured neutrino energy spectrum and the predicted spectrum, which could be caused by neutrino oscillations in certain regions of neutrino oscillation parameter space.

The most recent results, from the entire salt phase of SNO [39], indicate that the NC flux is approximately three times larger than the CC flux, as shown in Figure 2.6. This suppression of the CC flux relative to the NC flux indicates that neutrinos are changing flavor, which could be explained by MSW neutrino oscillations. The total neutrino flux from the sun measured with the NC interaction agrees with the predictions of the SSM, a major triumph for nuclear physics and astrophysics. The




Figure 2.6: Flux of muon and tau neutrinos versus flux of electron neutrinos. The CC, NC, and ES fluxes observed by SNO and the ES flux observed by Super-Kamiokande [31] are shown as filled bands, and their intercepts with the axes represent $\pm 1\sigma$ uncertainties. The total ${}^8\text{B}$ solar neutrino flux predicted by the SSM [13] is indicated by the dashed lines. The point and the contour lines represent ϕ_e from the CC flux and $\phi_{\mu\tau}$ from the NC-CC difference with 68%, 95%, and 99% confidence level contours. The non-zero value of $\phi_{\mu\tau}$ provides strong evidence for solar neutrino flavor transformation. Plot is from [39].

results from the salt phase of SNO are in agreement with previous results from the pure heavy water phase [40–42], but improve on those results by providing a model-independent measurement of the CC spectrum and increased NC precision.

Some confirmation of SNO’s results comes from the KamLAND experiment, which detects antineutrinos from nuclear reactors in Japan and surrounding countries. The particular neutrino oscillation parameters that best explain the combined solar neu-

trino results agree remarkably well with the oscillation parameters extracted from the KamLAND data, as shown in Figure 2.7. This agreement is a good confirmation of the neutrino oscillation model, particularly since KamLAND and SNO are dramatically different experiments, with SNO measuring matter-enhanced oscillations of solar neutrinos and KamLAND measuring vacuum oscillations of reactor anti-neutrinos.




Figure 2.7: Neutrino oscillation parameters. The plot is from [43] and shows KamLAND data, as well as a comparison to solar neutrino data. The agreement between the best-fit points from KamLAND and from solar neutrino experiments is a good confirmation of the neutrino oscillation model.

2.2 Motivation for the NCD Phase

For the NCD phase of the SNO experiment, which began in December 2004 and will continue through December 2006, the salt was removed and an array of ^3He proportional counters was installed in the heavy water to capture neutrons liberated in the NC interaction. The capture cross-section for thermal neutrons on ^3He is 5330 barns, about seven orders of magnitude larger than the capture cross-section on deuterium. Thus a sparse array of ^3He proportional counters occupying 0.1% of the heavy water volume is sufficient to provide a neutron-capture efficiency of 26%, giving a NC signal of about 3.3 events per day in the NCDs. The NCD array does not interfere dramatically with the light collection in the PMTs, blocking only about 10% of the light. A drawing of the SNO detector with the NCD array installed is shown in Figure 2.8.

The neutron-capture signal in the NCD array is read out through a separate data acquisition system from the Cherenkov light signal observed with the PMTs. The NC systematics in the NCD phase are largely independent from the pure heavy water and salt phases of the SNO experiment because the detection system is completely different. In essence, the NCD phase of SNO will be a different experiment providing an independent measurement of the NC flux.

Statistical separation of the NC signal from the PMT signals is unnecessary in the NCD phase of SNO, breaking the statistical correlations between the NC and CC signals. The NC signal contribution to the light collected in the PMTs is reduced compared to previous phases of SNO, since the capture efficiency on deuterium in the NCD phase is only about 18%. The NC flux measured in the NCDs will calibrate the NC contribution to the PMT signal and allow it to be subtracted from the CC and ES signals. This allows the best possible determination of the CC energy spectrum. One of the goals of the NCD phase of SNO is to search for distortions in the CC




Figure 2.8: An artist's drawing of the SNO detector with the NCD array partially installed. The AV and neck are visible, as well as part of the PMT array and part of the NCD array.

spectrum relative to the expected energy spectrum of ^8B solar neutrinos, which would provide confirmation that matter-enhanced neutrino oscillations are responsible for solar-neutrino flavor change.

In the NCD phase, the NC signals are distinguished on an event-by-event basis, so any time variations in the solar neutrino flux can be followed separately in NC and CC/ES. The eccentricity of the earth's orbit causes a periodic time variation that is visible in the solar neutrino flux. Other possible sources of time variations in the neutrino flux have been proposed, such as the 22-year sunspot cycle. Evidence in Super Kamiokande data for a 9.43-y^{-1} periodicity in the solar neutrino flux has been claimed [44] but SNO has not seen evidence for this periodicity in the pure-heavy-water or salt phases [45]. The topic of periodicity in the solar neutrino flux is an interesting and controversial one and the ability to follow time variations separately in the NC and CC/ES fluxes is an advantage of the NCD phase.

The SNO experiment will make precision measurements of solar neutrino fluxes and day-night asymmetries in the NCD phase. The errors on the NC flux, the CC flux, and the day-night asymmetry from the pure-heavy-water and salt phases are shown in Table 2.1, along with the expected errors on these quantities in the NCD phase.

Table 2.1: The uncertainty on the NC flux, the CC flux, and the day-night asymmetry from the pure-heavy-water and the salt phases, and the expected errors on these quantities in the NCD phase.

Uncertainty	D ₂ O	Salt	NCD (expected)
NC (σ_{NC}/NC)	12%	8%	~6%
CC (σ_{CC}/CC)	6%	6%	~4%
Day-Night (σ_A)	5%	7%	~5%

The NCD phase will provide SNO's best measurement of the CC/NC ratio, which constrains the solar neutrino mixing angle. While KamLAND is likely to provide the best constraint on Δm^2 , the best experimental constraint on the mixing angle will come from SNO's measurement of the CC/NC ratio. Figure 2.9 shows the expected improvement in the measurement of the solar neutrino mixing angle after the NCD phase of SNO is complete.




Figure 2.9: Currently allowed MSW region and the expected improvement in the mixing-angle constraint after the NCD phase of SNO. The shaded areas are the 90%, 95%, 99%, and 99.73% confidence-level allowed regions from solar neutrino data and KamLAND. The dashed lines are iso-CC/NC contours for the SNO experiment. The solid lines are the expected $\pm 1\sigma$ uncertainties on the CC/NC ratio from SNO after the NCD phase. Figure adapted from [46].

2.3 NCD Signals

The NCD array consists of 36 strings of ^3He -filled and 4 strings of ^4He -filled proportional counters deployed vertically on a one-meter grid in the SNO detector. Each string, one of which is shown in Figure 2.10, is between 9 and 11 meters in length. The total length of the NCD array is 398 meters. Each NCD string consists of three or four individual NCDs, with a single read-out through a cable at the top of the string. An individual NCD is a two-inch-diameter nickel tube with a length of 2, 2.5, or 3 meters. The NCDs are filled with a mixture of 85% partial pressure of ^3He or ^4He and 15% partial pressure of CF_4 gas, to a total pressure of 2.5 atmospheres. Each individual NCD is an independent gas volume, but the outer walls are welded together and electrical connections are made between the three or four NCDs in a string.

The basic operation of the NCDs is like that of all proportional counters. An energetic charged particle ionizes the gas inside the detector, creating a number of electron-ion pairs. The number of these primary ionization pairs is dependent on the energy loss of the particle in the gas. The electrons are accelerated by application of high voltage to the anode of the proportional counter, which is the central wire in the typical cylindrical geometry. The electrons are accelerated to sufficient kinetic energy to produce secondary ionization in the avalanche region near the anode. The number of secondary electron-ion pairs is proportional to the number of primary pairs, and this proportionality constant is the gas gain of the counter. The movement of the electron-ion pairs in the proportional counter induces an electrical signal on the anode which is proportional to the energy of the original ionizing particle.

The NCDs detect neutrons via the $^3\text{He}(n,p)^3\text{H}$ reaction, which has a Q-value of 764 keV. To conserve energy and momentum, the proton and triton are always emitted back-to-back, and always carry off 573 keV and 191 keV of kinetic energy, respectively.

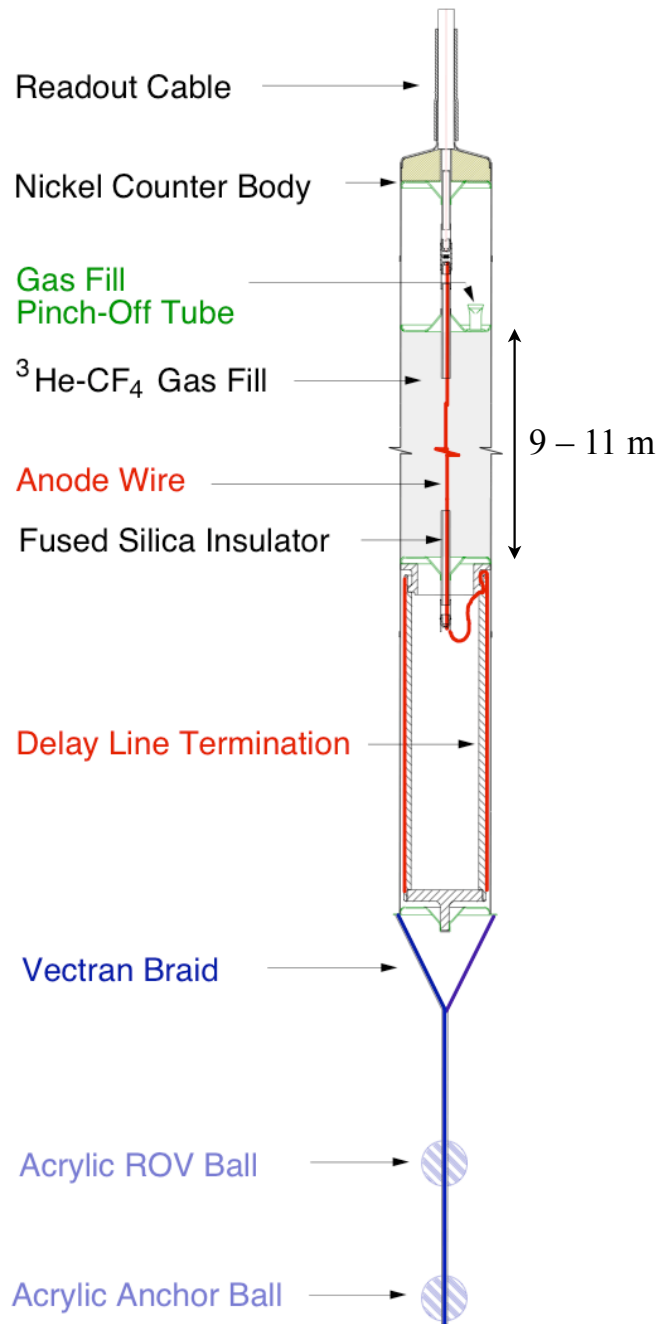


Figure 2.10: Diagram of an NCD string. The counter body actually consists of three or four independent gas volumes welded together with a continuous electrical connection along the anode. The anchor balls at the bottom of the string allow it to be attached to the SNO AV. Adapted from [47].

This energy is deposited in the proportional counter as the proton and triton ionize the gas. The primary electrons accelerate toward a central anode wire that is at 1950 V. This voltage is sufficient to cause an avalanche as the electron nears the wire, but low enough that the NCD essentially remains in the proportional regime. Thus the total charge collected on the anode wire during the event is nearly proportional to the energy deposited in the gas by the initial ionizing particles. Deviations from proportionality are caused by the positive ions created in the avalanche altering the electric field near the wire, which is known as the space-charge effect.

The NCDs have been deployed on a one-meter grid in the AV, providing x- and y-position resolution on that scale. It is desirable to have z-position resolution on the same scale. Typically, z-position is determined in cylindrical proportional counters by using double-ended readout, which is undesirable for the NCDs because the additional cables would interfere with the light detected by the PMTs and increase the radioactive backgrounds. Z-position determination can be achieved with single-ended readout by measuring the time delay between the initial pulse and the reflection off the bottom of the NCD string. In order to resolve this reflection when the neutron capture occurs near the bottom of the string, a 90-ns round-trip delay line is added to the bottom of each string.

In a ^3He proportional counter neutron-capture event, if either the proton or the triton hits the wall before losing all its energy to the gas, then less than the full 764 keV will be deposited in the gas. An energy spectrum from a ^3He proportional counter has a distinct peak at 764 keV. In addition, the spectrum has a shoulder at 573 keV, from the triton hitting the wall and only the proton ionizing the gas, and at 191 keV, from the proton hitting the wall and the triton entering the gas. An idealized spectrum from a ^3He proportional counter that illustrates these features is shown in Figure 2.11.

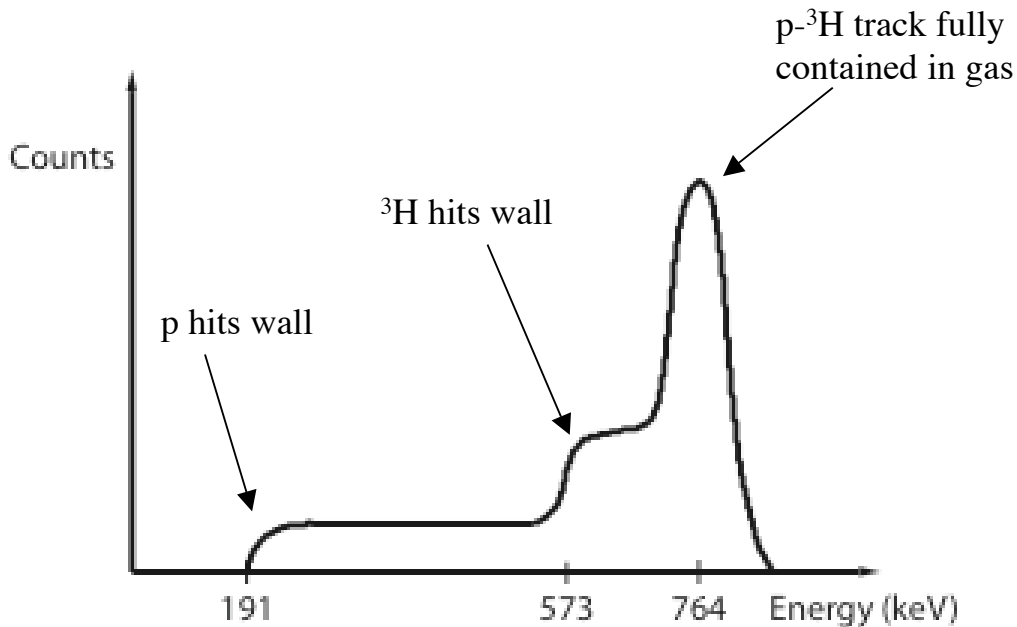


Figure 2.11: An idealized neutron spectrum from a ^3He proportional counter that illustrates the 764-keV peak and the 573-keV and 191-keV shoulders from the triton and proton striking the wall. Figure is from [48].

Neutron energy spectra taken from the NCD array look very much like the idealized spectrum, as shown in the energy histogram of AmBe neutron source data in Figure 2.12. NCD neutron spectra exhibit a clear peak at 764 keV. The shoulder at 191 keV is also visible, although it is sometimes partially obscured by low-energy noise and high-voltage discharge events. The shoulder at 573 keV and the low-energy edge of the 764-keV peak are broadened by space charge effects.

The NCDs are optimized to detect neutrons liberated in the NC reaction in SNO, but are capable of detecting other types of signals as well. Any thermal neutron can capture on ^3He , producing a characteristic back-to-back proton-triton track. In addition, any charged particle can ionize the gas and produce a signal. Minimum

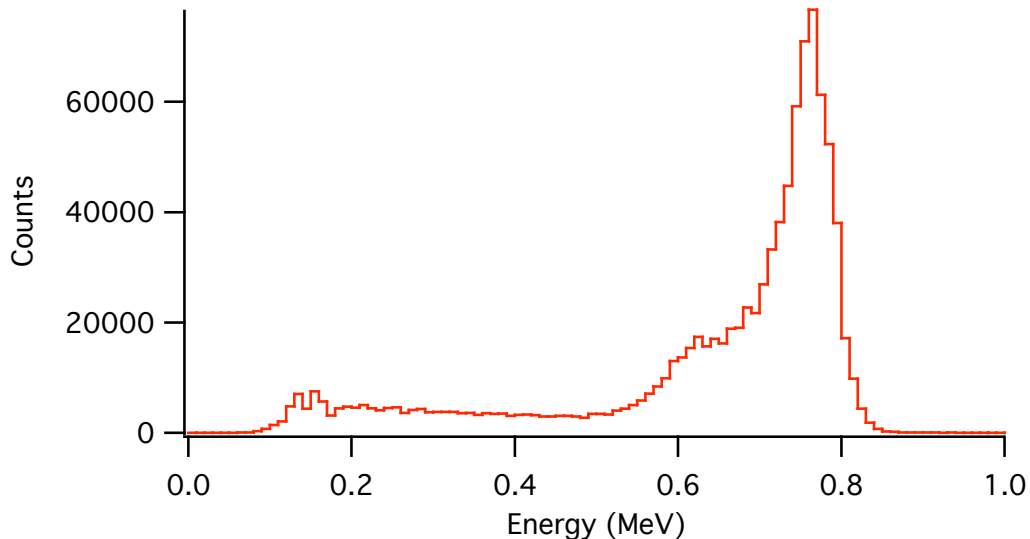


Figure 2.12: An energy histogram of AmBe neutron source data from a scan of the entire NCD array. The essential features of a ^3He proportional counter are evident, but the shoulder at 573 keV is somewhat washed out due to space charge effects and some low-energy noise events partially obscure the shoulder at 191 keV.

ionizing particles are not a serious source of background because they rarely deposit enough energy in an NCD to mimic a neutron-capture event. The most important background is alpha particles from radioactive decays, which are discussed in detail in chapter 6. An alpha particle from outside an NCD cannot travel through the nickel wall, so only alphas produced in the NCD itself are a source of background.

In addition to physics events such as neutron captures and alpha backgrounds, there are instrumental backgrounds in the NCD array. These come from many different sources, both internal and external to the NCD array. A significant source of instrumental background in the NCD array is high-voltage breakdown events within the NCDs themselves, known as microdischarge. In addition, RF noise pickup from external sources such as the PMT electronics can be a problem for the NCD array.

Another source of instrumental backgrounds is the response of the NCDs to seismic activity and blasting in the mine or to other activity in the SNO detector such as circulation of the heavy water or high-voltage breakdown in the PMTs. Instrumental backgrounds are eliminated by a suite of data-cleaning cuts implemented prior to the high-level analyses. More information about these data-cleaning cuts is available in Appendix A.

2.4 NCD Data Acquisition

The data acquisition (DAQ) system for the NCD array consists of two primary paths, each optimized to serve a different function. The current collected on the anode wire from an event in an NCD string is read out from the top end of the string through a cable that also provides high voltage to the string. The signal is sent through a current preamplifier and then to a multiplexer (MUX), where it is split into the two paths. On one data path, shaping circuitry followed by a sampling analog-to-digital converter acquires integrated current signals proportional to energy. There is a Shaper/ADC on each channel. The Shaper/ADC channels can easily handle the kHz data rates expected from a supernova burst, when background rejection through pulse shape analysis is not as important due to the short timescale.

On the other data path the time profile of the current reaching the anode wire is digitized, allowing the use of pulse shape analysis techniques to reject backgrounds. If the pulse passes the digitization threshold, it is logarithmically amplified and sent through a 300-ns delay and a summing junction to an oscilloscope where it is digitized. The purpose of the MUX is to allow many strings to be digitized with just two oscilloscopes, each with four channels. The summing junction optimizes the alternation between the two scopes so if one is busy the next event can be sent to the other one by appropriate routing of an external trigger to the scopes. The logarithmic amplifier

maintains acceptable signal-to-noise, despite the factor of 150 difference between the amplitude of the largest expected physics events and the smallest. The form of the equation governing the logarithmic amplifier is $V_{out} = a \cdot \log(1 + V_{in}/b) + c$.

A digitized event from an NCD consists of 15,000 points of one nanosecond each, covering a 15-microsecond time interval. The first 1.5 microseconds (1500 points) occur prior to the trigger and only contain the baseline level of noise from the oscilloscope. Then the MUX trigger occurs and the next 300 ns are the delay as the signal travels through the delay cable in the MUX to the scope. The event itself takes up the rest of the digitization time, although typically the amplitude returns very close to the baseline value within about 5 microseconds.

The triggers for the two data paths are independent and very different, so many spurious background events trigger only one of the two paths. The MUX trigger for the digitization path is based on a very fast integration of the current signal from the NCD string, so it is essentially a current trigger. Events with very low amplitude will not trigger the digitization path. These include certain types of unusual discharge events that have a long duration and low amplitude, as well as the lowest-amplitude neutron events. More information about optimizing the trigger thresholds in the digitization path can be found in Section 4.6. The trigger for the Shaper/ADC path is essentially a charge trigger. The signal is continually shaped and integrated over a six-microsecond time period – if it exceeds a certain threshold then the Shaper/ADC triggers and records the shaped, integrated charge of the event. Very narrow events with only a small amount of charge will not trigger the Shaper/ADCs, even though the event may have an amplitude large enough to trigger the digitization path. In general, candidate physics events are required to have triggered both data paths to be considered analyzable, except in the case of a supernova, when only the Shaper/ADC data would be used.

It is very important to understand the effects of the electronics on the NCD data. To that end, electronics calibrations are performed weekly by injecting a pulser signal into each NCD preamplifier and recording it as if it were an actual event. This signal injection closely approximates the situation encountered by a real signal, but does not exactly mimic it since the signal input to the preamp sees a different impedance than the pulser input. The form of the injected pulse is well known, so the effects of the electronics can be determined. The pulser signal is injected using a pulser distribution system that allows the output from one programmable pulser to be sent to any combination of the 40 NCD strings or the spare electronics channels.

All data from the NCD system are read out and recorded using the Object-oriented Real-time Control and Acquisition (ORCA) software, a data acquisition package that was custom-designed and written at the University of Washington. Thresholds, gains, and other hardware settings can be controlled and monitored through ORCA. ORCA interfaces with SNO Hardware Acquisition and Real-time Control (SHARC), the custom-built University of Washington software used to acquire data from the SNO PMT array. This allows the operator to use only SHARC for routine running of the entire SNO detector, rather than having to use two different DAQ systems. All important operating parameters of the NCD system can be monitored through SHARC, although advanced changes to NCD hardware settings require use of the ORCA interface.

Chapter 3

NCD CONSTRUCTION AND DEPLOYMENT

3.1 *NCD Fabrication*

One of the primary considerations in the NCD design was minimization of radioactive backgrounds introduced into the SNO detector, particularly ^{238}U and ^{232}Th . Decay products in the ^{238}U and ^{232}Th chains can photodisintegrate deuterium, producing neutrons that are indistinguishable from the NC signature. In order to limit the NCD array's photodisintegration-neutron production rate to less than 2% of the NC signal, the NCD bodies must contain less than 24 picograms-per-gram ^{238}U and less than 2 picograms-per-gram ^{232}Th ¹. Commercial proportional counters could not be used, since they did not meet these stringent radioactivity limits. To reach the required cleanliness, the NCDs were custom designed and built by the University of Washington and Los Alamos National Laboratory. The assembly of the NCDs took place at the University of Washington in Seattle. A brief description of the NCD construction is provided here, but more information can be found in [49] and [47].

The NCD bodies were made from ultra-pure, chemical-vapor-deposited (CVD) nickel. Nickel was chosen both for its strength and for its ability to participate in the chemical reactions involved in the CVD process. Under pressure, at about 50°C, nickel combines with carbon monoxide to make gaseous nickel carbonyl ($\text{Ni}(\text{CO})_4$). At

¹These are the construction specification used for the NCD array, but the gamma intensities used in calculating them were not correct so they do not produce 2% of the NC signal. An updated calculation is presented in Section 8.1.

about 250°C nickel carbonyl decomposes again into nickel and carbon monoxide. The nickel is precipitated onto an aluminum mandrel to make the nickel tubes for the NCD bodies. This CVD process allows the nickel to be separated from impurities, since no other elements form carbonyls reversibly at these temperatures. It is known that both uranium and thorium can form carbonyls [50, 51], so it is plausible that trace amounts of these impurities remained in the nickel throughout the CVD process and continue to contaminate the nickel, although at greatly suppressed levels compared to natural nickel. The ^{238}U and ^{232}Th content of most metals in the Earth's crust is on the order of 1 – 10 parts per million. By using the CVD process, the ^{238}U and ^{232}Th content of the NCD nickel was reduced to a few parts per trillion, an improvement of about 10^6 .

During the construction of the NCDs, the nickel tubes were manufactured faster than the NCDs were assembled. In order to reduce radioactivity caused by cosmic activation of the nickel, it was desirable to store the tubes in an underground location until they were ready to be assembled. A suitable location was found in a tunnel near Index, Washington, less than two hours from Seattle, which had originally been excavated as a test of tunnel-boring machinery for the Superconducting Supercollider. When the first group of tubes were brought back to Seattle from Index, it was found that they had very high rates of 5.3 MeV alphas from ^{210}Po decay. The nylon bags in which the NCDs were stored at Index had not been sealed, and subsequent measurements in the tunnel indicated a radon level of greater than 900 pCi/l [49]. Radon daughters such as ^{210}Pb and ^{210}Po had plated onto the inner and outer surfaces of the nickel tubes. The event rate near the 5.3 MeV peak was about 10^5 events/m²·day (the deployed NCD array has a surface area of 63.5 m²), with several thousand events per day falling into the energy region of the neutron-capture signal [49]. This rate was unacceptably high, so cleaning methods were devised to remove

the contamination from the surface of the nickel.

The nickel tubes were electropolished and acid etched to remove contaminants such as aluminum from the mandrel and radon daughters from air. Electropolishing was necessary to remove ^{210}Po adhering to the nickel tubes that could not be removed by acid etching alone because polonium displaces nickel into acidic solution. Although only about half of the tubes were stored at Index, there is some radon present in all air, so the same cleaning procedures were applied to almost all the tubes, except that the tubes from Index were prewashed in an acetic acid solution to remove surface ^{210}Pb . For the electropolish, the tubes were placed in a solution of 1 M sulfuric acid and heated to 130°F , and a current of 155 A was applied. The outsides of the tubes were electropolished for 45 seconds, removing $2\ \mu\text{m}$ of nickel, and the insides were electropolished twice for a total of nine minutes, removing a total of $20\ \mu\text{m}$ of nickel. More material was removed from the insides of the tubes because surface alphas on the outside cannot penetrate the nickel to produce a signal in an NCD.

After the electropolish, the tubes were acid etched at room temperature in a solution of hydrofluoric and nitric acid for 10 minutes, then transferred to a nitric acid solution for one minute, and finally to two weak acetic acid rinses for one minute each. The acid etch removed $2\ \mu\text{m}$ of nickel from each surface of the tubes. The hydrofluoric acid component was included to remove adhering granules of aluminum oxide from the mandrel on which the nickel tubes were formed. The electropolish and acid etch together removed a total of $26\ \mu\text{m}$ of nickel from the tubes, which had originally been about $360\ \mu\text{m}$ thick. This wall thickness is sufficient to support the NCDs against the pressure of the D_2O once deployed, but care had to be taken not to collapse the tubes when evacuating the NCDs prior to gas fill.

CVD nickel endcaps with a fused-silica high-voltage feed-through were laser welded into the ends of each NCD. The fused-silica feed-through extends out of the NCD

by 2.5 cm, allowing electrical contact between NCD anode wires to be made at their conducting copper tips. The silica also insulates the anode wire from the nickel bodies of the NCDs. The fused-silica feed-through is coated internally with pyrolytic graphite and extends 2.5 cm into the NCD, producing a multiplication-free region at either end of the NCD to reduce the effects of the distorted electric field there. One endcap on each NCD is slightly flared so it can slip over the straight end of another NCD, allowing structural attachment of the NCDs into strings. The connection between the endcaps of two NCDs is shown in Figure 3.1.

Each NCD was strung with a 50- μm -diameter copper anode wire. Copper is an excellent material for the anode, since it has good electrical properties and can be made to have very low levels of radioactive impurities, less than ten picograms per gram [52]. It was unsuitable for the NCD bodies because of the softness and porosity of electrodeposited copper, and due to the amounts of radioactive lead introduced with that much copper. The optimal anode wire thickness comes from the balance of several factors. A thinner anode wire allows for increased gas gain, but increases the risk of breakage. A thicker anode wire decreases space-charge effects, but requires higher voltage to achieve the same gas gain. Higher voltage is problematic because the rate of micro-discharge events increases with increased operating voltage.

At the upper end of each NCD, the endcap also has a copper tube through which the NCD was filled with 2.5 atmospheres of gas. The gas used is a 85:15 mixture (by pressure) of helium and carbon tetrafluoride (CF_4). Most of the NCDs were filled with ^3He for neutron capture, but about 8% of them were filled with ^4He instead for systematic studies. Addition of CF_4 provides several advantages over pure helium. The heavier CF_4 helps to minimize wall effects by providing additional stopping power without capturing neutrons. In addition, the CF_4 serves as a quenching gas so that if an atom in the gas is excited rather than ionized, the photon emitted when it decays




Figure 3.1: The endcap regions of two NCDs. The endcap, endcap wall, and everything contained within that region is often collectively referred to as the endcap. A coupler (described in Section 3.2) is used to connect the copper tips of the two fused-silica feed-throughs together to make the electrical connection.

back to its ground state can be absorbed in ionization, vibrational or rotational degrees of freedom by the CF_4 .

Gas purity is very important to reduce backgrounds from tritium contamination. The ^3He used for the NCDs was obtained from the decay of tritium, and prior to purification, had about 1 mCi/STP-l (4×10^7 Bq/STP-l) of tritium in the form of water vapor (HTO) and some free hydrogen (HT). This amount would produce 4.4×10^8 decays/s in a 2-meter NCD. While the β -decay of ^3H is quite low-energy (18.6 keV, with 6 keV deposited in the NCD gas on average), pile-up of tritium decays could produce events with enough energy to cause a background to the neutron-capture signal. The ^3He was purified to reduce the tritium level below 2.7 nCi/STP-l (100 Bq/STP-l). This limit represents a 1% probability that a tritium decay will occur in a 10- μs event integration time, making pile-up a negligible background.

Maintaining consistent gas mixture and pressure is crucial to achieving uniform gas gain throughout the NCD array. In order to reduce the tritium contamination of the ^3He and achieve uniform gas fills for the entire NCD array, a detailed gas fill procedure was used. Up to 7 NCDs could be filled at once on a manifold. For each set of NCDs, the appropriate amount of ^3He was introduced into holding tanks after passing through a cooled charcoal filter. The ^3He was then circulated through a SAES getter and a cooled charcoal trap for at least three hours. The proper amount of CF_4 was added to the holding tanks and the mixture was circulated through a different SAES getter for at least 15 minutes. The purification steps removed water vapor and free hydrogen, which are sources of tritium, as well as air and other contaminants. Finally the purified mixture was introduced into the NCDs.

After gas fill, but before the copper fill tube was sealed off, a low-energy spectrum was obtained from the NCD itself to measure the tritium contamination in the gas. Data were taken for 10 minutes at 2225 V from one NCD in the group just filled to

verify that the tritium content of the ^3He was below 2.7 nCi/STP-l. The gas gains of the filled NCDs were also measured before they were sealed to ensure uniformity of the gas fill. For ^3He NCDs an energy spectrum was taken with a ^{252}Cf neutron source to verify the gas gain. Data were taken at 1600 V for 10 minutes from all the NCDs as well as from a standard NCD with known gain. The positions of the 764-keV neutron-capture peaks on each NCD were recorded, then the NCDs were switched systematically to different electronics channels and another spectrum was taken. By comparing the newly-filled NCDs to the standard NCD and by comparing the data taken on the different electronics channels, it was possible to remove the effects of the electronics and determine the gas gain of each NCD. If the neutron peak position on an NCD was within 3% of the standard and the full-width at half-maximum (FWHM) resolution was better than 5%, then the NCD was declared to have passed the neutron source test.

The fill procedure for ^4He was different because it was undesirable to contaminate the holding tanks with ^4He and because tritium removal was unnecessary, so less purification was required. The ^4He was mixed with the CF_4 in the NCDs themselves, rather than in the holding tanks. The CF_4 was passed through the CF_4 SAES getter on its way to the tanks, but the 99.9999% purity ^4He was not purified at all. For the ^4He NCDs, the neutron source test was meaningless since no neutrons are captured, but it is still desirable to measure the gas gain and resolution of the NCD. To do this, an ^{241}Am gamma source was used and data were taken for 10 minutes at 1600 V with the amplifier gain raised by a factor of five. If the position of the Compton edge from the 59-keV gamma scattering in nickel indicated the correct gas gain and if the FWHM resolution was less than 16%, then the NCD passed. NCDs that passed the tritium test and either the neutron or ^{241}Am test had the copper fill tube sealed and were packed for shipment to the SNO laboratory.

3.2 NCD Storage and Handling

As the NCDs were completed in Seattle, they were transported to Sudbury and brought underground to the control room of the SNO lab for storage. Initially, commercial shipping companies were used to ship the NCDs from Seattle to Sudbury, but NCDs entrusted to a commercial company were broken. As a result, the NCDs were instead driven to Sudbury by UW graduate students and staff. Between April 1998 and November 2002, a total of nine trips were made to transport the 296 NCDs to Sudbury, with between seven and 63 NCDs in each load.

When the NCDs arrived underground, they were tested to make sure that they had survived the shipment. Each NCD was visually inspected for broken fused-silica feed-throughs or other obvious signs of damage. The resistance was measured to verify that the anode wire was still intact. Then a temporary cable was connected to the NCD and it was brought up to voltage. A neutron spectrum was taken with a ^{252}Cf neutron source. The presence of a neutron peak with good resolution and a reasonable rate and peak position indicated that the NCD was not leaking gas. NCDs that passed these tests were placed in a rack for storage. All the NCDs arrived safely in the SNO lab, but a few were broken during subsequent testing and installation, and had to be returned to Seattle for repair. The most common failure mode was breakage of the fused-silica feed-through, resulting in loss of the gas, although broken anode wires and pinhole leaks were other causes for repair.

General NCD-handling procedures focused on maintaining cleanliness. The NCDs were stored in nylon bags to prevent dust accumulation on the nickel. To minimize dust deposition, both ends of the bags were to remain sealed at all times, either with a heat-sealer or by folding the end of the bag and securing it with a binder clip for a temporary seal. If data were being taken from an NCD, the end of its bag was opened to allow the cable connection to be made, then the bag was clipped shut

around the cable. While most people were very diligent about sealing NCD bags, it was not uncommon to come across an NCD with its bag unsealed. Gloves were always worn when touching a bagged NCD, and if the bag was open, the NCD could only be touched with clean stainless steel tools, not with gloved hands.

An overhead cable tray allowed cables to be routed from the NCD electronics to the west end of the NCD rack, so data could be taken from the NCDs to characterize the array prior to deployment. Because the number of electronics channels was limited, and most of the temporary NCD cables were only long enough to reach to the west end of the rack, it was only possible to take data from a subset of the NCD array at a time. To take data from an NCD, a temporary cable connection was made using a cable connector, a CVD nickel tube about 5 cm long. The side of the connector that mated to the NCD was designed to mimic the endcap of another NCD. The other end of the cable connector was closed off with a copper plate to provide RF shielding and had a SHV (safe high voltage) connector in the center for the cable connection. A coupler was placed on the feed-through inside the connector to make the electrical connection, then the nickel body was slipped inside the flared end of the NCD. A CVD nickel sleeve was slipped over the joint and tightened with a hose clamp. Care was taken that the hose clamp only touched the nickel sleeve, not the NCD itself, so no contaminants could rub off onto the NCD.

It was possible to take data from more than one NCD on the same electronics channel by connecting multiple NCDs together. It was desirable to do this since data taken from multiple NCDs more closely approximated the final array configuration of three or four NCDs connected together in a string. In addition, this configuration took full advantage of the limited number of available electronics channels. Another advantage to taking data from multiple NCDs connected together was that it gave people practice at making connections between NCDs prior to the high-pressure environment

of deployment, although there was a risk of breaking a fused-silica feed-through while making a connection.

Three different types of couplers were used to take data from the NCDs in the control room, two of which were reused when the NCD array was deployed. Temporary copper couplers consisted of a spring of wire that slipped onto the cable connector or NCD feed-through. A copper plate on the end of the spring rested against the copper tip of the feed-through to make the electrical connection. Nickel couplers (shown in Figure 3.1) were used to connect individual NCDs into complete deployed strings. These couplers were twists of CVD nickel, with a loop at either end that could be opened wide enough to slip over the feed-through by pinching a set of ears on the loop using a pair of tweezers. Couplers with a resistance of 325Ω (to match the characteristic impedance of the NCD string to the NCD cable) were used to connect the uppermost NCD in each deployed string to its cable, as shown in Figure 3.2. These resistive couplers were made out of a ring of teflon with a very fine wire strung back and forth around one half and connected to CVD nickel loops that slipped over the NCD feed-through and the cable bell connector. Attaching a nickel coupler or a resistive coupler was quite challenging since tweezers had to be used at all times, any sideways force on the feed-through could break it, and there was not much room to work in the gap between endcaps.

3.3 Pre-deployment Welding

The NCDs arrived underground in sections that were 2, 2.5, or 3 meters long, since anything longer would not fit into the SNO transport rail cars. Installed, each NCD string is 9 to 11 meters long and consists of three or four of these sections. Each section is an independent gas volume, with electrical feed-throughs at each end to allow the sections to be connected together into strings. Electrical connections were




Figure 3.2: The cable bell and its attachment to the upper NCD in a string by means of a resistive coupler. Figure is from [47].

made between the sections in a string, then the nickel endcaps were welded together using a very clean laser welding process. To weld together two NCD sections, a 1024-nm Nd-YAG pulsed laser was directed at the overlap between the NCDs and rotated relative to the NCDs, creating a series of overlapping craters melted into the nickel that fused the two NCDs together. During welding, the NCDs were held in a stainless-steel weld fixture that was custom-built at the University of Washington. The weld fixture, shown in Figure 3.3, contained a filtration system that prevented nickel dust from escaping, and was hermetically sealed during welding to prevent the possibility of anybody being injured by the laser.

NCD deployment was done in two stages, pre-deployment and deployment, in order to minimize the time that the SNO detector was offline. During the pre-deployment welding phase, the NCDs were welded together into sections that did not exceed 5.5 m in length. Longer NCD sections could not be positioned vertically over the neck of the AV for deployment, due to limited overhead clearance above the neck of the vessel. At this time, the delay lines and the high-voltage/readout cables were welded on as well. The welding was done with the weld fixture in a horizontal configuration on a weld bench in a part of the SNO lab where it did not interfere with operation of the SNO detector. Pre-deployment welding was performed by teams of about four people, which consisted of a master welder, a quality-assurance record keeper, and one or more assistants. When everything went smoothly, a weld could be completed in about 45 minutes, but a failed weld or a poor fit between the endcaps could increase the required time considerably. Pre-deployment welding was performed from January 30 – February 18, 2003, at which point it was interrupted due to a failure of the laser used for the welding. The welding was finished from April 2 – 21, 2003, although some repairs were necessary during the summer of 2003.

Of the 196 required welds, 148 could be done during pre-deployment welding,




Figure 3.3: The fixture used for welding and leak testing the NCDs. Figure is from [53].

greatly reducing the SNO detector's down-time during deployment itself. Figure 3.4 shows which of the welds were performed during pre-deployment welding. Three different types of weld were done during pre-deployment welding: 40 delay-line welds, 68 NCD-to-NCD welds (36 in the lower half of the string and 32 in the upper half), and 40 cable welds. The nickel cable bells and delay-line housings were designed to connect to the endcaps of NCDs, so the welding process was nearly identical to welding two NCDs together. Because the lower end of the delay-line housing has a lip around it for attachment of the anchor balls, it could not be pulled through the weld

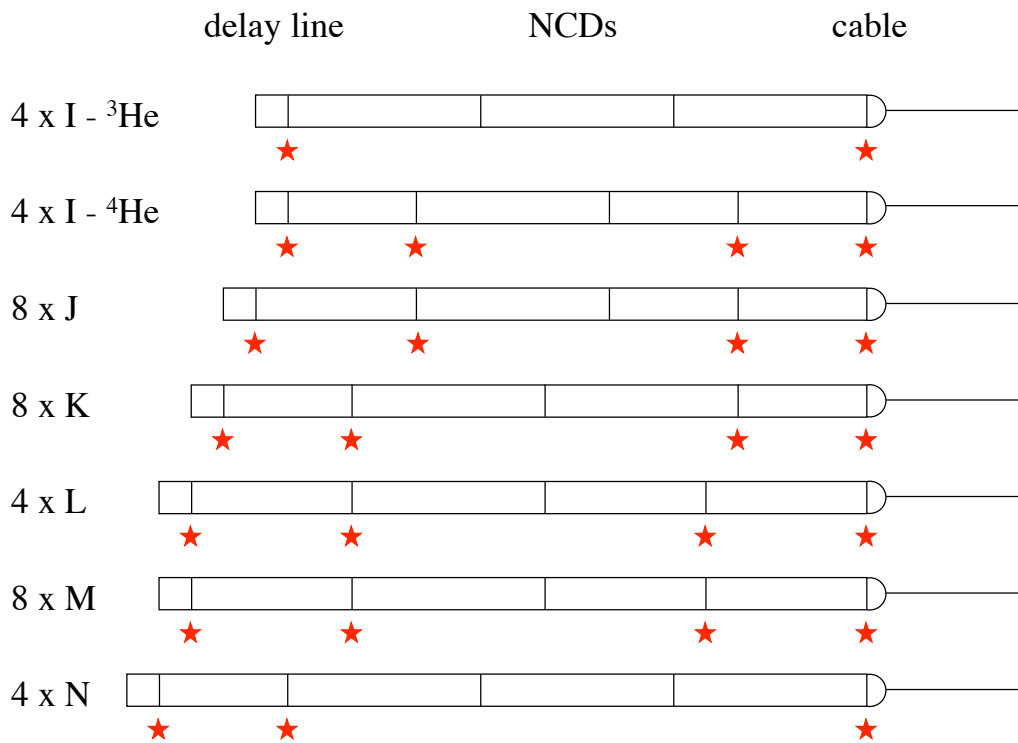


Figure 3.4: The configuration of each NCD string. The string type (which corresponds to its length), fill gas (^3He unless otherwise noted), and number of strings of this type are shown to the left of each string schematic. The welds marked with a star were performed during pre-deployment welding.

fixture, so a different weld fixture configuration was required to perform delay-line welds. Thus all of the delay-line welds were performed first, then the welder was reconfigured to perform the other welds. A delay line and its housing are shown in Figure 3.5.

One of the logistical complications of pre-deployment welding involved locating and organizing the NCDs that were to be welded into each string. Data were being taken from the NCD array in the control room up until just a few days prior to pre-deployment welding. During this time, the NCDs were not in the arrangement in which they would be deployed, and most of them were connected together with temporary NCD-to-NCD electrical connections. When it was time to weld a given NCD to another NCD, a cable, or a delay line, the NCD had to be located on the storage rack and brought to the weld bench. Before moving the NCD, its temporary connections had to be undone carefully, without breaking the fused-silica feed-throughs. Once welding of a given NCD segment was completed, it was brought back to the storage rack and put into a location that would make it easy to locate for deployment. This often involved shuffling around the other NCDs on the rack to make room for the segment that had just been welded.

For each weld, the NCDs to be welded together were first identified in the control room storage rack, disconnected from any other NCDs, and brought out to the welding corridor. The air quality was carefully monitored in the welding corridor throughout pre-deployment welding and measures were taken to minimize the dust to which the NCDs were exposed. The nylon bags were removed from the NCDs and they were placed on rollers that allowed the NCDs to slide along the weld bench and also rotate during welding. Time-domain reflectometry (TDR) measurements were made of the NCDs before and after they were connected to verify the integrity of the electrical connections. The electrical connection between the individual NCDs was made, then




Figure 3.5: The delay line connected to the bottom of a NCD string. Figure is from [47].

the entire segment was slid into the correct position in the weld fixture. The edge of the flared NCD endcap was located using a video camera in the weld fixture. The NCD segment was moved until crosshairs on the video screen were lined up next to the edge of the flare, so that the weld would be on the region of overlap between the two NCD pieces.

Just prior to welding, about 10 cm^3 of ^4He tracer gas was injected into the endcap region. Silicone cuffs were inflated around the NCDs to hold them in position. During pre-deployment welding, with the welder in the horizontal configuration, the NCD rotated and the weld fixture remained stationary. The peak laser power used in welding was 1.3 kW. The pulse rate was 3.0 Hz, each pulse was 8.5 ms long, and a total of 1065 pulses were used for one weld. During the first 15 pulses the laser power was ramping up and during the last 55 pulses the power was ramping down, thus these pulses were designed to overlap each other. Each weld took approximately six minutes to complete.

No ^3He can escape through a faulty weld between NCD sections, but heavy water could leak into the space between the NCDs, which contains the high-voltage connections. It was crucial to ensure that the weld is leak tight, so a series of three leak tests was performed after each weld. First, a ^4He sniffer was held near the weld joint and rotated for one complete revolution around the weld. Any change in the detected ^4He rate was an indication of a large leak. Leaks of this size were rare, but did occur occasionally when the nickel walls were too thin and the laser penetrated the wall completely, leaving a visible hole into the endcap region.

For the medium leak test, the NCD segment was moved until the weld was centered between the silicone leak-test cuffs, which were inflated. The leak-test volume was connected to the leak detector and evacuated. The fractional pressure difference from the ambient mine pressure was recorded as a measure of the vacuum in the leak-test

volume. This vacuum reading was nominally 130%, since the ambient mine pressure is about 30% above standard atmospheric pressure. Then the leak-test volume was opened to the region around the new weld. A dramatic pressure increase in the leak-test volume to a pressure-difference reading below 40% indicated a leak. If the pressure reading was acceptable, then the pressure was monitored for five minutes. If the pressure increased by more than 2% per minute or if the ^4He leak rate was significantly above background, then the weld failed the medium leak test. Interpretation of the results of this leak test could be complicated by the silicone cuffs sealing poorly on the nickel tubes, which allowed ambient air to seep into the leak-test volume even if the weld did not leak.

Finally, if the weld passed the medium leak test, then the ^4He leak rate from the weld region was monitored using the leak detector. A good weld had a leak rate of less than 10^{-8} mbar-l/s. If a weld failed a leak test, one of two actions was performed, depending on how the leak was detected and how large it was suspected to be. In most cases, the weld was repeated immediately by putting down another weld line just next to the failed one. The overlap on a typical NCD joint was large enough to allow three welds to be performed side-by-side with good overlap. For a large leak, particularly if it was found using the medium leak test method that involved pumping on the weld joint region, there was the possibility that none of the ^4He tracer gas was still in the weld joint region. Without tracer gas in the endcap region, leak testing the second weld would be ineffective, so simply redoing the weld was not an option. In these cases the weld was cut apart, the end was reflared, and the welding process was started all over again. This was a very time-consuming process, but certainly preferable to deploying a leaking NCD string.

After the leak check, a final TDR was taken of the welded segment. Then the segment was cleaned before being placed in a new nylon bag for storage until de-

ployment. The cleaning was done with a specially designed teflon vacuum head that encircled an NCD and slid along the length of the NCD. The vacuum head was attached to a clean-room vacuum and a compressed nitrogen line so that dry nitrogen gas was blown onto the NCD through the vacuum head, while the vacuum sucked up any dust that was dislodged. This process was designed to remove any dust that had accumulated on the NCD while it was exposed to air. Finally, the NCD was returned to the control room and placed in the correct location to be easily found for deployment.

Data were taken from each newly-welded segment to verify the electrical connection and to search for high-voltage discharging or other problems. The goal for the deployed NCD system is to have the MUX trigger rate below 0.3 Hz, to maximize livetime in the digitizing oscilloscopes, which take about 0.7 seconds to read each event. If discharging or other problems were evident in the control room, where elevated MUX thresholds were used to discriminate against the ambient noise, then the rates would likely be too high for production data-taking. Welded sections with high discharge rates were identified prior to the start of deployment, and these sections were repaired during the summer of 2003. During the repairs, two NCDs were replaced, six delay-line welds were redone, seven cable welds were redone, and nine NCD-to-NCD welds were redone.

In many cases, the discharges were caused by blow-through from the welding process. If the laser power was slightly too high or the nickel being welded was slightly too thin, some of the nickel from the inside wall could be vaporized by the welder. This could occur without the weld actually penetrating the inner wall, so the weld would still be leak-tight. However, the vaporized nickel dust could settle on high-voltage regions and create discharging when the NCD was brought up to operating voltage. This was especially likely in the case of delay-line welds, since the

weld is directly over the delay line, where the high-voltage wire is very close to the nickel wall of the delay-line housing. Other potentially faulty welds were the cable welds, since the cable bells had been made from scrap pieces of CVD nickel that were too thin to use for construction of actual NCDs.

The suspect welds on segments with high discharge rates were cut apart and the endcap regions were carefully cleaned. Data were taken from the NCDs with temporary connections to ensure that the problem had indeed been caused by the weld, not by the NCDs themselves. Then the weld was redone, with greater attention paid to the prevention of blow-through. The wall thicknesses of the weld regions were measured and an attempt was made to align the thickest part of one NCD with the thinnest part of the other. In addition, backing rings were introduced to prevent weld dust from settling on the couplers or other high-voltage regions. These were thin CVD nickel rings that went just inside the straight end being welded, effectively adding another layer of nickel to the weld region. This greatly reduced the chance of blow-through reaching the high-voltage region. These refined welding procedures were adopted for deployment.

3.4 Deployment Hardware

A significant amount of specialized equipment was required to deploy the NCD array into the SNO detector. Figure 3.6 shows the clean room above the SNO detector (known as the DCR), the AV, and the NCD deployment hardware. For the duration of deployment, this hardware was installed in the DCR or in the AV itself. All deployment hardware was made of clean materials, such as stainless steel or acrylic, that are suitable for use in the DCR or insertion into the SNO detector. Equipment that contacted the heavy water was tested prior to use to ensure that no contaminants could leach into the heavy water. Most of the deployment equipment was custom-




Figure 3.6: The AV with the equipment used for NCD deployment. Figure is from [53].

designed and built at the University of Washington.

When the SNO detector was built, acrylic anchor points for the NCDs were installed on the bottom of the AV on a one-meter grid. These anchor points are sockets that are concave downwards and that open towards the center of the AV. At the bottom of each NCD string is an anchor, consisting of two acrylic anchor balls. The lower anchor ball hooks into an anchor point and is held in place by the buoyancy of the NCD string pulling the ball upwards in the socket. Once anchored, the NCD string, which has positive buoyancy, naturally floats vertically in the heavy water.

A special stainless steel deployment plate, shown in Figure 3.7, was installed over the neck of the AV for NCD deployment. The welding fixture was attached to the deployment plate in a vertical configuration on an arm that allowed it to swing out of the way when it was not in use. The global-view and neck-view cameras and their controls were mounted to the deployment plate. There were four view ports that could have gloves fitted to them if needed for manipulating NCD parts during deployment. The haul-down crankset, the boathook to manipulate the NCD cables, and a feed-through for the ROV umbilical were also mounted on the deployment plate. The NCD insertion port was at the bottom of a well about 0.5 m deep that provided extra overhead clearance above the insertion port so longer NCDs could be turned vertically for deployment. This allowed longer segments to be welded in pre-deployment welding and minimized the number of welds required during deployment.

The NCD segments were positioned vertically over the deployment plate with the help of a pulley system known as the skyhook. A pulley was attached to the ceiling of the DCR and a long cable was strung through it. At one end of the cable was a cradle that attached to a clamp that could be placed around an NCD. Counterweights could be attached to the other end of the cable to provide upward force on the NCD in the skyhook, preventing it from falling into the AV. The skyhook allowed the NCD




Figure 3.7: The plate installed over the neck of the AV for NCD deployment. Figure is from [53].

to be positioned vertically for insertion into the AV in a controlled manner without the deployment crew having to use ladders.

NCDs were lowered into the AV using a haul-down system. This consisted of two plastic floats on a Vectran polymer fiber line. The lower float was anchored to a special anchor point on the bottom of the AV directly below the insertion port. The upper float, known as the shuttle float, was raised and lowered on the Vectran lines by turning the haul-down crank, which was attached to the deployment plate. The lower anchor ball of an NCD could be inserted into a socket on the shuttle float, and

would remain engaged in the socket as long as upward force was applied to the anchor ball.

Each NCD string was installed in its anchor point using a remotely-operated submersible vehicle (ROV), shown in Figure 3.8. The ROV was custom-built by Deep Ocean Engineering for clean operation in heavy water, and was extensively leach tested prior to deployment. An acrylic plate with a socket is mounted on the front of the ROV to hold the upper anchor ball on an NCD string. The ROV has two ballast tanks that could be purged with nitrogen gas or filled with 20 – 25 kg of D₂O to make the ROV heavy enough to support the buoyancy of an NCD string, which varies by about 25% depending on the string length. The ROV is maneuvered by three pairs of thrusters controlled by two joysticks on a hand-held control unit. The hand-held control unit also adjusts the trim of the ROV, which is the amount of automatic thrust provided by the vertical thrusters, and controls the solenoid valve that opens the ballast tanks.

There were two cameras that helped the ROV operator to see where the ROV was going. One was mounted on the ROV itself and provided close-up views of the ROV socket and the anchor balls and anchor points. Another, global-view, camera was mounted on the deployment plate and extended down into the neck of the AV, providing views of the entire interior of the AV. The displays for these two cameras were side-by-side on a control console that also contained some other controls for the ROV system. The ROV pilot would typically sit in front of the control console with the hand-held control unit on his or her lap, making it simple to access all the important information and controls for the ROV.

Prior to deployment, all ROV pilots participated in a week-long training session at Los Alamos National Laboratory. The training was carried out with the ROV immersed in a pool about six meters on a side that contained NCD anchor points into




Figure 3.8: The remotely operated vehicle (ROV) used to deploy the NCDs.

which the pilots deployed mock NCD strings. During the training, the pilots became adept at maneuvering the ROV and practiced appropriate responses to potential emergency situations.

The final pieces of deployment hardware aided in securing the NCD cables to

prevent them from getting tangled in the manipulator lines used for positioning calibration sources in the SNO detector. The NCD cables attach to the bottom and top of the AV neck to help route them up through the neck in an orderly fashion. Each cable goes straight up from the top of the NCD string until it contacts the upper hemisphere of the AV, then it follows the curve of the AV to the bottom of the neck. A rigid plastic quarter-circle known as a cable bend attaches the cable to the cable ring, an acrylic ring mounted just inside the bottom of the AV neck. The cable bend forces the cable to make a 90° bend and lie flat against the surface of the AV. At the top of the neck each cable is secured in a metal ring known as the cable comb to ensure that there is no slack cable in the neck. Plastic cable bullets attached to each NCD cable slip into indentations on the top of the cable comb, allowing for upward tension to be applied on the cable in the neck. Above the cable comb the end of the cable emerges sideways through a hermetic feed-through to the NCD preamplifiers, which are located in a ring above the top of the neck. During deployment a long pole with pegs on the bottom, known as the boat hook, was used to attach the cable bend to the cable ring. A third camera, the neck-view camera, provided views of the neck in order to aid in this process.

3.5 Deployment Procedures

NCD deployment began December 2, 2003 and the last string was deployed on February 12, 2004. However, repairs were required after that, so deployment was not completed until April 21, 2004. Deployment was performed by teams of up to seven people, which consisted of a team leader, a master welder, a quality-assurance record keeper, a ROV pilot, a DAQ expert, and one or more assistants. While each team member had some definite roles, team members would often trade tasks and do whatever needed the most attention at a given moment. Deployment of an NCD string was

quite complicated and typically took several hours, although problems at any stage in the procedure could increase the time required to perform a deployment. Detailed deployment procedures (23 pages) and quality assurance sheets (seven pages) were adhered to and are summarized here.

The NCD string to be deployed was first located in the control room storage rack. Temporary connections between unwelded sections were disconnected. The nickel coupler was left on the lower segment for reuse when the string was welded. Each section was identified by the name written on the nylon bag and scribed on the end of the NCD to verify that the correct sections were being deployed. The nylon bags were wiped with deionized water to remove any dust prior to the NCDs being moved into the DCR and placed on a storage rack or rail on the south DCR wall. Once the NCD segments were in the DCR, the ends of the nylon bags were opened, the anchor assembly was attached to the bottom of the NCD string, the thicknesses of all the weld regions were measured to avoid blow-through, and the endcaps were cleaned with dry nitrogen gas. A CVD nickel backing ring was inserted inside the bottom of the straight end of the upper NCD segment to prevent nickel dust from getting onto the center conductor of the NCD string and causing high-voltage discharge problems. TDR measurements gave a good measure of the length of the segment and the position of any connections within the segment, such as previously welded NCD sections and cable and delay-line welds.

The skyhook was used to position the NCD vertically for insertion into the deployment port. Even with the skyhook, four people were required to safely raise a long NCD segment into the vertical orientation, since care had to be taken not to hit the deployment plate or the edges or top of the cupola. Once the NCD segment was hanging vertically from the skyhook above the deployment plate, the lower end of the NCD was positioned in the insertion port. The insertion port was only opened if the

particle count rate in the DCR was below 500 counts per minute. Care was taken that there were two safety clamps on the NCD at all times to ensure that it could not fall into the AV. The lower anchor ball was inserted into the shuttle float using one of the glove ports on the deployment plate. The skyhook was still attached to the upper end of the NCD segment and additional counterweight was added to ensure that an upward force was being applied to the NCD to keep the anchor ball engaged in the shuttle float. Then the NCD segment was lowered into the D₂O using the haul-down. Once the lower NCD segment had been lowered to the welding position, the insertion port clamp was closed and the skyhook was removed. Another TDR measurement was made to ensure the segment had not been damaged during positioning.

Next, the upper NCD segment was positioned for welding. If there were three segments, as in the case of the N strings and the ³He I strings, then the following procedure was performed first for the middle segment, then for the upper segment. The NCD segment was raised vertically using the skyhook until it was above the top of the weld fixture, which was rotated away from the insertion port at this time. The segment was carefully lowered into the weld fixture until the bottom edge was just below the bottom of the weld fixture and two silicone cuffs were inflated to hold the NCD in place. The weld fixture was then raised until it could clear the lower NCD segment and swung into the weld position above the insertion port. Then the weld fixture was lowered until the straight end at the bottom of the upper NCD segment slid into the flared top of the lower segment.

The weld joint was located on the video screen and the weld fixture was rotated 360° to verify that the weld joint was in the correct position. The engagement and spacing of the two NCD segments were checked. Then the spacing selector knob on the weld fixture was placed in the connect position and tweezers were used to connect the coupler. A TDR measurement was made of the connected string. Just prior to

welding, the spacing selector was moved to the inject position and about 10 cm³ of ⁴He tracer gas was injected into the gap using a syringe, then the selector was put in the weld position. The weld fixture access door was closed, flow of the weld cover gas was started, and welding proceeded for about six minutes. During deployment welding with the weld fixture in the vertical configuration, the weld fixture rotated around the stationary NCD. After the weld was completed, the leak check procedure described in Section 3.3 was followed to verify the weld.

Once the weld passed the leak check, the cuffs were deflated and the string was hauled down until only a few inches protruded through the insertion port, which was then clamped shut. The weld fixture was raised to its maximum elevation and swung out of the way. The cable was fed through the weld fixture, and the bag was removed carefully to avoid getting the cable dirty. Another TDR was performed on the welded string to ensure that the electrical continuity was still intact. The cable bend and cable bullet were installed on the cable in the correct positions, which had been calculated beforehand. The locations were measured from the dry end of the cable using a tape measure with the cable stretched along the storage rail. The U-shaped opening on the cable bend was oriented upwards (towards the dry end of the cable) in order to engage it in the cable ring. Four small cable ties were used to attach the cable bend to the cable. The rounded portion of the cable bullet was pointed downwards in order to engage it in the cable comb and provide upward tension on the cable in the neck. A small hose clamp was used to attach the cable bullet to the cable. A cable tether was attached to the SHV connector on the end of the cable to prevent the end of the cable from falling into the heavy water, and the other end of the tether was tied off securely. Then the string was hauled all the way down to the bottom of the AV.

The buoyancy of the ROV was adjusted if necessary so the ROV could lift off the bottom of the AV using the thrusters with the trim set at 50%. The ROV was flown

to the upper anchor ball on the NCD string and engaged the ball. The ROV trim was set to zero and downward and forward thrust was applied to keep the anchor ball engaged. The gas valve on the control console was set to flood and the solenoid valve on the hand-held control was opened. Flooding the ballast tanks took about three minutes. During the first minute of flooding, the downward thrust being applied by the ROV was slowly decreased, since the increased weight of the ROV was enough to keep the anchor ball engaged. Once the lower anchor ball dropped down in the shuttle float anchor socket, the ROV was supporting the full buoyancy of the NCD string. The solenoid valve was closed and the ROV was backed away from the haul-down and allowed to settle gently on the bottom of the AV near the haul-down. If necessary, the ballast tanks were purged a bit until the ROV was just slightly negatively buoyant.

At this point the neutron source tests were performed. The NCD readout cable was connected to the test preamplifier and a 1850 V potential was applied. The shuttle float, containing an embedded AmBe neutron source, was brought up to the top of the neck, then lowered a predetermined number of turns of the haul-down crank to place it near the middle of the topmost string segment. Data were taken for five minutes to verify that the NCD produced a good neutron spectrum in the Shaper/ADCs. Then the shuttle float was lowered to the center of the next NCD in the string and the data-taking was repeated. In this way, each of the three or four NCDs in the string was verified. Once the neutron source tests were completed, the voltage was returned to zero, the preamplifier was disconnected, and the tether was reconnected to the NCD cable.

Now the string was ready to be flown to its anchor point. The optimal route to the appropriate anchor point was discussed by the deployment team, taking care to avoid tangling the ROV umbilical or wrapping the NCD cable around anything. The ROV was flown to the anchor point, with one team member moving and zooming the

global view camera to continually guide the ROV in the appropriate direction. Other anchor points along the way acted as signposts. Care was taken to avoid running into the AV with the ROV or flying too high and hitting the top of the NCD on the AV. Once the anchor point was reached, the CVD nickel nameplate mounted on the anchor point was checked to verify its identity.

Figure 3.9 shows the ROV inserting an NCD string into an anchor point. To do this, the ROV was aligned so it was pointing radially outwards facing the anchor point. The lower anchor ball was inserted in the anchor point and the trim was set




Figure 3.9: The ROV inserting an NCD into an anchor point. Another anchor point is shown in the background. Figure is from [53].

all the way positive to engage the anchor ball. The control console valve was turned to purge and the solenoid valve was opened. After about three minutes of purging, bubbles emerged from the ballast tanks that were visible on the ROV camera. At around this time the ROV rose up off the upper anchor ball, indicating that the tanks were fully purged. The ROV was flown up and back to disengage from the upper anchor ball and the NCD was left in the anchor point. The solenoid valve was closed and the control console gas valve was switched off. The ROV was flown back to the bottom center of the AV near the haul-down, taking care not to tangle the umbilical, then the thruster power was turned off.

The last step of NCD deployment was installing the cable in the correct spots on the cable ring and cable comb. This was often the most time-consuming and difficult part of deploying an NCD, since visibility was poor and maneuvering the boat-hook was not easy. The global view and neck view cameras were adjusted to locate the cable bend. One of the boathook pegs was engaged in the hole in the cable bend, with the curved side of the cable bend near the boathook pole. This procedure was usually quite difficult and the boathook operator was often assisted by one or more spotters and possibly someone manipulating the cable through a glove port to position the cable bend more optimally. To install the cable bend, the slot on the cable bend was engaged around the cable ring. The cable bend had to be on the correct part of the cable ring, above the NCD string and just below the designated spot on the cable comb. Then a glove port was used to place the cable bullet in its spot in the cable comb. Tension was applied to the cable to lift the bullet above the cable comb, but too much tension would slide the cable through the bullet and the cable bend, causing the top of the NCD string to tilt inwards.

The NCD preamplifiers are located in a ring above the top of the AV neck in an enclosure known as the preamp raceway. The dry end of each NCD cable emerges

into the preamp raceway through a compression fitting. Once the cable bend and cable bullet of a newly deployed NCD string were secured, the cap of the compression fitting was removed. A thin metal rod with a hook at the end was inserted through the compression fitting to hook the cable tether. The tether was used to pull the end of the cable out through the compression fitting just enough to attach the SHV connector to the preamplifier. The compression-fitting cap was installed, clamping the cable into position, and the preamplifier was attached. The high-voltage supply, signal output, and pulser input cables for that preamplifier were also attached and the connection was tested by taking data. Finally, heat shrink was applied over the connection to ensure that no light would enter the AV through the elastomer seal of the compression fitting.

3.6 Difficulties during Deployment

The positions of the 40 deployed NCD strings are shown in Figure 3.10. In general, the NCDs were deployed from the outermost ring of I strings in to the innermost N strings. However, due to high-voltage discharging on some of the strings that required repair, the order was not strictly followed. Several of the strings were removed from the AV after deployment, by dismounting the string and bringing it back to the neck then cutting one or more welds, so that welds could be redone or NCDs replaced in an attempt to mitigate the discharging problems. The original deployment order was intended to reduce the chance of striking an already deployed NCD string with the ROV. During the repairs it became clear that the ROV pilots were skilled enough to fly between deployed strings without bumping into them.

When deployment began, NCD sections with dramatic discharge problems had already been repaired. However, more problems arose during deployment. Several strings developed discharge problems after deployment that had not been evident in

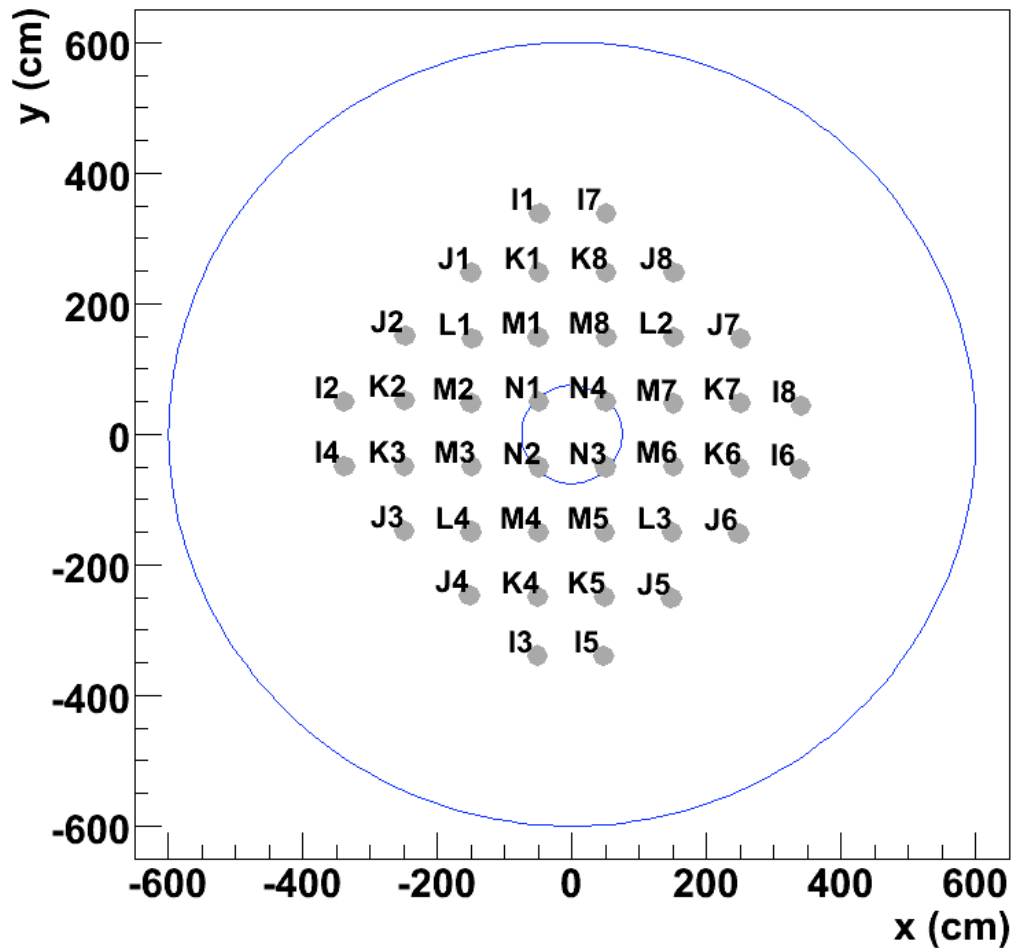


Figure 3.10: The positions of each NCD string in the 12-meter diameter AV, with the 1-meter diameter neck also shown. The string names consist of a letter/number designation with the letter denoting strings of the same length and distance from the center of the AV, and the number increasing (in most cases) in a counterclockwise fashion. Strings I2, I3, I6, and I7 are filled with ^4He . Figure is from [54].

the control room despite the precautions taken to avoid blow-through on the deployment welds. One hypothesis is that when the NCDs were turned from their horizontal storage configuration to the vertical deployed configuration, nickel dust from previous

welds was dislodged and settled on high-voltage regions. In addition, the discharge behavior of some strings is somewhat intermittent, so problems sometimes crop up seemingly randomly. Also, the noise level in the control room was quite high, so some problems could not be seen until the NCD strings were in a much quieter environment after deployment. Because of discharges and a variety of other problems, deployment took several months longer than originally planned.

These deployment difficulties are summarized in Table 3.1, which includes deployment problems, but not performance problems that arose after deployment. The performance of the NCD array after deployment is discussed in Section 4.5. Table 3.1 includes the string number associated with each string name. To simplify analysis, each NCD string is assigned a number from 0 to 39. The numbers are based on the position of each string's preamplifier, which are numbered clockwise around the top of the AV neck, thus N4 is string 0 and N1 is string 39. In general, string numbers rather than string names are used when analyzing NCD data.

The first problem discovered during deployment was the realization that the cable bends slid along the cable as they were being attached to the cable ring. This meant that the carefully measured position of the cable bend along the cable length was no longer correct after the slide. To correct this, the deployment procedure was modified to assume some sliding and the cable bend position was adjusted accordingly. Having slightly too much cable below the cable bend was not a serious problem, but having too little was, because it could cause the NCD to tilt inwards. This potential problem was identified after the first two strings (J3 and J4) were deployed, so first J3 then J4 were removed from their anchor points, placed on the shuttle float and hauled up until their cable bends could be repositioned, and then they were redeployed. After deployment was finished, laser range measurements of the NCD strings were performed to verify that they were all vertical.

Table 3.1: A summary of the problems that arose with each string during NCD deployment. The NCD strings are listed by both name and string number.

String Name	String Number	Problem
I1	36	—
I2	30	Deployed three times due to discharges
I3	20	Deployed into I4 anchor point
I4	29	Deployed into I3 anchor point
I5	19	—
I6	10	—
I7	3	Deployed into I8 anchor point; anchor assembly 3.8 cm longer
I8	9	Deployed into I7 anchor point
J1	35	—
J2	32	Redeployed due to leak that caused current draw
J3	26	Cable bend slipped
J4	23	Cable bend slipped
J5	16	—
J6	13	—
J7	6	Deployed five times due to discharges
J8	4	—
K1	37	—
K2	31	—
K3	28	—
K4	21	—
K5	18	Leak into endcap region
K6	11	—
K7	8	—
K8	2	Anchor assembly 2.5 cm longer
L1	34	—
L2	5	Deployed into L4 anchor point
L3	15	Deployed three times due to discharges
L4	24	Deployed into L2 anchor point
M1	38	—
M2	33	—
M3	27	—
M4	22	—
M5	17	—
M6	12	Deployed twice due to discharges
M7	7	Lost control of ROV during deployment
M8	1	—
N1	39	Deployed three times due to discharges
N2	25	Anchor assembly 1.3 cm longer
N3	14	—
N4	0	—

During the deployment of the I8 string it became clear that it was impossible to insert the anchor ball into the anchor point, due to the presence of pipes on the floor of the AV that are used for circulating the heavy water. The pipes are not exactly centered between the I1 and I8 anchor points, so it was possible to deploy the I1 string, but the corner of the ROV kept bumping into the pipes when trying to deploy I8. The easiest solution was to modify the anchor ball assembly on the NCD string that would be deployed in the I8 anchor point so there was a larger space between the upper and lower anchor balls. Thus the ROV could fly above the pipes but the anchor ball could still reach the anchor point. Modifying the anchor ball assembly on I8 would have required cutting two welds to get the entire string out of the AV. While a decision was made about what to do, it was decided to temporarily deploy I8 in a different anchor point, thus I8 was placed in the I7 anchor point.

When the difficulty with the I8 anchor point was discovered, all three of the other ^3He I-strings had already been deployed, but none of the ^4He I-strings had yet been deployed. The decision was made to switch I7 and I8 and to switch I3 and I4. The result was that the ^4He strings were on the anchor points just to the right of each cardinal direction, rather than being on both the anchor points in the west and east. This provided the advantage of greater symmetry in neutron capture in the detector, but decreased the redundancy if one ^4He string failed, since they were further apart in this configuration. The repositioned strings are still referred to by their names, not by the names of the anchor points in which they are located.

The modified anchor assembly on the I7 string (in the I8 anchor point) is 3.8 cm longer than a typical anchor assembly. The anchor assembly on the K8 string was also lengthened by 2.5 cm, in order to prevent a similar problem when deploying it next to the pipes. The N2 string was deployed last and the haul-down system was directly in front of its anchor point, so the N2 string had to be deployed with the ROV behind

the anchor point, rather than in front of the anchor point. In order to facilitate this unusual deployment, the anchor assembly on the N2 string was lengthened by 1.3 cm.

Another switch was performed near the end of deployment because some strings were undergoing repairs and it was desirable to minimize the amount of maneuvering around deployed strings that was necessary for the last deployments. The L2 string was deployed into the L4 spot, since the rest of the strings in that quadrant were ready to deploy, but L4 and strings in L2's quadrant were still being repaired. This allowed the deployment of the last five strings to proceed faster than if L2 and L4 had not been switched.

A near-disaster occurred during the deployment of the M7 string. The string was on the ROV and the neutron calibration had just finished. The ROV pilot was purging the ballast tanks a little so the ROV could lift off and fly the string to its anchor point. More purging than usual was necessary, and the ROV pilot was starting to wonder if something was amiss. The upward thrusters were briefly engaged to see if the ROV would lift off the floor of the AV, and the ROV began to ascend rapidly. The ROV pilot responded quickly and correctly by flooding the ballast tanks, setting the trim to full negative, switching the ROV into boost mode and thrusting full down, and moving to the center of the AV so the NCD string would go up in the neck rather than striking the upper hemisphere of the AV. The top meter or so of the NCD string emerged above the water surface in the neck before the ROV became heavy enough to compensate for the buoyancy. Then the ROV began to sink rapidly, striking the lower haul-down float and dislodging it from its anchor point.

The likely cause of the incident was that the corner of the ROV was stuck under the haul-down float, preventing it from lifting off until it had been significantly overpurged. Then when the thrusters were applied, the ROV dislodged from the haul-down float and ascended rapidly. Quick maneuvering prevented the NCD string

from striking the AV or the bottom of the deployment plate, but the haul-down float was dislodged during the rapid descent of the ROV. Several days of deployment were lost reinstalling the haul-down float and rethinking deployment procedures to ensure that a similar incident could not occur again. Initial tests of the M7 string indicated that it had not been hurt by the incident. Later it was observed that the rate of M7 discharge events above 9 MeV that pass the currently existing data-cleaning cuts is about 0.3 events per day, about as high as the rest of the array combined. This may or may not be related to some damage sustained during the incident.

The problems that added the most time to deployment were high-voltage discharges and other problems requiring repairs on deployed strings. Seven of the 40 strings were deployed more than once because welds were redone or NCDs replaced to try to alleviate discharge problems. Sometimes the entire string was removed from the AV, but other times just one of the upper welds was cut and redone, if, based on pulse reflection timing, the location of the discharge was thought to be well known. The worst string was J7, which was deployed a total of five times. L3, N1, and I2 were deployed three times each and M6 was deployed twice.

Leaking welds also added time to NCD deployment. J2 was removed from the AV and later redeployed because a leaking weld allowed heavy water into the endcap region and caused the string to draw current when ramped up to voltage. K5 turned out to have a small leak in the cable-middle NCD that allowed gas from the endcap region to enter the NCD, poisoning it and causing the gas gain to drift slowly downwards over the space of many months. A series of ^3He leak tests were performed on the relevant section with the upper half of the NCD pulled up through the insertion port on the deployment plate. No ^3He was seen, so it was concluded that the leak was only into the endcap region and K5 was redeployed without being repaired.

Despite these complications, the NCD array was deployed successfully, and on

April 21, 2004 the ROV was removed from the heavy water. Figure 3.11 shows the view down the neck of the acrylic vessel before the ROV was removed from the heavy water. The four innermost NCD strings (the N strings) are visible below the neck. All the NCD cables can be seen around the edges of the neck. The pipes at the top of Figure 3.11 are part of the water circulation system. The global-view camera (above the ROV) and the haul-down system (lower left) have not yet been removed from the AV. Some of the PMTs are clearly visible in the background.




Figure 3.11: The ROV, viewed through the neck of the AV. Also visible are four NCD strings, the NCD cables, the global-view camera and the haul-down system, as well as PMTs in the background.

Chapter 4

NCD CHARACTERIZATION AND OPTIMIZATION

4.1 *Cool-down Data*

After construction in Seattle, the NCDs were transported to Sudbury and brought underground as quickly as possible to maximize the amount of time for cool-down of cosmogenically activated ^{56}Co . The ^{56}Co is produced by spallation neutron interactions with ^{58}Ni and ^{60}Ni in the NCDs, and it decays via electron capture and subsequent gamma emission with a 77-day half-life. Approximately 33% of ^{56}Co decays result in a gamma above the photodisintegration threshold of deuterium [55], thus ^{56}Co constitutes a neutron-producing background in SNO. If the NCD array had been on surface long enough to reach ^{56}Co saturation and were then to have been deployed into SNO with no underground cool-down time, five photodisintegration neutrons per day would be produced by ^{56}Co decays [49].

The final shipment of NCDs arrived underground in Sudbury in November 2002. Production data-taking in the NCD phase did not begin until November 2004, so the entire NCD array was underground for at least two years prior to deployment. If the entire array had been brought underground in the final shipment, then the initial ^{56}Co neutron production rate in SNO at the start of the NCD phase would have been 7×10^{-3} neutrons per day. This constitutes an upper limit to the actual rate, since most of the array was brought underground well before November 2002. The average time underground prior to the commencement of deployment activities was almost four years, or about 18 half-lives of ^{56}Co .

The cool-down phase provided time for data to be taken from the NCDs in the underground environment, which allowed for characterization of NCD backgrounds in the absence of cosmic rays. In addition, the NCDs and their read-out systems were tested in the active mining environment where production data would actually be taken. This proved to be useful for testing electronics and cables and debugging the data acquisition system. The cool-down data were dominated by neutrons produced by (α,n) reactions in the rock walls of the mine, and by alphas from intrinsic NCD radioactivity. Due to the noisy environment of the control room, the MUX thresholds could not be set as low as was desired for production running. Because of this, the digitized data from the cool-down phase were not representative of production data. Nonetheless, the Shaper/ADC data from the cool-down phase were analyzed and several major tests of the NCD system were carried out during the cool-down phase.

During April and May of 2002 each of the nearly 300 NCDs that were underground at that point was tested with a ^{252}Cf neutron source to verify that no damage had occurred since the NCD was initially brought underground. The array had been handled during its time underground and breakage was possible. In addition, pinhole leaks could have developed that might not have been noticed during the initial test but that might result in an observable gas gain shift after several months or years. When these tests were performed, construction of the NCD electronics had not yet been completed, so there were only 30 available electronics channels. The available cables weren't long enough to reach to NCDs in the far end of the cool-down rack, so the NCDs to be tested in each batch had to be moved to the west end of the rack.

After moving the 30 NCDs to be tested and carefully attaching cable connectors and SHV cables to each of them, the voltages were raised to 1850 V and a ^{252}Cf neutron source was placed nearby. With only five minutes of data it was possible to see a distinct neutron peak in the Shaper/ADC energy histogram, and thirty minutes

of data gave enough statistics to allow a fit to the neutron peak. The preamplifiers and other electronics were consistent to better than 5% from channel to channel, so it was possible to verify that the NCD gas gain was approximately correct simply by requiring the neutron peak position to fall near bin 150 in the Shaper/ADC. Only three NCDs out of the entire array failed the test. These were B6-250A-1, which had a broken anode wire, and K6-200A-1 and M7-250-3, which had pinhole leaks. Over the course of the testing period, K6-200A-1 was observed to be losing gas. The first time K6-200A-1 was tested, the gas gain was too high, but within a week it had dropped too low and then the signal disappeared completely. This is the behavior that is expected from a leaking NCD, since the gas gain varies inversely with the gas pressure until the pressure gets extremely low.

Oscillatory noise pickup in the high-voltage/readout cables was also studied using cool-down data. Many oscillatory noise events integrate to zero in the Shaper/ADCs, but asymmetric events can integrate to very large values. Initially, cool-down data were being acquired using single shielded cables (96% braid) and the oscillatory noise rate in the Shaper/ADCs was 0.2 events/cable-day. Double shielded cables (98% braid and 96% braid) were built and the noise rate was reduced to less than 0.01 events/cable-day in the Shaper/ADCs, which improved the quality of the data taken during the cool-down phase. The NCD cables that were deployed are only shielded, with a single woven shield, not braid. During tests performed prior to deployment, the oscillatory noise rate on the real cables was 0.06 events/cable-day in the Shaper/ADCs. After deployment, efforts were made to reduce the noise pickup in the NCD cables, and oscillatory noise events are no longer seen in the Shaper/ADCs. It is also easy to discriminate cable noise from physics events using the digitized data, so oscillatory noise is not a problem for the deployed array.

The data rates in the NCD array were monitored carefully during the cool-down

phase. At one point the rate nearly tripled, from around 7 Hz total Shaper/ADC rate in the entire NCD array to almost 20 Hz. This sudden increase was correlated to a calibration of the SNO detector that was being done at that time using gammas from ^{16}N beta decay. The ^{16}N was produced by $^{16}\text{O}(n,p)^{16}\text{N}$ using ^{16}O in the form of CO_2 and neutrons from a deuterium-tritium (D-T) generator. The D-T generator was located in a pit in the SNO lab about 15 meters from the NCD rack and was shielded by a meter of concrete. When the power in the generator was turned down, the rate seen by the NCDs decreased within a matter of seconds, demonstrating that neutrons from the generator were being detected by the NCD array. The total neutron flux seen by the NCDs corresponded to a dosage of 4×10^{-4} Sievert/year just outside the D-T generator shielding, comparable to a typical surface background dosage [56].

From May 2002 until NCD deployment began in December 2003, all changes to the configuration of the NCDs in the control room were recorded. Most of the NCD array was connected together into full strings at this time. The configuration changed often, since it was desirable to make optimal use of the available electronics. As more electronics channels were built and installed underground, more of the NCD array came online. Even with the initial limited number of available electronics channels it was possible to collect high-statistics data from the entire array by rotating which NCDs were taking data every few months. These data were used to study internal backgrounds in the NCDs.

4.2 NCD Selection

The original size of the NCD array was based on the solar neutrino flux deduced from previous experiments, rather than the full SSM flux observed via the NC channel in the SNO experiment. Hence, a smaller, optimized array could be used in the NCD phase that would block less of the light going to the PMTs. A significant NC

measurement could still be made in a reasonable amount of time, and a better CC spectrum could be extracted than with the larger NCD array. The original NCD array was designed to consist of 96 strings, ranging from 4 m to 11 m in length. Each string would contain two, three, or four individual NCDs, so the original array would have required 300 NCDs. The array actually deployed in the NCD phase contains 40 NCD strings, ranging from 9 m to 11 m in length, and required 156 NCDs.

It was desirable to deploy the NCDs with the lowest individual levels of intrinsic radioactivity, in order to reduce the backgrounds introduced into SNO. These were chosen by analyzing Shaper/ADC data from the cool-down phase. Data from three different datasets were taken between June and December 2002 and almost all of the NCDs were represented in at least one of these datasets. Each dataset contained between one week and three months of data.

One problem with this data was high rates of discharge events on some channels, caused by poor connections between NCDs or between an NCD and a cable connector. In addition, only Shaper/ADC data were used in the analysis and no data cleaning was attempted, so discharge events and other types of instrumental noise events were counted along with the physics events in the Shaper/ADCs. Because of these issues, some NCDs were rejected by the selection analysis due to instrumental backgrounds, rather than high alpha rates. The use of three different datasets with the NCDs in different configurations allowed for some of these instrumental problems to be detected because they were only present in one of the datasets. However, some of the NCDs were only taking data in one of the three datasets, so this cross check could not be used to reject instrumental backgrounds for all of the NCDs. Since most of the data were taken with the NCDs connected into strings, rates on a given string were assumed to be representative of each NCD in the string.

The criteria used in NCD selection included good gas gain and resolution, and

low levels of alpha activity, particularly from uranium and thorium chain decays. For each NCD, the number of events in four different energy ranges were counted: neutron captures up to 1 MeV, wall effect alphas from 1 to 4.5 MeV, ^{210}Po peak alphas from 4.5 to 6 MeV, and uranium and thorium chain alphas above 6 MeV. In the neutron-capture energy range, the neutron peak was fit with a Gaussian and the mean and width of the fit were also recorded. For an NCD to be selected for deployment, the neutron peak position was required to be close to the average of the entire array. Additional selection criteria were good neutron peak resolutions and low numbers of counts in the ranges above 1 MeV. In particular, the range above 6 MeV was required to have very low counts, since this range contains only uranium and thorium chain alphas, not ^{210}Po alphas. Since uranium and thorium chain decay products can photodisintegrate deuterium, they are much more problematic than ^{210}Po .

An independent analysis that only used the third of the three datasets was performed by another graduate student, using slightly different cut values for the various criteria [57]. The NCDs that passed both the selection analyses were chosen for deployment. In order to have enough selected NCDs, some were chosen that were not taking data during the third dataset, but that passed the selection in one or both of the other two datasets. After choosing enough NCDs to make the reduced array, about six spares of each length were also selected that could be used if an NCD were damaged during welding or deployment.

Once the NCDs were selected, their string assignments were specified by trying to match the gains as closely as possible. To a large extent, NCDs that had originally been intended to go into the same string were left together, since they had typically been filled with gas from the same batch and so had very similar gas gains. When that was not possible, the neutron peak position was used as a measure of the gas gain. This was not completely accurate, since variations in preamplifiers could introduce

gain variations that were not due to the gas fill.

This optimized NCD configuration¹ is shown in Table 4.1. The NCDs in each position are listed, from the lowest (anchor) position up through the uppermost (cable) position. The table also shows the alpha rates in events per day measured during the cool-down phase for each string in this configuration. These rates were obtained by summing the rates determined in the NCD selection analysis for each NCD in the string. If a NCD acquired data in more than one dataset, all datasets that did not have high discharge rates were averaged. For comparison, the alpha rates measured for the original 96-string configuration are also shown in the table. In the original configuration, each string consisted of the NCDs whose identifiers began with the string name. Since not all of these NCDs were built and because a few did not have data taken from them during cool-down, some of the original strings do not have alpha rates listed.

The alpha rate of the strings in the optimized configuration is lower than that of the original 96-string configuration, except for strings 2, 3, and 18. String 3 is a ^4He string that was originally intended to be a ^3He string, and strings 2 and 18 are both ^3He strings that were originally intended to be ^4He strings. Because of the shortage of ^4He NCDs and because the optimized array required different numbers and lengths of ^4He NCDs than the original array, a direct comparison of the alpha rates of strings that changed fill gas between the original and optimized arrays is not meaningful.

4.3 Deployment Source Tests

During deployment, after each NCD string was welded and lowered to the bottom of the AV, a neutron source test was performed to ensure that each section of the

¹As discussed in Section 4.4, the final deployed array is somewhat different than this optimized array, due to repairs.

Table 4.1: The NCDs in each string in the optimized array, along with the alpha rates in events per day measured in the cool-down phase. For comparison, the alpha rates of the original 96-string configuration are also shown, where available. A * marks the ^4He strings.

String	#	Anchor	Anchor-middle	Cable-middle	Cable	Opt. Rate	Orig. Rate
I1	36	I7-300A-1	I7-300-2	—	I7-300C-3	2.01	6.87
I2*	30	A4-200A-1	K8-300-3	A1-200C-2	A4-200C-2	5.9	10.6
I3*	20	A5-200A-1	G6-300C-3	A1-200A-1	A5-200C-2	6.16	16.94
I4	29	H4-300A-1	H4-300-2	—	H4-300C-3	2.57	—
I5	19	I5-300A-1	I5-300-2	—	I5-300C-3	4.31	4.31
I6*	10	K8-200A-1	K8-300-2	K5-200A-1	K8-200C-4	7.7	17.38
I7*	3	K1-200A-1	K1-300-3	A8-200A-1	A8-200C-2	637.71	2.01
I8	9	I8-300A-1	I8-300-2	—	I8-300C-3	5.33	5.33
J1	35	F2-250A-1	F2-300-2	F7-200C-3	F2-200C-3	1.24	200.71
J2	32	J2-250A-1	J2-300-2	J2-200-3	J2-200C-4	3.35	3.35
J3	26	J3-250A-1	E4-300-2	E4-200A-1	E4-200C-3	3.28	4.35
J4	23	J4-250A-1	J4-300-2	J4-200-3	J4-200C-4	2.12	2.12
J5	16	F3-250A-1	F3-300-2	D7-200-2	F3-200C-3	1.44	26.46
J6	13	J6-250A-1	J6-300-2	J6-200-3	J6-200C-4	2.43	2.43
J7	6	J7-250A-1	J7-300-2	J7-200-3	J7-200C-4	2.12	2.12
J8	4	F6-250A-1	F6-300-2	F4-200C-3	F6-200C-3	3.21	55.44
K1	37	A3-200A-1	G1-300-2	G1-300C-3	A3-200C-2	2.59	970.47
K2	31	K2-200A-1	K2-300-2	K2-300-3	K2-200C-4	3.22	3.22
K3	28	A7-200A-1	C1-300A-1	C1-300C-2	A7-200C-2	1.35	12.69
K4	21	E5-200A-1	C3-300A-1	C3-300C-2	E5-200C-3	2.81	—
K5	18	E6-200A-1	E6-300-2	F7-300-2	E6-200C-3	2.84	1.31
K6	11	K6-200A-1	K6-300-2	K6-300-3	K6-200C-4	13.46	13.46
K7	8	K7-200A-1	H2-300-2	H2-300A-1	K7-200C-4	8.18	21.16
K8	2	D6-200-2	G6-300-2	C2-300A-1	D6-200C-3	15.64	8.6
L1	34	B1-250A-1	B1-300C-2	G6-250A-1	F4-250A-1	1.5	—
L2	5	B3-250A-1	B3-300C-2	G1-250A-1	B6-250A-1	1.98	14.34
L3	15	L3-250A-1	L3-300-2	L3-250-3	L3-250C-4	0.85	0.85
L4	24	B4-250A-1	B4-300C-2	G3-250A-1	B8-250A-1	2.26	—
M1	38	M1-250A-1	M1-300-2	M1-250-3	M1-250C-4	1.29	1.29
M2	33	B5-250A-1	B5-300C-2	G7-250A-1	D6-250A-1	2.07	4.48
M3	27	M3-250A-1	M3-300-2	M3-250-3	M3-250C-4	1.11	1.11
M4	22	M4-250A-1	M4-300-2	M4-250-3	M4-250C-4	1.26	1.26
M5	17	M5-250A-1	M5-300-2	M5-250-3	M5-250C-4	0.96	0.96
M6	12	B7-250A-1	B7-300C-2	D7-250A-1	F7-250A-1	2.01	17444.87
M7	7	M7-250A-1	M7-300-2	M7-250-3	M7-250C-4	1.32	1.32
M8	1	M8-250A-1	M8-300-2	M8-250-3	M8-250C-4	6.94	6.94
N1	39	F8-200C-3	N1-300-2	N1-300-3	N1-300C-4	1.44	26.53
N2	25	N2-200A-1	N2-300-2	N2-300-3	N2-300C-4	1.9	1.9
N3	14	N3-200A-1	N3-300-2	N3-300-3	N3-300C-4	1.19	1.19
N4	0	D7-200C-3	H1-300-2	H1-300A-1	H1-300C-3	3	17.22

string was still operating correctly. During the test, the ROV was parked, holding the NCD string, about 30 cm from the haul-down. The shuttle float, with an encapsulated AmBe source, was raised to the top of the neck and then lowered a predetermined amount to place it approximately in the center of the uppermost section of the string. After five minutes of data-taking, a clear neutron spectrum was visible in the Shaper/ADCs. The centroid and width of the neutron peak were calculated using ORCA's built-in centroid-calculating routine. Data-taking was stopped and the shuttle float was lowered to the center of the next NCD section and five minutes of data were taken there. The process was repeated until a neutron spectrum had been recorded for each section of the NCD string.

These neutron source tests served several important purposes. First, they verified that each NCD section was electrically connected and recording data. With the source only about 30 cm from the center of one NCD section, the neutron capture was dominated by the nearest NCD section, so a dead section would have resulted in a drastically-decreased data rate. Second, the neutron source tests allowed clean neutron spectra to be collected from each of the NCD sections with the source nearby. After installation of the NCD array, the closest a source can approach most NCDs is 0.5 meter, and for some strings the source is even further away. The relative gains of each NCD section in the array were well determined by the deployment neutron source tests, since the same preamplifier and electronics channel was used for each test. In addition, the gain measured during the deployment source tests could be compared to data taken prior to deployment in order to verify that no damage that might cause a gain shift had occurred during deployment. The deployment source tests also served to establish a baseline gain for each NCD section to which subsequent data could be compared.

Because an AmBe source emits gammas as well as neutrons, calibration of the

^4He NCDs was possible as well. In an AmBe source, the decay of ^{241}Am produces an alpha, which has energies of 5.49 MeV, 5.44 MeV, or 5.39 MeV in the dominant branches. The alpha then initiates an (α, n) reaction with ^9Be , producing ^{12}C and a neutron. Typically, AmBe sources are used as neutron sources, but ^{241}Am is a gamma source as well. In 36% of the ^{241}Am decays the alpha is accompanied by a 59-keV gamma. By raising the operating voltage of the NCD to 2150 V and raising the gain of the Shaper/ADCs as well, it was possible to see the Compton edge from electrons in the NCD nickel scattered by these gammas. This was done for the ^4He strings to provide a calibration of their gains and resolutions.

The centroid and width of the neutron peak for each ^3He NCD are shown in Figure 4.1. The centroid and width of the Compton edge for each ^4He NCD are also shown. The average width (1σ) of the entire array from these tests is 8.6 Shaper/ADC bins. For a given fill gas, the centroids are fairly consistent, and the consistency within a given NCD string is typically quite good. One notable exception is caused by a leaking NCD in the cable-middle position of string 18. Gas from the endcap region is slowly poisoning the NCD and lowering the gas gain. Since deployment, the neutron peak position on that section has drifted down from around bin 100 to around bin 30 over one year. Two other strings that have slight gain mismatches are 20 and 35. String 20 is a ^4He string, and good gain matching was more difficult to accomplish for the ^4He strings than for the ^3He strings due to the limited number of ^4He NCDs and the difficulty calibrating them. The gain matching on string 35 was not very good when its optimized configuration was determined, perhaps because of bad luck with the electronics gains of the components that went into it.

The conditions under which the neutron calibrations were performed were not completely stable. The NCD electronics were being optimized during deployment, so occasional changes were made that impacted the data. The most dramatic of these

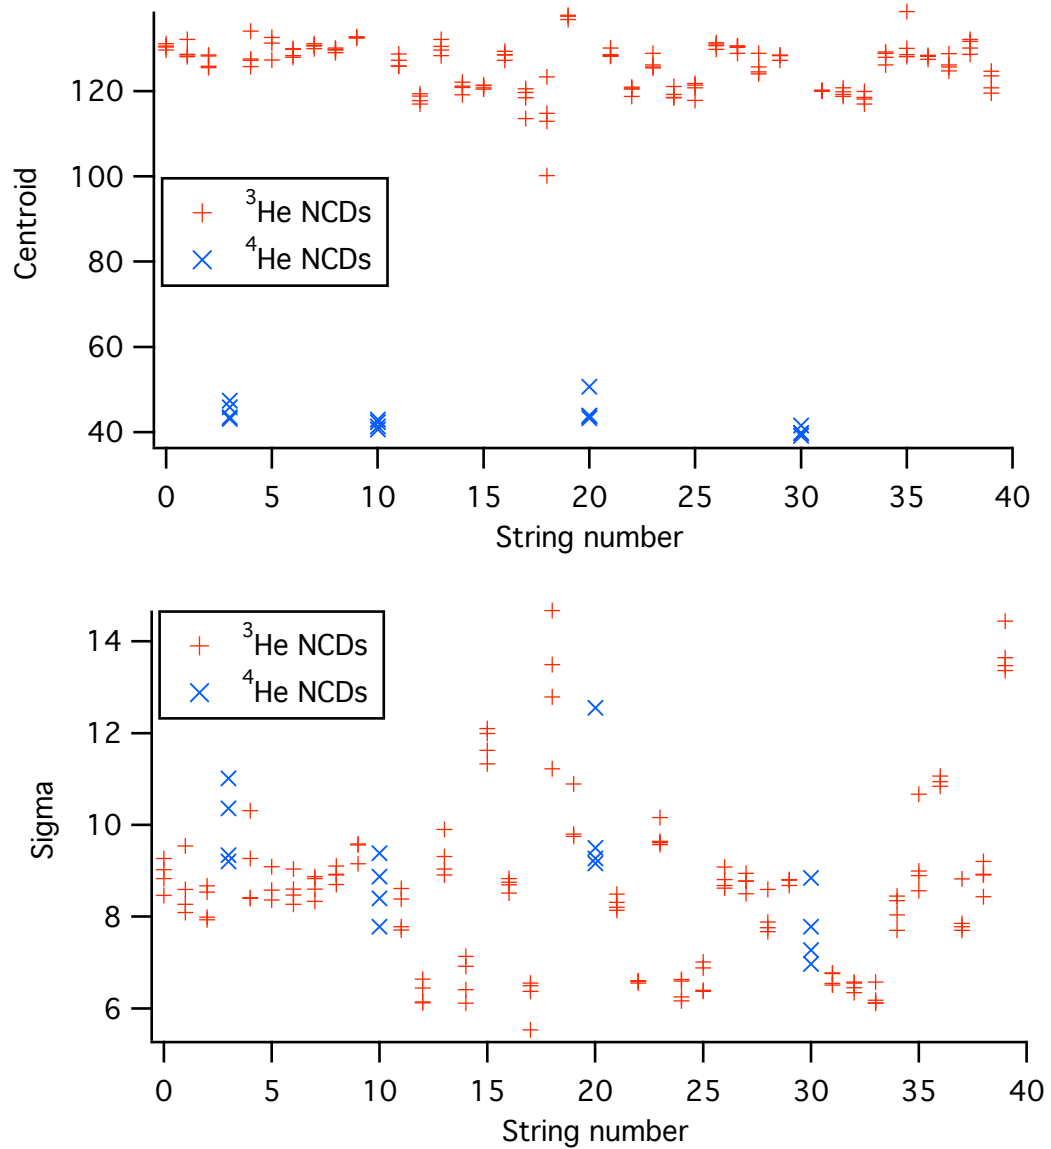


Figure 4.1: The centroid (upper panel) and 1σ width (lower panel) of the neutron peak (for ^3He NCDs) or the Compton edge (for ^4He NCDs) obtained from calibration with an AmBe source during the deployment.

changes was the replacement of the high-voltage digital-to-analog converter (DAC) board on January 20, 2004. With the new high-voltage DAC board, the measured gains were slightly lower than prior to the replacement. The neutron calibration data on strings 0, 6, 12, 14, 15, 17, 18, 22, 24, 25, 31, 32, 33, and 39 were taken after the DAC replacement. As evident in Figure 4.1, all of these strings except for 0 and 6 are slightly lower than the rest of the strings in the array.

The operating voltage of the NCD array was changed from 1850 V to 1950 V on April 1, 2004, near the end of deployment. This increase in voltage increased the gas gain by about a factor of two, but the Shaper/ADC gain was decreased to make the neutron peak fall in approximately the same place. The amplitude of the smallest neutron events was increased, allowing them to be digitized with somewhat higher MUX thresholds, which decreased deadtime because fewer noise events were being digitized. The disadvantage of raising the voltage was that the rate of high-voltage discharging tends to increase at higher voltage. Since noise triggers were dominating the trigger rate, and since discharges are relatively easy to remove based on pulse shape characteristics, increasing the voltage improved the overall data quality. Another disadvantage of higher voltage is that space-charge effects are enhanced, which tends to worsen the detector resolution. The neutron calibration data on strings 18 and 39 were taken with 1950 V, unlike the rest of the strings. The widths in Figure 4.1 on strings 18 and 39 are larger than the rest of the array, as expected from space-charge effects.

4.4 Quality Control and Repairs

Starting during pre-deployment welding, data were taken from welded NCD sections in the control room to monitor their performance. Many of the sections behaved as desired, producing energy spectra with clean neutron peaks from ambient neutrons

from the rock walls, low rates of higher-energy alpha events, and little low-energy noise. Some NCD sections, however, had high-voltage discharge problems introduced during the welding process. Repairing these problems caused significant delays during deployment, but allowed the NCD analysis experts to become very familiar with different types of NCD pathologies and their likely causes.

Very high rates of discharging were evident in the overall trigger rate and many types of discharging events produced distinctive spectra in the Shaper/ADCs. However, it was difficult to detect low-rate discharging behavior that was observable only in the digitized data. The analysis method that evolved in response to this problem was to hand-scan the digitized scope traces from a suspect NCD section. Data taken overnight were usually scanned first thing in the morning, prior to that day's welding. Typical data rates on a single NCD section produced several hundred to a few thousand events that had to be scanned. The graduate students and postdocs doing the hand scanning soon became quite adept at identifying neutrons, alphas, and discharge events very quickly, and could typically scan two or three events per second.

The MUX thresholds used in the control room were higher than those used for the deployed array, so many discharges did not trigger the control-room system but were problematic after a string was deployed. Because of this, and perhaps also because some discharging was intermittent or caused when the NCD was turned vertically during deployment, some discharging strings were impossible to identify until they were deployed. Thus it was sometimes necessary to undeploy a string and carry out repairs on it to eliminate discharging behavior.

By examining the pulse shape of a digitized discharge event, it was sometimes possible to make an educated guess as to the location of the discharge. Discharge events are typically quite narrow, which allows a clean reflection signal to be seen

in many cases. By measuring the timing between the initial and reflected pulse, it was often possible to determine which weld in the NCD string was responsible for the discharge. This sometimes allowed repairs to be performed without cutting a good weld and fully removing a deployed string from the AV, if the bad weld was at the cable bell, for example. Sometimes, however, this position identification did not work because the reflections were either absent or misleading.

During pre-deployment welding and deployment, many welds were cut and redone in an attempt to reduce discharge rates in the NCD array. In some cases NCDs had to be replaced, either because the discharges were suspected to be inside the NCD or because the weld had been cut too many times and there was not enough endcap wall extending beyond the tip of the fused-silica feed-through to safely connect the NCD to another NCD without risking breakage of the feed-through. Another reason NCDs were replaced during pre-deployment and deployment was due to breakage of the feed-through. In all, 27 NCDs in 14 different strings were replaced between the start of pre-deployment welding and the end of deployment. The final deployed array configuration is shown in Table 4.2, along with the impact these replacements had on the alpha rates in each string based on cool-down data.

4.5 Pathological Strings

Despite all of the careful repair work, several strings continue to have discharge problems during production data-taking in the NCD phase. Table 4.3 summarizes the behavior of strings in the NCD array with discharging problems that are apparent during data-taking. These are the easiest problems to identify, since they occur at relatively high rates, and data-cleaning cuts have already been developed that remove these discharges when they appear in the digitized data path. The digitization system can only handle rates of up to 0.3 Hz without the deadtime becoming appreciable

Table 4.2: The NCDs in each string in the final array, along with the alpha rates in events per day measured in the cool-down phase. The difference between the rate on each string in the final array and in the optimized array is also shown. A * marks the ^4He strings.

String	#	Anchor	Anchor-middle	Cable-middle	Cable	Final Rate	Rate Diff.
I1	36	I7-300A-1	I7-300-2	—	I7-300C-3	2.01	0
I2*	30	A4-200A-1	S4-272A-1	S4-227C-2	A4-200C-2	6.3	0.4
I3*	20	A5-200A-1	G6-300C-3	A1-200A-1	A5-200C-2	6.16	0
I4	29	H4-300A-1	H4-300-2	—	H4-300C-3	2.57	0
I5	19	I5-300A-1	I5-300-2	—	I5-300C-3	4.31	0
I6*	10	K8-200A-1	K8-300-2	K5-200A-1	K8-200C-4	7.7	0
I7*	3	K1-200A-1	K1-300-3	A8-200A-1	A8-200C-2	637.71	0
I8	9	I8-300A-1	I8-300-2	—	I8-300C-3	5.33	0
J1	35	F2-250A-1	F2-300-2	F7-200C-3	F2-200C-3	1.24	0
J2	32	J2-250A-1	B6-300C-2	K1-200C-4	J2-200C-4	2.21	-1.15
J3	26	J3-250A-1	E4-300-2	E4-200A-1	E4-200C-3	3.28	0
J4	23	J4-250A-1	J4-300-2	J4-200-3	J4-200C-4	2.12	0
J5	16	B2-250A-1	F3-300-2	D7-200-2	F3-200C-3	1.37	-0.07
J6	13	J6-250A-1	J6-300-2	J6-200-3	J6-200C-4	2.43	0
J7	6	D4-200C-3	E7-300-2	D1-200-2	B6-250A-1	2.48	0.36
J8	4	F6-250A-1	F6-300-2	A2-200C-2	F6-200C-3	3.08	-0.13
K1	37	A3-200A-1	G1-300-2	G1-300C-3	A3-200C-2	2.59	0
K2	31	K2-200A-1	K2-300-2	K2-300-3	K2-200C-4	3.22	0
K3	28	A7-200A-1	C1-300A-1	C1-300C-2	A7-200C-2	1.35	0
K4	21	E5-200A-1	C3-300A-1	C3-300C-2	E5-200C-3	2.81	0
K5	18	E6-200A-1	E6-300-2	F7-300-2	E6-200C-3	2.84	0
K6	11	K6-200A-1	N4-300C-4	N4-300-3	K6-200C-4	15.14	1.68
K7	8	K7-200A-1	H2-300-2	H2-300A-1	K7-200C-4	8.18	0
K8	2	D6-200-2	G7-300C-3	C2-300A-1	D6-200C-3	16.4	0.77
L1	34	B1-250A-1	B1-300C-2	G6-250A-1	F4-250A-1	1.5	0
L2	5	D5-250A-1	B2-300C-2	G1-250A-1	M2-250C-4	2.29	0.31
L3	15	G4-250A-1	H3-300C-3	L3-250-3	L3-250C-4	1.48	0.63
L4	24	B4-250A-1	B4-300C-2	G3-250A-1	B8-250A-1	2.26	0
M1	38	M1-250A-1	M1-300-2	M1-250-3	M1-250C-4	1.29	0
M2	33	B5-250A-1	B5-300C-2	G7-250A-1	M2-250A-1	2.36	0.29
M3	27	M3-250A-1	M3-300-2	M3-250-3	M3-250C-4	1.11	0
M4	22	M4-250A-1	M4-300-2	M4-250-3	M4-250C-4	1.26	0
M5	17	M6-250C-4	F4-300-2	M5-250-3	M5-250C-4	2.13	1.17
M6	12	D4-250A-1	B7-300C-2	D7-250A-1	F7-250A-1	2.79	0.78
M7	7	M7-250A-1	M7-300-2	M7-250-3	M7-250C-4	1.32	0
M8	1	M8-250A-1	M8-300-2	M8-250-3	M8-250C-4	6.94	0
N1	39	A2-200A-1	I3-300A-1	I3-300C-3	N1-300C-4	8.36	6.93
N2	25	N2-200A-1	N2-300-2	N2-300-3	N2-300C-4	1.9	0
N3	14	N3-200A-1	N3-300-2	N3-300-3	N3-300C-4	1.19	0
N4	0	D7-200-2	H1-300-2	H1-300A-1	H1-300C-3	3	0.01

(above a few percent), so discharging on a single string becomes a deadtime problem when the rate is above 0.1 Hz. One tactic that is effective for at least a few months is to temporarily raise the voltage on a discharging string to 2150 V for a day or two, which seems to burn off whatever is causing the discharges. Despite this treatment, there are times when the MUX rate on a string is too high and its MUX threshold has to be raised to bring the deadtime down to acceptable levels.

Table 4.3: The NCD strings that have discharging problems at high enough rates that they are noticeable during data-taking. Data cleaning cuts have been developed to remove these events.

String	#	Shaper/ADC rate	MUX rate	Stability	Deadtime
I3	20	~ 1 Hz	$\ll 0.1$ Hz	variable	no
I7	3	$\ll 0.1$ Hz	$\ll 0.1$ Hz	during blasting	no
L1	34	up to 0.4 Hz	up to 0.4 Hz	variable	sometimes
M3	27	~ 0.1 Hz	~ 0.1 Hz	stable	no
M8	1	~ 0.5 Hz	0	few month period	no
N1	39	up to 0.2 Hz	up to 0.2 Hz	occasional	sometimes
N4	0	up to 25 Hz	up to 0.2 Hz	variable	sometimes

There are also some strings with discharge problems at much lower rates (a few events per day), so they are not noticeable during data-taking or from a cursory look at the data. These problems are only apparent when cleaned data is carefully analyzed and the rate or energy distribution on a particular string is found to be different from what is expected. Table 4.4 summarizes these problems. Once the development of data-cleaning cuts for the NCD phase is complete, these events will presumably be removed from the analyzed datasets. The most problematic of these strings is J3, which has a class of events that fall into the neutron-capture energy region and look

somewhat similar to neutron-capture events. The ion tail is less apparent than it should be on a physics event and there are too many of them to be neutrons, but it is not yet clear how they will be removed.

Table 4.4: The NCD strings that have discharging problems at lower rates that are only noticeable because they pass currently existing data-cleaning cuts.

String	#	Energy	Description
J1	35	overflow bin	large discharges
J3	26	neutron region	look similar to neutrons
J5	16	near neutron region	look like discharges
M7	7	across entire spectrum	look like discharges

4.6 MUX Threshold Optimization

After the NCD array was deployed, several months were spent commissioning the SNO detector with the NCDs installed. Part of the commissioning involved setting thresholds to optimize data quality in the NCD phase. The NCD MUX thresholds determine the minimum amplitude of a pulse that will trigger the MUX and in turn be digitized by one of the two oscilloscopes (assuming they are not both busy). Since scope data is required for an analyzable solar neutrino candidate event, optimizing the MUX thresholds is quite important. The goal is to set the threshold such that the trigger rate is as low as possible but the fraction of neutron captures detected is as high as possible. If the thresholds are too low, then unacceptable levels of scope deadtime are introduced during neutrino data-taking. Scope deadtime causes the random loss of a fraction of all analyzable neutron-capture events. If the thresholds are too high, then the lowest-amplitude neutron-capture events will not be recorded.

These low-amplitude events are the neutron captures that are easiest to tell apart from alphas, due to their pulse shape characteristics. Preliminary studies have shown that about half of all neutron-capture events are distinguishable from alphas by pulse shape characteristics, so losing two neutron-capture events due to scope deadtime is roughly equivalent to losing one due to the MUX threshold being too high.

A simple calculation gives the approximate scope deadtime expected for a given MUX rate. By pulsing the University of Washington test system at 100 Hz for 300 seconds, 428 scope events were recorded on a single oscilloscope, indicating a maximum digitization rate of 1.4 Hz, or 0.7 seconds per sample. With two scopes, as in the actual DAQ system, the deadtime can be calculated as the probability of getting three MUX events within the 0.7 seconds it takes for the first scope to be read out. For a given average rate R and per-event deadtime D (in this case 0.7 seconds), the rate of missed events due to the deadtime is approximately R^3D^2 . For a MUX rate of 0.2 Hz, this gives 0.0039 events missed per second, so 2.0% of the events are missed. For a 0.3 Hz rate, 4.4% of the events are missed, and for 0.4 Hz 7.8% are missed.

As part of optimizing the MUX thresholds it was crucial to know what threshold triggers on all neutron-capture events and what fraction are missed for a given threshold above that. A threshold test was performed on string 14 on September 29 – 30, 2004 using the weak AmBe neutron source (the neutron-capture rate from this source is 2 Hz in the NCD array when the source is in the center of the AV). The MUX thresholds on string 14 ranged from $0 \times 3F$ to $0 \times 4E$ in hex (63 to 78 in decimal). The MUX thresholds were set to FF (turning off that channel) for the rest of the NCD array, except string 39, whose threshold remained the same as a cross-check. The quantity of interest was the threshold cut fraction,

$$Q = (S - M/L)/S, \tag{4.1}$$

where S is the number of Shaper/ADC neutron-capture events (defined by bins 30 –

145), M is the number of those events with correlated MUX events, and L is the MUX livetime fraction (measured with the MUX livetime counters). Therefore, Q is the fraction of neutron-capture events that are missed in the MUX due to the threshold. The Shaper/ADC livetime is 99.95% at a data rate of 2 Hz, so this deadtime can be neglected. The uncertainty on Q is given by the statistics in the events missed (due to threshold and to deadtime), and can be written $\pm\sqrt{S - M}/SL$. Figure 4.2 shows the measured threshold cut fraction Q and its associated uncertainty for each threshold. For thresholds below 0×44 (decimal 68), the threshold cut fraction is zero within uncertainties, indicating that it is sufficiently low to trigger on all neutron-capture events.

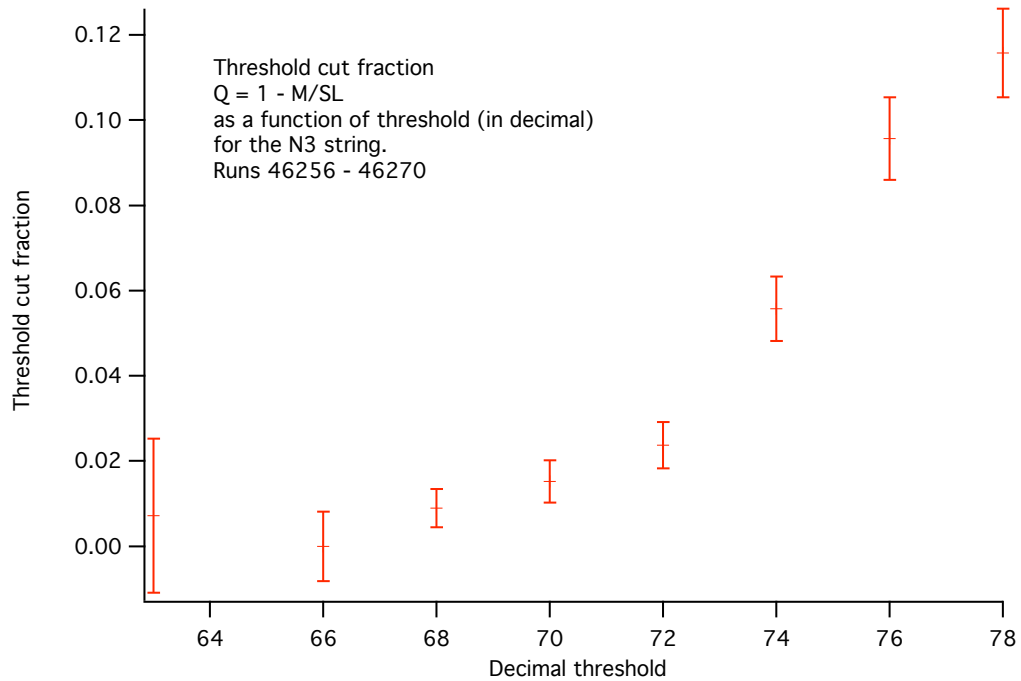


Figure 4.2: The threshold cut fraction on string 14 at different MUX thresholds.

The uniformity of the MUX thresholds across the array can be tested by pulsing

each string with square waves of decreasing amplitude until the channel no longer triggers. If all channels stop triggering at the same amplitude, then the thresholds are set uniformly. Channels that stop triggering at larger amplitudes should have their thresholds lowered from the setting when the test was performed, and ones that stop triggering at smaller amplitudes should have their thresholds raised. Figure 4.3 shows the results of this test with the “standard noisy” thresholds used prior to optimization [54]. The amplitudes are in uncalibrated units, since NCD electronics gains are not accounted for. Most channels stopped triggering at amplitudes between 17.1 and 19.5 (the dashed lines in Figure 4.3). Strings 12, 20, 28, and 36 were lower than average because they had their thresholds intentionally lowered to aid in noise studies. Strings 4, 19, and 38 were higher than average, indicating that their thresholds should be lowered.

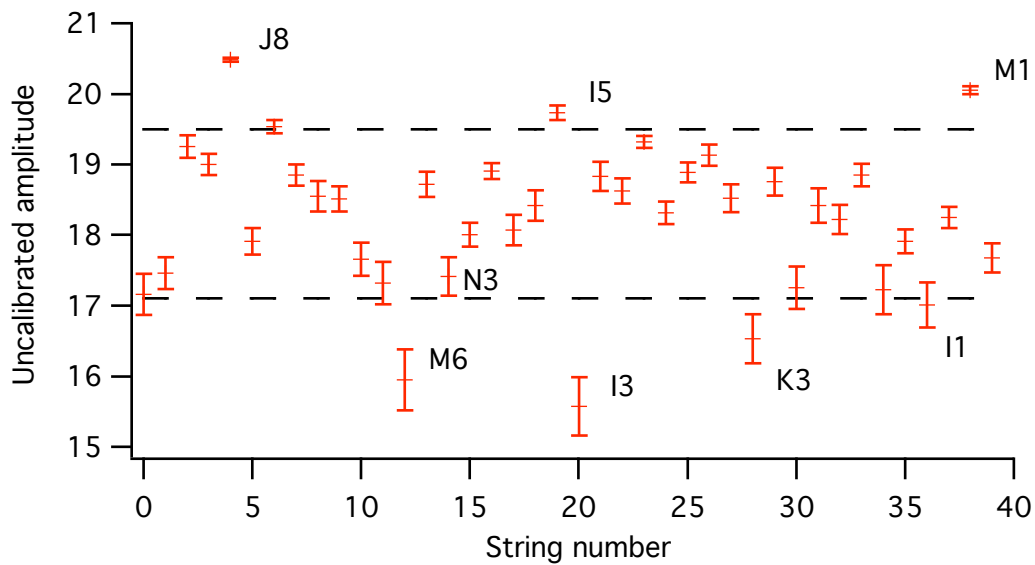


Figure 4.3: The amplitude at which each string stopped triggering, with the MUX thresholds set to the “standard noisy” values.

Even if the amplitudes in Figure 4.3 were uniform and low enough to trigger on all neutrons, the MUX thresholds would not necessarily be optimized, since noise triggers could cause unacceptable levels of deadtime. During September 18 – 20, 2004, a series of neutrino runs were taken with five different sets of MUX thresholds to investigate noise rates. One dataset used the “standard noisy” thresholds used prior to September 2004. The “quiet” thresholds are the same as the “standard noisy” thresholds, except the noisy strings have had their thresholds raised (by one least-significant bit on strings 28 and 36 and by two bits on strings 12 and 20). The other three sets of thresholds are derived from the “quiet” thresholds by raising or lowering the threshold on each channel by a fixed number of bits. The decimal threshold on each string in each set of thresholds is shown in Table 4.5. For each set of thresholds, the average MUX rate is shown in Figure 4.4. The two strings with the highest MUX rates, even in the “plus 1” thresholds, are 34 and 27, which both have discharging problems that are largely unaffected by these subtle changes to MUX thresholds.

The best verification of the MUX thresholds comes from MUX-Shaper/ADC correlations, such as those for the central AmBe runs on July 17, 2004 (runs 44631 – 44632) [58]. There were Shaper/ADC and MUX deadtime problems with the DAQ system at that time, but there is strong evidence that the deadtimes were correlated, so the MUX-Shaper/ADC correlation information should be useful despite the deadtime. All the correlations were above 92% and only four were below 96%. The lowest correlation factor was on string 18, which has low gain on one leaky NCD. Two of the other three low correlations were on strings 4 and 38, which had very high threshold amplitudes and low MUX rates. The other low correlation was on string 2, which had a threshold amplitude of just over 19 and a fairly low noise rate.

Based on the MUX rates, the threshold amplitudes, and the MUX-Shaper/ADC correlations, an optimized threshold was determined for each MUX channel. One

Table 4.5: The decimal thresholds on each string in each set of thresholds used in the MUX rate tests.

#	Name	“Standard Noisy”	“Quiet”	“Plus 1”	“Minus 1”	“Minus 2”
0	N4	81	81	82	80	79
1	M8	59	59	60	58	57
2	K8	73	73	74	72	71
3	I7	69	69	70	68	67
4	J8	85	85	86	84	83
5	L2	68	68	69	67	66
6	J7	65	65	66	64	63
7	M7	72	72	73	71	70
8	K7	75	75	76	74	73
9	I8	65	65	66	64	63
10	I6	60	60	61	59	58
11	K6	77	77	78	76	75
12	M6	79	81	82	80	79
13	J6	79	79	80	78	77
14	N3	70	70	71	69	68
15	L3	77	77	78	76	75
16	J5	74	74	75	73	72
17	M5	60	60	61	59	58
18	K5	76	76	77	75	74
19	I5	74	74	75	73	72
20	I3	50	52	53	51	50
21	K4	52	52	53	51	50
22	M4	74	74	75	73	72
23	J4	77	77	78	76	75
24	L4	70	70	71	69	68
25	N2	68	68	69	67	66
26	J3	78	78	79	77	76
27	M3	51	51	52	50	49
28	K3	60	61	62	60	59
29	I4	66	66	67	65	64
30	I2	60	60	61	59	58
31	K2	90	90	91	89	88
32	J2	46	46	47	45	44
33	M2	75	75	76	74	73
34	L1	74	74	75	73	72
35	J1	58	58	59	57	56
36	I1	59	60	61	59	58
37	K1	68	68	69	67	66
38	M1	73	73	74	72	71
39	N1	68	68	69	67	66

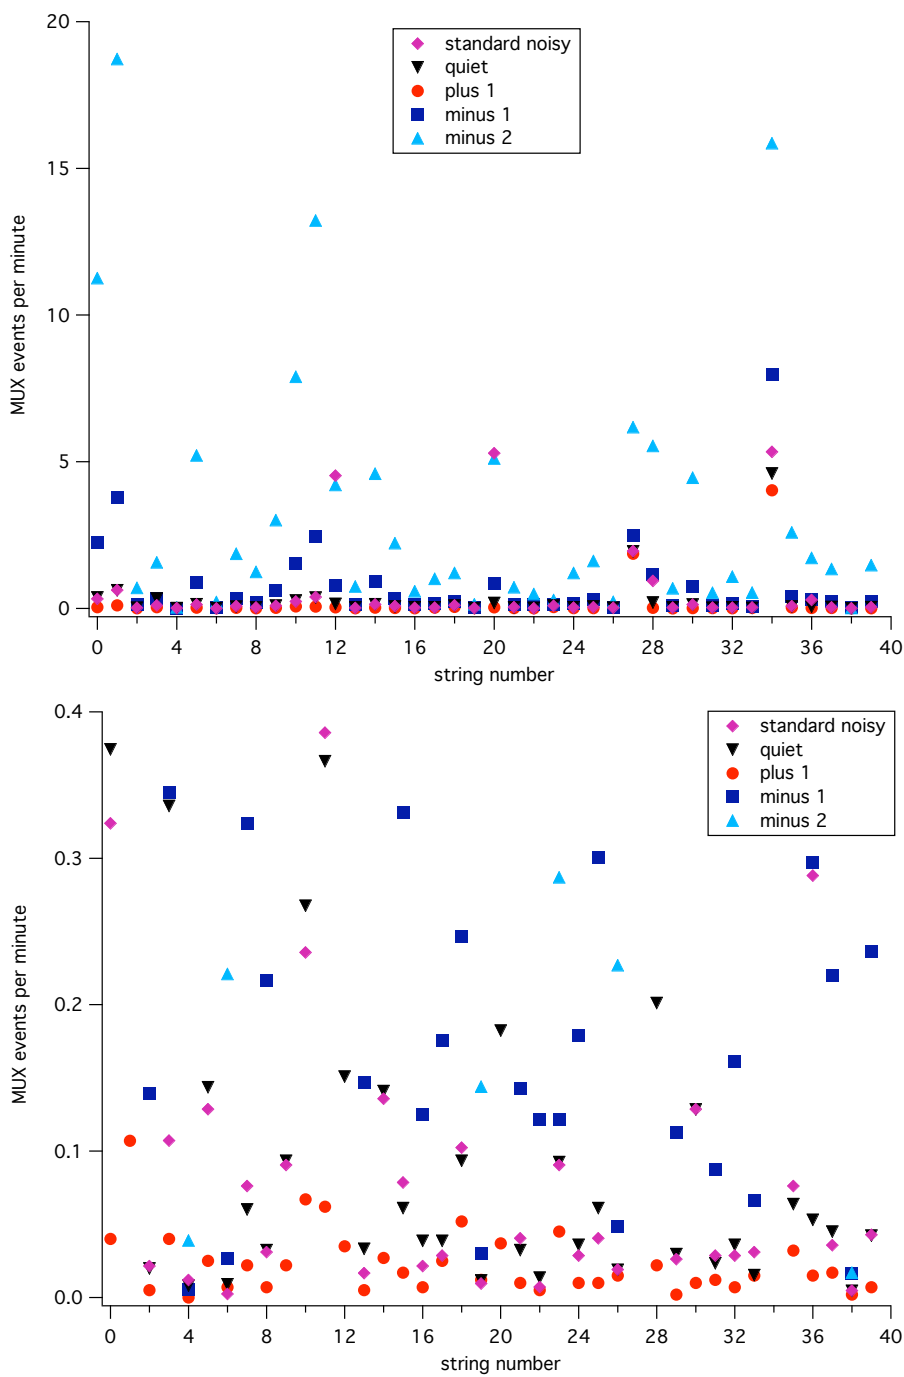


Figure 4.4: The MUX rate on each NCD string with different sets of MUX thresholds. The top panel shows the full scale and the bottom panel is zoomed in on lower rates, which are desired for optimized data-taking.

goal was to keep the overall MUX rate during neutrino running below 0.3 Hz. The two discharging strings contribute 0.109 Hz to the overall rate, with string 34 providing 0.077 Hz and string 27 providing 0.032 Hz. This leaves 0.2 Hz that can be distributed among the other 38 strings and still stay below 0.3 Hz total, so 0.3 events per minute per string is allowable. Another goal of the optimization was an overall MUX-Shaper/ADC correlation of better than 97%, based on the AmBe data analysis. Any strings with correlations worse than 97% had their thresholds lowered, if possible. Table 4.6 shows the new “optimized” thresholds for each string. The expected total MUX rate with these optimized thresholds is 0.217 Hz during neutrino running, based on the MUX rate data collected September 18 – 20, 2004.

Based on these “optimized” thresholds and the results of the string 14 threshold test, some preliminary conclusions can be drawn. To first order, the noise rates are similar if the absolute amplitude of the thresholds are set to be the same, which is required to make Q the same for all the strings. A threshold of 68 on string 14 (which brought Q down to zero) would raise its trigger rate to 0.08 Hz. If all 40 of the NCD strings had trigger rates of 0.08 Hz, then the overall trigger rate in the MUX system would be over 3 Hz, which would saturate the oscilloscopes. At the “optimized” thresholds (70 for string 14), the trigger rate should be acceptable at less than 0.25 Hz in the entire array (leading to scope-deadtime losses of about 3% of the events), and the threshold cut percentage should be about 1.5%.

After a significant amount of data had been collected with the “optimized” MUX thresholds, including neutron source data, the MUX thresholds were further refined based on similar data as in the initial optimization. For the most part, the overall data rates on each string were in good agreement with the predicted rates. The threshold amplitudes analysis was repeated using the “optimized” thresholds and all of the threshold amplitudes fell between 17.1 and 19.5. For the MUX-Shaper/ADC corre-

lation studies, the threshold cut fraction had to be used because the Shaper/ADC deadtime had been greatly reduced and was no longer correlated to the MUX deadtime. Figure 4.5 shows the threshold cut fraction Q for each string for the weak AmBe source runs taken on Sept 28 – Oct 3, 2004. Ideally Q would be zero for all the strings, indicating that the thresholds are low enough to trigger on all the neutrons.




Figure 4.5: The threshold cut fraction for each string with “optimized” MUX thresholds. Strings 3, 10, 20, and 30 are not shown because they are the ^4He strings.

If the MUX thresholds were lowered enough to lower the threshold cut fraction on all the strings to zero, then the MUX rate would increase significantly. Having the threshold cut fraction below 2% on all strings is an acceptable compromise. This keeps the MUX rate around 0.25 Hz, where the scope deadtime is reasonably small. From Figure 4.5, the only strings whose error bars are completely above 2% are the leaking string 18 and strings 1, 4, 7, 12, 16, 33, 35, and 38. Thus, to bring the

threshold cut fractions below 2%, the thresholds on those eight strings were lowered by one bit each. When the threshold on string 7 was lowered from 71 to 70, its rate went up to 0.04 Hz, which was unacceptably high, so its threshold was left at 71. The rate on the rest of these strings was less than 0.02 Hz at the lowered MUX thresholds. After these changes, the predicted threshold cut fraction on all strings (except for string 18, and also string 7, since its threshold was not changed) was below 2% and the predicted total MUX rate was below 0.3 Hz. These “reoptimized” thresholds are shown in Table 4.6, along with the “optimized” thresholds.

Based on an analysis of data taken with the AmBe source over the weekend of October 15 – 17, 2004, the threshold cut fraction for each string with the “reoptimized” thresholds are shown in Figure 4.6. The threshold cut fraction for the leaking string 18 was $10.3\% \pm 0.1\%$, off the scale of the graph. The threshold cut fractions for the rest of the strings are below 2% within error bars, except for strings 1, 21, and 34. Strings 1 and 34 both experience low-energy discharge events that trigger the Shaper/ADCs below bin 40. Some of these events do not trigger the MUXes and they account for the slightly high threshold cut fractions on those two strings. String 21, however, is a mystery. In earlier and later data the threshold cut fraction on string 21 was in line with the rest of the array, so the high threshold cut fraction during this set of neutron calibrations was anomalous. The overall threshold cut fraction for all 36 of the ^3He strings is also plotted on Figure 4.6. It is $1.83\% \pm 0.05\%$, even including the very high threshold cut fraction on string 18.

The “reoptimized” MUX thresholds are all quite consistent, as evidenced by updated threshold amplitude measurements that indicate all the amplitudes are between 17.75 and 19.25, much tighter than prior to optimization. The threshold cut fraction is quite low in the array, indicating that all of the thresholds are acceptable. If all the thresholds were lowered by one or two hex bits, the threshold cut fraction could




Figure 4.6: The threshold cut fractions on each string with “reoptimized” MUX thresholds. Strings 3, 10, 20, and 30 are not shown because they are the ^4He strings.

be reduced dramatically below 2%. The cost of doing this would be an increased MUX rate, which would lead to increased scope deadtime. A test done on October 24, 2004 indicated that lowering the thresholds by one bit would decrease the scope livetime from 98.49% to 83.81%. The advantages of decreasing the threshold cut fraction would be completely outweighed by the negative impact of this large deadtime. In addition, the MUX rate is more likely to go up than down (if more strings develop discharging problems), and there is a bit of leeway in the scope deadtime to handle such an occurrence with the “reoptimized” thresholds.

Table 4.6: “Optimized” MUX thresholds, set on Sept 22, 2004, and “reoptimized” thresholds, set on October 14, 2004.

#	Name	MUX	“Optimized”		Change from “Quiet”	“Reoptimized”		Change from “Optimized”
			hex	decimal		hex	decimal	
0	N4	3/4	0×51	81	0	0×51	81	0
1	M8	2/3	0×3C	60	1	0×3B	59	-1
2	K8	1/4	0×48	72	-1	0×47	71	-1
3	I7	2/4	0×44	68	-1	0×44	68	0
4	J8	0/9	0×53	83	-2	0×52	82	-1
5	L2	1/2	0×44	68	0	0×44	68	0
6	J7	2/2	0×3F	63	-2	0×3F	63	0
7	M7	0/8	0×47	71	-1	0×47	71	0
8	K7	3/1	0×4A	74	-1	0×4A	74	0
9	I8	1/3	0×41	65	0	0×41	65	0
10	I6	0/5	0×3C	60	0	0×3C	60	0
11	K6	2/1	0×4D	77	0	0×4D	77	0
12	M6	1/1	0×51	81	0	0×50	80	-1
13	J6	3/3	0×4E	78	-1	0×4E	78	0
14	N3	2/1	0×46	70	0	0×46	70	0
15	L3	0/6	0×4C	76	-1	0×4C	76	0
16	J5	1/1	0×49	73	-1	0×48	72	-1
17	M5	3/2	0×3B	59	-1	0×3B	59	0
18	K5	0/7	0×4B	75	-1	0×4B	75	0
19	I5	2/9	0×48	72	-2	0×48	72	0
20	I3	3/9	0×34	52	0	0×34	52	0
21	K4	1/5	0×33	51	-1	0×33	51	0
22	M4	2/8	0×49	73	-1	0×49	73	0
23	J4	0/4	0×4B	75	-2	0×4B	75	0
24	L4	3/1	0×45	69	-1	0×45	69	0
25	N2	0/2	0×43	67	-1	0×43	67	0
26	J3	2/7	0×4C	76	-2	0×4C	76	0
27	M3	1/10	0×33	51	0	0×33	51	0
28	K3	3/10	0×3D	61	0	0×3D	61	0
29	I4	0/3	0×41	65	-1	0×41	65	0
30	I2	1/8	0×3C	60	0	0×3C	60	0
31	K2	2/5	0×59	89	-1	0×59	89	0
32	J2	3/5	0×2D	45	-1	0×2D	45	0
33	M2	0/0	0×4A	74	-1	0×49	73	-1
34	L1	2/6	0×4A	74	0	0×4A	74	0
35	J1	1/6	0×3A	58	0	0×39	57	-1
36	I1	3/8	0×3B	59	-1	0×3B	59	0
37	K1	0/1	0×43	67	-1	0×43	67	0
38	M1	3/7	0×47	71	-2	0×46	70	-1
39	N1	1/11	0×43	67	-1	0×43	67	0

Chapter 5

NCD PULSE SHAPE ANALYSIS TECHNIQUES

5.1 Pulse Shape

The expected NC neutron-capture signal in the NCDs is about three events per day, compared to the alpha background of just over 250 events per day, about 22 of which are in the energy range where neutron-capture events are expected. Because of the small signal-to-background ratio, identification of neutron-capture events through pulse shape analysis is essential for signal extraction. The purpose of digitizing the NCD data is that the pulse shapes can be used to separate neutron-capture events from backgrounds, particularly alphas.

The pulse shape of a NCD event is mainly determined by the profile of the primary ionization. The rate of energy loss for charged particles in an absorber is known as the stopping power and is defined as the differential energy loss divided by the differential path length. Stopping power is described by the Bethe formula [48]:

$$S = -\frac{dE}{dx} = \frac{4\pi e^4 z^2}{m_0 v^2} N Z \left[\ln \frac{2m_0 v^2}{I} - \ln \left(1 - \frac{v^2}{c^2} \right) - \frac{v^2}{c^2} \right], \quad (5.1)$$

where v and ze are the velocity and charge of the ionizing particle, N and Z are the number density and atomic number of the absorber atoms, and m_0 and e are the electron mass and charge. I is a parameter representing the average excitation and ionization potential of the absorber, which is experimentally determined for each element. For nonrelativistic particles only the first term in the brackets is significant. The terms within the brackets vary slowly with energy, so for a given particle in a

given absorber $-dE/dx$ goes like $1/v^2$, or inversely with particle energy. Because of the z^2 dependence, a particle with the greatest charge will have the largest energy loss in the same absorber. Media with high atomic number and high density make the best absorbers, because of the NZ dependence.

The stopping power as a function of distance follows a characteristic shape known as the Bragg curve, shown in Figure 5.1. At first, the stopping power increases along the track, since the particle is losing energy and slowing down, so on average it spends more time in a given region and can interact more strongly. Near the end of the track, however, the effective charge of the positive ion decreases as it picks up an electron cloud, so the stopping power decreases sharply. The energy at which the maximum of the stopping power curve occurs is known as the Bragg peak, and is a constant for a given ionizing particle in a given medium, regardless of the initial energy.




Figure 5.1: Energy loss along an alpha track. Stopping-power curves are shown for a single particle and for a parallel beam of identical particles. Figure is from [48].

Thermal neutron capture on ^3He produces a back-to-back proton and triton with 573 keV and 191 keV of energy, respectively. The stopping-power curve for the proton and triton created in a neutron-capture event is shown in Figure 5.2. For a proton in the NCD gas mix, the Bragg peak occurs at about 100 keV, so the peak is evident in neutron-capture events. However, for a triton in the NCD gas, the Bragg peak is at 300 keV, above the initial energy of the triton produced in the neutron-capture reaction. Thus, the ionization profile of the triton in a neutron-capture event is sharply decreasing along the track length. The stopping power of a 191-keV triton is greater than that of a 573-keV proton, so the back-to-back track has a discontinuity at the point where the neutron capture occurred. That discontinuity, combined with the Bragg peak of the proton, creates a doubly-peaked ionization profile that makes




Figure 5.2: Stopping-power curve for the proton-triton pair created in a neutron-capture event. The neutron capture occurs at the origin, with the proton moving to the right and the triton moving to the left. Figure is from [59].

neutron-capture events quite distinctive.

The primary ionization occurs quite rapidly, within several nanoseconds. The primary electrons then begin to drift towards the anode wire under the influence of the electric field. This drift can take up to 3.5 microseconds for an electron that starts near the wall of the NCD. The radial position as a function of drift time can be derived from the electron drift velocity and is shown in Figure 5.3. While the velocity is not constant with radius, it is smoothly varying, and in general the time structure of the electrons as they near the wire mimics the primary ionization density profile.




Figure 5.3: The radial position as a function of drift time. Figure is from [59].

If the primary ionization track is perpendicular to the wire with the triton pointing towards the wire, for example, then the pulse shape looks similar to the stopping power curve in Figure 5.2, as shown in the left panel of Figure 5.4. There will be an initial rapid increase in electron density reaching the wire from the triton ionization, then a

drop as the proton ionization begins to reach the wire, followed by a gradual increase to the proton Bragg peak, then a rapid drop near the end of the proton ionization track. If the primary ionization track is parallel to the wire, however, then all the primary electrons are created at about the same distance from the wire and the time structure of the electrons reaching the wire will not reflect the ionization profile, as shown in the right panel of Figure 5.4.




Figure 5.4: Neutron-capture events where the proton-triton track is perpendicular (left) and parallel (right) to the anode wire. The recorded NCD pulses are negative-going, as shown here.

As the primary electrons approach within a few tens of microns of the anode wire, they begin to undergo multiplication. The very strong electric field near the wire accelerates the electrons enough that they ionize the gas near the wire, and these secondary electrons further ionize the gas, creating an avalanche of electrons. At the voltages where the NCDs are operated, the gas multiplication is nearly the same for all primary electrons, so the total charge is roughly proportional to the initial energy deposited by the primary ionization. The voltage is high enough, however, that the proportionality is affected by the influence of the positive ions created in the avalanche.

The positive ions alter the electric field near the wire and reduce the multiplication. This space-charge effect depends on the density of the primary ionization and the track orientation, thus it can affect the proportionality of the NCDs.

The collection of avalanche electrons on the wire happens quite rapidly and makes only a small contribution to the current pulse. The actual signal on the NCDs is caused primarily by the slow drift of the positive ions away from the wire. The positive ions, which are roughly proportional to the initial energy deposition, induce a current on the wire as they drift away. This drift can take up to a millisecond, so in order to capture the full duration of an event, recording times of hundreds of microseconds would be required. This is impractical, partly for deadtime reasons, but mainly because of the low signal-to-noise ratio, and is not done with the NCDs, which digitize each current pulse for only 13.5 microseconds. Fortunately, the majority of the current is near the beginning of the event, while the ions are quite close to the wire and the image charge changes more quickly.

In order to understand the structure of a current pulse from an NCD event, it is helpful to think of the primary ionization track as consisting of many short segments. The time structure of the event is determined by the distance of each segment from the wire, and therefore the time it takes to drift to the avalanche region. The avalanche from a given primary electron creates a sharp spike with a long decay time, as the positive ions slowly drift to the wall. The recorded events consist of a superposition of these spikes and decays from the avalanches created by each of the primary electrons as they reach the wire, as illustrated in Figure 5.5.

The NCD electronics also impact the shapes of NCD pulses. The preamplifier and other components each act as effective RC circuits. Thus the electronics chain smoothes features of the original pulse, making it more difficult to extract information about the ionizing particles, because some details of the pulse shape are washed out.




Figure 5.5: Simulated ion tails from many individual primary electrons. This figure depicts positive-going ion tails, but the recorded NCD pulses are actually negative-going. Figure is from [60].

Prior to analyzing digitized events from the NCD array, many of the effects of the electronics are removed from the pulses. Electronics calibrations allow the electronics transfer function to be determined for each channel. This transfer function can then be used to remove the effects of the electronics from each event. The primary purpose is to undo the logarithmic amplification and return to a linear pulse shape. This process is known as delogging and is applied to each digitized event prior to analysis. Details about the electronics transfer functions and the delogging process will be available in [61].

Another feature in the NCD pulses is the reflection off the bottom end of the NCD string. The timing of the reflection can be used to obtain information about the position of an event along the axis of the NCD. In order to increase the separation between the direct pulse and the reflected pulse, there is a 90-ns round-trip delay line at the bottom of each NCD string. In general, the reflection is visible as a distinct

peak for sharp pulses that occur near the top of an NCD string, but the reflected peak cannot be resolved and only causes a subtle inflection in the majority of pulses.

5.2 Pulse Width

The current default NCD pulse shape analysis methods are based on using the pulse duration or related parameters to discriminate neutron-capture events from alphas. The track length for a 764-keV alpha is 3.3 mm, which is shorter than the 10.2-mm track length for a proton-triton pair of the same energy since energy loss scales as z^2 of the ionizing particle. For certain track orientations, the recorded neutron-capture pulses are wider than the widest possible alpha event of the same energy. Thus the duration of a pulse can be used as a parameter for discriminating neutron captures from alphas. A simple technique involves producing two-dimensional phase-space plots of a pulse duration parameter and energy. The region in phase space that contains neutron captures but no alphas is referred to as the background-free neutron region. The rest of the neutron-capture events are so narrow that they are indistinguishable from wall-effect alphas.

A very simple pulse duration parameter is the full width at some fraction of the maximum pulse amplitude, such as 10%. The delogged current pulse is used to determine this pulse width. The pulse width algorithm finds the maximum amplitude of the pulse and then defines the start and end of the pulse as the points where the pulse amplitude is 10% of that maximum amplitude. The algorithm steps back towards the start of the pulse from the maximum-amplitude point until the pulse first crosses below 10% of the maximum amplitude, then steps forwards towards the end of the pulse from the maximum-amplitude point until the pulse first crosses below 10% of the maximum amplitude. The time difference between those two 10% crossings is the pulse width. Different amplitudes can also be used instead of 10%.

A plot of pulse width vs. energy for simulated NCD events is shown in Figure 5.6. These are ideal events, without noise or electronics effects. The width is defined as the time between the first deviation from zero amplitude and the return to zero amplitude. For real data this is not the best definition of pulse width because noise and possible non-zero baselines make zero-crossings unreliable. Figure 5.6 clearly illustrates how a pulse duration parameter can be used to separate neutron-capture events from alpha backgrounds. About half of the simulated neutron-capture events fall into the background-free region above the maximum alpha pulse width.




Figure 5.6: Simulated alpha and neutron capture width vs. energy. Simulation is from [62].

Another pulse duration parameter that has been used to analyze NCD data is the risetime of the integrated pulse, defined as the amount of time it takes for a

given fraction of the total charge of the event to be deposited. The default risetime algorithm finds the time at which the integral of the digitized scope trace reaches 50% of its full charge. First the start time of the pulse is determined by stepping backwards from the point of maximum pulse amplitude to the first zero crossing. If no zero crossing is found, then point 1500 is used as the starting point, since that is the point at which the MUX trigger occurs. The risetime method integrates the scope trace from the start time to the end of the trace and stores that integral value as the total charge. Then it steps along the integrated pulse from the start time until it reaches the point where the charge is 50% of the total charge. The time difference between that point and the start position is assigned as the risetime of the pulse. Separation similar to what is shown in Figure 5.6 is also possible for simulated data with risetime of the integrated current profile instead of pulse width. When used with real data, however, a flaw in the risetime method is that if the post-pulse baseline is slightly below zero, then the charge obtained by integrating all the way to the end of the pulse will be lower than the true value, since the pulses are negative-going. The point at which the charge reaches 50% of that will then be too far along in the pulse.

5.3 Proof-of-principle Pulse Width Analysis

In December 2004 an analysis workshop was convened in Seattle to run through a pulse duration vs. energy analysis as an exercise to demonstrate that there were no major problems with the NCD data. The livetime of the dataset used was less than a week, so the statistical error on the number of extracted neutrons was very large, but it was intended as a proof-of-principle exercise, not a rigorous analysis. The analysis performed at the workshop used a 50% risetime parameter on the integrated pulse, which had the problems discussed in Section 5.2. The analysis presented here is a variation on that analysis, using the width at 10% of the maximum amplitude as the

pulse duration parameter.

Figure 5.7 presents a pulse width vs. Shaper/ADC energy plot for cleaned neutrino data from selected runs between 48184 – 48298 taken during November 26 – December 3, 2004. The pulse width plotted here was defined as the time between crossings of 10% of the maximum amplitude. There is a significant overabundance of events on string 26, the widest of which is shown in Figure 5.8 (a). These events are a class of background that has not yet been explained. The widest event on a ^4He string






Figure 5.7: Pulse width vs. energy plot for neutrino data taken November 26 – December 3, 2004. The events above the horizontal line fall into the background-free neutron region used in this analysis.


is shown in Figure 5.8 (b). There are a few neutron-capture candidates with larger widths than any of the events on string 26 or any of the ^4He strings. A possible, very simplistic, background-free region based on this plot would be all events with a pulse width greater than 2000 ns in the appropriate energy region. This excludes all events on string 26 and on the ^4He strings, and gives a total of four neutron-capture candidates in the background-free region, which are shown in Figure 5.8 (c) – (f).




(a) Run 48269, ID 589156.




(b) Run 48229, ID 515776.




(c) Run 48283, ID 209520.



(d) Run 48225, ID 295849.



(e) Run 48227, ID 1679550.



(f) Run 48226, ID 964632.

Figure 5.8: Pulse shapes of various NCD events. The event in (a) is the widest event on string 26, with a pulse width of 1912 ns. The event in (b) is the widest event on a ^4He string, with pulse width of 1702 ns. The events in (c) – (f) are the four neutron candidates with pulse widths above 2000 ns. The horizontal axes are in nanoseconds and the vertical axes are in uncalibrated current units.

The efficiency of the background-free region was calculated by counting the fraction of all events that fell in the background-free region during a neutron source run. Above a pulse width of 2000 ns the efficiency is 11.5%. Thus the four candi-

date neutrons in the 2000 ns background-free region correspond to 35 ± 17 neutrons captured in the NCDs. In order to determine the number of neutrons produced in SNO, the overall capture efficiency of 26.4% must be accounted for. Thus the 2000-ns background-free region gives 132 ± 66 neutrons produced in this dataset. The livetime for this dataset was 137.28 hours, so the observed rate is 23 ± 12 neutrons produced per day. The measured neutron-production rate in the salt phase of SNO was 12.6 neutrons per day [63]. Background neutrons have not been accounted for here, but are expected to contribute fewer than one neutron per day. The result of 23 ± 12 neutrons per day is consistent with the salt phase result, given the low statistics.

There is a great deal of similarity in pulse shape between the neutron-capture candidate in Figure 5.8 (e) and the event on string 26 in Figure 5.8 (a) and the ^4He string event in Figure 5.8 (b). Indeed, it is plausible that the neutron candidate in Figure 5.8 (e) is actually an alpha event or a discharge similar to what is seen on string 26. This casts doubt on the validity of this simple pulse duration analysis technique. At this time, development of data-cleaning cuts to remove discharges and other instrumental backgrounds is not yet complete. Until robust data-cleaning cuts have been finalized and pulse shapes are more completely understood, it is extremely difficult to make a good measurement of the NC flux with a pulse duration background-free region technique.

5.4 Other Pulse Shape Analysis Techniques

Other pulse shape analysis techniques that utilize the pulse shape more fully are likely to yield superior results. A pulse shape fitting routine is a very promising technique to distinguish neutron-capture events from alpha backgrounds. A rudimentary pulse shape fitter was developed at the University of Washington [59] based on semi-analytic functions that describe the expected signal from neutron captures and alphas. A

likelihood that an event is a neutron capture can be assigned based on the goodness-of-fit with the expected neutron-capture and alpha signals. The full pulse shape information can be used in a fit, rather than just extracting two parameters from each pulse as in the pulse duration background-free region technique. In the fitting technique, a likelihood of being a neutron capture is assigned for each event, but in the background-free region technique no ambiguity on the event identity is allowed, leading to the necessity of very aggressive cuts.

The pulse shape fitter involves a simulation of both neutron-capture and alpha events, which provides the functional forms for the fits. Only four physics parameters are required to describe any event. For a neutron-capture event these parameters are the radial position of the neutron capture with respect to the wire, the z-position of the event along the wire (needed to correctly include the reflected pulse), and two angles that describe the direction of the proton-triton track. Alpha events originate in or on the wall, so the parameters are the two angles describing the track direction, the z-position, and the energy with which the alpha leaves the inner surface of the wall of the NCD. Using known stopping powers, drift velocities, and electron multiplications for the NCDs, as well as electronics response functions, the pulse shape associated with a given simulated event can be determined. Once the functional forms describing pulse shapes have been simulated to the desired level of detail, these functional forms can be used to fit real data to differentiate neutron-capture events from alphas. A neutron-capture event generated by this simulation is shown in Figure 5.9, along with the track orientation.

This fitter is quite rudimentary, but more refined fitting techniques are being developed by other members of the SNO collaboration. All of these fitters determine a neutron-capture likelihood and an alpha likelihood for each event. One method uses a grid fit with a database of simulated pulses to find the best-fit parameters in a




Figure 5.9: A simulated proton-triton pulse shape (top) and the end view (bottom left) and side view (bottom right) of the track the produced it. Only the solid lines represent the actual proton-triton track; the dotted line is the chord of the NCD that the track falls along, but the particles range out before travelling the full length of the chord.

neutron-capture model and an alpha model. The goodness-of-fit parameters for the neutron-capture model and alpha model are compared and used to separate neutron-capture events from alphas [62]. Another technique uses matrices that turn a simple track geometry into a full simulated pulse shape. By inverting the matrices, real events can be simplified into basic descriptions of alpha or neutron-capture geometry

and energy deposition [64]. Another fitter uses fully-analytic descriptions of pulse shapes and attempts an analytic fit [65]. The published analysis from the NCD phase will likely be based on one or more of these fitters.

Chapter 6

ALPHA MODELING

6.1 *Alpha Sources*

Both ^{238}U and ^{232}Th have gammas in their decay chains, following the beta decays of ^{214}Bi and ^{208}Tl , respectively, that are above the 2.2-MeV photodisintegration threshold of deuterium. Neutrons produced by photodisintegration are indistinguishable from the neutron signature of the NC process in SNO. Thus it is important to know precisely how much ^{238}U and ^{232}Th are present in the SNO detector in order to quantify this photodisintegration-neutron background and subtract it from the measured NC flux. Impurities in the nickel walls of the NCDs dominate, since the nickel constitutes the vast majority of the mass of the NCDs. In order to limit this background, the NCD construction specifications called for less than 24×10^{-12} grams of ^{238}U per gram of nickel (< 24 ppt) in the NCD array and less than 2×10^{-12} grams of ^{232}Th per gram of nickel (< 2 ppt). The alphas from the uranium and thorium decay chains can provide a measurement of the ^{238}U and ^{232}Th in the NCDs, and thus of the photodisintegration-neutron background introduced by the NCDs.

The NCDs are sensitive to alphas from the uranium and thorium decay chains if the alphas are able to reach the active NCD volume. The highest energy alpha in the uranium and thorium decay chains is the 8.8-MeV ^{212}Po alpha in the ^{232}Th chain, which has a range of $20 \mu\text{m}$ in nickel. Thus, the NCDs are only sensitive to alphas in about the inner 5% of the nickel walls, since alphas produced deeper in the walls cannot escape the nickel into the gas.

shape discrimination techniques can separate neutrons from these alpha backgrounds. Knowledge of the total alpha rate aids in these analyses by providing an estimate of the number of alphas that fall into the neutron energy region.

Several assumptions must be made when modeling these alpha impurities. Due to the complicated chemistry of the CVD process used to purify the nickel, it cannot be assumed that the uranium and thorium decay chains are in equilibrium. A likely scenario breaks equilibrium at radium in both chains¹. In this case, the bottom of the uranium chain (^{226}Ra decay and below) would be suppressed to the extent that ^{226}Ra was removed during the CVD process. This activity would grow back slowly, following the 1600-year half-life of ^{226}Ra . In the thorium chain, ^{228}Ra and ^{224}Ra are shorter-lived, so the disequilibrium scenario would be different. It would only take a short time for 3.66-day ^{224}Ra to reestablish the lower part of the chain, but then this activity would decay away with the 1.9-year half-life of ^{228}Th before being replenished by 5.8-year ^{228}Ra .

The ^{238}U and ^{232}Th are assumed to be uniformly distributed throughout the nickel walls of the NCDs. This assumption could be violated by inclusions of other materials in the nickel. It is known that during construction of the NCDs, small pockets of aluminum oxide from the mandrels used in the CVD process sometimes adhered to the nickel tubes. The thickness of the oxide layer on the mandrel was 38 μm , so these inclusions could be up to 38 μm deep. The cleaning procedures applied to the NCD tubes included a hydrofluoric acid etch that was intended to remove aluminum oxide [67], so it is likely that the majority of the inclusions were removed. However, aluminum typically has ^{238}U and ^{232}Th impurities at the 1 – 10 parts per million level, considerably higher than the 2 – 24 parts per trillion specification set for the NCD nickel. So any aluminum inclusions remaining in the nickel are likely to appear as hot

¹Preferential removal of other elements from the chains would allow equilibrium to be restored in one or both chains in a matter of days.

spots on or in the NCD walls.

In addition to the alphas from uranium and thorium chain bulk contamination, there are 5.3-MeV alphas from ^{210}Po surface contamination that was introduced during construction. Prior to NCD construction, the nickel tubes were stored in an underground location that was high in radon. Progeny of ^{222}Rn , such as ^{210}Pb and ^{210}Po , adhered to the surface of the nickel. Electropolishing removed much of the contamination, but a distinct line at 5.3 MeV still remains in energy spectra taken from the NCDs. The shape of this line indicates that the ^{210}Po is primarily on the surface of the nickel walls, so the alphas always enter the NCD gas with the full 5.3 MeV, unlike the bulk uranium and thorium chain alphas that are produced deeper in the walls. There is a lower-energy tail below the 5.3-MeV ^{210}Po line that arises from alphas that strike the NCD wall before depositing all of their energy in the gas.

6.2 Alpha Energy Simulation

Much can be learned about the alpha-producing impurities in the NCD array from their energy spectra. Because the highest-energy thorium chain alpha is more energetic than the highest-energy uranium chain alpha, some separation of ^{232}Th and ^{238}U contamination is possible. In addition, because the ^{210}Po is a surface contaminant, the distinct 5.3-MeV line provides a good determination of the level of ^{210}Po contamination. In order to understand the NCD energy spectra, the energy distributions of bulk uranium chain, bulk thorium chain, and surface ^{210}Po alphas were simulated to produce probability density functions (PDFs) for comparison with data.

In the simulation, the uranium and thorium chains were broken according to a disequilibrium model where radium was preferentially removed. The details of equilibrium-breaking do not have a huge impact on the simulated spectra because all of the alphas in the upper part of the chains have similarly low energies (below 5.5

MeV). The alphas in the upper parts of both chains (from ^{232}Th , ^{238}U , ^{234}U , and ^{230}Th decays) were combined into one PDF since separating them based on their spectra is difficult. This separation is also not needed to understand the photodisintegration backgrounds that come from near the bottom of the chains. A second PDF was produced with the five alphas in the bottom part of the uranium chain (from ^{226}Ra , ^{222}Rn , ^{218}Po , ^{214}Po and ^{210}Po decays), and a third PDF contained the five alphas in the bottom part of the thorium chain (from ^{228}Th , ^{224}Ra , ^{224}Rn , ^{216}Po , and ^{212}Bi or ^{212}Po decays).

The uranium and thorium chains were simulated using Igor Pro version 5.00 [68], a commercially-available analysis software package. Alphas were generated following the branching ratios and energies from the uranium and thorium chains [49], shown in Table 6.1. One million decays of each portion of each chain were simulated. The primary ^{238}U , ^{232}Th , ^{226}Ra , and ^{228}Th atoms were uniformly distributed in the nickel walls of the NCDs. To reduce computation time, the primaries were only distributed to a depth of $22\ \mu\text{m}$, since no alphas produced deeper in the walls can escape the nickel to reach the active detector volume. A randomly-generated number was used to set the depth of the primary atom and then all of the alphas (three in the upper part and five in the lower part of the uranium chain and one in the upper part and five in the lower part of the thorium chain) from subsequent decays in the chain were started at the same depth. Different randomly-generated numbers drawn from isotropic distributions ($\cos\theta = [-1, 1]$ and $\phi = [0, 2\pi]$) determined the trajectory of each alpha in the chains.

The coordinate system used in the simulation has its origin at the position where the alpha originates in the NCD wall, as shown in Figure 6.2. The y-axis points towards the anode wire, the z-axis is parallel to the anode wire, and the x-axis is in the mutually perpendicular direction to form a right-handed coordinate system. The

Table 6.1: The energies and branching ratios of alphas in the ^{238}U and ^{232}Th decay chains. Only branching ratios over 1% are shown. In the ^{232}Th chain, ^{212}Bi either alpha-decays to ^{208}Tl with a 36% branching ratio or beta-decays to ^{212}Po with a 64% branching ratio. Thus, the branching ratios of the alphas shown in ^{212}Bi and ^{212}Po together add up to 100%. Data is from Appendix 1 of [49].

Decay Chain	Element	Branching Ratio	Energy (MeV)	Branching Ratio	Energy (MeV)
^{238}U upper	^{238}U	77%	4.196	23%	4.149
	^{234}U	72%	4.777	28%	4.723
	^{230}Th	76.5%	4.688	23.5%	4.621
^{238}U lower	^{226}Ra	94.5%	4.784	5.5%	4.602
	^{222}Rn	>99%	5.490	–	–
	^{218}Po	>99%	6.002	–	–
	^{214}Po	>99%	7.687	–	–
	^{210}Po	>99%	5.305	–	–
^{232}Th upper	^{232}Th	77%	4.011	23%	3.957
^{232}Th lower	^{228}Th	73.1%	5.423	26.9%	5.340
	^{224}Ra	95%	5.686	5%	5.449
	^{220}Rn	>99%	6.288	–	–
	^{216}Po	>99%	6.778	–	–
	^{212}Bi	10.1%	6.090	25.9%	6.051
	^{212}Po	64%	8.785	–	–

angle θ is defined relative to the positive z-axis and increases toward the xy-plane and the angle ϕ is defined relative to the positive x-axis and increases toward the positive y-axis in the xy-plane. To calculate the energy loss of the alpha in the nickel wall of the NCD, the wall is assumed to be a flat plane that extends infinitely in the x- and z-directions. The approximation of the wall as a flat plane is valid, since the range

of a 5-MeV alpha is less than $10 \mu\text{m}$, while the radius of the NCD is 2.54 cm. The line s is the distance travelled in the wall by an alpha that starts at depth a moving in the direction defined by the angles θ and ϕ . The projection of the line s into the xy -plane is $p = s \cdot \sin \theta$. The projection of p onto the y -axis is $a = p \cdot \sin \phi$. Thus the distance traveled to escape the wall is $s = a / (\sin \theta \cdot \sin \phi)$.




Figure 6.2: Simulated geometry describing an alpha escaping from the NCD wall. Due to the very short range of the alphas in the nickel wall, the curved wall can be approximated as a flat plane.

Once the initial depth and trajectory of a given alpha was determined in the simulation, the energy loss of the alpha in the nickel was calculated. To do this, the distance that the alpha must travel to exit the nickel and enter the detector volume was calculated from the initial depth and direction, then the amount of energy it would lose while traveling that distance was calculated. This was done using range tables from Stopping and Range of Ions in Matter (SRIM) [69], for alphas in nickel with a

density of 8.9080 g/cm^3 . The range of an alpha with the appropriate initial energy was found from the SRIM table, then the distance the alpha had to travel through the nickel was subtracted from that range, giving the final range of the alpha. The energy corresponding to that final range was found from the SRIM table, giving the alpha's initial energy in the active volume. If the range was too short for the alpha to escape from the wall or if the initial direction of the alpha was deeper into the wall, then the event energy was recorded as zero.

The distributions of energies with which simulated alphas enter the active NCD volume are shown in Figure 6.3. Each amplitude step in the distribution corresponds to a different initial alpha energy in the decay chain. The shape of the distribution below the initial alpha energy reflects the shape of the energy loss curve, which is shown in Figure 6.4. The density of energy deposition is greatest near 1 MeV, de-




Figure 6.3: The energy with which simulated alphas leave the NCD wall and enter the active volume. In the simulation, the alphas are uniformly distributed in the nickel walls of the NCD.

creasing gradually above 1 MeV and sharply below 1 MeV. This effect can be seen in the shape of the distribution below the lowest energy alpha line at 4 MeV in Figure 6.3. For a given initial alpha energy and a uniform depth distribution, the number of alphas that lose a given amount of energy in the wall is directly related to the energy-deposition density at that energy.




Figure 6.4: The stopping power, or energy loss, for an alpha in nickel as a function of the alpha energy.

Once the alpha's energy as it left the NCD wall was known, the simulation checked to see if the alpha hit another part of the wall before depositing all its energy in the detector volume. The same coordinate system was used as in the simulation for alphas in the wall, with the curvature of the walls now taken into account. Since the range of alphas in nickel is so small, the approximation was made that the alpha had not moved while escaping from the wall, so the origin could be shifted to the point where the alpha emerged from the wall without appreciably changing either of the angles.

Figure 6.5 shows a side view and an end view of an NCD to illustrate the distance d that the alpha travels in the detector volume before hitting another part of the wall. The projection of d into the xy -plane is given by $b = d \cdot \sin \theta$. Applying the law of cosines to the triangle in the right panel of Figure 6.5 gives $b = 2 \cdot r \cdot \sin \phi$. Thus $d = 2 \cdot r \cdot \sin \phi / \sin \theta$ is the distance the alpha travels in the detector volume in terms of the NCD radius $r = 2.54$ cm and the angles θ and ϕ .




Figure 6.5: Simulated geometry describing an alpha traveling through the NCD gas. The left panel shows a side view of an NCD and the right panel shows an end view.

SRIM range tables were again used to calculate the energy deposition, now with alphas in the NCD gas mix. The NCD gas is 85% ^3He and 15% CF_4 by number, which corresponds to atomic percentages of 53.12% helium, 9.38% carbon, and 37.5% fluorine. The density of the NCD gas used in the calculation was 1.76×10^{-3} g/cm³, which was calculated with 85% ^3He and 15% CF_4 and a pressure of 2.5 atmospheres. A Bragg correction of -4.31% was applied to account for the difference between molecular

CF₄ and atomic carbon and flourine [69]. Based on the alpha's direction of travel, the distance to the NCD wall was calculated. If that distance was greater than the alpha's range in the NCD gas mix, then the full energy with which the alpha entered the detector volume was deposited. If the range was longer than the distance to the NCD wall, then the distance was subtracted from the initial range to obtain the final range, and the corresponding energy was found from the range tables. The energy deposited in the active volume was then the energy with which the alpha entered the detector volume less the energy with which it struck the wall.

A randomly-generated number was added to each energy to simulate the effects of detector resolution. A good approximate form for resolution is $\sigma_t = \sqrt{(\sigma_d\sqrt{E})^2 + (\sigma_e)^2}$, where σ_t is the total resolution, σ_d is the contribution from the detector itself at an energy E , and σ_e is the contribution from the electronics. The electronics resolution was determined to be 8 keV by fitting a Gaussian to the peak produced in one Shaper/ADC channel by a mono-energetic pulser input. The detector resolution was determined from neutron source data taken with the NCD array. For each string, the Shaper/ADC bin of each event was converted into energy, then all the strings were summed together into one histogram. A Gaussian was fit to the neutron peak, giving $\sigma \sim 21$ keV as the total resolution at 764 keV (thus $\sigma \sim 2.8\%$ for the neutron peak). Using the resolution equation above, the detector contribution was determined to be $\sigma_d \sim 0.022 \sqrt{\text{MeV}}$, so the resolution factor added to each simulated event energy E (in MeV) was randomly drawn from a Gaussian distribution with $\sigma = \sqrt{(0.022\sqrt{E})^2 + (0.008)^2}$. The distributions of energies deposited in the NCD by simulated uranium and thorium chain alphas are shown in Figure 6.6.

As a check on the simulation geometry, the probability of simulated alphas escaping from the wall was compared to an analytic calculation. Since the coordinate system used in the simulation is not the natural one, a different coordinate system




Figure 6.6: The energy that simulated alphas deposit in the NCD gas. Energy losses as the alphas leave the nickel walls are simulated, as well as the energy lost due to the alphas striking another part of the wall prior to depositing the full energy in the gas. In addition, detector resolution effects are included in the simulation.

was used in the analytic calculation. In the calculation, the z -axis points towards the wire so that θ defines a cone whose surface is the paths of all particles that travel the same distance to escape the wall. In this coordinate system ϕ does not affect the distance that an alpha must travel to escape from the wall. Monoenergetic alphas were assumed to be uniformly distributed to a depth equal to their range r . The energy loss was assumed to be linear, so an alpha lost the same amount of energy going a given distance at the end of its path as at the start. The angle $\theta = \cos^{-1}(a/s)$, where a is the depth and s is the distance the particle travels to escape from the wall. The probability of escape is given by the fraction of all alphas that need to travel a distance less than their range, or the probability that $s < r$. For a given θ , this probability can be expressed as the fraction of solid angle occupied by the cone defined

by $\theta = \cos^{-1}(a/r)$, which is the cone with $s = r$ as the surface. The solid angle subtended by the cone is $\int_0^{2\pi} \int_0^\theta \sin \theta \, d\theta \, d\phi = 2\pi(1 - \cos \theta)$. As a fraction of the total solid angle 4π , this is $(1 - \cos \theta)/2$. Since θ itself depends on a , which can vary from 0 to r , the probability of escape is given by $\int_0^r (1 - \cos \theta[a])/2 \, da = \int_0^r (1 - a/r)/2 \, da = r/4$. Thus 25% of alphas that start uniformly distributed to a depth equal to their range escape the wall. When 6-MeV alphas distributed to a depth equal to their range of $11.84 \mu\text{m}$ were simulated with linear energy loss, 250,317 out of one million escaped into the active detector volume, in good agreement with the analytic calculation.

6.3 ^{210}Po Alpha Simulation

Initially, the ^{210}Po energy simulation was the same as the uranium and thorium simulations, except all the alphas were generated on the inner surface of the NCD wall, so no energy was lost in the nickel, and all the alphas had an initial energy of 5.3 MeV. This produced poor agreement with the data, since the energy of the ^{210}Po peak in the simulation was higher than in the data. Based on the hypothesis that nuclear recoils or diffusion could move ^{210}Po (or its 22-year parent ^{210}Pb) into the nickel walls, non-surface alphas were simulated. Energy distributions were produced with the alpha origins distributed with different depth profiles up to $3 \mu\text{m}$ deep in the nickel. None of these distributions agreed very well with the data, since they broadened the ^{210}Po peak more than was observed in the data. The proposed mechanisms for non-surface ^{210}Po were not very convincing either, since the distances travelled by the recoiling heavy nuclei are less than 20 nm and it not known if ^{210}Pb will readily diffuse through nickel.

Much better agreement with the data was achieved when the simulation was expanded to include space-charge effects, which occur when the electric field near the anode wire is modified by the presence of positive ions produced in the avalanche.

The modification of the electric field reduces the recorded energy for a given event from its true energy. The reduction can depend on several factors that differ from event to event, such as the density of the primary ionization track and the track orientation. In general, ionizing particles with higher charge produce denser primary ionization tracks and will be more affected by space charge than lower-charge ionizing particles. Thus, for the NCDs, alpha events have larger space-charge effects than neutron-capture events. This means that the average energy of a ^{210}Po alpha event, for example, will be lower than expected if neutron-capture events are used to calibrate the energy scale.

The effect of space charge on neutron captures and ^{210}Po alphas simulated with the SNOMAN NCD pulse shape simulation is shown in Figure 6.7. Adding space charge to the simulation shifts the maximum charge of neutron-capture events down by 1.9% from its value without space charge, but for ^{210}Po alphas the downward shift due to space charge is 5.4%. There are several parameters in the space-charge model, the most important of which are a constant associated with the mobility of the positive ions in the NCD gas, and the gas gain, for which a default value of 100 was used. Adjusting these parameters has a significant impact on the energies and pulse shapes of simulated events. Not all of these parameters have been measured precisely for the NCDs, and a rigorous study optimizing these parameters to produce good agreement between simulated pulses and data has not yet been conducted. For this simulation, default parameter values were used, and no attempt has been made to quantify the systematic uncertainties associated with the space-charge model.

Since space-charge effects require using a full pulse shape simulation rather than the simple energy simulation used previously, the SNOMAN NCD pulse shape simulation was used to produce a simulated ^{210}Po energy distribution, shown in Figure 6.8. Alphas with 5.3 MeV of initial energy were produced on the inner surface of the NCD

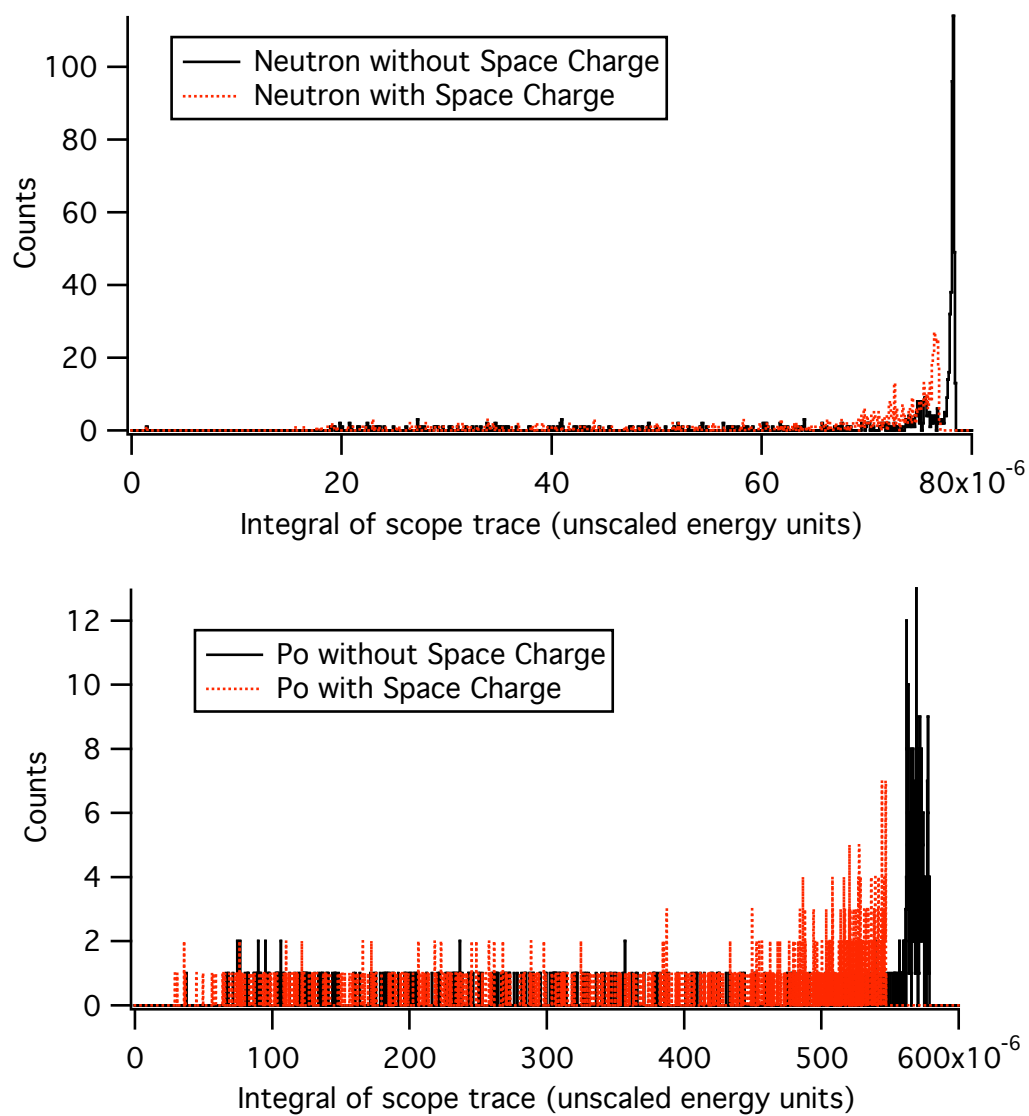


Figure 6.7: The integrated charge deposited in the NCD by simulated neutron captures (upper panel) and ^{210}Po alphas (lower panel), with and without space charge




Figure 6.8: The energy deposited in the NCD by simulated ^{210}Po alphas, with initial energy 5.3 MeV. The alphas are all assumed to be on the inner wall of the NCD. Space charge effects are included in the simulation. A detector resolution of $0.03\sqrt{E}$ is included.

wall and the full pulse shape simulation, including reflections and space charge, was performed. Each simulated event was saved as it left the NCD (before entering any of the electronics) and then integrated to determine its total charge. The energy was calibrated using another SNOMAN simulation without space charge included, so that the upper edge of the energy distribution without space charge was at 5.3 MeV. After this energy calibration was applied to the space-charge simulation, detector resolution was added. The resolution factor added to each energy was \sqrt{E} times a random number drawn from a Gaussian distribution with $\sigma = 0.03^2$.

²This form neglects the energy-independent electronics resolution and assumes that all the resolution comes from energy-dependent detector effects. At higher energies, it is a good approximation to the correct form used in Section 6.2, but the approximation breaks down below about 130 keV, where the electronics resolution contribution is comparable to the detector resolution. This low-energy regime is not used in the energy spectrum fitting, so the difference is negligible.

6.4 Pulse Width Simulation

Some of the pulse shape analysis techniques described in Chapter 5 can also be used for determination of alpha backgrounds. Alphas that deposit less than their full energy in the active volume of the NCD can do so for two different reasons. Either the alpha loses some energy in the nickel wall before entering the active volume, or else the alpha strikes another part of the wall after having deposited some energy in the gas. Bulk alpha energy loss is dominated by the first mechanism, but for surface alphas only the second of these energy-loss mechanisms applies. Given that the range of a 5-MeV alpha in the NCD gas is about the same as the NCD radius, if the alpha is to strike another part of the wall before stopping in the gas, it cannot be traveling very perpendicular to the wire. Because of this geometry effect, low-energy alphas that originate on the surface tend to be less perpendicular to the wire than those that originate in the bulk. Thus, in energy ranges below the full initial alpha energy, the pulse widths of bulk alphas are expected on average to be wider than those of surface alphas.

In order to perform a width vs. energy pulse shape analysis to estimate the relative amounts of bulk and surface alpha contamination, the alpha energy simulation in Section 6.2 was expanded to also include the width of simulated NCD events. For the pulse width vs. energy analysis, the uranium and thorium decay chains were assumed to be in equilibrium, unlike in the energy spectrum simulation in Section 6.2. As in the ^{210}Po simulation in Section 6.3, the energy resolution used here was drawn from a Gaussian with $\sigma = 0.03\sqrt{E}$. To determine the width of each simulated event, the point of closest approach to the wire and the furthest point from the wire were found for each event. The electron drift time from each of these two points to the wire were found from the curve shown in Figure 5.3. The difference between the drift time from the closest point to the wire and the drift time from the furthest point to the wire

gives the width of the pulse.

Since all the alphas originate on or in the walls, the furthest point from the wire is always the inner surface of the wall at a distance $r = 2.54$ cm from the wire. The alpha travels along the chord d that is $b = d \cdot \sin \theta$ when projected into the xy-plane, as was shown in the left panel of Figure 6.5. The alpha may or may not lose all of its energy before hitting the wall. The point on the chord that is closest to the wire is halfway along the chord at $d/2 = r \cdot \sin \phi / \sin \theta$. If the alpha travels more than this distance along the chord before stopping, then it passes through the point closest to the wire. In this case, the distance of closest approach is simply $w = r \cdot \cos \phi$, as shown in the left panel of Figure 6.9. If, however, the alpha loses all of its energy after traveling a distance $g = l / \sin \theta$ that is less than half the distance along the chord,




Figure 6.9: Simulated geometry illustrating the point of closest approach of an alpha to the wire. The left panel shows the distance of closest approach w when the alpha travels more than $d/2 = r \cdot \sin \phi / \sin \theta$ before losing all of its energy and the right panel shows the distance of closest approach v if the alpha stops before traveling $d/2 = r \cdot \sin \phi / \sin \theta$.

then it will not reach the point along the chord that is closest to the wire. It will instead be a distance v from the wire at its closest approach. Using the law of cosines on the triangle in the right panel of Figure 6.9 gives $v = (l^2 + r^2 - 2 \cdot l \cdot r \cdot \sin \phi)^{1/2}$ as the distance of closest approach in this case.

Pulse width vs. energy plots for simulated bulk uranium chain, bulk thorium chain, and surface ^{210}Po alphas are shown in Figure 6.10. The width distributions of the bulk uranium chain and bulk thorium chain alphas are very similar, as they should be since the initial distribution of the ^{238}U and ^{232}Th in the nickel is the same. The surface ^{210}Po looks very different, however, and is confined to much narrower widths in the energy range below the 5.3-MeV peak. This very clear separation between the regions occupied by bulk and surface alphas indicates that the use of a pulse width

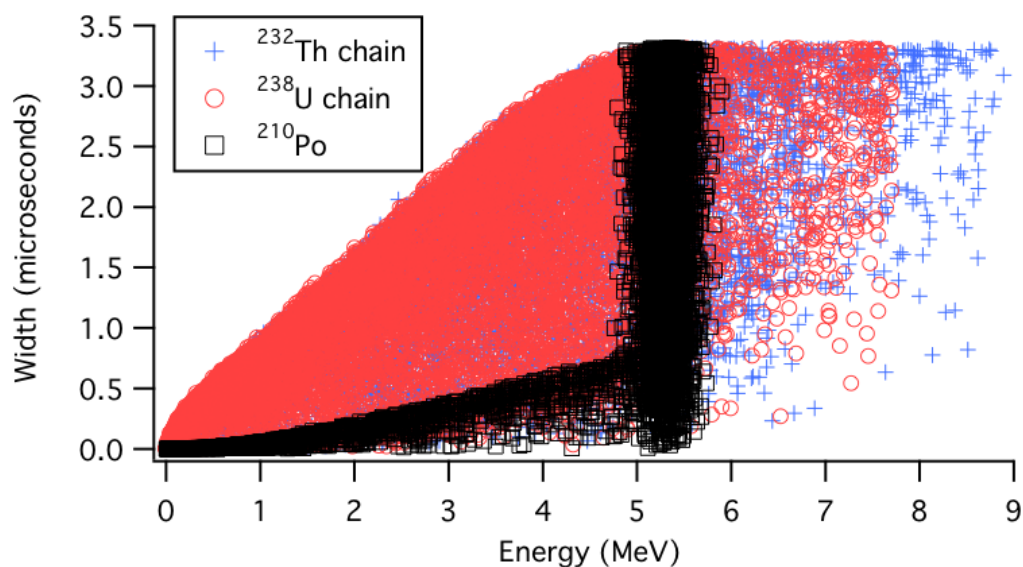


Figure 6.10: Pulse width vs. energy plot for simulated bulk uranium chain, bulk thorium chain, and surface ^{210}Po alphas. For clarity, this plot only contains 1% of the simulated events.

parameter along with the event energy should provide a good distinction between bulk and surface alpha contamination.

6.5 Time Coincidence Simulation

Another technique that is complementary to the alpha energy spectrum analysis takes advantage of the short-lived alpha emitters in the uranium and thorium decay chains. In the uranium chain, the half life of ^{218}Po is only 3.05 minutes, thus the alphas from the decay of ^{222}Rn and ^{218}Po will be correlated in time. In the thorium chain, the half life of ^{220}Rn is only 55.6 seconds and the half life of ^{216}Po is only 0.15 second. Thus it is possible to observe the triple coincidence between the alphas from ^{224}Ra , ^{220}Rn , and ^{216}Po . In this analysis, a thorium chain triple was defined as two events within 111.2 seconds of each other, followed by an event within 0.3 seconds, all of which deposited at least 0.1 MeV. This energy threshold is approximately the hardware threshold on the Shaper/ADCs. Coincident pairs were also studied with time windows of 111.2 seconds and 0.3 seconds. The energy threshold for the 0.3-second doubles was also 0.1 MeV, but the threshold for the long doubles was set to 1 MeV to avoid accidental coincidences between low-energy events that occur at relatively high rates in the data. The coincidence times were chosen to be two half-lives of the thorium chain decays involved.

The alpha energy simulation described in Section 6.2 was modified to simulate the detected energies and time differences of the two alphas involved in the coincidence in the uranium chain and the three alphas in the triple coincidence in the thorium chain. The depths of the simulated alphas were uniformly distributed to $22\ \mu\text{m}$, as before. The time between decays was drawn randomly from a distribution of $t = -\ln(x/\lambda)/\lambda$, which is simply the inverse of the equation governing the timing of a radioactive decay, $f(t) = \lambda e^{-\lambda t}$. Thus, the time between decays has a 50% probability of being less than

$T_{1/2}$. The values of the decay rate λ are $3.78769 \times 10^{-3} \text{ s}^{-1}$ for ^{218}Po , $1.24667 \times 10^{-2} \text{ s}^{-1}$ for ^{220}Rn , and 4.62098 s^{-1} for ^{216}Po .

For the uranium chain coincidence, 10 million decays were simulated. 11.88% of the 5.490-MeV ^{222}Rn alphas deposit some energy in the gas, and 13.43% of the 6.002-MeV ^{218}Po alphas deposit some energy. The rest of the alphas lose all of their energy in the walls. Since the range of a 6-MeV alpha in nickel is only $11.8 \mu\text{m}$, about half the simulated alphas have no chance of escaping regardless of their orientation. Half of the remaining alphas will have their initial direction away from the gas. And only about half of the remaining alphas will make it out of the wall, due to their track orientation and initial depth. In the simulation, 9.93% of the ^{222}Rn alphas deposit at least 1 MeV in the gas, and 11.47% of the ^{218}Po alphas deposit at least 1 MeV. 3.50% of the time, both of the alphas deposit at least 1 MeV. This fraction is considerably larger than if the 9.93% and the 11.47% were simply multiplied together, because they are not independent since the starting depth for both decays is the same. The time between decays is less than or equal to 111.2 seconds 34.38% of the time. The chance of a long-double coincidence where both alphas deposit at least 1 MeV and the time between them is no greater than 111.2 seconds is $1.205\% \pm 0.003\%$. The time between decays is less than or equal to 0.3 seconds 0.11% of the time, and the chance of the short double coincidence is less than 0.005%.

For the thorium chain triple coincidence, 10 million decays were simulated. 12.01% of the 5.686-MeV or 5.449-MeV ^{224}Ra alphas deposit at least 0.1 MeV in the gas, 13.99% of the 6.288-MeV ^{220}Rn alphas deposit at least 0.1 MeV, and 15.60% of the 6.778-MeV ^{216}Po alphas deposit at least 0.1 MeV. 1.71% of the time, all three alphas deposit at least 0.1 MeV in the gas. About 25% of the events are lost to each timing cut, as expected since the coincidence time cuts are set to two half-lives. 75.00% of the events pass the 111.2-second requirement for the first decay and 74.99% pass the

0.3-second requirement for the second decay. The chance of a triple coincidence where all three alphas deposit at least 0.1 MeV and the time differences are within 111.2 seconds and then 0.3 seconds is $0.961\% \pm 0.003\%$.

Both long and short doubles from the thorium chain were also simulated. Exactly two of the three alphas from the thorium chain triple are seen within 111.2 seconds of each other with energies above 1 MeV $8.909\% \pm 0.009\%$ of the time. The two detected alphas are ^{224}Ra and ^{220}Rn 31.35% of the time, ^{224}Ra and ^{216}Po 20.31% of the time, and ^{220}Rn and ^{216}Po 48.34% of the time. The efficiency to see thorium chain doubles is much higher than the efficiency to see uranium chain doubles because the relevant half-lives in the thorium chain are considerably shorter, and because there are two short-lived alpha emitters, so if one alpha is missed there is still a chance to observe a coincidence. The short double is dominated by ^{220}Rn and ^{216}Po , which are seen within 0.3 seconds with energies above 0.1 MeV $3.678\% \pm 0.006\%$ of the time. Short doubles from the other two pairs of thorium chain alphas are seen 0.02% of the time.

System deadtimes also impact the time coincidence analysis. The MUX deadtime is about 1 millisecond after each event, which is non-negligible for the 0.15-second half-life decay in the thorium chain triple coincidence and for the 0.3-second double. 0.46% of the 0.15-second half-life decays were lost to this deadtime. Since there are two scopes in the system, the scope deadtime is only a problem if both scopes are busy when an event occurs. If the ^{220}Rn decay fell within about 0.75 seconds of the previous event, then both scopes would be busy when the ^{216}Po decay occurred. 0.93% of the thorium chain triples were lost to this deadtime. The other situation that causes scope deadtime is an accidental coincidence that busies one of the scopes in conjunction with an actual coincidence. Since the overall scope livetime is 96.6%, a 3.4% efficiency cut was taken on all coincidences to account for this scenario. Taking these deadtimes into account gives a uranium long-double efficiency of $1.164\% \pm$

0.003%, a thorium long-double efficiency of $8.606\% \pm 0.009\%$, a thorium short-double efficiency of $3.557\% \pm 0.006\%$, and a thorium triple efficiency of $0.916\% \pm 0.003\%$.

The rate of accidental coincidences can be calculated for each string by $r_d = r_s^2 t$, where r_s is the average singles event rate and t is the size of the coincidence window. This can be thought of as the fraction of time that lies within t of a preceding pulse (given by $r_s t$) multiplied by the average event rate (r_s) [48]. For a triple coincidence, $r_t = r_s^3 t_1 t_2$, where t_1 and t_2 are the sizes of the two coincidence windows. These equations are valid if $r_s t \ll 1$. The total number of accidental coincidences is simply this coincidence rate multiplied by the livetime.

Chapter 7

ALPHA BACKGROUND MEASUREMENTS**7.1 Datasets**

The commissioning and open datasets were used in the analyses presented in this chapter. The commissioning dataset was taken between June 5, 2004 and November 15, 2004, while the recommissioning of the SNO detector with the NCDs was still underway. The runlist¹ contained 423 runs spanning runs 43466 – 47886. The open dataset was taken between November 27, 2004 and January 3, 2005, when production-quality data were being taken, but no blindness constraints were yet imposed. The runlist² contained 126 runs between runs 48220 – 48673. For the pulse width analysis in Section 7.2, a reduced open dataset of 116 runs³ was used because there were problems processing the digitized data in some runs.

The livetime of each dataset was determined using a Perl script [70] to obtain the length of each run from the PMT system. The times of the first and last events in the run are used to calculate the length of each run, which provides a good approximation to the true run length since the event rate in the PMTs is about 20 Hz. Assuming that the length of each run is 30 minutes (the minimum length of analyzed runs) and that there is a ± 0.05 second uncertainty on the timing of the start and the end of

¹The runlist was `commissioning_ncd_p9.runlist`, with runs 43448, 43450, 45651, 47002, 47608, 47609, 47610, 47611, 47612, 47642, and 47643 removed due to processing problems.

²The runlist was `ncd_open_b8_p1.runlist`, with runs 48302, 48595, and 48650 removed.

³The same runlist, `ncd_open_b8_p1.runlist`, was used, but runs 48302, 48308, 48489, 48493, 48496, 48527, 48530, 48552, 48579, 48605, 48607, 48619, and 48650 were removed.

the run, the uncertainty due to this timing method is 0.0056%. The total runtime for the commissioning dataset was 79.25 days, the open dataset runtime was 27.42 days, and the reduced open dataset runtime was 25.32 days.

A livetime correction was applied to account for the 20 seconds removed after each muon by the muon-follower cut⁴. This correction assumed an average muon rate of 64.6 ± 0.9 muons per day [71] inside the PMT support structure (PSUP). This is the relevant volume since any muons inside the PSUP will be seen by the PMTs and will trigger the muon-follower cut. This correction removed 1.185 ± 0.017 days from the commissioning dataset, 0.410 ± 0.006 days from the open dataset, and 0.379 ± 0.005 days from the reduced open dataset. The livetime uncertainty due to this correction is 18 seconds per day, or 0.021%. The other burst cuts removed very small amounts of time from the datasets (less than a minute) so no correction was applied.

Deadtimes in the NCD system have a significant impact on the livetime. An alternative method of calculating the livetime uses the NCD livetime clocks. There is one livetime clock for the overall NCD system, plus additional clocks for the Shaper/ADCs, the MUXes, and the digitizing oscilloscopes. In order to compare the different livetime calculations, the livetime of a subset of the open data was calculated with these clocks as well as with the method based on the PMT timing. This comparison is shown in Table 7.1. Since cleaned data that required correlated Shaper/ADC and scope events were used in the analyses presented here, the scope deadtime is relevant. To account for this, the livetime of each dataset was reduced by the 3.399% scope deadtime.

The NCD clock records the time at which the run actually starts and stops, which should produce a similar result to the PMT timing method. The 0.3435% difference in the livetime between the PMT timing method and the NCD livetime clock is counted

⁴For published SNO analyses the livetime removed by the data-cleaning cuts is corrected for on an event-by-event basis, but here an average correction was used instead for simplicity.

Table 7.1: The top row shows the livetime of a subset of the open data, calculated with the PMT timing and with each of the NCD livetime clocks. The middle row shows the difference between the livetime calculated with the PMT timing and with each of the NCD clocks, as a percentage of the PMT livetime. The bottom row shows the percent deadtime of each component of the NCD system.

	PMT	NCD	Shaper/ADC	MUX	Scope
Livetime (hours)	583.9486	585.9543	585.8546	585.0602	566.0377
Diff. from PMT (%)		0.3435	0.3264	0.1904	-3.0672
NCD deadtime (%)			0.017	0.153	3.399

as a systematic uncertainty. After correcting for the muon-follower cut and the scope deadtime and applying this additional systematic uncertainty, the final livetimes for the commissioning, open, and reduced open datasets are 75.41 ± 0.26 days, 26.10 ± 0.09 days, and 24.09 ± 0.08 days, respectively.

Data-cleaning cuts were applied to remove non-physics events based on pulse shape and on high-rate bursts. The cuts used were: retrigger, muon-follower short, NHIT burst, missed muon-follower short, Shaper/ADC burst, MUX burst, reverse fork, fork, NCD oscillatory noise, flat trace, narrow pulse, NCD pulser, and Shaper/ADC-overflow follower⁵. These data-cleaning cuts are described in Appendix A. Figure 7.1 shows the impact of these cuts on energy spectra from the commissioning and open datasets. The analyses in this chapter were performed before the production data-cleaning cuts for the NCD phase of SNO were finalized, so some discharge events are still present in the cleaned data. To reduce the impact of these discharge events on the analyses, certain strings were not analyzed, as shown in Table 7.2.

The Shaper/ADC values were converted into energies for use in the analyses presented here. For the ^3He strings, this conversion was done using the peak and edge

⁵The data-cleaning mask cut an event if `GetDamnID()&0xC8440001` or `GetDamnID1()&0x1BE`.

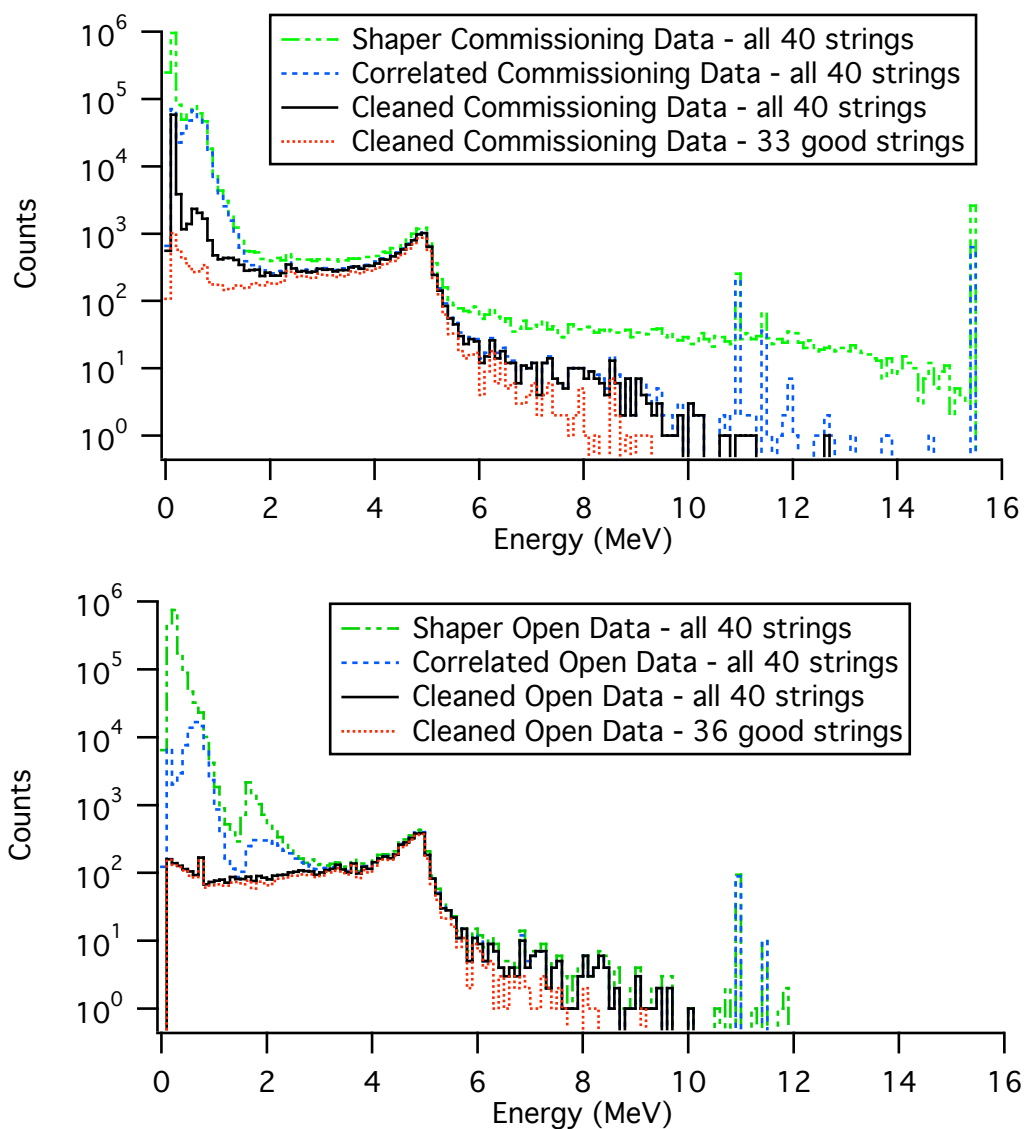


Figure 7.1: Shaper energy spectra of the commissioning and open datasets, showing the impact of data cleaning. For each dataset the raw Shaper/ADC spectrum is shown, along with the spectrum of Shaper/ADC events that have correlated scope events, the Shaper/ADC events with correlated scope traces that pass the data-cleaning cuts, and finally the spectra from the 33 or 36 strings chosen for analysis. The top figure is for the commissioning dataset and the bottom figure is for the open dataset.

Table 7.2: The strings that were not analyzed in the commissioning and open datasets, and the analyzed surface area (in m^2) in each dataset. Also shown is the event rate in each dataset with all 40 strings and the difference between that rate and the rate on the analyzed strings. The rates are given in $\text{events}/\text{m}^2\cdot\text{day}$.

	Removed Strings	Area	Rate	Difference
Commissioning	3, 7, 16, 18, 20, 27, 34	52.51	3.56 ± 0.04	0.27 ± 0.06
Open	3, 7, 18, 20	57.37	3.60 ± 0.06	0.06 ± 0.08
Reduced open	none	63.52	3.58 ± 0.06	—

from neutron source data as a calibration. The neutron-capture peak should be at 764 keV and the low-energy edge of the neutron-capture distribution (caused by protons hitting the NCD wall) should be at 191 keV. The Shaper/ADC bin where these two features fall was measured with source data and used, along with the intercept from electronics calibrations, to determine a linear conversion formula from Shaper/ADC bin to energy for each NCD string [54]. When this formula is used to convert Shaper/ADC bin to energy, the 5.3-MeV ^{210}Po peak falls at 4.9 MeV. Electronics calibrations indicate that the Shaper/ADCs and the preamplifiers are linear, but the position of the ^{210}Po peak shows that the response of the NCDs is not linear over this energy range, primarily due to space-charge effects.

No neutrons capture in the ^4He strings, so the neutron-capture peak cannot be used to determine the Shaper/ADC bin to energy conversion. Instead, conversion formulas for the ^4He strings were based on the ^{210}Po peak. In order to align the ^4He strings' spectra with the ^3He strings' spectra, the energy of the ^4He strings was calibrated to put the ^{210}Po peak at 4.9 MeV. For the width vs. energy analysis, the energy of all the strings was scaled upwards by $5.3/4.9$ to put the ^{210}Po peak at 5.3 MeV since the pulse width simulation did not include space charge.

7.2 Width vs. Energy

As shown in Section 6.4, pulse shape analysis techniques can provide discrimination between bulk and surface alphas. A width vs. energy pulse shape analysis was performed on the reduced open dataset using the full width at 20% of the maximum amplitude and the calibrated energy from the Shaper/ADCs. The 10% width parameter described in Section 5.2 was not used because of concerns that the width was measured so close to the baseline noise level that noise or the ion tail was altering the width of some events. Figure 7.2 shows the width vs. energy for each correlated,




Figure 7.2: The full width at 20% of the maximum amplitude for cleaned events from the open dataset. The lines delineate regions where bulk alphas, surface alphas, and the full-energy ^{210}Po peak can be found. There were 68 narrow discharge events (pulse width below 300 ns) that were removed from this data, 25 from the surface alpha region and 43 from the high-energy bulk-alpha region.

cleaned NCD event. The pulse width simulation described in Section 6.4 does not include space charge, ion tail, reflections, noise, or the electronics response function, most of which are likely to widen the pulses. Another significant difference between the data and the simulation is that the simulation gives the full width at zero amplitude, not at 20% amplitude. Despite these differences, the simulation shown in Figure 6.10 qualitatively resembles the data shown in Figure 7.2.

This data has been divided into regions dominated by bulk alphas, surface alphas, the full-energy ^{210}Po peak, and high-energy bulk alphas. The bulk-alpha region is between 1 and 4.5 MeV above a diagonal line $W = 600 + 1100 \cdot E/4.5$, where E is in MeV. The surface-alpha region spans the same energies but falls below the diagonal line. There are bulk alphas in the surface-alpha region, but no surface alphas should fall into the bulk-alpha region. The ^{210}Po peak is between 4.5 and 6 MeV, and only bulk uranium and thorium chain alphas can be found above 6 MeV. The counts in each region from the simulation described in Section 6.4 and from this data are shown in Table 7.3.

Table 7.3: The percentage of simulated bulk ^{238}U chain, bulk ^{232}Th chain, and surface ^{210}Po alpha events that occupy the bulk, surface, ^{210}Po peak and > 6 MeV regions. Also shown are the numbers of events that occupy each of those regions from data. The percentages do not total to 100% because events that fall below 1 MeV are not counted. 68 discharge events were removed from the counts shown here.

	Simulation			Data
	Bulk U Chain	Bulk Th Chain	Surface ^{210}Po	
Bulk	34.5%	29.0%	0%	777
Surface	28.3%	25.5%	16.5%	2164
^{210}Po peak	15.6%	22.0%	76.7%	2392
> 6 MeV	4.9%	9.1%	$\sim 0\%$	123
Total				6346

The bulk-alpha region could be contaminated by surface alphas that partially ionize the gas in the non-multiplying end of the NCD and partially ionize in the normal multiplication zone. A study of this effect was performed using two short (25-cm) end-effect NCDs made in part from uncleaned nickel tubes. One end-effect NCD had intentionally high surface ^{210}Po contamination on the 7.5 cm at each end, and the other had high ^{210}Po surface contamination on the 10 cm in the center. Analysis of data from these end-effect NCDs gives an effective cross section of $1.86 \times 10^{-3} \text{ m}^2/\text{NCD}$ for surface end-effect events to occupy the bulk-alpha region. Based on the 2392 ^{210}Po peak data events, 11 end-effect events are expected to contaminate the bulk-alpha region in this dataset. The end-effect analysis used the risetime of the integrated pulse, not the pulse width, thus an uncertainty of ± 11 events is taken to account for systematic differences between the pulse duration parameters.

The data events above 6 MeV and in the bulk-alpha region can be used to determine the number of uranium and thorium chain alphas, assuming the chains are in equilibrium. If all 123 events above 6 MeV were from the thorium chain, then the 9.1% efficiency gives a total of 1346_{-127}^{+128} thorium chain events, and if they were all from the uranium chain, then the 4.9% efficiency gives a total of 2508_{-239}^{+242} uranium chain events. If all 766 events in the bulk-alpha region (with the 11 ± 11 end-effect events subtracted) were from the thorium chain, then the 29.0% efficiency gives a total of 2646 ± 104 thorium chain events, and if they were all from the uranium chain, then the 34.5% efficiency gives a total of 2220_{-86}^{+87} uranium chain events. These numbers have been converted into areal rates in Figure 7.3, where the diagonal bands represent the possible uranium and thorium chain alpha rates determined from the bulk-alpha region and the region above 6 MeV. The intersection of the diagonal bands gives the best solution for the uranium and thorium chain alpha rates.

To calculate the intersection point of the two bands, two equations can be written




Figure 7.3: The alpha rate in alphas/ $\text{m}^2\cdot\text{day}$ determined from the bulk-alpha region and the region above 6 MeV. The horizontal-axis intercept of each diagonal band gives the thorium chain alpha rate if all the alphas were from thorium chain decays, and the vertical-axis intercept gives the uranium chain alpha rate if all the alphas were from uranium chain decays. In actuality there is a combination of thorium chain decays and uranium chain decays, and the intersection of the two bands gives the best solution for the thorium chain and uranium chain alpha rates, also shown with the total statistical and systematic uncertainty.

and solved for the unknown uranium and thorium chain contributions using the data events in the bulk-alpha region and above 6 MeV as well as the efficiencies from Table 7.3. The equations are $0.091 \cdot Th + 0.049 \cdot U = 123$ and $0.290 \cdot Th + 0.345 \cdot U = 766$, which have the solution $Th = 280$ and $U = 1984$. There is a statistical uncertainty on the number of counts in each region and a systematic uncertainty on each efficiency that goes into this analysis. A Monte Carlo was used to calculate the impact of these uncertainties on the extracted uranium and thorium chain values. To calculate the statistical uncertainty, values for the numbers of counts in the bulk and above 6 MeV regions were drawn from Gaussian distributions with the appropriate means and

standard deviations, then the equations were solved for Th and U . This was done one million times, then Gaussian peaks were fit to the resulting Th and U distributions to extract the means and standard deviations. To calculate the systematic error, this process was repeated with the numbers of counts held at their central values, but the values for the efficiencies drawn from appropriate Gaussian distributions. Converting into areal rates gives $0.18 \pm 0.22 \pm 0.01$ thorium chain alphas/m²·day and $1.30 \pm 0.22 \pm 0.01$ uranium chain alphas/m²·day, where the first errors are statistical and the second are systematic. The statistical errors are large because of the limited amount of data in the reduced open dataset used. These results are also shown in Figure 7.3, where they fall in the center of the intersection of the two diagonal bands.

The thorium chain result is within 2σ of zero, so it is appropriate to use the prescription of Feldman and Cousins [72] to create asymmetric error bars that exclude non-physical negative amounts of thorium. This technique assumes that many identical experiments measuring the thorium chain alpha rate would produce Gaussian-distributed results. The thorium chain result was divided by its 1σ uncertainty to produce a measured mean of 0.846σ . 68.27% confidence intervals are given in Table X of [72] for the mean of a Gaussian constrained to be non-negative. These were interpolated to give 0.168σ and 1.846σ as the lower and upper limit of the confidence interval. Multiplying these bounds by σ gave 0.036 and 0.399 as the lower and upper bounds on the number of thorium chain alphas/m²·day. Asymmetric statistical errors were constructed by finding the difference between the best-fit value and these bounds, giving $0.18_{-0.15}^{+0.22}$ thorium chain alphas/m²·day. Because the livetime uncertainty from Section 7.1 and the simulation efficiency error are so small, adding them in quadrature with the statistical error leaves the error unchanged, so the final results are $0.18_{-0.15}^{+0.22}$ thorium chain alphas/m²·day and 1.30 ± 0.22 uranium chain alphas/m²·day.

This analysis assumes that the uranium and thorium decay chains are in equilib-

rium, which is not very likely given the complicated chemistry involved in the CVD process. A disequilibrium scenario would produce events in the bulk region from the upper parts of the uranium and thorium decay chains without adding events to the region above 6 MeV, since that is only accessible to high-energy alphas that occur near the bottoms of the decay chains. The horizontal-axis and vertical-axis intercepts of the bulk-alpha band in Figure 7.3 would shift to smaller values, since there would be a reduction in the number of bulk-alpha events that are attributed to sources in equilibrium with the events above 6 MeV. Thus, the point of intersection of the two bands would be at larger thorium chain values and smaller uranium chain values than calculated here, although these values would only represent the lower parts of the chains in the disequilibrium scenario. It is impossible to quantify this shift using only this width vs. energy analysis, but the other two analyses presented here are able to produce quantitative results even in disequilibrium scenarios.

7.3 Energy Spectrum Fits

The energy ranges used in NCD selection (described in Section 4.2) gave only crude estimates of the uranium chain, thorium chain, and ^{210}Po contributions since the goal was to choose the cleanest NCDs, not to quantify the impurities. To improve on that analysis, simulated energy distributions were produced as described in Section 6.2 and 6.3, then smoothed, truncated below 1 MeV, and normalized to unity to produce PDFs for fitting. PDFs of the lower part of the uranium chain, the lower part of the thorium chain, the upper parts of the chains, and ^{210}Po alphas were fit to the commissioning and open data using a binned maximum-likelihood fit in ROOT [73]. The fits were done from 1 MeV to 9 MeV in 10 keV bins, with 1000 events as the initial guess for each parameter. The data and fits are shown in Figure 7.4, the fit residuals are shown in Figure 7.5, and the results are summarized in Table 7.4.

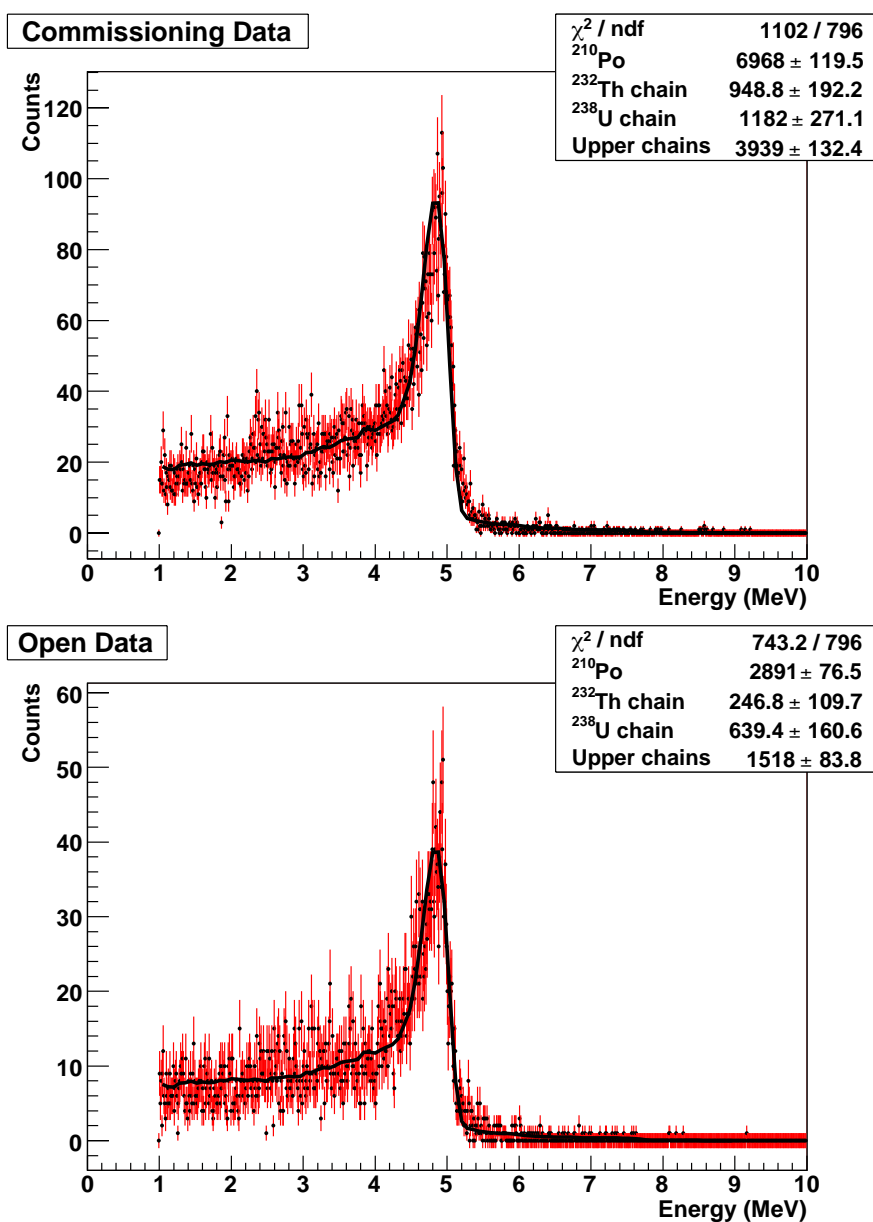


Figure 7.4: Shaper energy histograms of the commissioning and open datasets, with fits to the data. The top figure is for the commissioning dataset and the bottom figure is for the open dataset. For the commissioning dataset $\chi^2 = 1102$ for 796 degrees of freedom and the likelihood parameter $-2 \cdot \log(L) = 3722.5$. For the open dataset $\chi^2 = 743.2$ for 796 d.o.f. and $-2 \cdot \log(L) = 2754.88$. The χ^2 was calculated using all bins from 1 – 9 MeV, and assigning an error of ± 1 to any bins with zero events in them.




Figure 7.5: Residuals for the fits shown in Figure 7.4. The residual in each bin is defined as the data value minus the fit value, divided by the data error. The top figure is for the commissioning dataset and the bottom figure is for the open dataset.

Table 7.4: The results from the energy spectrum fits for the commissioning and open datasets. The fit results for the number of events from ^{210}Po , the lower part of the thorium chain, the lower part of the uranium chain, and the upper parts of the chains are shown for each dataset, along with the corresponding areal rates. Also shown is the difference between the commissioning and open dataset results, in units of sigma.

	^{210}Po	Lower Th Chain	Lower U Chain	Upper Chains
Commissioning:				
Number	6968 ± 119.5	948.8 ± 192.2	1182 ± 271.1	3939 ± 132.4
Rate ($\alpha/\text{m}^2\cdot\text{day}$)	1.76 ± 0.03	0.24 ± 0.05	0.30 ± 0.07	0.99 ± 0.03
Open:				
Number	2891 ± 76.5	246.8 ± 109.7	639.4 ± 160.6	1518 ± 83.8
Rate ($\alpha/\text{m}^2\cdot\text{day}$)	1.93 ± 0.05	0.16 ± 0.07	0.43 ± 0.11	1.01 ± 0.06
Difference	-2.9σ	0.9σ	-1.0σ	-0.3σ

The discrepancy between the commissioning and open dataset fit results can be attributed to the difference in the rate of cleaned events above 1 MeV in the two datasets, which is due to the removal of some strings from each dataset. The 13,040 events on the 33 strings used in the commissioning dataset correspond to a rate of 3.29 ± 0.04 events/ $\text{m}^2\cdot\text{day}$, whereas in the open dataset the 5296 events on the 36 strings used correspond to 3.54 ± 0.06 events/ $\text{m}^2\cdot\text{day}$, a 3.3σ difference. As a fraction of the total number of events in the commissioning dataset, the best-fit values are $53.4\% \pm 0.9\%$ ^{210}Po , $7.3\% \pm 1.5\%$ lower thorium chain, $9.1\% \pm 2.1\%$ lower uranium chain, and $30.2\% \pm 1.0\%$ upper chains. These uncertainties are based only on the errors from the fits, assuming no uncertainty on the total number of events. For the open dataset, the fractions are $54.6\% \pm 1.4\%$ ^{210}Po , $4.7\% \pm 2.1\%$ lower thorium chain, $12.1\% \pm 3.0\%$ lower uranium chain, and $28.7\% \pm 1.6\%$ upper chains. This fractional lower thorium chain contribution to the open dataset is 1.0σ lower than the

fractional lower thorium chain contribution to the commissioning dataset, and the other fractional contributions are in even better agreement between the two datasets.

Since it is known that anodized aluminum from the mandrels used in making the nickel tubes for the NCDs could be deposited on the inner surface of the NCDs, uranium and thorium chain alphas in aluminum oxide (Al_2O_3) were also modeled. A simulation like the one described in Section 6.2 was used, with the impurities uniformly distributed to a depth of $38 \mu\text{m}$ in aluminum oxide with density of 3.98 g/cm^3 . The uranium and thorium chains were assumed to be in equilibrium in this simulation. The simulated uranium chain and thorium chain alpha energy spectra were used to produce PDFs that were then fit to the data, along with the ^{210}Po PDF used before. The χ^2 was 1923 for the commissioning dataset and 817.6 for the open dataset for 797 degrees of freedom. The likelihood parameter $-2 \cdot \log(L)$ was 4089.02 for the commissioning dataset and 2885.76 for the open dataset. These poorer fits indicate that the model with the impurities in nickel is in better agreement with the data than the model with the impurities in aluminum oxide. The aluminum oxide model fits are not dramatically worse than equilibrium nickel model fits, however, so these results do not rule out the hypothesis that at least some of the uranium and thorium impurities originate in aluminum oxide.

Fits were also performed with equilibrium surface uranium and thorium chain PDFs included along with equilibrium bulk uranium and thorium chain and surface ^{210}Po PDFs. The reduced χ^2 and the likelihood parameter were smaller when the surface PDFs were included, but the best-fit values for the surface uranium and thorium chain PDFs were significantly negative. For the commissioning dataset, the surface uranium chain contribution was -504.7 ± 40.9 events and the surface thorium chain contribution was -430.8 ± 28.0 events. For the open dataset, the surface uranium chain contribution was -108.8 ± 23.9 events and the surface thorium

chain contribution was -132.7 ± 16.5 events. As these values are clearly unphysical, the fits were repeated with the surface uranium and thorium chain contributions fixed to 100 events each. This produced χ^2 and likelihood parameter that were significantly larger than without the surface uranium and thorium chain PDFs included in the fits. These results indicate that the amount of surface uranium and thorium is very small.

This energy spectrum fitting analysis is dependent on the energy calibration, but care was taken to account for space-charge effects in the simulation used for this analysis, so the discrepancy between data and simulation is small. The ^{210}Po peak falls at 4.9 MeV in both the data and the ^{210}Po simulation with space charge included. To test the sensitivity of the fit results to the energy calibration, the fits were repeated with a floating multiplicative shift. The fit results were very close to those found with the energy calibration held fixed and the best-fit values for the multiplicative shift indicate that the energy calibration is within 0.5% of the best-fit value. The systematic uncertainty associated with the difference between the fits with the energy calibration fixed and allowed to float is $^{+0.005}_{-0.013}$ ^{210}Po alphas/m²·day, -0.0001 lower thorium chain alphas/m²·day, -0.06 lower uranium chain alphas/m²·day, and $+0.05$ upper chain alphas/m²·day.

The two different datasets were combined by taking a weighted average given by $\Sigma(x_i/\sigma_i^2)/\Sigma(1/\sigma_i^2)$, where only the uncertainties from the fits are used in the σ_i . The uncertainty of the average is given by $1/\sigma^2 = \Sigma(1/\sigma_i^2)$, which is smaller than the individual fit uncertainties, as it should be, since more data are being used. This is added in quadrature to the systematic uncertainties from the livetime (discussed in Section 7.1) and the energy calibration to produce the final uncertainty. The final results are 1.80 ± 0.03 ^{210}Po alphas/m²·day, 0.22 ± 0.04 lower thorium chain alphas/m²·day, $0.34^{+0.06}_{-0.08}$ lower uranium chain alphas/m²·day, and $1.00^{+0.06}_{-0.03}$ upper chain alphas/m²·day.

The use of two datasets in this analysis allows a study of the time stability of the ^{210}Po rate. The half-life of ^{210}Po is 138 days. The midpoint of the open dataset was 113 days later than the midpoint of the commissioning dataset, so the majority of a half-life of ^{210}Po elapsed between the two datasets. Thus, unless the ^{210}Po is supported by ^{210}Pb , its decay should be evident. Yet the ^{210}Po rate in the later dataset is actually slightly higher than in the earlier dataset, indicating that the ^{210}Po is supported by ^{210}Pb . In this case, the reduction in the rate of ^{210}Po events should follow the 22-year half-life of ^{210}Pb , which cannot be extracted on the available time scale.

7.4 Time Coincidence

A time coincidence analysis based on efficiencies calculated from the simulation in Section 6.5 was performed on the combined open and commissioning datasets. This analysis assumes that the uranium and thorium impurities are uniformly distributed in the nickel NCD walls. Three different analyses were performed: a study of triple coincidences from the thorium chain, a study of coincident pairs from the thorium chain with a short time difference of less than 0.3 seconds, and a study of coincident pairs from both the uranium chain and the thorium chain with a longer time difference of less than 111.2 seconds.

A triple coincidence was defined as two events within 111.2 seconds of each other, followed by another event within 0.3 second, all with deposited energy above 0.1 MeV. There were 26 observed triples that fit these criteria. The expected number of accidental coincidences was calculated for each string according to $r_t = r_s^2 t_1 t_2 / L$, where r_s is the singles rate on that string, $t_1 = 0.3$ seconds and $t_2 = 111.2$ seconds are the first and second coincidence windows, and L is the livetime of the dataset. Subtracting the 0.023 expected accidentals from the observed triples and converting to an areal rate yields 4.8×10^{-3} triples/m²·day. Based on the simulation-determined

efficiency of 0.916% to see a triple coincidence from bulk thorium activity, the detected thorium impurity level is 0.52 ± 0.10 decays/m²·day.

Several tests were performed to understand the coincident events further and to verify that their source was short-lived alpha emitters in the thorium decay chain. The energy distributions of the events in the triple coincidences is shown in Figure 7.6. The energies range from a maximum energy down to zero energy, indicating that the source is primarily a bulk impurity, allowing the alphas to lose energy in the nickel walls. The maximum energy is in agreement with the initial energies of the alphas in the coincidence. The ²²⁴Ra alpha that comes first in the coincidence has an initial energy of 5.686 MeV or 5.449 MeV; the ²²⁰Rn alpha that is second has 6.288 MeV; and the ²¹⁶Po alpha that is last in the coincidence has an energy of 6.778 MeV. The




Figure 7.6: The energy distribution of the first, second, and third event in the 26 triple coincidences. The three distributions have been offset on the vertical axis for clarity.

pulse shapes of most of the events in the triple coincidences were verified by eye to be good candidate alpha pulse shapes. In addition, the distribution of coincident events across the NCD array was studied and was found to be fairly uniform, as shown in Figure 7.7.




Figure 7.7: The number of doubles (within 0.3 second) and triples on each string. Strings 3, 7, 18, and 20 were not analyzed, so they are not shown here.

Alpha pairs with a time difference of less than 0.3 seconds also provide information about the thorium impurity. There were 89 doubles with time differences within 0.3 seconds where both alphas deposited above 0.1 MeV of energy. A accidentals estimate for each string yields 0.9 expected accidentals within a 0.3-second coincidence window. The distribution of the 0.3-second doubles across the NCD array is shown in Figure 7.7. The energy distributions of these doubles is shown in Figure 7.8. Like the triples spectra, these energy distributions indicate a bulk impurity. The 0.3-second doubles should be closely related to the triples, since the short doubles are almost

all from ^{220}Rn decay. Based on the efficiency of 3.557% to see a 0.3-second double from simulated bulk thorium chain activity (neglecting 0.3-second doubles from the uranium chain because the efficiency is so low), the detected thorium impurity level is 0.45 ± 0.05 decays/ $\text{m}^2 \cdot \text{day}$. This is in good agreement with the thorium level inferred from the triples.




Figure 7.8: The energy distribution of the first and second events in the 89 0.3-second doubles. The two distributions have been offset on the vertical axis for clarity.

Analysis of doubles that fall within a 111.2-second time window is complicated by the appreciable number of accidental coincidences in a time window that large. Because there was a relatively high rate of singles below 1 MeV, 111.2-second doubles in that energy range were dominated by accidental coincidences, so a 1-MeV threshold was used for the long-doubles analysis. There were 444 doubles above 1 MeV within 111.2 seconds. The total expected number of accidental coincidences is 236.5 ± 14 ,

leaving 207.5 ± 25 candidate doubles. The distribution of candidate doubles across the array is shown in Figure 7.9. It is uniform, with the exception of string 30, which has an excess compared to the rest of the array. This string also had more triples than any other string in the array, and an excess of 0.3-second doubles. This may indicate that this string is somewhat high in thorium and possibly uranium, which could be uniform or concentrated in a small region near the inner surface of the wall.




Figure 7.9: The number of coincidences (within 111.2 seconds) on each string, corrected for accidental coincidences. The error bars shown on each point are statistical only and include the uncertainty on the number of observed coincidences and on the accidentals. Strings 3, 7, 18, and 20 were not analyzed, so they are not shown here.

The energy distribution of the 111.2-second doubles is more complicated than the triples or the 0.3-second doubles, since the long doubles can come from both the uranium chain and the thorium chain and because there is a significant number of accidental coincidences. In order to understand the doubles energy distribution, energy

spectrum fits were performed using PDFs describing the expected energy distributions of accidental coincidences, the first and second event in thorium chain doubles, and the first and second events in uranium chain doubles. The PDF for the accidentals was based on the singles data above 1 MeV. The PDFs for the uranium chain doubles were created using the energy simulation described in Section 6.2, but only including the ^{222}Rn and ^{218}Po alphas and requiring both alphas to be detected above 1 MeV. This requirement creates some correlation between the energies of the two alphas in the coincident pair, because they are started at the same depth and the probabilities of being detected above 1 MeV are related to the initial depth. The PDFs for the thorium chain doubles were similar to the uranium chain doubles PDFs, but all three doubles that can be produced by the alphas in the thorium chain triple were simulated. Then the PDFs were created by weighting each pair according to the probability that it is the pair observed as a double in 111.2 seconds, which is 31.35% for ^{224}Ra and ^{220}Rn , 20.31% for ^{224}Ra and ^{216}Po , and 48.34% for ^{220}Rn and ^{216}Po .

Simultaneous ROOT maximum-likelihood fits were performed to the energy distributions of the first and second events in the doubles, and are shown in Figure 7.10. The same PDF for accidentals was used for both fits. The energy distribution of the first event in each double was fit with the PDFs created from the first event in simulated bulk uranium and thorium chain coincidences, and the distribution of the second event in each double was fit with the PDFs created from the second event in the simulated coincidences. The number of events from each source (accidentals, thorium chain, and uranium chain) was required to be the same in the first-event fit as in the second-event fit.

Even though maximum-likelihood fits were performed, χ^2 is still a useful estimator of the goodness of fit. Calculating χ^2 using only bins containing a non-zero number of events yields $\chi^2 = 126.2$ for 111 degrees of freedom, corresponding to a probability of

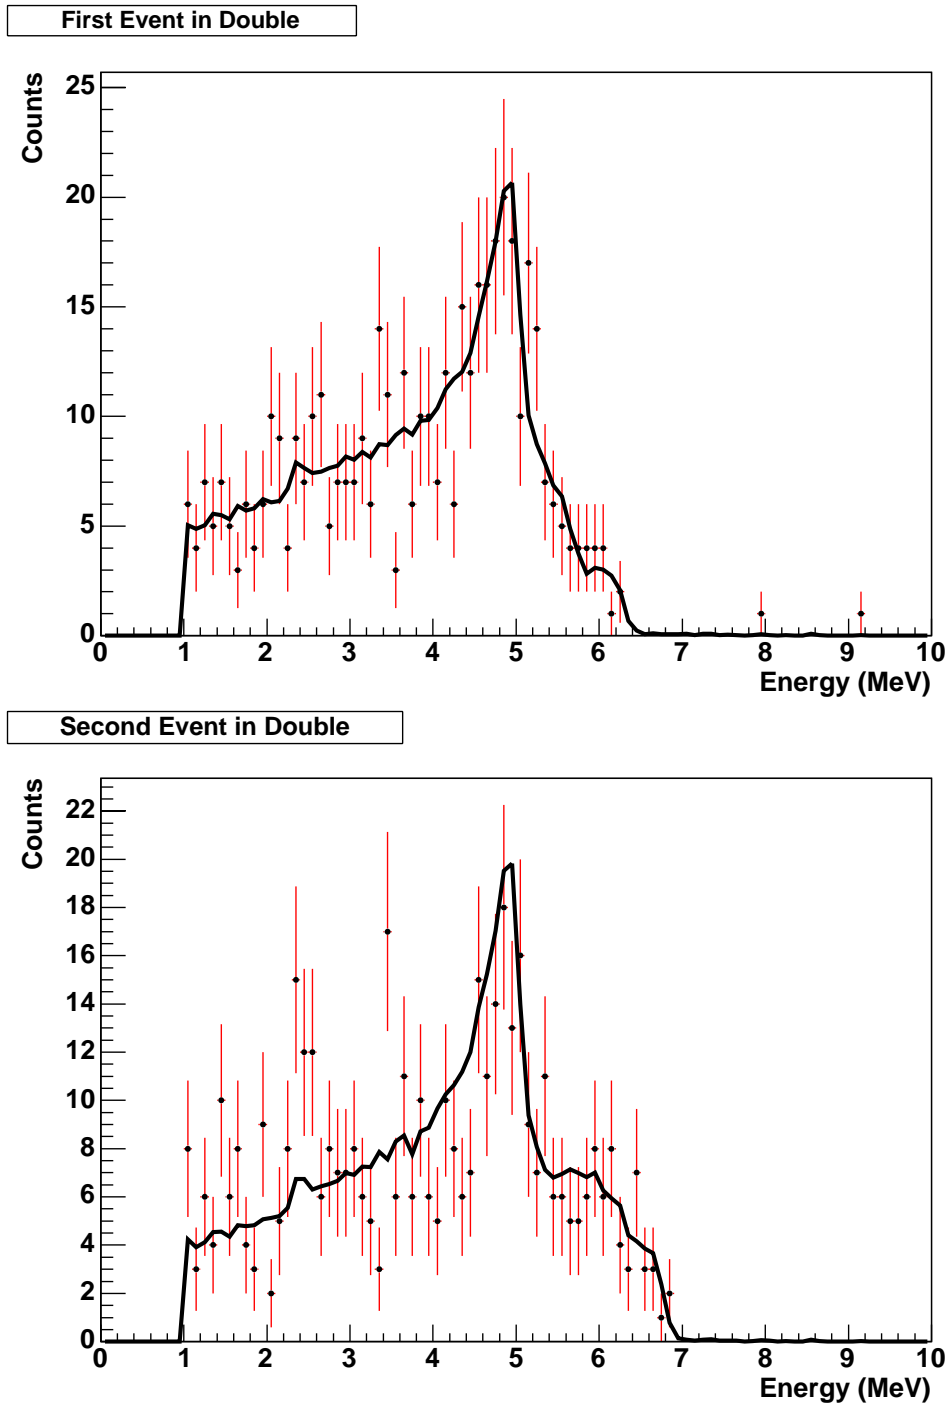


Figure 7.10: The energy distribution and fit of the first (top panel) and second (bottom panel) events in the 444 double coincidences.

15.4%. Because χ^2 can introduce bias in the case of data with fewer than five events in some bins, the Poisson-likelihood chi-square estimator [74] was also used. This estimator is defined as $\chi_{\lambda,p}^2 = 2 \sum [y_i - n_i + n_i \ln(n_i/y_i)]$, where y_i are the values of the best-fit function in each bin i , and n_i are the numbers of data events in each bin. This estimator yields $\chi_{\lambda,p}^2 = 123.81$ when only bins containing a non-zero number of events are used. Both of these estimators indicate reasonably good fits, and the uniformity of the residuals shown in Figure 7.11 is also evidence that the model is sensible. The agreement between the data and the model shows that the assumption that bulk uranium and thorium chain decays are the sources of the coincident alphas is a good one. This supports the qualitative conclusion from the triples and 0.3-second doubles that the coincident alphas come from a bulk impurity, not a surface impurity.

The fit results are 211 ± 25 accidental coincidences, 201 ± 27 thorium chain coincidences, and 32 ± 33 uranium chain coincidences. The uranium result is close to zero, so the prescription of Feldman and Cousins [72] was used to create asymmetric error bars that exclude non-physical negative amounts of uranium, giving 32_{-25}^{+33} uranium chain doubles. This accidental coincidence result is consistent with the 236.5 ± 14 accidentals expected from the singles rate on each string. Based on the thorium impurity calculated from the 26 observed triples, 244 ± 48 thorium chain doubles are expected, consistent with the fit result. Based on the thorium calculated from the 89 observed 0.3-second doubles, 213 ± 23 thorium chain doubles are expected, also consistent with the fit result. The fits were also performed with the thorium contribution constrained to the results from the triples and the 0.3-second doubles. When the thorium was constrained to 244 ± 48 events, the best-fit uranium result was 31.92 ± 31.12 , almost identical to the unconstrained fit. When the thorium was constrained to 213 ± 23 events, the best-fit uranium result was 31.9 ± 32.8 , again almost identical to the unconstrained fit. This is a good indication of the consistency between the results of

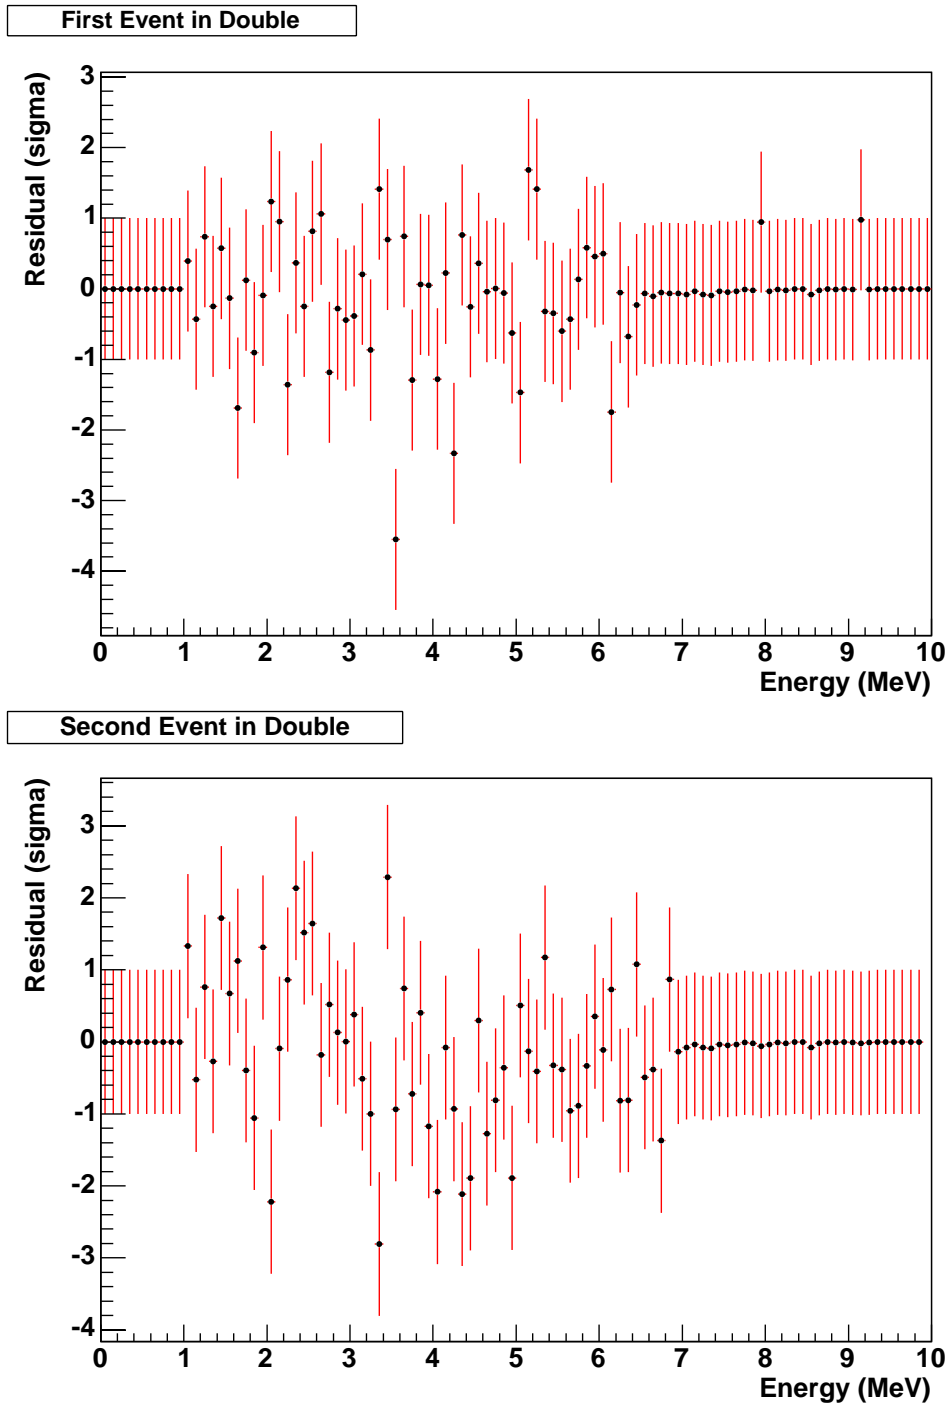


Figure 7.11: The residuals from the fits shown in Figure 7.10. The residual in each bin is defined as the data value minus the fit value, divided by the data error.

the triples, the 0.3-second doubles, and the long-doubles fit.

The simulation in Section 6.5 indicates that the efficiency for a bulk thorium decay chain to produce a coincident pair within 111.2 seconds is 8.606% and the efficiency for a bulk uranium decay chain to produce a coincident pair within 111.2 seconds is 1.164%. Converting the fit results into impurity levels using these efficiencies yields 0.43 ± 0.06 thorium decays/m²·day, which agrees with the results from the triples and the 0.3-second doubles, and $0.50_{-0.39}^{+0.53}$ uranium decays/m²·day. These results, like the results from the triples and 0.3-second doubles, assume that the impurities are uniformly distributed throughout the bulk of the nickel walls. This assumption is strongly supported by the energy distributions of the coincident events.

The three measurements of the thorium impurity level are based on the same data, so they cannot be combined as if they were independent. The thorium contribution and the uranium contribution are somewhat correlated in the long-doubles fits, with a correlation coefficient of -0.686 . It is preferable to use the thorium results from the triples or 0.3-second doubles, since they are essentially uncorrelated with the uranium. The statistical uncertainty on the 0.3-second doubles is smaller, so that was chosen as the best measurement of the thorium impurity level from this coincidence analysis. The difference between the three thorium measurement techniques is taken as a systematic uncertainty of $_{-0.03}^{+0.07}$ thorium decays/m²·day. This uncertainty is due in part to possible deviations from uniform bulk thorium, since the triples analysis and the long-doubles analysis with a 1-MeV energy threshold sample to shallower maximum depths (10.97 μm and 10.94 μm , respectively) than the 0.3-second doubles analysis (12.62 μm).

Several systematic uncertainties impact this analysis. There is a small fractional uncertainty of less than 0.3% from the simulation statistics used to determine the efficiencies. The livetime uncertainties discussed in Section 7.1 also contribute a frac-

tional uncertainty of 0.3%. The long-doubles analysis depends slightly on the energy calibration because of the 1-MeV threshold used, leading to a systematic uncertainty. The energy calibration used in this analysis was based on the neutron-capture peak, and placed the ^{210}Po peak at 4.9 MeV. Shifting the data up by 5.3/4.9, to put the ^{210}Po peak at 5.3 MeV, adds 17 more doubles to the dataset. Assuming that the breakdown of uranium and thorium coincidences is the same in these additional doubles as in the 444 that were analyzed, this produces a systematic uncertainty of +0.02 uranium decays/m²·day. The statistical uncertainty was added in quadrature with these systematic uncertainties and the spread between different techniques for the thorium result. The final results are $0.45_{-0.05}^{+0.08}$ thorium decays/m²·day and $0.50_{-0.39}^{+0.53}$ uranium decays/m²·day. These results were calculated with the assumption of uniform bulk impurities, and the data indicate excellent consistency with that model.

7.5 Combined Alpha Background Results

In order to compare the results of these three analyses, they must all measure the same quantity. The width vs. energy analysis results are in alphas/m²·day over the entire energy range, since the simulation efficiencies gave the fraction of all events that fell into a given region. This analysis assumed that the uranium and thorium decay chains are in equilibrium. The energy spectrum fit analysis results are in alphas/m²·day above 1 MeV, since the PDFs used in the fit were normalized over that energy range. This analysis assumed a particular disequilibrium scenario, where equilibrium was broken by removing radium from the decay chains. The time coincidence analysis results are in decays/m²·day, since the simulation efficiencies gave the chance of seeing a double or triple coincidence from one decay chain. This analysis makes no assumption about equilibrium since the short-lived alpha emitters that produce the coincident events are in equilibrium with each other in any scenario.

The width vs. energy results can be converted into decays/m²·day by using efficiencies from the simulation in Section 6.4. Based on one million simulated thorium decay chains, the efficiency to get an alpha depositing some energy in the NCD active volume from one thorium chain decay is 82.0%. For the uranium decay chain the efficiency is 93.4%. Applying these efficiencies to the final width vs. energy results gives $0.22^{+0.26}_{-0.18}$ thorium decays/m²·day and 1.39 ± 0.23 uranium decays/m²·day. Similarly, the energy spectrum fit results can be converted into decays/m²·day by using efficiencies from the simulation in Section 6.2. The efficiency to get one alpha above 1 MeV from one decay of the lower part of the thorium chain is 64.2%. For the lower part of the uranium chain the efficiency is 55.9%. Applying these efficiencies to the final energy spectrum fit results gives 0.34 ± 0.06 lower thorium chain decays/m²·day and $0.60^{+0.10}_{-0.15}$ lower uranium chain decays/m²·day. A small systematic uncertainty is introduced by this conversion, and has been included in the quoted uncertainties.

The results of all three analyses are shown in Figure 7.12. The time coincidence analysis result is 1.4σ larger than the energy spectrum fit result for the lower part of the thorium chain and is 0.2σ smaller than the energy spectrum fit result for the lower part of the uranium chain. The width vs. energy analysis assumed that the decay chains are in equilibrium, so it is not meaningful to compare its results to the results of the energy spectrum fits. However, the width vs. energy results can be compared to the time coincidence analysis results, which do not make any assumption about the equilibrium scenario of the decay chains. The time coincidence results of $0.45^{+0.08}_{-0.05}$ thorium decays/m²·day and $0.50^{+0.53}_{-0.39}$ uranium decays/m²·day are 0.9σ higher and 1.5σ lower, respectively, than the width vs. energy results. This is consistent with the expectation that disequilibrium would produce a larger thorium chain value and a smaller uranium chain value in the width vs. energy analysis.

The final results from the three different analyses and the errors impacting these




Figure 7.12: Comparison of the results from the energy spectrum fitting analysis, the time coincidence analysis, and the width vs. energy analysis. The smaller error bars for each measurement show the statistical uncertainties and the larger bars show the total uncertainty. In some cases the uncertainties are completely dominated by the statistical uncertainties, so the two sets of error bars cannot be resolved. For the width vs. energy analysis the results are interpreted as limits since the analysis assumed secular equilibrium.

analyses are summarized in Table 7.5. In general, the statistical uncertainties are dominant over the systematics, particularly for the uranium impurity level extracted from the time coincidence analysis. Livetime uncertainties and simulation efficiency uncertainties are very small. Energy calibration uncertainties, although fairly small, are significant for the energy fits. These uncertainties are likely to be reduced in the future. The difference between the three coincidence analysis methods used to extract the thorium impurity level is a significant systematic uncertainty for that analysis, although the results of the different methods are consistent with each other within the statistical errors.

Table 7.5: The results from the time coincidence analysis and the energy spectrum fitting analysis, including uncertainties. The sources of the systematic uncertainties are shown in the lower half of the table.

	Coincidence		Energy fits			
	Th	U	Lower Th	Lower U	^{210}Po	Upper
	decays/m ² ·day		decays/m ² ·day		$\alpha/\text{m}^2\cdot\text{day} > 1 \text{ MeV}$	
Central value	0.45	0.50	0.34	0.60	1.80	1.00
Stat. unc.	± 0.05	$^{+0.53}_{-0.39}$	± 0.06	± 0.10	± 0.03	± 0.03
Total sys. unc.	$^{+0.066}_{-0.026}$	$^{+0.019}_{-0.002}$	$^{+0.001}_{-0.001}$	$^{+0.002}_{-0.103}$	$^{+0.008}_{-0.014}$	$^{+0.052}_{-0.003}$
Livetime	± 0.002	± 0.002	± 0.001	± 0.002	± 0.006	± 0.003
Sim. eff.	± 0.001	± 0.001	± 0.001	± 0.001	—	—
Energy calib.	—	+0.019	-0.001	-0.103	$^{+0.005}_{-0.013}$	+0.052
Diff. meth.	$^{+0.066}_{-0.026}$	—	—	—	—	—

These three analyses take advantage of pulse shape, energy, and timing in order to gain as much information as possible about the sources of the alphas detected in the NCD array. The width vs. energy analysis presented here is statistics limited and only effective in an equilibrium model. The energy spectrum fitting analysis has the advantage of being the only one of these three that provides a quantitative measure of the ^{210}Po contamination of the NCDs. Because this analysis is not statistics-limited it was also possible to perform similar energy spectrum fits to data from each NCD string individually, which is described in Appendix B. The time coincidence analysis presented here is quite powerful, since it uses the energy spectra of coincident events to complement the timing information. This analysis could be improved by incorporating the timing information directly into the fits, rather than just setting a coincidence window. In addition, the use of z-position information to verify that the alphas in a

coincidence originate in the same location could provide a powerful constraint.

The results of the energy spectrum fit analysis and the time coincidence analysis were combined by taking an average weighted by the total errors, giving $0.39_{-0.06}^{+0.07}$ lower thorium chain decays/m²·day and $0.59_{-0.16}^{+0.13}$ lower uranium chain decays/m²·day. The errors quoted here are weighted averages of the individual errors. The uranium result is dominated by the energy spectrum fit analysis, since it has the smallest uncertainties, but the thorium result is a more uniform combination of the two analyses. The ²¹⁰Po result from the energy spectrum fits is 1.80 ± 0.03 ²¹⁰Po alphas/m²·day above 1 MeV. The result from the energy spectrum fits for upper chain alphas is $1.00_{-0.03}^{+0.06}$ alphas/m²·day above 1 MeV.

Chapter 8

IMPLICATIONS OF ALPHA BACKGROUNDS

8.1 *Converting Alpha Rate to Impurity Level*

It is useful to convert the measured bulk-alpha rates presented in the previous chapter to impurity levels in ppt (parts per trillion or picograms per gram) or $\mu\text{Bq}/\text{kg}$, in order to compare these results to the construction specifications and to previous measurements. These conversions will be done with reference to ^{238}U and ^{232}Th , despite the fact that a disequilibrium scenario was assumed in some of the analyses presented here and was supported by the data. This is convenient because the construction specifications and most previous measurements refer to ^{238}U and ^{232}Th and implicitly assume equilibrium. Because the activities of importance for photodisintegration background production are near the bottom of the chains, a decay of the lower part of a chain is the same as a decay of a complete chain in terms of photodisintegration backgrounds. Table 8.1 summarizes the calculation used to convert the measured decays/ $\text{m}^2\cdot\text{day}$ into ppt and $\mu\text{Bq}/\text{kg}$. Using these conversion factors, the measured impurity levels in the NCDs are $2.8_{-0.8}^{+0.6}$ ppt or 35_{-10}^{+8} $\mu\text{Bq}/\text{kg}$ ^{238}U and $5.7_{-0.9}^{+1.0}$ ppt or 23 ± 4 $\mu\text{Bq}/\text{kg}$ ^{232}Th .

A goal for the NCD array was to limit the ^{232}Th and ^{238}U impurities in the NCDs so the photodisintegration backgrounds they each produced were less than 1% of the NC signal of 12.6 neutrons per day. Based on SNOMAN simulations [75], one neutron is produced for every 482.84 ± 4.9 ^{208}Tl decays and one neutron is produced for every 32051 ± 907 ^{214}Bi decays. The specific activity of ^{232}Th is 4057 Bq/g and the branching

Table 8.1: Outline of the calculation used to convert decays/m²·day into ppt (upper part of table) and $\mu\text{Bq/kg}$ (lower part of table).

	²³⁸ U	²³² Th
Nickel density (g/cm ³)	8.908	
Simulated depth (μm)	22	
Simulated areal mass density of nickel (g/m ²)	196	
Areal mass density of 1 ppt impurity (g/m ²)	1.96×10^{-10}	
Molar mass	238.05	232.04
Areal number density of impurity (atoms/m ²)	4.96×10^{11}	5.09×10^{11}
Half life (years)	4.468×10^9	1.405×10^{10}
Decay rate for 1 atom (decays/day)	4.25×10^{-13}	1.35×10^{-13}
Areal rate for 1 ppt (decays/m ² ·day)	0.211	0.069
Decay rate for 1 atom (decays/s)	4.92×10^{-18}	1.56×10^{-18}
Atoms for 1 Bq	2.03×10^{17}	6.40×10^{17}
Impurity for 1 $\mu\text{Bq/kg}$ (g/kg)	8.04×10^{-11}	2.46×10^{-10}
Parts per trillion for 1 $\mu\text{Bq/kg}$	0.080	0.246

ratio to ²⁰⁸Tl is 35.94%, so one neutron per day is produced by 3.833 μg of ²³²Th. The specific activity of ²³⁸U is 12437 Bq/g and the branching ratio to ²¹⁴Bi is 100%, so one neutron per day is produced by 29.827 μg of ²³⁸U. The deployed NCD array is 398 meters long and consists of one-inch radius nickel tubes with density 8.908 g/cm³ that are about 334 μm thick, so the total mass of the nickel walls is about 189 kg¹. The goal of having the photodisintegration-neutron production rate be less than 1% of the NC signal rate is met if the ²³²Th impurity level is less than 2.6 ppt and the ²³⁸U impurity level is less than 19.9 ppt.

The measured ²³⁸U impurity is well below this limit, but the measured ²³²Th im-

¹This calculation neglects the mass of the endcaps, but includes the endcap walls.

purity is above its limit. Since the measured ^{238}U impurity level is only 2.8 ppt, the ^{238}U photodisintegration-neutron production rate is only $0.14\%_{-0.04\%}^{+0.03\%}$ of the NC signal. Since the measured ^{232}Th impurity level is 5.7 ppt, the ^{232}Th photodisintegration-neutron production rate is $2.22\%_{-0.34\%}^{+0.41\%}$ of the NC signal. The total photodisintegration-neutron background from ^{238}U and ^{232}Th combined is $2.4\%_{-0.3\%}^{+0.4\%}$ of the NC signal, slightly above the goal of 2%. The capture efficiency in the NCDs for photodisintegration neutrons produced by activity in the NCDs is 50.0%, whereas the efficiency for NC neutrons to capture in the NCD array is 26.4%. These capture efficiencies are different because of the different geometric distributions of the neutrons from the two sources. Thus the neutrons captured in the NCDs from NCD impurity photodisintegration will number $4.5\%_{-0.6\%}^{+0.8\%}$ of the detected NC signal.

8.2 Comparison to Previous Measurements and Detectors

Alpha background results from early data taken from the first few NCDs brought to the underground SNO lab were presented in [49]. The analysis techniques used were a pulse risetime vs. energy analysis and a time coincidence analysis, similar to those presented here. The results were < 1 ppt ^{232}Th and < 9.5 ppt ^{238}U , although the uncertainties were quite large and not fully quantified. The estimated fractional uncertainties given in [49] were 700% on the ^{232}Th measurement and 20% on the ^{238}U measurement. Given the large estimated ^{232}Th uncertainty on the previous results, they are not in disagreement with the results presented here.

In addition, samples of materials used in the construction of the NCD array were radioassayed by Radiochemical Neutron Activation Analysis (RNAA) or by direct gamma counting, yielding a result of 1 ± 1 ppt ^{232}Th in the nickel [49]. The ^{232}Th result presented here is 3.6σ above the radioassay results. This discrepancy could indicate that the inner surface layer to which these alpha analyses are sensitive contains

a greater concentration of impurities than the rest of the nickel. However, the discrepancy could simply indicate a systematic difference between the two measurement techniques that was not accounted for.

The measured ^{238}U and ^{232}Th impurity levels in the NCD array are much lower than the levels in commercially-available ^3He proportional counters or even any ^3He proportional counters designed for previous low-background experiments. Typical low-background ^3He proportional counters are concerned only with alpha contamination of the neutron-capture signal, not also with photodisintegration-neutron production, as in the case of the NCDs. Thus the region of concern for contamination is only the inner tens of microns of the proportional counter wall, so applying a low-background coating to the inner wall is an effective background-reduction technique.

A Soviet group [76] coated the inner surface of stainless-steel ^3He proportional counters with 60 microns of an “organoflourine compound having an activity 10 times less than the steel” and observed rates of about 100 alphas/ $\text{m}^2\cdot\text{day}$. A group at University of California at Irvine [77] prepared copper-coated stainless-steel ^3He proportional counters that reduced the rate from 3000 alphas/ $\text{m}^2\cdot\text{day}$ with untreated stainless steel to 750 alphas/ $\text{m}^2\cdot\text{day}$ with the electrorefined copper coating. It is not clear whether either of these rates are only for the neutron-capture energy region of interest or for the entire alpha energy spectrum. In either case, the NCD rates of 3.95 alphas/ $\text{m}^2\cdot\text{day}$ across the entire spectrum and 0.35 alphas/ $\text{m}^2\cdot\text{day}$ in the neutron-capture energy range are orders of magnitude below the rates achieved by these previous detectors.

It is initially somewhat surprising that better results were not achieved by the Irvine group, given that purity levels of a few ppt have been achieved with electrorefined copper. The ^{76}Ge double beta decay experiment carried out by Pacific Northwest Laboratory (PNL) and the University of South Carolina (USC) [78] achieved

dramatic reduction in the gamma lines seen in a germanium detector by electrorefining the copper used in their detector assembly. The intensity of the remaining ^{212}Pb line indicates that the purity of the copper is better than 4 pg/g ^{232}Th [79]. If the Irvine group achieved this purity level, then their detected rate should have been below 1 alpha/m²·day from thorium chain activity. However, tracer experiments indicate that lead is deposited along with the copper in the electrorefining process [79], so the majority of the alphas seen by the Irvine group were probably from ^{210}Po decays supported by 22-year ^{210}Pb in the copper.

The challenge of the NCD radiopurity requirements was that two different goals had to be simultaneously met: very low ^{238}U and ^{232}Th impurity levels to limit the photodisintegration-neutron background and also very low rates of alpha backgrounds in the NCDs themselves. These requirements overlap somewhat, but not completely, as shown by the Irvine proportional-counter example. The ^{232}Th impurity level achieved for the CVD nickel in the NCDs is comparable to the level achieved by the PNL-USC group with electrorefined copper. The total alpha rate in the NCD array is less than 4 alphas/m²·day across the entire alpha energy spectrum, orders of magnitude below the total alpha rates seen in the Soviet and Irvine low-background ^3He proportional counters. Given these impressive comparisons with other efforts, it is safe to say that the SNO Neutral Current Detectors are the world's best low-background ^3He proportional counters.

8.3 Statistical Estimate of Alpha Sensitivity in the NCD Phase

The analyses presented in Chapter 7 were based on up to about 100 days of data. By the end of the NCD phase of SNO, considerably more data will be available to study alpha backgrounds, and the statistical uncertainties on these analyses will be reduced. The NCD phase began in early December 2004 and will continue through

the end of December 2006, for a total of about 750 days of data-taking. The livetime fraction in the NCD phase so far has been nearly $2/3$. If this livetime fraction remains unchanged for the entire NCD phase, then the total livetime of the NCD phase will be over 450 days. This includes the open dataset used in these analyses, but not the commissioning dataset, since it is unlikely that any published analysis will be based on the latter.

For the time coincidence analysis, the result was based on combining the commissioning and open datasets, which had a total exposure of $5457 \text{ m}^2\cdot\text{days}$. If 450 days of data were available with the entire 63.52 m^2 of the array analyzable, then there would be 136 triples instead of 26, 466 0.3-second doubles instead of 89, and 2326 111.2-second doubles instead of 444. When the doubles energy spectrum fitting analysis in Section 7.4 is repeated assuming 2326 long doubles, the uncertainties from the fits are about 43% of the fit uncertainties found in the analysis in Section 7.4, very close to what would be expected from increasing the amount of data by a factor of 5.3 with \sqrt{N} statistics. The final statistical uncertainty on the time coincidence analysis from the entire NCD phase should be ± 0.02 thorium decays/ $\text{m}^2\cdot\text{day}$ and ± 0.23 uranium decays/ $\text{m}^2\cdot\text{day}$.

Based on \sqrt{N} statistics, the statistical uncertainty on the energy fit analysis should decrease to 44% of its current value. The statistical uncertainty and a systematic uncertainty from correlations between the different PDFs in the fits are combined into the uncertainty that ROOT calculates with the fit results and are difficult to separate. If this entire uncertainty is taken as a statistical uncertainty and reduced to 44% of its current value, then statistical uncertainties from the entire NCD phase would be ± 0.03 lower thorium chain decays/ $\text{m}^2\cdot\text{day}$, ± 0.05 lower uranium chain decays/ $\text{m}^2\cdot\text{day}$, ± 0.01 upper chain alphas/ $\text{m}^2\cdot\text{day}$ above 1 MeV, and ± 0.01 ^{210}Po alphas/ $\text{m}^2\cdot\text{day}$ above 1 MeV. The actual improvement may not be this large because the correlations between

the PDFs may not decrease with increased statistics.

Since a smaller dataset was used for the width vs. energy analysis than for the time coincidence and energy spectrum fitting analyses, it would improve even more if the full NCD phase dataset were analyzed. However, this analysis is not useful except as a cross-check because it assumes an equilibrium scenario, so it will not be considered here.

The potential statistical sensitivity of the full NCD phase is shown in Table 8.2, assuming 450 days of livetime and the central values found in Section 7.5. The table also shows the current fractional systematic uncertainties, many of which are smaller than the projected statistical uncertainties. It is desirable to reduce the remaining systematic uncertainties to below the statistical uncertainties, if possible.

Table 8.2: The projected fractional statistical uncertainty on the time coincidence and energy spectrum fitting alpha background analyses, assuming 450 days of data from the NCD array. The current fractional systematic uncertainties are also shown. The central values assumed in these calculations are 0.39 lower thorium chain decays/m²·day, 0.59 lower uranium chain decays/m²·day, 1.80 ²¹⁰Po alphas/m²·day above 1 MeV, and 1.00 upper chain alphas/m²·day above 1 MeV.

	Lower Th	Lower U	²¹⁰ Po	Upper
Coincidence (proj. stat.)	±5.4%	±38.8%	—	—
Coincidence (sys.)	+16.9% -6.7%	+3.3% -0.4%	—	—
Energy fits (proj. stat.)	±7.0%	±7.6%	±0.6%	±1.3%
Energy fits (sys.)	±0.3%	+0.3% -17.3%	+0.4% -0.8%	+5.2% -0.3%

8.4 Alpha Projection into the Neutron Energy Region

One of the main purposes of studying alphas in the NCDs is to understand the extent to which they are a background to the neutron-capture signal. The full-energy

neutron-capture signal falls at 764 keV, with a wall-effect tail down to 191 keV. Accounting for detector resolution, a conservative estimate of the region of interest for neutron-capture signals is 155 – 800 keV. In the simulations in Sections 6.2 and 6.3, the efficiency for detected lower thorium chain alphas to be seen in 155 – 800 keV is $7.93\% \pm 0.04\%$. For alphas from the lower part of the uranium chain, the efficiency is $8.88\% \pm 0.05\%$. The results for the upper parts of the chains are expressed in alphas/m²·day above 1 MeV, so the relevant ratio is of the simulated alphas in 155 – 800 keV to those above 1 MeV, giving an efficiency of $16.4\% \pm 0.1\%$. The ²¹⁰Po results are also expressed in alphas/m²·day above 1 MeV, and the efficiency is $6.9\% \pm 0.4\%$.

Converting the measured uranium and thorium impurity levels from decays/m²·day to alphas/m²·day and applying these efficiencies gives 0.023 ± 0.004 lower thorium chain alphas/m²·day, $0.035^{+0.008}_{-0.010}$ lower uranium chain alphas/m²·day, $0.164^{+0.011}_{-0.006}$ upper chain alphas/m²·day, and 0.12 ± 0.01 ²¹⁰Po alphas/m²·day expected between 155 keV and 800 keV. For the entire 63.52-m² NCD array, this corresponds to $1.5^{+0.3}_{-0.2}$ lower thorium chain alphas/day, $2.2^{+0.5}_{-0.6}$ lower uranium chain alphas/day, $10.4^{+0.7}_{-0.4}$ upper chain alphas/day, and 7.9 ± 0.6 ²¹⁰Po alphas/day detected in the energy range where they will be a background to the neutron-capture signal.

The expected NC neutron-capture rate in the NCDs is about 3.3 neutrons/day, so this background of $21.9^{+1.1}_{-1.0}$ alphas/day is about seven times the signal. With the current uncertainty on the alpha measurement, a statistical subtraction of the alpha background from the neutron-capture signal would produce about 30% uncertainty on the NC rate. It is interesting to note that the uranium chain alphas, which are much less problematic than thorium chain alphas as a photodisintegration background, are a larger contributor to the alpha background in the neutron-capture energy range, both in terms of their absolute number and their current uncertainty. This is because the efficiency for uranium chain alphas to fall into the neutron-capture energy range

is higher than for thorium chain alphas, and also because the systematic uncertainties on the uranium chain alpha measurement are larger than on the thorium chain alpha measurement. There are two good handles on the thorium chain alphas: the triple coincidence and the fact that they can have higher energy than any other alpha source. For the uranium chain alphas there are no such identifiers, so they are more difficult to measure. The upper chain alphas and the ^{210}Po alphas are even larger contributors to the alpha background in the neutron-capture energy range, and can only be measured using energy spectrum techniques.

It is also possible to define a reduced neutron-capture energy window with an improved signal-to-background ratio, since the majority of the neutron signal occurs near the peak at 764 keV. In the energy range from 530 – 830 keV, the expected signal is larger than the uncertainty that would be obtained by doing a statistical subtraction of the alpha background. Using this energy range would produce 0.6 ± 0.1 lower thorium chain alphas/day, $0.9_{-0.3}^{+0.2}$ lower uranium chain alphas/day, $4.4_{-0.2}^{+0.3}$ upper chain alphas/day, and $3.3_{-0.3}^{+0.4}$ ^{210}Po alphas/day. The neutron signal is only reduced by 14% by using this smaller energy range. The background of 9.3 ± 0.5 alphas/day is about three times the neutron signal.

The projected statistical uncertainties calculated in Section 8.3 from the full NCD phase are minimized if the time coincidence analysis is used for the lower thorium chain and the energy fitting analysis is used for the other alpha sources. By making a 5% measurement of the lower thorium chain rate, an 8% measurement of the lower uranium chain rate, a 1% measurement of the upper chain rate, and a 1% measurement of the ^{210}Po rate, it would be possible to measure the total number of alpha events in the 530 – 830 keV energy range to ± 0.35 alphas/day. With this uncertainty, a statistical subtraction of the alpha background from the neutron signal would provide a 12% measurement of the NC rate.

8.5 Statistical Alpha Subtraction

While a statistical subtraction is clearly not the method that will be used to extract the NC flux in the NCD phase, performing one at this early stage of the analysis of NCD phase data provides a good verification that there are no unknown problems with the data. The dataset used for this analysis was the cleaned, combined commissioning and open datasets (101.51 days of livetime) with nine NCD strings removed. The strings that were removed were all four of the ^4He strings (3, 10, 20, and 30), since they are not sensitive to neutrons, plus strings 16, 18, 26, 27, and 34. String 18 was removed because the gain difference caused by the leaking NCD makes analysis difficult, strings 16 and 26 have discharge events near the neutron region that are not removed by the existing data-cleaning cuts, and strings 27 and 34 have discharges, particularly in the commissioning dataset, that are not completely removed by data cleaning. The remaining 31 strings total 312.5 m of length, or 49.87 m² of surface area. The energy range from 530 – 830 keV was chosen for this analysis to maximize the signal-to-background ratio.

Based on the individual-string energy spectrum fits described in Appendix B, the nine strings not analyzed here account for 59.3% of the total number of lower thorium chain alphas in the energy region from 530 – 830 keV, 49.7% of the lower uranium chain alphas between 530 – 830 keV, 43.4% of the upper chain alphas in that energy range, and 26.4% of the ^{210}Po alphas. These fractions are larger than would be expected if the alpha backgrounds were uniformly spread throughout the array, which is why the individual string fits are necessary to understand the distribution of each alpha source on each string. For example, string 10 was found to be particularly high in ^{210}Po alphas and upper chain alphas, string 3 is high in lower uranium chain alphas, and strings 18 and 20 are high in lower thorium chain alphas.

Figure 8.1 shows the energy spectrum of the data and a spectrum produced by




Figure 8.1: The energy spectrum of data (upper) from 31 strings analyzed in the statistical subtraction analysis, and a simulated plot (lower) showing the expected spectrum.

simulation. The simulated spectrum uses normalized alpha spectra from lower thorium chain alphas, lower uranium chain alphas, upper chain alphas, and ^{210}Po alphas, each weighted according to the measured alpha rates from Chapter 7. For the lower thorium chain, there was a measured alpha rate of 18.4 alphas/day, of which 40.7% was on the 31 analyzed strings. For the lower uranium chain, 50.3% of the 24.8 alphas/day were on the analyzed strings. For the upper chains, the measured rate was 81.0 alphas/day and 56.6% of that was on the 31 analyzed strings. For the ^{210}Po alphas, 73.6% of the 126.9 alphas/day were on the analyzed strings. The neutron spectrum used to produce the simulated spectrum was from a neutron source calibration, and it was normalized to 12.63 neutrons/day times 312.5/362 to account for the reduced array used². The neutron and alpha spectra were summed together, then multiplied by the livetime of 101.51 days to agree with the normalization of the data.

The data and simulated spectra are qualitatively quite similar, although there are some features in the data that are absent from the simulation. There are low-energy events in the data that are primarily caused by a few strings (1, 7, 15, 17, and 29) that had an excess of low-energy noise not removed by data cleaning in this dataset. The high-energy data events above 9 MeV are primarily caused by large discharges on string 7, which are present in this dataset but will ultimately be removed by data cleaning. The shape of the ^{210}Po peak is somewhat different in the data and the simulation, and the data exhibits a small feature around 2.5 MeV that is not present in the simulation. The peak shape and the 2.5-MeV feature are believed to be caused by space-charge effects, which need to be refined in the simulation. This is currently an active area of simulation efforts by several members of the SNO collaboration working on pulse shape models and fitting techniques.

Using the best results of the analyses in Chapter 7, and the projections in Section

²The total length of the ^3He NCDs in the full array is 362 m, but only 312.5 m are analyzed here.

8.4, the number of alphas in the 31 analyzed strings in the 530 – 830 keV energy range was calculated to be 578_{-31}^{+34} . There were 875 events in 530 – 830 keV in this dataset on these 31 strings. This leaves 297_{-43}^{+45} events attributed to neutrons, or $2.9_{-0.4}^{+0.5}$ neutrons per day detected in the NCD array. The capture efficiency for the entire array for NC neutrons is 26.4%. Since a reduced portion of the array is analyzed here, this efficiency is scaled down by the fraction of the ^3He NCDs that are used in this analysis ($312.5/362$) to 22.8%. Applying this capture efficiency gives $12.8_{-1.9}^{+2.0}$ neutrons per day produced in the heavy water. This is in excellent agreement with the 12.6 neutrons produced per day that was measured in the salt phase of SNO. This analysis has not accounted for background neutrons, which would reduce the neutron production rate attributed to NC interactions.

As a cross-check, the analysis was performed separately for the commissioning dataset and the open dataset, still using the same 31 strings. The results of $12.6_{-2.0}^{+2.1}$ neutrons produced per day in the commissioning dataset and 13.5 ± 2.9 neutrons produced per day in the open dataset are consistent with each other within the statistical uncertainties. This indicates that the result is stable over the several-month time span between the two datasets and that there is not a large systematic uncertainty associated with problems in either dataset.

A statistical subtraction of this sort has very limited sensitivity compared to the pulse shape analysis techniques that will be used to produce published NC flux results from the NCD phase of SNO. However, this analysis has provided a good verification of the data that is being taken in the NCD phase. The fact that this statistical subtraction is even possible (albeit on a subset of the NCD array), given that data cleaning is not yet complete, is a testament to the high quality of the NCD data and the power of having an array of 40 essentially independent detectors. In addition, it would be very easy to get an incorrect measure of the alpha contribution to the

neutron energy range if the wrong number of alphas were attributed to lower uranium and thorium chain decays, upper chain decays, and ^{210}Po decays. This is because the alpha rates are primarily measured above 1 MeV and projected down into the neutron energy range, and these projections are different for each alpha source. The agreement between the result of this subtraction analysis and the result from the salt phase of SNO provides confidence in the measured alpha rates, and their breakdown into the different alpha sources. Thus, while this subtraction analysis should not be taken as an accurate measurement of the NC flux from the NCD phase of SNO, it provides a very important verification of the NCD data quality and indicates that the measured alpha rates from each source are reasonable.

Chapter 9

CONCLUSIONS

Solar neutrinos provide insight into both solar physics and neutrino properties. As John Bahcall wrote when proposing solar neutrino detection in 1964 [25], “neutrinos, with their extremely small interaction cross sections, can enable us to *see into the interior of a star* and thus verify directly the hypothesis of nuclear energy generation”. Interest in neutrinos as a means for studying the sun was augmented by interest in the puzzling properties of neutrinos themselves, as it became clear that neutrinos change flavor while they propagate. The Sudbury Neutrino Observatory has provided concrete evidence of solar neutrino flavor transformation and confirmation of solar models by detecting solar neutrinos via the neutral-current interaction. In the third phase of the SNO experiment, the detector’s capabilities have been enhanced by the addition of the Neutral Current Detectors, an array of ^3He proportional counters that detect neutrons liberated in neutral-current interactions of neutrinos with deuterium.

The main consideration in the design of the NCD array was the reduction of radioactive impurities that are able to create backgrounds that mimic neutron-capture signals in the NCDs. A variety of innovative materials and techniques were used to construct and install the NCDs, and extensive testing and optimization was carried out before, during, and after deployment of the NCD array to ensure data quality. Data-taking in the NCD phase of SNO began in December 2004 and will continue through the end of 2006. The performance of the NCD array, the amount of maintenance required to take useable data from the array, the livetime fraction in the NCD

phase, and the overall data quality have been very good.

Photodisintegration induced by ^{208}Tl and ^{214}Bi gammas is the primary source of neutron backgrounds to a neutral-current flux measurement in SNO. In the previous phases of the SNO experiment, thorium and uranium impurities in the D_2O have been the dominant source of photodisintegration neutrons, and this background has been quantified through a combination of ex-situ and in-situ measurements. The ex-situ measurements involve periodic assays of the D_2O to count decays of thorium and uranium progeny. The in-situ measurements use the Cherenkov signal produced by low-energy beta and gamma events from ^{208}Tl and ^{214}Bi decays to determine the rate of these decays. This in-situ technique can also be used to measure the additional photodisintegration background introduced by the NCD array. This measurement using the PMTs will be complemented by measurements of the NCD thorium and uranium impurities made with the NCDs themselves.

The primary background to the 764-keV NCD neutron-capture signal is alpha decays that do not deposit their full energy in the active volume of the NCD. A variety of pulse shape analysis techniques are under development to distinguish these alpha backgrounds from neutron-capture signals. A neutron capture produces a back-to-back proton-triton pair, which can create a distinct current signal in the proportional counter compared to the single ionizing particle in an alpha event. Pulse shape parameters such as the duration of the current pulse, or fitting techniques that compare pulse shapes to simulated pulses, can be used to distinguish neutron-capture events from alpha backgrounds. At this point, the data-cleaning cuts for NCD data have not been finalized, and understanding the impact of the electronics and other factors on NCD pulse shapes is not yet complete. As a result, robust pulse shape analysis techniques are still under development and a thorough analysis separating neutron-capture events from alpha backgrounds is not yet possible.

Understanding the alpha backgrounds in detail will be a great benefit when the pulse shape analysis of the NCD data is carried out. With the preliminary data-cleaning cuts it is easier to study alpha backgrounds than neutron captures because the higher-energy part of the spectrum is less impacted than the neutron-capture signal region by instrumental backgrounds that are not removed by data cleaning. One of the main goals of the analyses presented here is to quantify the alpha background expected in the neutron-capture energy range. Another goal is to verify that the impurity levels are within design specifications and will not cause a large photodisintegration-neutron background in SNO.

Three different analysis techniques were used to measure and characterize the alpha backgrounds in the NCD array. A pulse width vs. energy analysis was used to separate bulk alphas in the nickel walls of the NCDs from surface alphas on the interior surface of the walls, and to separate thorium chain alphas from uranium chain alphas. Unfortunately, this analysis was only valid under the assumption of secular equilibrium in the uranium and thorium decay chains, which proved to be a poor assumption. An energy spectrum fitting method was also used to identify ^{210}Po alphas and uranium and thorium chain alphas, and to quantify the disequilibrium in the decay chains. This method used simulations of bulk, disequilibrium uranium and thorium chain activity, and surface ^{210}Po activity to produce energy PDFs that were fit to NCD data. In addition, a time coincidence analysis that took advantage of short-lived alpha emitters in the thorium and uranium decay chains was used to measure the thorium and uranium impurities in the NCDs. Detailed simulations of alpha events in the NCDs were required to carry out these analyses. These three analysis techniques complement each other and together they provide the capability to make a good measurement of the alpha impurities in the NCD array.

The surface ^{210}Po alpha rate across the entire energy spectrum, averaged over the

63.52-m² NCD array, is 2.00 ± 0.03 alphas/m²·day. Additionally, bulk thorium and uranium chain activity produces $0.29_{-0.04}^{+0.05}$ alphas/m²·day from the lower part of the disequilibrium thorium decay chain, $0.39_{-0.11}^{+0.09}$ alphas/m²·day from the lower part of the uranium decay chain, and $1.28_{-0.04}^{+0.08}$ alphas/m²·day from the upper parts of the two decay chains. These rates were measured primarily above 1 MeV and simulations were used to extrapolate the results to include the range below 1 MeV. The energy distribution of the alpha events indicate that the amount of surface uranium and thorium is negligibly small. These alpha rates are orders of magnitude lower than have been achieved in previous low-background ³He proportional counters.

The thorium and uranium alpha rates were converted to impurity levels, under the assumption that the alphas detected in the NCDs, which probe only about 20 μm into the nickel walls, are representative of the bulk activity of all the nickel. A disequilibrium scenario was assumed where the uranium and thorium decay chains were broken by preferential removal of radium. The decay rates from the lower parts of the chains containing the photodisintegration-producing gammas were used, but the calculated impurity levels were referenced to ²³²Th and ²³⁸U. The activity from the lower part of the thorium chain is equivalent to what would be produced by $5.7_{-0.9}^{+1.0}$ picograms of ²³²Th per gram of nickel, and the activity from the lower part of the uranium chain is equivalent to $2.8_{-0.8}^{+0.6}$ picograms of ²³⁸U per gram of nickel. These impurity levels are comparable to what has been achieved in other materials such as electrorefined copper. The photodisintegration-neutron background produced by these impurity levels is 2.4% of the NC neutron production rate. Because the photodisintegration neutrons originate from the NCDs, their capture efficiency is enhanced relative to NC neutrons, so the background neutron-capture rate is 4.5% of the NC neutron-capture rate in the NCDs.

By using simulation to project the measured alpha rates below 1 MeV, the total

alpha rate in the neutron-capture energy range from 155 keV to 800 keV was calculated to be $21.9_{-1.0}^{+1.1}$ alphas per day in the NCD array, 6.6 times the signal expected from neutrons produced by neutral-current interactions. With the statistical sensitivity from the entire NCD phase and a modest reduction in systematic errors to make them comparable to the expected statistical errors, this alpha background should be able to be measured to about 4%. A statistical subtraction performed using the available data in an optimized energy range from 530 – 830 keV, where the signal-to-background ratio is 0.3, produced a result of $12.8_{-1.9}^{+2.0}$ neutrons produced per day in SNO, in agreement with the measured neutron production rate from previous phases of the SNO experiment. This statistical subtraction has limited sensitivity compared to the pulse shape analysis techniques that are under development to separate the neutron-capture signal from alpha backgrounds. These techniques are likely to allow about half of the neutrons to be positively identified in a background-free manner, so statistical separation will only be used as a secondary technique for the remaining portion of the neutron-capture signal.

The NCD phase goal of creating a robust array of low-background ^3He proportional counters capable of taking data for several years while submerged in ultra-pure heavy water has been met. While pulse shape analysis techniques to separate neutron-capture signals from alpha backgrounds have not been fully developed, preliminary work indicates that the separation techniques will be successful. The alpha background studies presented here indicate that the alpha rates in the NCD array are in line with expectations, and will not pose a significant problem to the successful extraction of the neutral-current signal from the NCDs.

BIBLIOGRAPHY

- [1] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, *Science* **124**, 103 (1956).
- [2] R. Davis, *Phys. Rev* **97**, 766 (1955).
- [3] R. Davis, UNESCO Conference Proceedings, Paris (1958).
- [4] E. Majorana, *Nuovo Cimento* **14**, 171 (1937).
- [5] G. Danby, J.-M. Gaillard, K. Goulios, L. M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger, *Phys. Rev. Lett.* **9**, 36 (1962).
- [6] M. L. Perl, G. S. Abrams, A. M. Boyarski, M. Breidenbach, D. D. Briggs, F. Bulos, W. Chinowsky, J. T. Dakin, G. J. Feldman, C. E. Friedberg, et al., *Phys. Rev. Lett.* **35**, 1489 (1975).
- [7] DONUT Collaboration, *Phys. Lett. B* **504**, 218 (2001).
- [8] G. Gamow and E. Teller, *Phys. Rev.* **49**, 895 (1936).
- [9] T. D. Lee and C. N. Yang, *Phys. Rev. Lett.* **4**, 307 (1960).
- [10] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, *Phys. Rev.* **105**, 1413 (1957).
- [11] F. J. Hasert, S. Kabe, W. Krenz, J. von Krogh, D. Lanske, J. Morfin, K. Schultze, H. Weerts, G. H. Bertrand-Coremans, J. Sacton, et al., *Phys. Lett. B* **46**, 121 (1973).
- [12] F. J. Hasert, S. Kabe, W. Krenz, J. von Krogh, D. Lanske, J. Morfin, K. Schultze, H. Weerts, G. H. Bertrand-Coremans, J. Sacton, et al., *Phys. Lett. B* **46**, 138 (1973).
- [13] J. N. Bahcall, A. M. Serenelli, and S. Basu, *Astrophys. J. Lett.* **621**, 85 (2005).
- [14] H. A. Bethe, *Phys. Rev.* **55**, 434 (1939).
- [15] L. C. Stonehill, J. A. Formaggio, and R. G. H. Robertson, *Phys. Rev. C.* **69**, 015801 (2004).

- [16] N. B. Gove and M. J. Martin, Nucl. Data Tables **10**, 205 (1971).
- [17] I. Iben, K. Kalata, and J. Schwartz, Astrophys. J. **155**, 551 (1967).
- [18] R. B. Firestone, *Table of Isotopes*, 8th ed. (Wiley, 1996).
- [19] E. E. Salpeter, Aust. J. Phys. **7**, 373 (1954).
- [20] A. V. Gruzinov and J. N. Bahcall, Astrophys. J. **490**, 437 (1997).
- [21] J. N. Bahcall, M. H. Pinsonneault, and S. Basu, Astrophys. J. **555**, 990 (2001).
- [22] A. Kurylov, M. J. Ramsey-Musolf, and P. Vogel, Phys. Rev. C **67**, 035502 (2003).
- [23] F. L. Villante, Phys. Lett. B **460**, 437 (1999).
- [24] J. N. Bahcall, URL <http://www.sns.ias.edu/~jnb/>.
- [25] J. N. Bahcall, Phys. Rev. Lett **12**, 300 (1964).
- [26] R. Davis, Phys. Rev. Lett. **12**, 303 (1964).
- [27] R. Davis, D. S. Harmer, and K. Hoffman, Phys. Rev. Lett. **20**, 1205 (1968).
- [28] J. N. Abdurashitov, J. Exp. Theor. Phys. **95**, 181 (2002).
- [29] W. Hampel, J. Handt, G. Heusser, J. Kiko, T. Kirsten, M. Laubenstein, E. Pernicka, W. Rau, M. Wojcik, Y. Zakharov, et al., Phys. Lett. B **447**, 127 (1999).
- [30] T. Kirsten and GNO Collaboration, Nucl. Phys. B (Proc. Suppl.) **118**, 33 (2003).
- [31] S. Fukuda, Y. Fukuda, M. Ishitsuka, Y. Itow, T. Kajita, J. Kameda, K. Kaneyuki, K. Kobayashi, Y. Koshio, M. Miura, et al., Phys. Lett B **539**, 179 (2002).
- [32] V. N. Gribov and B. M. Pontecorvo, Phys. Lett. B **28**, 493 (1969).
- [33] L. Wolfenstein, Phys. Rev. D **17**, 2369 (1978).
- [34] S. Mikheev and A. Y. Smirnov, Sov. J. Nucl. Phys. **42**, 913 (1985).
- [35] SNO Collaboration, Nucl. Instrum. Meth. A **449**, 172 (2000).
- [36] R. C. Allen, H. H. Chen, P. J. Doe, R. Hausammann, W. P. Lee, X. Q. Lu, H. J. Mahler, M. E. Potter, K. C. Wang, T. J. Bowles, et al., Phys. Rev. D **47**, 11 (1993).

- [37] S. D. Ellis and J. N. Bahcall, Nucl. Phys. A **114**, 636 (1968).
- [38] H. H. Chen, Phys. Rev. Lett. **55**, 1534 (1985).
- [39] SNO Collaboration, nucl-ex/0502021 (2005).
- [40] SNO Collaboration, Phys. Rev. Lett. **87**, 071301 (2001).
- [41] SNO Collaboration, Phys. Rev. Lett. **89**, 011301 (2002).
- [42] SNO Collaboration, Phys. Rev. Lett. **89**, 011302 (2002).
- [43] KamLAND Collaboration, Phys. Rev. Lett. **94**, 081801 (2005).
- [44] P. A. Sturrock, Astrophys. J. **605**, 568 (2004).
- [45] SNO Collaboration, hep-ex/0507079 (2005).
- [46] A. Bandyopadhyay, S. Choubey, S. Goswami, S. Petcov, and D. Roy, Phys. Lett. B **608**, 115 (2005).
- [47] K. M. Heeger, Ph.D. thesis, University of Washington (2002).
- [48] G. F. Knoll, *Radiation Detection and Measurement*, 2nd ed. (Wiley, 1989).
- [49] M. C. Browne, Ph.D. thesis, University of Washington (1999).
- [50] M. Zhou, L. Andrews, J. Li, and B. E. Bursten, J. Am. Chem. Soc. **121**, 12188 (1999).
- [51] M. Zhou, L. Andrews, J. Li, and B. E. Bursten, J. Am. Chem. Soc. **121**, 9712 (1999).
- [52] T. J. Bowles, P. J. Doe, A. Hime, R. G. H. Robertson, T. C. Spencer, P. M. Thornewell, J. B. Wilhelmy, and J. F. Wilkerson, *Neutral-Current Detection in the Sudbury Neutrino Observatory*, SNO Internal Document (1992).
- [53] J. F. Amsbaugh, private communication.
- [54] K. Rielage, private communication.
- [55] E. Browne and R. B. Firestone, *Table of Radioactive Isotopes* (Wiley-Interscience, 1986).
- [56] A. L. Hallin, private communication.

- [57] G. A. Cox, private communication.
- [58] L. L. Kormos, private communication.
- [59] T. V. Bullard and M. W. E. Smith, *A Technique for Separating Alpha and Neutron Events in Neutral Current Detectors*, SNO Internal Document (2001).
- [60] T. V. Bullard, *A Method of Deconvoluting the Ion Tail from Proportional Counter Pulses*, SNO Internal Document (1999).
- [61] G. A. Cox, Ph.D. thesis, University of Washington (in preparation).
- [62] H. W. C. Tseung, private communication.
- [63] A. W. P. Poon, private communication.
- [64] N. S. Oblath, private communication.
- [65] C. Kraus, A. L. Hallin, and J. R. Leslie, private communication.
- [66] J. Farine, Ph.D. thesis, Neuchatel University (1996).
- [67] M. W. E. Smith, Ph.D. thesis, University of Washington (2002).
- [68] WaveMetrics, *Igor Pro*, Software package (2003).
- [69] J. F. Ziegler, *Stopping and Ranges of Ions in Matter*, Software package (2003).
- [70] M. Dunford, private communication.
- [71] N. Tagg, Ph.D. thesis, University of Guelph (2001).
- [72] G. J. Feldman and R. D. Cousins, *Phys. Rev. D* **57**, 3873 (1998).
- [73] R. Brun and F. Rademakers, *ROOT*, Software package (2005).
- [74] S. Baker and R. D. Cousins, *Nucl. Instrum. Meth.* **221**, 437 (1984).
- [75] G. McGregor, Ph.D. thesis, Oxford University (2002).
- [76] I. I. Gurevitch et al., Kurchatov Institute Preprint IAE-4986/2 (1989).
- [77] E. L. Pasierb, Ph.D. thesis, U. C. Irvine (1979).
- [78] R. L. Brodzinski, H. S. Miley, J. H. Reeves, and F. T. Avignone, *Nucl. Instrum. Meth. A* **292**, 337 (1990).

- [79] R. G. H. Robertson, *Electrorefining of Copper*, SNO-STR-91-071 (1991).
- [80] N. K. McCauley, Ph.D. thesis, Oxford University (2001).
- [81] H. Deng, private communication.
- [82] N. Tolich, private communication.
- [83] D. Horning, private communication.

Appendix A

DATA CLEANING

Prior to the analysis of NCD data, instrumental backgrounds are removed by a series of data-cleaning cuts. The trigger rate in the NCDs is about 0.25 Hz, but only a few hundred events per day are physics events caused by alphas or neutron captures. The majority of triggers are caused by noise fluctuations, high-voltage breakdown, response of the NCD system to external influences, or other instrumental effects. There are two primary types of data-cleaning cuts: cuts that remove high-rate bursts of data and cuts that remove specific types of events based on event characteristics.

High data rates can be caused by high-voltage breakdown in the SNO PMTs, blasting and rockbursts in the mine environment, or high-voltage discharging events in the NCD string itself. The only known physics events in the NCDs that could occur in bursts are neutrons following cosmic-ray muons, which are identified by the muon signal in the PMT array and removed from the data, and galactic supernovae, which would be identified by the supernova trigger system. Other periods of high data rates are caused by undesirable instrumental backgrounds and should be removed. Since the PMT system also experiences instrumental background bursts, code that can identify and flag bursts already exists in the SNO analysis software [80]. Thus, development of burst cuts for the NCDs simply required the determination of the appropriate type, size, and duration of a burst to be cut.

Studies of NCD event rate bursts were carried out on five neutrino runs taken in September 2004: 45632, 45877, 45880, 45893, and 45895. These runs were chosen

to represent normal data-taking as well as a range of potentially burst-inducing situations: blasts, PMT high-voltage breakdown, high-rate situations that caused the PMT data buffers to overflow, and high-energy muons. Since the PMT burst cuts are applied to the NCD data as well by default, their impact on NCD bursts was studied prior to the development of NCD-specific burst cuts.

A plot of the time between consecutive NCD Shaper/ADC events in 0.1 s bins showed a large excess in the first bin. Applying the NHIT burst cut¹ reduced the first bin from 237 events to 94 events but did not completely eliminate the excess. A closer look at these events indicated that the majority of them were bursts of multiple strings firing with the same time (within 0.1 μ s) and that many of them had MUX events as well. These events are easily cut by the PMT retrigger cut, which removes events within 5 μ s of a flagged first event. Another class of burst that contributed to the excess in the first time bin was events on the same NCD string following a Shaper/ADC overflow event (an event that is recorded in Shaper/ADC bin 2047).

The time between consecutive Shaper/ADC events was plotted for the same set of five runs with strings 1 and 27 removed because of discharges that occurred every few seconds on those strings. The cuts that were applied were the NHIT burst cut, the retrigger cut, the burst cut that removes bursts of at least 3 events in 1 ms, and the muon-follower cut that removes 20 seconds of data after a muon event. The expected exponential distribution of the time between Poisson-distributed events produced a good fit to the data (the χ^2 was 807.6 for 817 degrees of freedom), but the first 1.5 s had an excess of events. Looking in more detail at the events in the first 1.5 s indicated that many of them were bursts of three or more events in a few seconds, and that an unusual number of them were on string 20. There were also some alpha-alpha coincidences, as expected from 0.15-second half-life ^{216}Po in the thorium chain. In

¹The NHIT burst cut removes bursts of 6 or more events in a 4 s time window with at least 40 hit PMTs in each event.

order to avoid cutting the triple alpha coincidence in the thorium chain, the number of events in an NCD burst cut must be at least four.

Based on these studies, two NCD burst cuts were developed to be applied after the PMT NHIT burst cut and retrigger cut: a Shaper/ADC-overflow follower cut and a 0.1-s burst cut of five or more events. The Shaper/ADC-overflow follower cut removes a fixed amount of time before and after any event that falls into the overflow bin in the Shaper/ADCs. Before the event, 15 μs are cut, in order to remove the digitized event associated with that Shaper/ADC event, which typically falls 8 μs prior to the Shaper/ADC event. After the event, 5 ms are cut, to remove smaller discharge events on the same string that follow the large primary discharge. The 0.1-s burst cut is implemented as two separate cuts, one for Shaper/ADC events and one for MUX events. For both of the burst cuts, any event that falls into the burst cut window is removed, even if it is on another NCD string or in the PMT system. This makes it possible to account for deadtime introduced by the burst cuts.

The burst cuts were verified with 137.28 hours of data taken between November 26 and December 3, 2004. In this dataset there were 20 Shaper/ADC overflow events, 16 of them on string 35 and four on string 7. Twelve of these had between one and eight events following on the same string within a few milliseconds, while the other eight were not followed closely by additional events. In total, 89 Shaper/ADC events were removed by the Shaper/ADC-overflow follower cut in this dataset. There were nine Shaper/ADC bursts and one MUX burst removed by the 0.1-s NCD burst cut in this dataset. The sizes of the bursts ranged up to 50 events and they were spread across the NCD strings. The total livetime lost due to this cut was 1.1 s and a total of 168 Shaper/ADC events and 4 MUX events were cut.

In addition to the Shaper/ADC-overflow follower cut, the Shaper/ADC burst cut, and the MUX burst cut, there are a variety of data-cleaning cuts that use characteris-

tics of individual events to identify them as being caused by instrumental backgrounds. These cuts are being developed by two other members of the SNO collaboration, but will be briefly described here. The primary types of NCD instrumental backgrounds cut by these data-cleaning cuts are high-voltage discharging events, oscillatory noise events, and essentially flat events that trigger on very small noise excursions. Two independent sets of cuts are being developed so that they can be used to verify each other and to provide redundancy.

One set of data-cleaning cuts uses pulse shape characteristics of known classes of non-physics events to identify and reject instrumental backgrounds [81]. High-voltage discharges produce very narrow pulses that may exhibit a forked structure caused by the reflection of the initial discharge off the bottom end of the NCD string. Three separate cuts are used to remove these events, one for narrow pulses without a reflection (the narrow pulse cut), one for pulses where the initial pulse is larger in amplitude than the reflection (the fork cut), and one for pulses where the reflected pulse is larger than the initial pulse (the reverse fork cut). Other cuts remove oscillatory noise (the NCD oscillatory noise cut) and flat traces (the flat trace cut). These cuts were developed based on the types of events seen on the strings with high rates of instrumental backgrounds, and are quite effective at removing these types of events, but not as effective at other classes of backgrounds that may exist at lower rates in the NCD array. These data-cleaning cuts have already been implemented and were used in the analyses presented here, although they are still being refined.

The other set of data-cleaning cuts operates in the frequency domain, rather than the time domain [82]. Various properties of Fourier transforms of individual events are used to identify events as physics events or instrumental backgrounds. Many of the non-physics events that are removed by the other data-cleaning cuts can be removed by this technique, and it is more effective at certain classes of events than

is the time-domain technique. These cuts are still under development and were not used in the analyses presented here. For published analyses of NCD data, both sets of data-cleaning cuts, as well as the NCD burst cuts, will be used to provide the best possible removal of NCD instrumental backgrounds.

Because the NCD data-cleaning cuts are still under development, a rigorous study of the sacrifice and contamination of the cuts has not yet been performed. This makes it difficult to quantify the systematic errors introduced by the data-cleaning cuts, so no data-cleaning systematics were assumed in the analyses presented here. In addition, because the data-cleaning cuts are still preliminary, some discharge events are still present in the cleaned data. This necessitated the removal of certain strings from some analyses because the non-physics events on those strings could not be removed easily and interfered with the analysis. In addition, the incompleteness of the data-cleaning cuts makes pulse shape based separation of neutron-capture events from alpha backgrounds very difficult, so only a preliminary analysis could be performed here.

The full set of data-cleaning cuts used in the analyses presented here also includes the NCD pulser cut and some PMT data-cleaning cuts. The NCD pulser cut removes events on spare NCD channels, such as events injected by a random pulser used to help quantify the NCD deadtime. In addition to the retrigger cut and the NHIT burst cut already mentioned, the PMT cuts that were used were the muon-follower short cut and the missed muon-follower short cut. These two cuts are designed to remove neutrons that follow cosmic-ray muon events in SNO. The muon-follower short cut removes a 20-s time period after a muon is detected in the PMT system, but the muon itself is not removed from the data. The missed muon-follower short cut removes a 250-ms time period after events with at least 150 NHIT, since these large-NHIT events could be muons that were not reconstructed as muons by the PMT event fitters. A

detailed description of these PMT cuts, as well as the rest of the data-cleaning cuts developed for the PMT array, is available in [80].

Appendix B

STRING-BY-STRING ALPHA RATES

In addition to the energy spectrum fits done for the commissioning and open datasets with the array summed together, energy spectrum fits were done for each string individually. In order to increase the data available for these fits, the commissioning dataset and the open dataset were combined. The initial guesses used for the fits were 500 ^{210}Po events, 100 lower thorium chain events, 100 lower uranium chain events, and 200 upper chain events. Because negative values are non-physical, the best-fit values of each parameter were constrained to be greater than zero. The energy scale was allowed to float by introducing an additional multiplicative parameter to account for any discrepancies in energy calibration. Several of the strings had noise or discharges that were not removed with data cleaning, so limited energy ranges were used in the fits on those strings. These were strings 7 (1 – 6 MeV), 10 (1.1 – 9 MeV), 16 (2 – 9 MeV), 18 (4 – 9 MeV), 27 (1.2 – 9 MeV), and 34 (2 – 9 MeV).

The results for the individual string fits are summarized in Figures B.1, B.2, B.3, and B.4. Simulations were used to extrapolate the rates from the fit ranges used to the entire energy range from 0 – 9 MeV. Some of the fits gave best-fit values that were very close to zero, particularly for the lower parts of the thorium and uranium chains. The asymmetric 68.27% confidence intervals reported here are calculated using the prescription of Feldman and Cousins [72]. If the ratio of the best-fit value to σ was greater than 1.9, the errors simply become the original 1σ symmetric errors.

The ^{210}Po alpha rate obtained for the whole array by averaging the result from




Figure B.1: The rate of ^{210}Po alpha decays on each of the 40 strings in the NCD array, in units of alphas/ $\text{m}^2\cdot\text{day}$ across the entire energy range, along with 1σ error bars.

each string is 1.91 ± 0.03 alphas/ $\text{m}^2\cdot\text{day}$ across the entire energy range. This is 2.0σ below the combined result from the summed-string fits for the commissioning and open datasets extrapolated to the entire energy range, which was 2.00 ± 0.03 alphas/ $\text{m}^2\cdot\text{day}$. Several strings stand out in Figure B.1 as having high ^{210}Po rates, in particular string 10 (a ^4He string) and string 12. These strings, especially string 10, have been known to have high alpha rates. While the overall goal is to minimize the alpha rate in the NCDs as much as possible, a somewhat elevated ^{210}Po rate on a ^4He string provides more neutron-free alphas to use for testing pulse shape discrimination routines.

The lower uranium chain alpha rate obtained for the whole array by averaging the result from each string is $0.27^{+0.08}_{-0.04}$ alphas/ $\text{m}^2\cdot\text{day}$. This is 1.0σ below the combined summed-string fit result from both datasets extrapolated to the entire energy range,




Figure B.2: The rate of lower uranium chain alphas in each of the 40 strings in the NCD array, in units of alphas/m²·day across the entire energy range, along with 1 σ error bars.

which was $0.39^{+0.07}_{-0.10}$ alphas/m²·day. String 3 and, to a lesser extent, string 10 stand out in Figure B.2 as having particularly high lower uranium chain alpha rates. These are both ⁴He strings, and while additional lower uranium chain activity is not good because of the added photodisintegration background, it is far better to have it on a ⁴He string where it can be an aid in testing pulse shape discrimination techniques.

The lower thorium chain alpha rate obtained for the whole array by averaging the result from each string is $0.40^{+0.07}_{-0.03}$ alphas/m²·day. This is 2.7 σ above the combined result from the summed-string fits for both datasets extrapolated to the entire energy range, which was 0.25 ± 0.05 alphas/m²·day. The discrepancy is probably due to the fact that strings 18 and 20 were not included in the summed string fits, but were included in the average of the individual string fits. Strings 20 and 18 stand out in Figure B.3 as having a particularly high lower thorium chain rate. As with the lower




Figure B.3: The rate of lower thorium chain alphas in each of the 40 strings in the NCD array, in units of alphas/m²·day across the entire energy range, along with 1σ error bars.

uranium chain activity, it is preferable to have lower thorium chain activity on a ⁴He string such as string 20 rather than on a ³He string. The high best-fit value of lower thorium chain alphas on string 18 is probably due to the leaking NCD in that string, whose gain shift distorts the spectrum and can produce spurious effects.

The upper chain impurity level obtained for the whole array by averaging the result from each string is $1.43^{+0.14}_{-0.09}$ alphas/m²·day. This is 1.4σ above the extrapolated, combined result from the summed string fits for both datasets, which was $1.28^{+0.08}_{-0.04}$ alphas/m²·day. String 10 (a ⁴He string) stands out in Figure B.4 as having a particularly high rate of events from the upper parts of the uranium and thorium chains. There are no photodisintegration backgrounds associated with activity in the upper parts of the chains, and this additional activity on a ⁴He string provides more neutron-free alphas to use for testing pulse shape discrimination routines. The very

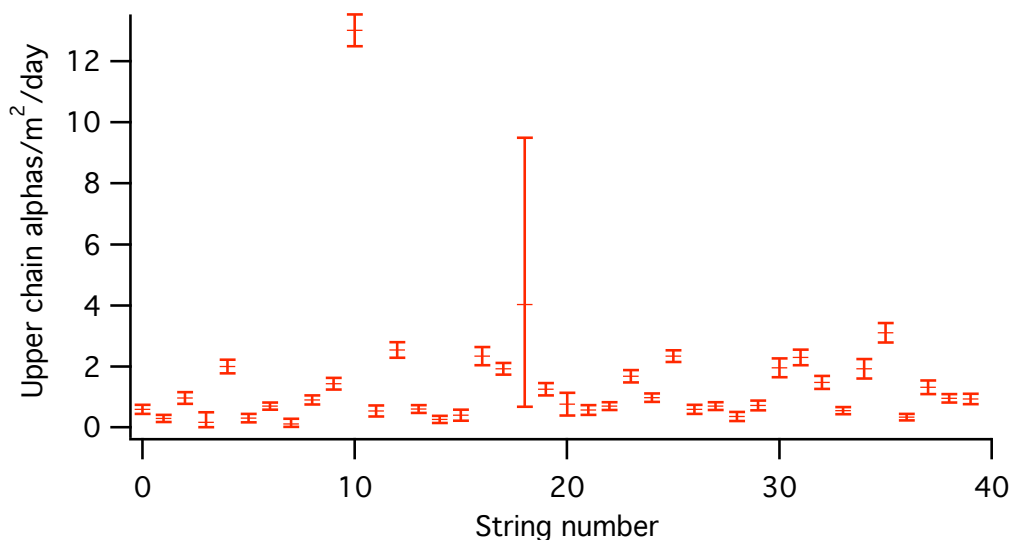


Figure B.4: The of alphas from the upper parts of the uranium and thorium chains in each of the 40 strings in the NCD array, in units of alphas/m²·day across the entire energy range, along with 1 σ error bars.

large uncertainty on string 18 is due to the fact that the reduced fit range that was necessary on this string because of the gain shift on the leaking NCD barely overlaps with the alpha spectrum from the upper chain activity.

The numerical results from the individual string fits are shown in Table B.1 for each string. The lower thorium chain, lower uranium chain, upper chains, and ²¹⁰Po alpha rates on each string in the NCD array are shown, in units of alphas/m²·day across the entire energy range. Also shown is the expected number of alphas per day that fall into an optimized neutron capture energy range from 530 – 830 keV, based on these alpha rates. This projection into the neutron capture energy range cannot be compared directly to the projected alpha rate used in the statistical subtraction analysis in Section 8.5 because that projection used the best results from all the

Table B.1: Alpha rates, in units of alphas/m²·day across the entire energy range, and the expected number of alphas per day in 530 – 830 keV. All uncertainties are 1 σ .

#	Name	Lower Th	Lower U	Upper	²¹⁰ Po	530 – 830 keV
0	N4	0.00 ^{+0.17} _{-0.00}	0.42 ± 0.17	0.60 ± 0.15	1.94 ± 0.15	0.18 ± 0.02
1	M8	0.53 ± 0.12	0.00 ^{+0.22} _{-0.00}	0.29 ± 0.11	0.38 ± 0.08	0.07 ^{+0.02} _{-0.01}
2	K8	0.30 ± 0.17	0.06 ^{+0.38} _{-0.06}	0.97 ± 0.19	2.88 ± 0.19	0.23 ^{+0.04} _{-0.02}
3	I7	0.52 ^{+0.50} _{-0.39}	2.61 ± 0.76	0.17 ^{+0.33} _{-0.16}	0.83 ± 0.17	0.21 ^{+0.07} _{-0.06}
4	J8	0.00 ^{+0.25} _{-0.00}	0.67 ± 0.27	2.00 ± 0.23	1.89 ± 0.17	0.28 ^{+0.04} _{-0.03}
5	L2	0.00 ^{+0.18} _{-0.00}	0.45 ± 0.19	0.31 ± 0.14	1.21 ± 0.12	0.11 ± 0.02
6	J7	0.15 ± 0.07	0.00 ^{+1.00} _{-0.00}	0.70 ± 0.11	0.64 ± 0.10	0.09 ^{+0.05} _{-0.01}
7	M7	0.00 ^{+1.25} _{-0.00}	0.78 ± 0.23	0.12 ^{+0.17} _{-0.10}	0.45 ± 0.10	0.08 ^{+0.06} _{-0.02}
8	K7	0.00 ^{+0.04} _{-0.00}	0.15 ± 0.08	0.91 ± 0.15	1.50 ± 0.14	0.15 ^{+0.02} _{-0.01}
9	I8	0.15 ± 0.07	0.00 ^{+0.25} _{-0.00}	1.44 ± 0.19	3.23 ± 0.20	0.24 ^{+0.03} _{-0.02}
10	I6	0.59 ^{+0.50} _{-0.42}	1.35 ± 0.74	13.01 ± 0.52	9.35 ± 0.38	1.48 ± 0.08
11	K6	0.31 ^{+0.36} _{-0.25}	0.15 ^{+0.48} _{-0.15}	0.54 ± 0.18	0.95 ± 0.12	0.11 ^{+0.04} _{-0.03}
12	M6	0.20 ^{+0.23} _{-0.16}	0.06 ^{+0.43} _{-0.06}	2.54 ± 0.25	5.83 ± 0.26	0.51 ^{+0.04} _{-0.03}
13	J6	0.00 ^{+0.05} _{-0.00}	0.13 ^{+0.11} _{-0.09}	0.61 ± 0.13	1.23 ± 0.15	0.11 ^{+0.02} _{-0.01}
14	N3	0.39 ± 0.16	0.00 ^{+0.68} _{-0.00}	0.26 ± 0.12	0.87 ± 0.10	0.09 ^{+0.04} _{-0.01}
15	L3	0.21 ^{+0.37} _{-0.19}	0.08 ^{+0.86} _{-0.08}	0.40 ± 0.18	0.71 ± 0.09	0.09 ^{+0.06} _{-0.02}
16	J5	0.32 ± 0.11	0.00 ^{+1.15} _{-0.00}	2.34 ± 0.29	2.89 ± 0.20	0.33 ^{+0.07} _{-0.03}
17	M5	0.14 ± 0.06	0.00 ^{+1.06} _{-0.00}	1.93 ± 0.19	2.15 ± 0.16	0.28 ^{+0.06} _{-0.02}
18	K5	2.57 ± 0.38	0.00 ^{+0.40} _{-0.00}	4.03 ^{+5.46} _{-3.34}	1.35 ± 0.21	0.55 ^{+0.33} _{-0.20}
19	I5	0.17 ^{+0.15} _{-0.12}	0.16 ^{+0.21} _{-0.13}	1.25 ± 0.20	2.41 ± 0.19	0.21 ^{+0.03} _{-0.02}
20	I3	4.53 ± 0.66	0.40 ^{+0.85} _{-0.39}	0.76 ± 0.37	0.80 ± 0.18	0.33 ^{+0.08} _{-0.06}
21	K4	0.57 ± 0.18	0.00 ^{+0.36} _{-0.00}	0.58 ± 0.16	2.19 ± 0.17	0.17 ^{+0.03} _{-0.02}
22	M4	0.00 ^{+1.33} _{-0.00}	0.11 ^{+0.07} _{-0.06}	0.70 ± 0.12	0.68 ± 0.10	0.10 ^{+0.06} _{-0.01}
23	J4	0.19 ± 0.10	0.00 ^{+0.23} _{-0.00}	1.68 ± 0.21	3.39 ± 0.21	0.29 ^{+0.03} _{-0.02}
24	L4	0.15 ± 0.07	0.00 ^{+0.11} _{-0.00}	0.98 ± 0.14	1.59 ± 0.14	0.17 ^{+0.02} _{-0.01}
25	N2	0.00 ^{+0.30} _{-0.00}	0.22 ^{+0.16} _{-0.15}	2.34 ± 0.19	2.35 ± 0.17	0.35 ^{+0.03} _{-0.02}
26	J3	0.00 ^{+0.12} _{-0.00}	0.42 ± 0.15	0.60 ± 0.15	0.87 ± 0.11	0.11 ± 0.02
27	M3	0.16 ± 0.07	0.00 ^{+0.58} _{-0.00}	0.70 ± 0.13	0.77 ± 0.09	0.11 ^{+0.03} _{-0.01}
28	K3	0.65 ± 0.23	0.07 ^{+0.41} _{-0.07}	0.36 ± 0.15	0.29 ± 0.08	0.08 ^{+0.03} _{-0.02}
29	I4	0.00 ^{+0.27} _{-0.00}	0.22 ^{+0.17} _{-0.15}	0.72 ± 0.16	1.64 ± 0.18	0.13 ^{+0.03} _{-0.02}
30	I2	1.04 ± 0.48	0.52 ^{+0.69} _{-0.43}	1.96 ± 0.31	1.08 ± 0.16	0.27 ^{+0.06} _{-0.05}
31	K2	0.36 ^{+0.26} _{-0.23}	0.27 ^{+0.39} _{-0.23}	2.30 ± 0.25	3.84 ± 0.22	0.40 ± 0.04
32	J2	0.00 ^{+0.06} _{-0.00}	0.25 ± 0.11	1.47 ± 0.21	3.52 ± 0.21	0.28 ± 0.02
33	M2	0.33 ± 0.14	0.00 ^{+0.28} _{-0.00}	0.55 ± 0.12	0.96 ± 0.11	0.11 ^{+0.02} _{-0.01}
34	L1	0.12 ^{+1.20} _{-0.12}	0.29 ^{+1.03} _{-0.29}	1.92 ± 0.32	2.51 ± 0.19	0.31 ^{+0.10} _{-0.04}
35	J1	0.38 ^{+0.35} _{-0.28}	1.10 ± 0.55	3.10 ± 0.32	2.65 ± 0.22	0.45 ^{+0.06} _{-0.05}
36	I1	0.03 ^{+0.03} _{-0.02}	0.00 ^{+0.10} _{-0.00}	0.34 ± 0.10	0.89 ± 0.11	0.06 ± 0.01
37	K1	0.64 ± 0.31	0.16 ^{+0.51} _{-0.16}	1.31 ± 0.22	1.56 ± 0.15	0.22 ^{+0.05} _{-0.03}
38	M1	0.00 ^{+0.14} _{-0.00}	0.15 ^{+0.11} _{-0.10}	0.96 ± 0.14	1.44 ± 0.13	0.16 ± 0.02
39	N1	0.90 ± 0.15	0.00 ^{+0.17} _{-0.00}	0.93 ± 0.17	1.68 ± 0.15	0.22 ± 0.02

analyses, and this one only uses the individual string fits.

An additional study was done to search for correlations between the ^{210}Po contamination measured on each string and the total time the nickel tubes in that string were exposed to the high-radon environment at Index [83]. No correlation was found, even for tubes that were never stored at Index. This result could indicate that the cleaning procedures were adequate to remove all excess ^{210}Po contamination from Index, leaving all the tubes at least as clean as they would have been with just the radon exposure from the air in the University of Washington laboratory where they were constructed. Given this interpretation, however, it is surprising that the ^{210}Po alphas rates are as high as they are.

VITA

Laura Stonehill has lived most of her life in Seattle, with brief stints in Alaska and Hawaii as a child. Laura became interested in physics as a high school senior at Lakeside School. She attended Pomona College as an undergraduate, where she earned a B. A. in Physics with a minor in Mathematics in 2000. After that, she entered the graduate program at the University of Washington, where she earned an M. S. in Physics in 2002 and a Ph. D. in Physics in 2005. If she had free time, Laura would enjoy reading with her cat Pandora, cooking, ceramics, and playing outside with her husband Minesh.